Two-Layer Panopticon: How the Chinese Government Uses Digital Surveillance to Prevent Collective Action*

Han Zhang[1]

[1]Assistant Professor, Department of Sociology and the Watson School of International and Public Affairs, Brown University

Abstract

Authoritarian regimes increasingly use digital surveillance to suppress collective action. Existing accounts emphasize how dictators use mass surveillance of citizens to gather information and deter mobilization, but overlook their continued reliance on human agents, whose shirking often undermines repression. We propose a two-layer Panopticon framework for digital surveillance. Dictators can directly surveil citizens. They can also surveil the frontline agents responsible for implementing repression, reducing shirking and improving prevention. We test this framework in China using

*This is the accepted version at *Social Forces*. The authors would like to thank Yongshun Cai, Xi Chen, Jingyang Huang, David Hendry, Junyan Jiang, Jennifer Pan, Gleb Papyshev, Matt Salganik, Bing Tian, Lai Wei, Hannah Wright, Xu Xu, Vivian Zhan, Dong Zhang, Yang Zhang, and Wenjuan Zheng for useful feedback on the draft. We would also like to thank participants at seminars at the University of Chicago, Brown University, and Georgetown University, as well as at the 2024 Annual Meeting of APSA. ChatGPT was used for proofreading. Direct correspondence to Han Zhang, Department of Sociology and the Watson School of International and Public Affairs, Brown University, Providence, RI 02912. Email: han_zhang4@brown.edu. Phone: +1 401-863-2766.

an original dataset of 51,611 government procurement contracts that capture digital workplace surveillance of agents alongside mass surveillance of citizens. We find that each layer independently reduces protest and that their interaction produces modestly reinforcing effects. Causal mediation analysis reveals an asymmetric mechanism: about one-third of the protest-reducing effect of citizen surveillance operates through increased oversight of agents, while agent-facing surveillance reduces protest directly. These results remain robust across dynamic panel models, instrumental variables, and alternative protest data. This article bridges and extends research on state repression, principal–agent problems in bureaucracy, and digital authoritarianism, offering new theoretical and empirical insight into how digital technologies strengthen the practice of authoritarian rule.

1 Introduction

Dictators have long hoped to shut down mobilization before it begins. But in reality, such prevention often fails, because dictators have historically depended on human agents—who frequently lack the capacity or motivation to act. Digital surveillance now provides vastly more information on citizens, bringing that dream closer to reality. Facial recognition, predictive policing, and social media monitoring are said to give rulers an unprecedented edge—allowing them to detect dissent before it surfaces, and suppress protest without resorting to overt violence (Dragu and Przeworski, 2019; Liu and Sullivan, 2021; Danneman and Ritter, 2014; Truex, 2019). Digital surveillance is thus often seen as a game-changer for authoritarian control. Many believe digital surveillance will eventually substitute for human agents in government preemption (Beraja et al., 2023).

This vision of preventive repression focuses on the information edge of the state over citizens, but overlooks a key reality: information does not act on its own. Most authoritarian regimes still rely on street-level agents—such as local police, grid workers, or cadres—not only to carry out protest-prevention tasks informed by digital surveillance data but also to gather human intelligence that machines cannot capture (Frantz et al., 2020). Yet dictators cannot observe whether these agents are truly patrolling, intervening, or reporting honestly. As a result, shirking persists. Even with precise digital data about where and when dissent may occur, protest prevention can fail if agents do not carry out the necessary interventions.

Our core argument is that dictators can apply digital surveillance on human agents to reduce shirking problems. With tools like real-time GPS tracking and automated reporting systems, rulers can detect whether agents are patrolling assigned areas, filing accurate reports, or responding to risks, and then reward or punish them accordingly. This form of workplace surveillance compels frontline agents to act more diligently in preventing protest. Moreover, this strategy is especially attractive to dictators, as it is easier to implement, lower-cost, and more publicly acceptable than population-wide mass surveillance. We therefore hypothesize that digital surveillance of agents enhances protest prevention, regardless of the existence of mass surveillance of citizens.

We propose a two-layer Panopticon to capture how authoritarian regimes use digital surveillance to suppress dissent. The citizen-facing layer targets citizens, monitoring them to detect and deter mobilization. The agent-facing layer surveils the frontline agents responsible for enforcing repression, ensuring that those who watch are themselves being watched. These two layers operate independently, each capable of reducing protest on its own, but may also reinforce one another when citizen-facing surveillance generates actionable intelligence and agent-facing surveillance ensures it is acted upon.

China provides an ideal setting to examine our two-layer Panopticon framework. The state possesses both the capacity and the will to suppress dissent, yet protest has long been a persistent challenge. Social unrest grew from an estimated number of 8,500 in 1990 to over 180,000 in 2013, despite the government allocating over $196 billion to "stability maintenance" in 2013 alone (Goebel and Ong, 2012; Lee and Zhang, 2013). Yet while decades of scholarship have focused on explaining why protests grew under a strong authoritarian regime (O'Brien and Li, 2006; Cai, 2010; Chen, 2012; Wright, 2018), new quantitative evidence shows that protests began declining around 2013 (Zhang and Pan, 2019; Chen, 2020; Goebel and Steinhardt, 2019). This shift calls for a new theoretical framework to explain why protest has begun to decline.

One promising explanation is China's use of mass digital surveillance to prevent protests. The government has built one of the most powerful domestic surveillance systems in the world (Chin and Lin, 2022), from camera networks (Xu, 2021) to social media monitoring (King et al., 2013) and biometric surveillance (Miao, 2024). It has also shifted its agenda from "stability maintenance" to "active prevention." (Ong and Luo, 2024; Truex, 2019). Indeed, Beraja et al. (2023) provided evidence that local governments' purchase of facial recognition software leads to declining protest activity .

However, this focus on citizen surveillance overlooks another development: Chinese governments use workplace surveillance to monitor their own frontline agents. In particular, grid workers—who play a key role in protest prevention—are now themselves subject to digital monitoring (Mittelstaedt, 2022; Chen and Greitens, 2022). Section 3 provides detailed background on grid workers.

To examine this dynamic, we construct an original measure of digital workplace surveillance: the number of local government procurement contracts for monitoring grid workers in each prefecture and quarter, published by the Chinese Ministry of Finance. We combine this with measures of mass surveillance on citizens, using data from Beraja et al. (2023). Our main dependent variable comes from CASM-China, the largest (N = 136,330) and most comprehensive dataset of Chinese offline social protests (Zhang and Pan, 2019), with additional data from Chen (2020) serving as robustness checks.

Using two-way fixed-effects regressions, we find that cities implementing digital surveillance to monitor grid workers experienced a statistically significant decrease in protest occurrences, after controlling for digital mass surveillance on citizens. Importantly, protests were not reduced in cities that only had grid management systems but lacked digital workplace surveillance, highlighting that monitoring frontline agents is essential for effective human-based prevention. We also find that about one-third of the impact of digital mass surveillance on protests is mediated through workplace surveillance of agents, whereas workplace surveillance has an independent direct effect in reducing protests. We also detect a modest negative interaction: protests declined most where both forms of surveillance were present, indicating that citizen- and agent-focused monitoring are mutually reinforcing.

To strengthen causal inference, we employed dynamic panel models, exogenous weather and policy shocks as instrumental variables, and extensive robustness checks to address concerns about endogenous relationships and unobserved confounders. We provide preliminary evidence that the reduction in protests stems from increased worker effort, rather than from surveillance creating a chilling effect on citizens.

This article bridges and extends three literatures: state repression of collective action, principal–agent problems in bureaucracy, and digital authoritarianism. For the literature on state repression, we build on recent work on preventive repression (De Jaegher and Hoyer, 2019; Truex, 2019; Dragu and Przeworski, 2019; Ong and Luo, 2024; Liu and Sullivan, 2021), but identify a neglected channel of authoritarian control: digital surveillance of street-level agents. We also provide the first empirical evidence demonstrating the effect of workplace surveillance on protest. These results also speak to the literature on bureaucracy. In general bureaucratic theory, classic studies see monitoring as intrinsically difficult and focus on incentive structures (Wilson, 1989; Brehm and Gates, 1999; Greitens, 2016; Dragu and Przeworski, 2019), but we show that digital workplace surveillance makes frontline effort observable and attributable in real time. Research on Chinese bureaucracy highlights a principal-supervisor–agent problem where supervisors sometimes collude with agents and sometimes issue instructions that diverge from principals' intent (Zhou, 2010), but we show that digital surveillance allows principals to monitor frontline agents directly, reducing reliance on intermediate supervisors and reshaping local governance. Finally, for the literature on digital authoritarianism, which tends to emphasize technology and citizen surveillance (King et al., 2013; Gohdes, 2020; Kendall-Taylor et al., 2021; Beraja et al., 2023), we shift attention to how dictators use these tools to monitor and discipline their own agents, thereby strengthening authoritarian rule.

2 Two-layer Panopticon

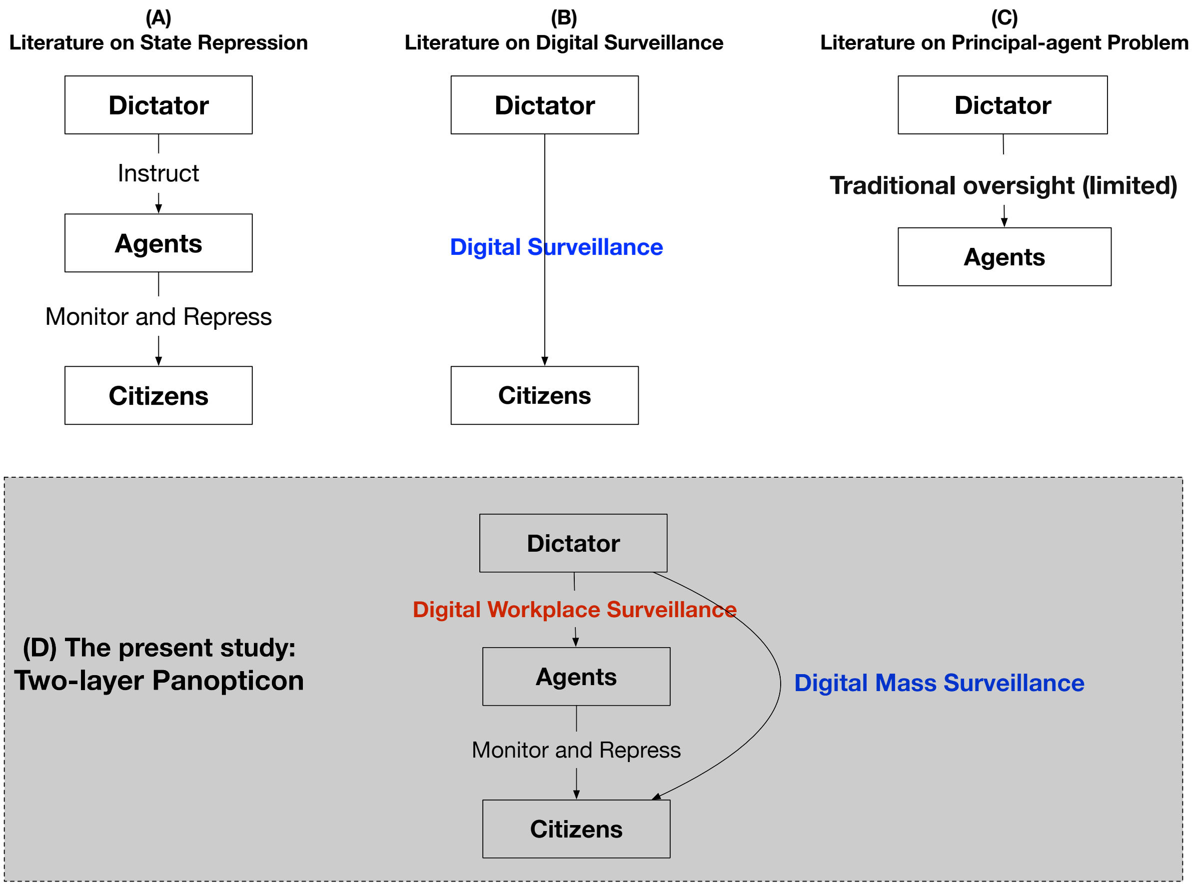

Figure 1: Illustration of the theoretical framework: the two-layer Panopticon. Panels A–C represent three distinct strands of literature that this study engages and extends. Panel D shows the theoretical framework proposed in this study, which bridges across these three strands of literature.

2.1 Preemptive repression and how it failed

This study focuses on preemptive repression—government efforts to stop mobilization before it occurs (Earl, 2003). In theory, governments prefer preemption because early intervention is often more efficient, less visible, and less risky than reactive repression (Davenport, 2005; Ritter and Conrad, 2016; Dragu and Przeworski, 2019; Truex, 2019). Preventive repression requires two steps: information collection and actual preemption. The first step is that the government must know who will mobilize, when they will mobilize, where they will gather, or in what ways they will protest. With such information, dictators can instruct their agents to use repression (torture, detainment, harassment, arrest; see Earl (2003); Dragu and Przeworski (2019); Liu and Sullivan (2021)) or cooptation ("buying off" potential protesters; see Gandhi and Przeworski (2006); Lee and Zhang (2013); Pan (2020)) to prevent collective action from happening in the first place.

Regardless of whether modern digital surveillance exists or not, dictators have always relied on street-level human agents for preemptive repression. Here, we define "dictators" as those who set up policy goals and issue instructions—such as central leaders or local officials—and "agents" as those who carry them out. Street-level agents include not only security personnel such as regular police, secret police, and military forces, but also non-security actors such as para-police, social workers, and members of state-sponsored civil groups who can help the state to monitor and co-opt dissidents in many authoritarian states (Mattingly, 2019). Agents can both collect information about mobilization and finish actual prevention tasks. Information collection efforts include patrolling, infiltration, eavesdropping, wiretapping, cultivating informants, and gathering local gossip (Mattingly, 2019; Greitens, 2016). After the information about mobilization has been collected, same or different agents will be ordered to prevent mobilization from happening. Panel A in Figure 1 shows this literature on state repression.

However, human-centered preventive repression often falls short in practice, due to failures at either the information-gathering or intervention stage—or both. Two limitations stand out. First, human agents face capacity constraints in their information gathering. Unlike machines, they need to sleep, rotate shifts, and can only cover limited ground at any given time. This limitation has motivated the rise of digital surveillance, which promises to overcome the coverage and granularity limitations of human monitoring. We develop this logic further in Section 2.2.

Second, human agents often shirk their responsibilities, both in information gathering and in actual intervention. They may not share the dictator's goals, and dictators often lack effective tools to monitor their performance—resulting in a classic principal–agent problem. As a result, even with perfect information about potential protesters, this information may fail to translate into repression if agents do not act. This weakness is critical because it undermines the link between information and repression. Yet it has received little attention in the literature, which often conflates dictators and their agents into a monolithic "state." In Section 2.3, we argue that dictators are well aware of this oversight failure and adopt workplace surveillance systems to strengthen oversight and discipline agents.

Together, these two digital strategies—citizen surveillance to overcome information capacity constraints and workplace surveillance to reduce shirking—form what we call a "two-layer Panopticon." Section 2.4 examines how these two layers work in combination.

2.2 Mass surveillance on citizens

Digital surveillance dramatically expands authoritarian state capacity by overcoming long-standing information constraints (Feldstein, 2021; Kendall-Taylor et al., 2021; Beraja et al., 2023; Tirole, 2021). Technologies like facial recognition and predictive policing enable real-time monitoring and early detection of dissent, allowing more targeted and efficient preemptive repression. Moreover, these tools may also create a chilling effect by making citizens fear they will be targeted if they mobilize. By enabling precision targeting and instilling fear of detection, digital surveillance deters dissent both actively and passively. These capabilities underpin what some scholars have described as a new era of algorithmic authoritarianism.

Hypothesis 1: *Digital surveillance of citizens reduces the likelihood of collective action.*

The implicit assumption of this literature, as illustrated in Panel B of Figure 1, is that digital surveillance allows dictators to directly observe and target citizens, reducing their reliance on street-level agents. Agents remain part of the system, but their informational role is weakened and, in some accounts, even replaced entirely by automated surveillance infrastructure.

Existing critiques of this "replacement view" primarily emphasize technological limitations: digital surveillance is porous, and human involvement is still needed to cover blind spots and make the best use of technology. While we agree with this technological critique, we focus on a distinct and largely overlooked organizational problem in the digital authoritarianism literature, one long recognized as fundamental in organizational sociology: dictators and agents often have misaligned goals. While dictators prioritize regime survival and seek to suppress mobilization, agents are motivated by routine bureaucratic incentives and rarely share the same sense of urgency.[1] As a result, dictators' orders are often not executed as intended.

2.3 Human-based preventive repression and principal-agent problems

One way to address such compliance problems is to increase monitoring of agents' actual efforts, but monitoring preventive repression is extremely difficult. These efforts include information-gathering activities (such as surveillance, patrolling, and informant recruitment) as well as preemptive actions (such as persuasion and intimidation). This information asymmetry creates moral hazard: because principals cannot verify whether these tasks are carried out effectively, agents often shirk their responsibilities.[2]

Another common approach is to use incentive structures (e.g., performance-based pay or punishment), which most bureaucracies already have in place. But their effectiveness requires accurately attributing outcomes to agent efforts, which is a different kind of monitoring challenge rarely overcome in preventive repression. When no protest occurs, leaders cannot tell whether unrest was genuinely prevented, agents shirked their duties and got lucky, or repression was so heavy-handed that it suppressed dissent temporarily, creating good paper records but fueling long-term backlash. When protests happen, attribution is unclear, and punishments are often arbitrary or collective. Faced with such persistent uncertainty, performance systems drift toward superficial compliance: supervisors focus on visible, auditable metrics such as completed forms and logged checkpoints, and agents learn to game those metrics, fabricating outputs rather than investing in substantive prevention of unrest.

These persistent monitoring failures—visualized in Panel C in Figure 1—create strong incentives for dictators solve such monitoring challenges. We argue that digital surveillance fundamentally enhances dictators' capacity to monitor agents (Panel D in Figure 1). First, various surveillance technologies now allow principals to directly observe agents' behaviors. GPS tracking through cellphones, wearable devices, or patrol cars enables precise monitoring of agents' locations and movements (de Brito and Ariel, 2017). Body-worn cameras help verify whether agents' actions align with established protocols, leading to significant reductions in misconduct (Ariel et al., 2015). Moreover, technologies initially designed to monitor citizens have been repurposed to monitor agents: fixed CCTV cameras in public spaces have been shown to reduce shirking among traffic police in India (Conover et al., 2023), while facial recognition systems in China are now widely used to track the presence, movements, and performance of street-level agents.[3]

Second, modern AI technologies help principals link outcomes to agent efforts. AI systems can analyze behavioral data to provide real-time performance assessments. Scores are automatically deducted based on algorithmic rules, and rewards or punishments can be administered immediately when violations occur (Sherman, 2013). By auto-generating verifiable performance records, these systems supersede traditional self-reports and enable dictators to administer rewards and punishments with greater precision.

Beyond improving monitoring capacity, these technologies make agent surveillance attractive to principals for cost and social acceptability reasons. Principals may prefer agent monitoring over mass citizen surveillance for practical reasons. First, it monitors a much smaller set of actors, making implementation more economical than surveillance of much larger civilian populations. Second, employee monitoring faces less social resistance than mass surveillance (Sewell and Barker, 2006). Workplace surveillance has become ubiquitous in all workplaces, in both private industries and public sectors (Ball, 2010; Levy, 2022). Because of this, employees generally accept that their work will be monitored by organizations in some way (Goold, 2003). Moreover, citizens widely support efforts to discipline corrupt or ineffective bureaucrats (Williams, 2021).

Hypothesis 2: *Digital surveillance on agents reduces collective action because it reduces the principal-agent problems prevalent in government preemption.*

2.4 Mediation and Interaction Between Citizen- and Agent-Facing Surveillance

Thus far we have examined two complementary ways that dictators deploy digital surveillance: they watch citizens directly, and they watch the street-level agents responsible for monitoring and suppressing collective action. Together, these constitute the two layers of our two-layer Panopticon framework, as we visualize in Panel D of Figure 1. We next discuss how these two layers mediate or moderate each other.

Regarding mediation effects, there is an important asymmetry between surveillance targeting citizens and surveillance targeting agents. Since agents can independently monitor and preempt mobilization, the impact of agent-facing surveillance is unlikely to be mediated by citizen-facing surveillance.

By contrast, citizen-facing surveillance systems often function primarily as informational tools. Their direct deterrent effect relies on psychological mechanisms—most notably, inducing fear among citizens that they are being watched (Foucault, 1977). However, if some organizers remain undeterred and begin to mobilize, these systems cannot themselves intervene: they cannot persuade, deter, or arrest. Enforcement ultimately depends on street-level frontline agents acting on the intelligence that surveillance systems produce. Therefore, the effectiveness of citizen-facing surveillance in reducing protest is partially mediated by whether frontline agents act on the information it produces—an outcome increasingly ensured by digital oversight of enforcement personnel.

There are at least two reasons why increased citizen-facing surveillance may lead governments to invest in agent-facing surveillance (the first stage of this mediation pathway). First, once governments gain more detailed intelligence about potential unrest, they face pressure to ensure that such information leads to concrete action. This creates both a stronger incentive and a clearer justification for investing in tools that monitor frontline agents and ensure they follow through, such as GPS-based tracking systems. Second, many citizen-focused surveillance tools—such as facial-recognition systems—can also be repurposed to track agent behavior: whether they patrol assigned areas, respond to alerts, or comply with operational protocols. The resulting metadata are increasingly fed into evaluation platforms that automatically flag underperformance and trigger disciplinary responses. Regardless of the specific mechanism, (working) agents are critically important for mass digital surveillance to reach its full potential.

Hypothesis 3a (mediation): *Protest-reducing effect of digital surveillance on agents is **not** mediated by the presence of digital surveillance on citizens, whereas protest-reducing effect of digital surveillance on citizens is mediated by the presence of digital surveillance on agents.*

Beyond mediation, these two forms of surveillance may also reinforce each other when deployed together, producing the strongest protest reduction. Citizen-facing technologies provide richer real-time information than what agents could collect by pure human monitoring efforts. Conversely, when frontline agents are themselves closely monitored, investments in citizen-facing surveillance yield greater returns because the intelligence they generate is more reliably translated into concrete preventive actions rather than being ignored or selectively enforced. This synergy suggests that:

Hypothesis 3b (moderation): *The combination of digital surveillance of citizens and agents will produce a stronger reduction in collective action than either mechanism alone.*

Not all regimes possess the capacity to implement both citizen- and agent-facing surveillance at scale. While some democracies employ workplace monitoring tools—such as CompStat dashboards and body-worn cameras—to address shirking and misconduct (Brayne, 2020), they face legal and normative constraints against extensive citizen surveillance. Even when such systems exist, they are typically designed for law enforcement or accountability rather than for preventing political protest. Many authoritarian regimes, despite their willingness, lack the administrative or technological resources necessary for large-scale deployment. Even if they do possess these capacities, achieving such synergy is challenging. Citizen- and agent-facing surveillance systems typically belong to separate bureaucracies, creating organizational silos that block information flow and undermine joint effectiveness. Even when both systems exist, integration depends heavily on overcoming entrenched principal–agent frictions: fragmented authority, competing bureaucratic incentives, and distorted performance metrics.

3 How Chinese Government Surveils Street-Level Agents

China therefore presents an especially informative case, combining significant political will, ample resources, and recent administrative reforms explicitly aimed at bridging the vertical and horizontal divides that historically exacerbated principal–agent problems.

Principal–Agent Problems in China's Bureaucracy China's bureaucracy combines horizontal fragmentation across departments with a multi-layered vertical hierarchy (Lieberthal and Oksenberg, 1988). Zhou (2010) characterizes this vertical structure as a three-tier "principal–supervisor–agent" system: the central government (principal) governs grassroots agents indirectly through intermediate local governments (supervisors). This structure creates serious monitoring challenges: Beijing often cannot observe what local officials actually do (Shirk, 1993), let alone the day-to-day actions of street-level personnel. To address these delegation problems, the central government designed incentive structures to motivate local officials, including fiscal contracting (Oi, 1999) and tournament-style promotion systems (Li and Zhou, 2005). Yet this multi-layered delegation also encourages two distinctive responses: one is collusion, where supervisors and their agents form protective networks to shield one another from inspections (Zhou, 2010); the other is "muddling through," where agents confronted with numerous and sometimes conflicting instructions from principals and supervisors adopt ad hoc responses (Zhou et al., 2013).

Taken together, these studies illustrate how China's fragmented authority structure exacerbates classic principal–agent problems, such as information asymmetry, metric gaming, collusive cover-ups, and ad hoc adaptation. These principal–agent problems have long been viewed as a key reason for China's frequent protests: gaps between central directives and local implementation create political opportunity for protesters to leverage (Cai, 2010; Chen, 2012).

Our study extends this literature on China bureaucracy in three ways. First, it examines a shift in delegated goals—from governance and economic development to preventive repression. Second, we highlight strengthened monitoring capacity as a solution to principal–agent problems: the central government increasingly bypasses intermediate supervisors to directly monitor frontline agents. Third, we focus on street-level workers—the lowest-tier agents who directly interact with citizens—while treating local and higher-level officials as principals. Preventive repression still relies heavily on these low-tech, labor-intensive street-level agents (Pei, 2024), yet they have received relatively little scholarly attention, with some exceptions (e.g., Scoggins and O'Brien (2016); Scoggins (2021) on police).

Grid Management for Integrated Local Control: Ambitions and Shortcomings

We offer an in-depth analysis of grid workers, China's frontline agents tasked with everyday social control under the *social grid management* system. Grid management aimed to consolidate previously fragmented street-level governance—historically split across multiple administrative agencies—by integrating their functions into a unified system of spatially defined units (Chen and Greitens, 2022; Mittelstaedt, 2022).[4] Each grid typically covers several hundred households and is assigned a dedicated full-time worker responsible for routine monitoring, conflict mediation, and incident response. Figure 2a visualizes this grid arrangement in a community in Fuzhou City.

To coordinate these units and integrate information flow, each county-level CCP Political and Legal Affairs Commission (PLAC) operates a centralized command center. These centers not only oversee grid workers, but also fuse inputs from other surveillance sources—including public security officers, CCTV networks, and facial recognition systems operated by police units. Drawing on these multi-channel inputs, supervisors in the command center assign grid workers a range of governance and coercive duties. Several of these are directly tied to preemptive repression, including:

- Regular visits to individuals on surveillance lists, such as petitioners, former convicts, and others flagged by local authorities
- Patrolling and reporting suspicious activity that may lead to mass incidents
- On-site "soft" interventions such as mediation, persuasion, and distributing compensation to aggrieved residents (Deng and O'Brien, 2013; Lee and Zhang, 2013).

Grid workers report to the command center and submit written completion reports, which are used by supervisors to assign monthly performance scores—typically on a 100-point scale tied to salary and bonus eligibility.

Yet despite these bureaucratic design, supervisors can hardly verify if these tasks were actually performed. Combined with low pay, minimal training, loose recruitment standards, and high turnover, these conditions create fertile ground for shirking and falsification.

Digital surveillance on grid workers Some local governments in China have adopted two distinct types of workplace surveillance to monitor grid workers and other street-level agents. Grid workers carry smartphones with real-time GPS tracking and are required to upload live photos and videos from incident scenes. To ensure diligence, smart platforms trigger random checks several times a day, demanding immediate submissions to verify presence and attention. Yet these systems are not foolproof: workers can sometimes delegate check-ins or falsify uploads—for example, having one person check in on behalf of the entire team. As a further check, CCTV and facial recognition systems—originally installed to monitor public behavior—can also incidentally capture whether agents appear in required locations during work hours.

These workplace surveillance records are integrated into a "smart" command platform, which uses AI-powered analytics to evaluate agent behavior in real time. These smart platforms—initially designed to surveil citizens—also turn their gaze on grid workers, automatically flagging unusual patterns in their movements and activities. For instance, if a worker remains stationary for too long or deviates from assigned routes, their phone receives immediate alerts to resume patrol. Separately, it cross-checks self-reports against GPS or CCTV traces. If a worker claims to have resolved an issue at a specific location, but GPS or CCTV data contradicts the report, the system flags the mismatch and reduces their wages.

Importantly, what makes the digital workplace surveillance effective is not harsher penalties, but the immediacy and certainty of enforcement. Our review of twelve provincial and municipal regulations from 2006 to 2024 shows that the formal penalty rules have remained largely unchanged: deductions of 5–10 points from a 100-point performance score, linked to salary and bonus eligibility. In the pre-digital era, such rules were enforced irregularly and often after-the-fact, typically in monthly reviews. Digital surveillance closes this gap by enabling real-time detection and automatic sanctions, turning symbolic accountability into enforceable discipline.

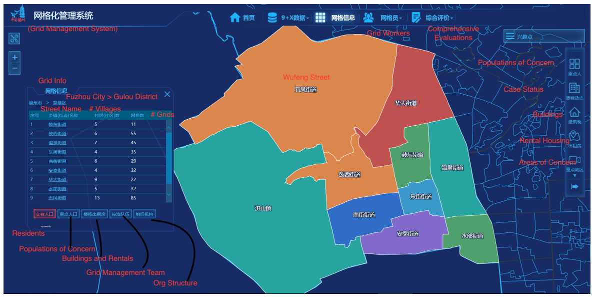

(a) Community divided into grids.

(b) Real-time GPS locations of grid workers.

(c) Mobile application for grid workers to report event scenes and solutions. The leftmost panel is the interface for grid workers to report emerging incidents to their superiors. The next panel shows possible categories for them to report regular information. The next panel shows their inbox that lists the todo tasks assigned by the system. The rightmost panel shows the detail of each task, indicating that they have to solve it within an hour.

(d) The purchase contract of the integrated grid management platform from company EGOVA in 2017

(e) The specific software EGOVA offered for the integrated platform, which includes panel a and b (for command center) and c (for grid workers).

Figure 2: Grid management system used in Fuzhou City, Fujian Province.

4 Data

4.1 Protest

Social movement studies in Western democracies have benefited significantly from protest event datasets (Earl et al., 2004), but historically most literature on collective action in China has taken a qualitative approach due to the lack of such data. Our analysis uses CASM-China, a recently constructed dataset containing 136,330 offline collective action events in China from 2010 to mid-2017 (Zhang and Pan, 2019).[5] CASM-China employs a two-stage deep learning algorithm to identify offline collective action events from 9.5 million Weibo (China's Twitter analog) posts containing at least one of 50 common protest-related words. The first stage removed irrelevant posts and kept posts that discussed offline protests or expressed online grievances. The second stage further selected posts discussing protests and merged them into events. While the online grievance posts were removed in the original CASM-China, they offer valuable data for our robustness checks. CASM-China is accurate[6] and comprehensive.[7]

4.2 Digital Surveillance on agents

Chinese local governments have considerable discretion on when and how to implement their grid management, and whether and when to add digital surveillance on their grid workers. To measure when local governments in China began implementing their grid management systems (with or without workplace monitoring), we utilized government procurement contracts. Since 2013, the Ministry of Finance has mandated local governments to report their purchases exceeding 600,000 RMB on the Ministry's website (*zhengfu caigou wang*).[8] We collected 51,611 contracts that contained the word "grid management". Research assistants then read and identified contracts that involved the purchase of two types of software or hardware. The first type is software aimed at collecting grid workers' real-time GPS locations and enabling the uploading of photos/videos/reports from incident scenes. The second type is software and hardware required to build "intelligence command centers" that analyze digital traces of grid workers, including deviations from their routes. Details of human coding is provided in Appendix A.1.

The Fuzhou Grid Management system (Figure 1) provides an example. Figure 2d displays EGOVA company's winning bid announcement for the grid management integrated system, valued at 7.16M RMB (approximately 1M USD). The detailed contract (highlighted in red) led us to EGOVA's website, where we found product demonstrations (first page shown in Figure 2e, with additional system details in Figures 2b to 2c).

From the procurement data, we counted the number of surveillance-related contracts each city government purchased, as a measure of surveillance intensity on grid workers.

4.3 Surveillance on citizens

For measuring digital surveillance of citizens, we use the facial-recognition procurement dataset from Beraja et al. (2023), which was compiled from the same Chinese Ministry of Finance portal used in our agent-surveillance measure but focuses on contracts explicitly purchasing facial-recognition cameras and software. Compared with monitoring grid workers, facial recognition procurement is more closely linked to China's mass surveillance programs—Skynet (Tianwang), Sharp Eyes (Xueliang), and Smart City reforms. We discuss these systems and their connections to our two digital surveillance measures in Appendix B.

4.4 Descriptive Trends and Summary Statistics

Table F.1 shows summary statistics of variables. The number of contracts for monitoring grid workers is only about 3.5% of the contracts for facial recognition systems monitoring citizens. However, this lower number is misleading because grid workers represent a much smaller population than citizens. In fact, the rate of surveillance contracts per person is 14 times higher for grid workers than citizens, if we make a reasonable assumption of one

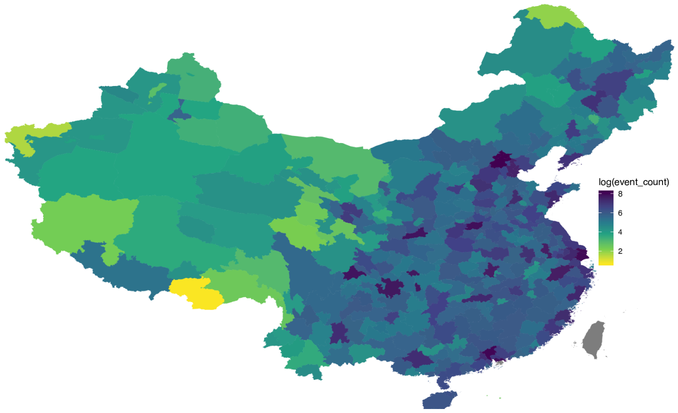

(a) Number of collective action events (logged) in CASM-China, by prefecture-level city.

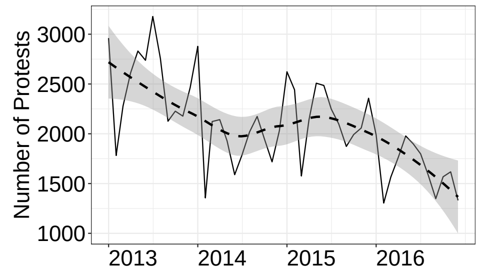

(b) Number of collective action events in CASM-China, by year

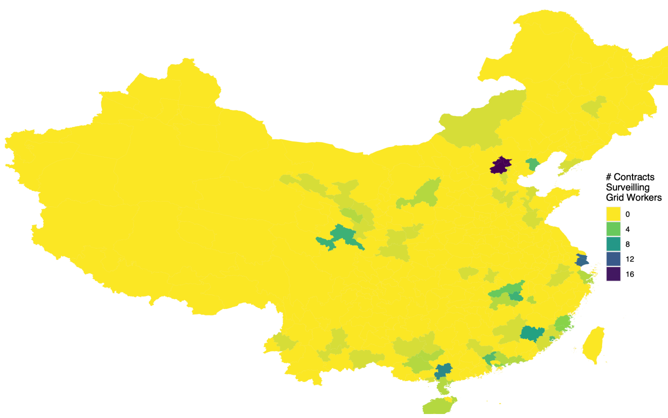

(c) Number of procurement contracts for grid worker surveillance, by prefecture-level city.

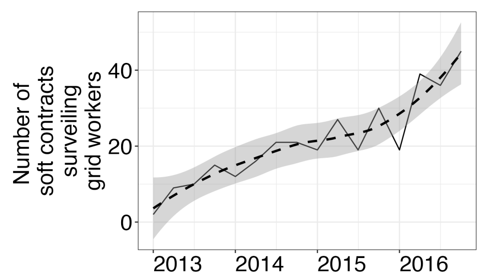

(d) Number of purchase contracts aims to surveil grid workers, by year.

Figure 3: Descriptive plots.

grid worker monitoring 500 citizens.[9] This suggests the heightened risk of being surveilled compared with citizens.

Figure 3a shows that the number of protests varied widely across cities, while Figure 3b documents a clear downward trend in protest activity over time. In contrast, surveillance capacity expanded: Figure 3c maps the spatial distribution of workplace surveillance contracts as of 2017, and Figure 3d shows their rapid temporal growth. Similar expansion patterns hold for facial recognition contracts (see Beraja et al. (2023)). Taken together, these descriptive patterns suggest an alignment with our theoretical expectations: as both citizen-facing and agent-targeted surveillance increased, protest activity declined. Our statistical analysis in the next section will formally see if it is the case.

5 Results

5.1 Digital surveillance of grid workers and citizens each reduces protests

City-quarter analysis Our analysis begins with a city-quarter fixed-effects model to examine how surveillance affects protest activities:

$$\text{Protest}_{i,t+1} = \alpha_i + \gamma_t + \beta \times \text{Surveillance on agents}_{it} + \delta \times \text{Surveillance on citizen}_{it} + \epsilon_{it}, \quad (1)$$

where $\text{Protest}_{i,t+1}$ represents the total number of protests in city i in the next quarter $(t+1)$, and the two surveillance measures are the number of contracts local government had to monitor agents, and citizens (facial recognition), respectively. Fixed effects help us remove confounding factors that are fixed at either city or quarter level.[10] Given that our outcomes are overdispersed event counts,[11] negative binomial models would seem natural. However, as is known among statisticians, standard (conditional) negative binomial models with fixed effects are not true fixed effects (Allison and Waterman, 2002). Although Allison and Waterman (2002) show simulation evidence that unconditional estimation (treating fixed effects as dummy variables) does not suffer from this bias, we lack a theoretical guarantee of its performance. Therefore, we present results from both conditional and unconditional negative binomial estimation. We also employ quasi-Poisson models which explicitly model overdispersion in outcomes, and are robust when fixed effects are included. Last, we use OLS to enable comparison with state-of-the-art two-way fixed effects panel models not yet available for count data.

The results are shown in Table 1 (top panel), which presents estimates from OLS, quasi-

	Panel A: City-Quarter Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts	−2.064*	−0.023**	−0.027*	−0.027**
	(0.859)	(0.007)	(0.012)	(0.010)
Facial Recognition Contracts	−0.517***	−0.004**	−0.005**	−0.005**
	(0.143)	(0.002)	(0.002)	(0.002)
Num. obs.	5312	5312	5312	5312

	Panel B: City-Year Fixed Effects			
	OLS	Poi	NB(conditional)	NB(conditional)
Workplace Surveillance Contracts	−6.687$^{\cdot}$	−0.021***	−0.014*	−0.014*
	(4.036)	(0.005)	(0.007)	(0.006)
Facial Recognition Contracts	−0.643**	−0.002**	−0.002*	−0.002*
	(0.221)	(0.001)	(0.001)	(0.001)
Population (M, log)	−11.175	−0.339	−0.163	−0.163
	(13.512)	(0.343)	(0.301)	(0.262)
GDP (B¥, log)	−7.238	−0.109	−0.172	−0.172
	(8.591)	(0.125)	(0.155)	(0.133)
Tax revenue (B¥, log)	−1.648	0.054	0.064	0.064
	(5.380)	(0.050)	(0.068)	(0.055)
Num. obs.	1325	1325	1325	1325

The outcome variable is the number of protests from CASM-China for each city-quarter. The main variable of interest is the number of workplace surveillance contracts designed to monitor grid workers. Standard errors are clustered at the city and quarter level.

Table 1: Effects of surveillance on protests using two-way fixed effects regression.

Poisson, and two variants of negative binomial models with fixed effects. We find that the number of contracts for monitoring grid workers significantly predicts a decline in protest numbers. As expected, facial recognition AI contracts also significantly predict protest reductions.[12] Overall, the evidence supports that monitoring human surveillants serves as an independent channel for preventive repression, distinct from direct citizen surveillance using facial recognition AI.

City-year analysis with control variables Although city-quarter data provides more observations, it has two limitations: potential seasonal variations and limited control variables at the quarterly level. We therefore run additional regressions at the city-year level with a more comprehensive set of controls. City-year data only has 4 years; given this short panel length, we examine contemporaneous rather than lagged effects. The bottom panel of Table 1 confirms our findings: contracts for monitoring grid workers significantly predict decreased protest numbers, even after controlling for observable city-year confounders.

Grid Management Alone Does Not Reduce Protest One may argue that the grid management system, as a new institution, would reduce protests alone, without digital surveillance. After all, the system is supposed to be a major overhaul of previous separate street-level social control systems (Chen and Greitens, 2022; Mittelstaedt, 2022). But as Section 2.1 explains, grid management is likely insufficient to reduce protests alone, due to either information capacity limitations or principal-agent problems.

We constructed a city-year dummy variable that equals one after the earliest year a city implemented grid management. We relied on multiple sources—newspapers and academic articles—to supplement procurement data (see Appendix A.2). Table 2 shows that grid management alone, without digital surveillance of agents, does not reduce protest activities. In other words, despite its institutional ambition, human agents alone are *unable* to prevent protests; they either require mass surveillance to provide more information or workplace monitoring by their superiors to ensure compliance.

Pre-trends checks and robust estimates Recent econometric research has highlighted potential issues with two-way fixed effects, as they can produce biased estimates by incorporating improper 2-by-2 comparisons. To address this concern, we employ the estimator developed by Abraham and Sun (2021), which eliminates these problematic comparisons. This approach also allows us to estimate pre-treatment effects, which should not exist if the parallel trends assumption holds.

Figure 4 presents these results, with the x-axis showing months since a city's first purchase of grid worker monitoring software. The analysis reveals three key findings. Prior to the first purchase, we find no statistically significant associations, supporting the parallel trends assumption. After cities implemented grid worker monitoring, there was a statis-

	OLS	Poi	NB(unconditional)	NB(conditional)
Implementation of grid management	−2.82	−0.00	−0.03	−0.03
	(4.22)	(0.04)	(0.05)	(0.04)
Population (M, log)	−8.28	−0.35	−0.16	−0.16
	(13.15)	(0.21)	(0.32)	(0.15)
GDP (B¥, log)	−10.35	−0.10	−0.18	−0.18
	(9.83)	(0.12)	(0.16)	(0.11)
Tax revenue (B¥, log)	−2.99	0.04	0.06	0.06
	(4.79)	(0.05)	(0.08)	(0.05)
Num. obs.	1325	1325	1325	1325

Two-way (city-year) fixed effects regression with a placebo outcome. The outcome variable is the number of protests from CASM-China for each city-year. The main variable of interest in this placebo test is the implementation of grid management. These are evidence by purchase of software or hardware for grid management command center, media reports, or academic articles, regardless of whether additional surveillance softward on grid workers were purchased. Column 1 presents results using OLS; column 2 uses quasi-Poisson regression, which accounts for overdispersion in event counts; column 3 employs unconditional negative binomial regression; and column 4 applies conditional negative binomial regression. Standard errors are clustered at the city level to account for potential heteroscedasticity and correlation within cities over time.

Table 2: Implementation of grid management on protests

tically significant reduction in protest numbers. Furthermore, the magnitude of reduction increased in subsequent quarters, suggesting that purchases had a cumulative impact on protest suppression.

Dynamic panel model If a place has experienced more protests before, its street-level agents are more likely to be monitored for their failure to prevent protests, which in turn leads to a future reduction in protests. In essence, past values of the outcomes may influence present values by affecting past independent variables. To address such endogeneity concerns, we use dynamic panel models. We first controlled for the impact of past protests (i.e., $\text{Protest}_{i,t}$) on future events ($\text{Protest}_{i,t+1}$) in Equation 1. Column 1 in Table F.3 shows that results are robust.

Directly adding a lagged dependent variable to a fixed-effects model introduces endogeneity, because the lagged outcome is correlated with the fixed effects—a problem known as Nickell bias (Nickell, 1981). To address this, we use the Arellano-Bond estimator, which differences the model to remove unit fixed effects and instruments the differenced lagged outcome using deeper lags of $\text{Protest}_{i,t}$ starting from $t-2$ (Arellano and Bond, 1991). The

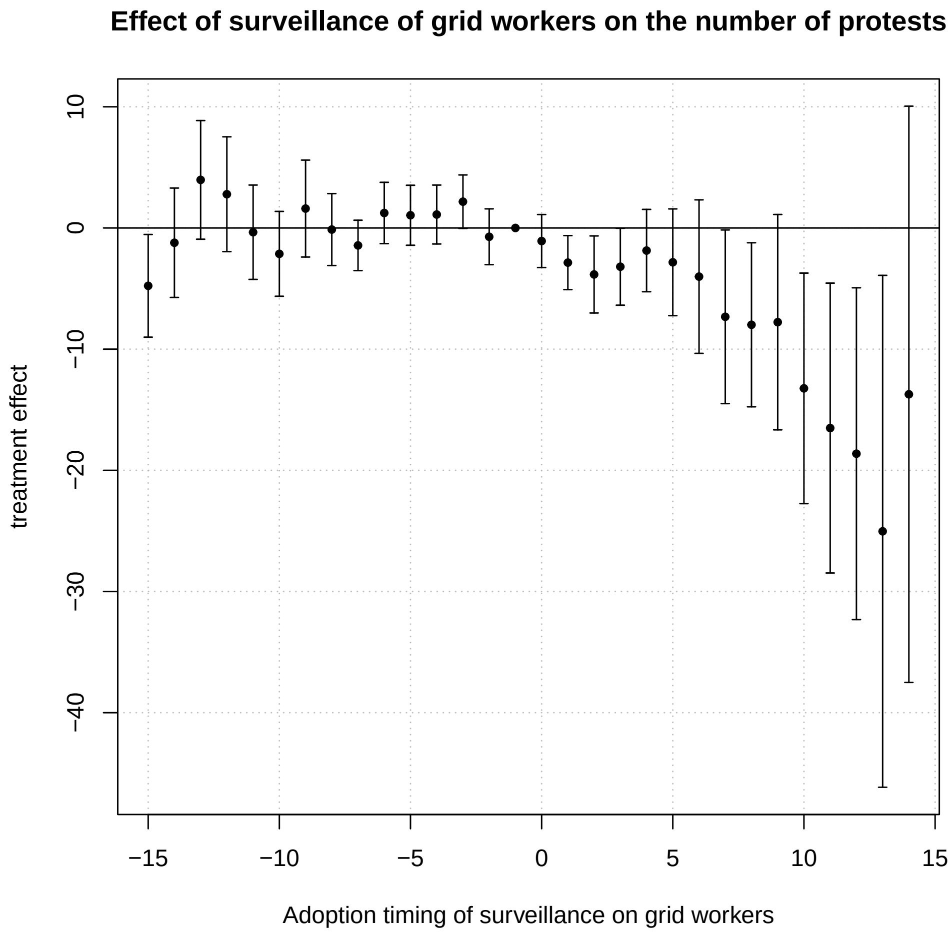

Figure 4: Quarters relative to the purchase of surveillance contracts monitoring grid workers (x-axis) on the number of protests (y-axis).

estimator uses a generalized method of moments (GMM) approach, and Column 2 in Table F.3 shows that both types of surveillance continue to reduce protests after accounting for this dynamic structure.

Exogenous conducive weather as IV The second way we deal with endogeneity is to rely on exogenous weather conditions conducive to protests. The key idea is that certain weather conditions often predict protest activities: bad weather discourages protests while good weather may encourage more protests, and it is clear that weather is exogenous to dictators (leaders), agents, and protesters. However, if digital workplace surveillance is effective, authorities can compel workers to maintain diligence during weather conditions conducive to protests. Consequently, we expect places with more surveillance contracts to show a weaker relationship between conducive weather conditions and protests. We built a conducive weather measure predictive of protests in each city-quarter based on Beraja et al. (2023) (details in Appendix D).[13] Also exactly following Beraja et al. (2023), we focus on how *accumulated* investment in digital surveillance of grid workers moderates the effects of exogenous weather shocks on protests. This is because *contemporaneous* number of contracts would be higher in areas that experience more political protests, which is exactly what we would want to isolate.

Table 3 confirms our expectations: First, conducive weather shows a positive correlation with protest numbers, showing the validity of our measures. Second, the interaction term reveals that digital surveillance of grid workers significantly weakens the relationship between conducive weather and unrest. This weather-based approach helps address concerns that unobserved city characteristics, such as leader's political will or city resources, might drive our results. If protest reduction were solely due to these unobserved city-level factors, we would expect the relationship between conducive weather and protests to remain consistent regardless of whether cities had implemented surveillance on their grid workers. This is because neither leader resolve nor city resources vary with weather conditions. However, we find that surveillance on grid workers offsets the weather-protest association. This finding is difficult to explain through unobserved city characteristics alone.

Exogenous policy shocks as IV Due to recent critiques of using weather as an instrument (Mellon, 2021), we also employ a policy-based instrument. We instrument workplace surveillance using the staggered rollout of the national Smart City pilot program. The Smart City program encouraged pilot cities to build a data-driven "systematic social governance" platform that could track the movements and tasks of grid workers alongside other frontline personnel, making investments in digital attendance and performance-monitoring tools more

	OLS	Poisson	NB(unconditional)	NB(conditional)
Conducive Weather	1.819^{**}	0.063^{*}	0.062	0.062^{*}
	(0.632)	(0.029)	(0.032)	(0.031)
Workplace Surveillance Contracts (cumulative)	-1.884^{***}	-0.021^{**}	-0.020^{**}	-0.020^{**}
	(0.568)	(0.007)	(0.007)	(0.007)
Workplace Surveillance Contracts (cumulative) * Conducive Weather	-0.230	-0.007^{*}	-0.006^{*}	-0.006^{*}
	(0.222)	(0.003)	(0.003)	(0.003)
Num. obs.	5312	5312	5312	5312

The outcome variable is the number of protests from CASM-China for each city-quarter. The main variable of interest is the interaction between conducive weather and the number of workplace surveillance contracts designed to monitor grid workers. Column 1 uses OLS; column 2 uses quasi-Poisson regression, which accounts for overdispersion in event counts; Column 3 uses unconditional negative binomial regression; and Column 4 uses conditional negative binomial regression. Standard errors are clustered at the city and quarter level.

Table 3: Effect of workplace surveillance on protest, instrumented by conducive weather, using two-way (city-quarter) fixed effects regression.

likely in participating municipalities (first-stage relevance). Secondly, the timing of pilot entry was centrally determined and unrelated to short-term protest dynamics in specific cities; each province usually will have some enrolled (exogeneity). Last, Smart City's main goal is to improve governance, not directly protest control, so their effect on protests operates primarily through expanded surveillance capacity. These conditions make the pilot of smart City likely to be an effective IV.

We use the three national Smart City pilot waves—December 2012, May 2013, and April 2015, and mark a city as treated from the first full quarter after it is selected (fewer than one-third of cities were selected by 2017). The indicator remains 1 thereafter; because this instrument provides only within-city time variation, we omit quarter fixed effects in the IV specifications to preserve that signal.

Table 4 shows the results. The IV estimates confirm that Smart City pilots significantly increased workplace surveillance, satisfying first-stage relevance. In the second stage, greater adoption of workplace surveillance is associated with a substantial and statistically significant decline in protest events. These effects remain robust when controlling for facial recognition contracts, suggesting that the observed impact is not driven by other concurrent surveillance technologies. Overall, the results support our core argument that digital monitoring of street-level agents can deter collective action.

	A. Without Facial Recognition		**B. With Facial Recognition**	
	First Stage	**Second Stage**	First Stage	**Second Stage**
Dependent Variable:	WS Contracts	**Protest Count**	WS Contracts	**Protest Count**
In pilot	0.074**		0.052*	
	(0.025)		(0.024)	
Workplace Surveillance Contracts		**-49.119****		**-56.908***
		(17.448)		(26.948)
Facial Recognition Contracts				0.402
				(0.442)
Fixed Effects				
City Code	Yes	Yes	Yes	Yes
Fit Statistics				
Observations	5,312	5,312	5,312	5,312
R^2	0.417	0.423	0.438	0.274
Within R^2	0.003	-2.638	0.038	-3.576

Clustered standard errors by city code in parentheses.

*Significance levels: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$*

The dependent variable is the number of protest events in each city-quarter, drawn from the CASM-China dataset. The key independent variable is the number of workplace surveillance contracts designed to monitor frontline grid workers. All models include city fixed effects. Standard errors are clustered by city code to account for serial correlation within units. The second-stage estimates represent the causal effect of surveillance contracts on protest activity, instrumented by a pilot policy rollout. Column set A excludes facial recognition contracts; column set B includes them as an additional control. The F-statistics for set A and B are 16.8 and 8.76, respectively.

Table 4: Effect of workplace surveillance on protest activity, instrumented by a staggered policy pilot for the Smart City program, with city fixed effects.

5.2 Mediation and interaction between two types of surveillance

As stated in Hypothesis 3a, we expect that citizen-targeted surveillance influences protests indirectly, by enhancing surveillance of street-level agents. In contrast, agent-targeted surveillance does not operate through citizen monitoring.

Using causal mediation analysis (Imai et al., 2010), we test both directions of the hypothesized relationship. Panel A of Table 5 shows that citizen-targeted surveillance reduces protest via increased grid-worker monitoring, with a consistently negative and highly significant ACME across OLS (–1.08), Poisson (–0.15), and negative binomial (–0.16) models ($p < 0.001$). This indirect pathway accounts for 35% (OLS), 28% (Poisson) and 26% (Negative Binomial) of the total effect of corresponding models. Conversely, Panel B of Table 5 shows that agent-targeted surveillance's impact on protest is virtually unmediated by citizen surveillance—the ACME is near zero, the mediated share only 5–7 percent, and the effect is not significant in the negative-binomial model.

	Panel B: Effect of Agent Surveillance on Protest, Mediated by Citizen Surveillance		
	OLS	Poisson	NB (Unconditional)
ACME	-1.09^{***}	-0.15^{***}	-0.16^{**}
ADE	-2.05^{***}	-0.38^{***}	-0.46
Total	-3.14^{***}	-0.53^{***}	-0.61^{*}
Prop. Mediated	0.34	0.28	0.26

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

	Panel B: Effect of Citizen Surveillance on Protest, Mediated by Agent Surveillance		
	OLS	Poisson	NB (Unconditional)
ACME	-0.03^{***}	-0.01^{***}	-0.01
ADE	-0.52^{***}	-0.07^{***}	-0.08^{***}
Total	-0.55^{***}	-0.08^{***}	-0.08^{***}
Prop. Mediated	0.05	0.07	0.07

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 5: Mediation Analysis.

Our mediation analysis reveals an asymmetric pattern: citizen surveillance influences protest in part through enhanced agent oversight, while agent surveillance can reduce protest directly. Even so, the two channels may still reinforce one another when implemented jointly, as posited in Hypothesis 3b. We test this by interacting citizen- and agent-level surveillance in our two-way fixed effect model. Table 6 presents the interaction effects at both the quarterly (top panel) and yearly (bottom panel) levels. OLS models consistently show statistically significant interaction terms between grid worker monitoring and facial recognition AI in both city-year and city-quarter specifications.

The count models reveal a more nuanced temporal pattern: city-year specifications show negative and statistically significant interaction coefficients in both Poisson and negative binomial models, while city-quarter specifications retain the negative direction but lose statistical significance. One plausible explanation is that while each surveillance channel independently reduces protests in the short term, their combined effect emerges only after sufficient time for institutional integration and interdepartmental coordination.

5.3 Robustness checks

We conducted a series of robustness checks to make sure our findings are robust.

Selection biases in cities A potential concern is that surveillance contracts may not be randomly distributed: cities with higher or lower protest risk may systematically differ in ways that also influence surveillance adoption. From raw descriptive statistics, cities that eventually sign surveillance contracts tend to be larger, wealthier, and have experienced more protests on average (Table F.2). However, these differences reflect underlying cross-city heterogeneity rather than causal selection. Once we control for time-invariant city characteristics (two-way fixed effects) and test whether last quarter's protest count predicts current quarter contract procurement, the relationship disappears (Table F.4). This pattern of higher unconditional means but no positive within-city effect suggests that cities' prior unrest does not systematically drive surveillance adoption.

	Panel A: City-Quarter Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts (sqrt)	4.434***	−0.017	−0.045	−0.045
	(0.892)	(0.032)	(0.046)	(0.040)
Facial Recognition Contracts (sqrt)	−0.756***	−0.019**	−0.018	−0.018*
	(0.185)	(0.007)	(0.011)	(0.009)
Facial Recognition (sqrt) * Workplace Surveillance (sqrt)	−3.194***	−0.011	−0.006	−0.006
	(0.249)	(0.007)	(0.010)	(0.010)
Num. obs.	5312	5312	5312	5312

	Panel B: City-Year Fixed Effects			
	OLS	Poi	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts (sqrt)	19.319**	0.017	0.020	0.020
	(5.872)	(0.036)	(0.034)	(0.033)
Facial Recognition Contracts (sqrt)	2.183$^{\cdot}$	0.004	0.003	0.003
	(1.186)	(0.010)	(0.011)	(0.008)
Facial Recognition (sqrt) * Workplace Surveillance (sqrt)	−7.323***	−0.013**	−0.013*	−0.013*
	(1.469)	(0.005)	(0.005)	(0.005)
Population (M, log)	−8.497	−0.318	−0.156	−0.156
	(13.654)	(0.377)	(0.320)	(0.153)
GDP (B¥, log)	−6.437	−0.091	−0.165	−0.165
	(7.648)	(0.143)	(0.158)	(0.106)
Tax revenue (B¥, log)	−2.578	0.052	0.060	0.060
	(4.672)	(0.067)	(0.079)	(0.050)
Num. obs.	1325	1325	1325	1325

The outcome variable is the number of protests from CASM-China for each city-quarter. The main variable of interest is the interaction between the number of workplace surveillance contracts designed to monitor grid workers and the number of facial recognition contracts designed to monitor citizens. Standard errors are clustered at the city and quarter level.

Table 6: Effects of surveillance on protests using two-way fixed effects regression; interaction effect to test the two-layer Panopticon framework.

Falsification tests: contracts irrelevant to surveillance of grid workers Some procurement contracts mention the Chinese term "grid management" (*wangge hua guanli*), but are unrelated to social grid managements, due to word polysemy. For instance, there are many contracts that are relevant to environmental grid management, specifically monitoring air pollution in spatial “grids”. The number of contracts unrelated to social grid management should have no effect on protests.[14] Consistent with this expectation, we find no significant association between these unrelated contracts and protest activity (Table F.5).

Outcome measures using other datasets We replicated our analysis using an alternative dataset constructed from Chinese-language newspaper articles covering 8,859 protests from 2010 to 2017 (Chen, 2020). All events in this dataset were manually verified, and they came from mass media, not social media, making it ideal to validate our findings. Table F.6 replicates the city-year two-way fixed effects regression from Table 1, showing similar results.[15]

Outcome measures using more precise predicted events Since CASM-China identified protests from social media posts, the subset of events identified from Weibo posts with both text and images should have higher classification accuracy compared to events that rely solely on text information. Table F.7 and F.8 replicate Table 1 and 6 with the subset of CASM-China that includes both text and images associated with them on social media. The conclusions are qualitatively similar.

Correcting statistical prediction errors We employed the statistical correction method proposed by Hopkins and King (2010) to adjust for potential false classifications in the raw CASM protest counts, producing a corrected protest measure as new outcome variable measures. Table F.9 and F.10 present the results that corrected event count measures. The results are qualitatively similar.

5.4 Mechanisms

Our primary theory posits that monitoring reduces grid workers’ shirking behavior. Alternatively, a chilling mechanism could explain the same results. Under this mechanism, grid workers may still shirk their duties in ways monitoring technology cannot easily detect. For instance, they might mechanically complete their patrol routes without paying adequate attention to their surroundings. Yet citizens, unable to distinguish between diligent and shirking grid workers, perceive increased risk and thus refrain from participating. We are not aware of any protester survey data in China on whether they feel restrained. Instead, we compare these mechanisms through two sets of indirect evidence based on predicted effects across protest types.

Target differences Under a chilling effect, citizens would strategically reduce more sensitive forms of protest, particularly those targeting the government. Under reduced shirking, we expect uniform reductions across protest types because grid workers have neither advance knowledge of protest characteristics nor instructions to prioritize specific types.[16] Once grid workers notice mobilization attempts, they have no reason to only report on anti-government protests while leaving other mobilizations untouched. We directly examined the difference between anti-government protests and protests against private companies as the outcome variable. We found no statistically discernible effect of surveillance measures on this difference (Table F.11).

Using a different formulation, we estimated separate regressions using anti-government and non-anti-government protests as outcomes. We then employed seemingly unrelated regression (SUR) to test whether surveillance measures have differential effects across the two models. The SUR results confirm that the coefficients of surveillance measures do not significantly differ between the two models using protests against/not against government as outcome variables ($p < 0.92$).

Overall, both approaches consistently show that digital surveillance on agents does not differentially affect protests based on their target. This uniform reduction pattern across targets aligns with the reduced shirking mechanism: monitoring appears to enhance overall agent effectiveness rather than inducing strategic responses from citizens.

Offline protest versus online dissent Grid workers monitor only physical spaces, having neither access to nor responsibility for *online* activities. Therefore, if digital surveillance works through reduced shirking, it should only reduce street protests. If it creates a broader chilling effect, letting citizens self-discipline, it should reduce both online and offline dissent.

To measure online grievances, we use the grievance posts from the first stage of CASM-China. These grievance posts, while initially filtered out during the second stage of CASM-China construction, provide critical information about online discontent (Zhang and Pan, 2019). We find that surveillance on grid workers has no effect on online grievances (Table F.12), even though offline protests decline and online and offline dissent are highly correlated (Pearson $\rho > 0.88$). These results support our mechanism that digital surveillance works through active prevention by monitored grid workers, rather than through a general chilling effect.

6 Conclusions And Discussions

Our two-layer Panopticon framework reveals two pathways through which dictators can use digital surveillance to reduce collective action. The conventional view emphasizes using digital technologies to monitor citizens, identify mobilization efforts, preempt protests, and deter potential participants. The new pathway proposed in this study shows that digital surveillance can also be used to monitor agents responsible for street-level prevention tasks and ensure they are working to fulfill dictators' goals rather than shirking. Theoretically, we identify a neglected channel of authoritarian control: digital surveillance of street-level agents; we challenge assumptions in bureaucratic control research by showing that digital tools make frontline effort directly observable and reduce reliance on intermediate supervisors; and we shift the focus of digital authoritarianism from citizen surveillance to the monitoring of state agents, highlighting the organizational foundations of authoritarian resilience.

Empirically, this study provides the first quantitative evidence on how AI-enhanced workplace surveillance of grid workers affects protest prevention. We find that surveilling agents and surveilling citizens are both independently effective in reducing protests. When both are present, protests decline even further, though this reinforcing effect emerges only after sufficient time for institutional coordination. Our causal mediation analysis further shows an asymmetric mechanism: about one-third of the protest-reducing effect of citizen-facing surveillance is indirectly transmitted through tighter oversight of grid workers. Together, these findings demonstrate how digital technologies reshape both state–society interactions and internal bureaucratic oversight.

We conclude by discussing some future research directions. Our two-layer Panopticon framework is the simplest starting point to disaggregate the state so that we can discuss how dictators surveil both bureaucrats and citizens. In real life, there are often multiple principals (such as officials at different levels of the government) and multiple types of agents (including grid workers, police, and secret police). Agency theory suggests that the coexistence of multiple principals and agents tends to increase the difficulty of monitoring activities (Waterman and Meier, 1998). This implies that in the absence of digital surveillance on agents, the severity of principal-agent problems is likely to become more pronounced, thus making our contribution even more relevant.

An important next step is to assess the broader applicability of our findings. Future work could examine how surveillance reshapes other protest dynamics, such as whether it shortens mobilization windows, reduces crowd size, or shifts tactics toward violence or online contention, and test the logic in contexts where only one surveillance layer is present (e.g., mass citizen monitoring without agent oversight, or the reverse). Such extensions face empirical challenges: once the state preempts earlier, the remaining protests become rarer and systematically different, introducing selection bias (Ritter and Conrad, 2016).

A two-layer Panopticon seems intimidating, but is not flawless. Sustaining the two-layer Panopticon is costly. Recent media reports suggest that after COVID, some local governments cannot afford to pay their grid workers,[17] raising concerns about the sustainability of the two-layer Panopticon.

Organizational factors also limit effectiveness. First, agent-facing monitoring can create adverse selection—pushing out more capable agents while retaining those with fewer outside options—and generate second-order principal–agent problems, as those tasked with operating surveillance systems can themselves shirk or manipulate data. Second, integrating citizen- and agent-facing systems is not automatic: in many regimes, they operate in separate bureaucracies, sometimes by deliberate design, and require sustained coordination as well as perceived legitimacy among frontline agents. In the short term, China mitigates many of these obstacles: grid workers have few outside options, automation reduces the need for human monitors, and authorities have invested in integrating agent- and citizen-facing systems (e.g., PLAC command centers). These features help explain why digital surveillance is effective in China so far, though whether similar results could emerge elsewhere or persist over time remains uncertain.

Last, although this paper focuses on top-down state behaviors, both agents and citizens may not only self-discipline and comply (thereby reducing protest) but also *resist* the system. When citizens perceive surveillance as excessive or intrusive—particularly when they are monitored by both state agents and neighbors—it can provoke resentment, subversion, or even backlash. Agents in authoritarian regimes rarely engage in overt resistance compared to their democratic counterparts, yet they often adopt subtle, everyday forms of opposition. Behaviorally, over-reliance on digital logs may encourage agents to find new ways to evade the gaze of workplace surveillance. Psychologically, it can make some feel doubly scrutinized —both as employees and as ordinary citizens—or see their experiential knowledge devalued (O'Brien, 2017; Brayne, 2020), generating resentment and disengagement. In extreme cases, dissatisfaction with the state has even pushed frontline agents to side with protesters, as seen during the collapse of some late-stage authoritarian regimes, including former Soviet states. Although a detailed investigation of these forms of resistance is beyond the scope of this article, they represent an important direction for future research on surveillance and contentious politics.

To conclude, while the two-layer Panopticon may appear to be a digital dystopia, it faces persistent vulnerabilities from fiscal, organizational constraints, and both human resistance. The long-term implications of such comprehensive surveillance systems for state-society relations remain an important area for future research.

Author Biography

Han Zhang is an assistant professor in the Department of Sociology and the Watson School of International and Public Affairs, Brown University. His research interests include political sociology, social movements, computational social science, and quantitative methodology. His research has been published in *Sociological Methodology*, *Sociological Methods and Research*, *Socio-Economic Review*, *Mobilization*, *The China Quarterly*, among others. He has won best paper awards from American Sociological Association and International Communication Association.

Data availability

The data collection procedures are available in the article and in its online supplementary material. It will be released at `https://osf.io/fy63j/overview?view_only=537af1a5b56c4c9e97fe64980a`

Endnotes

[1]Even within the category of "dictator," agents may receive conflicting instructions, as central leaders and local officials often pursue different policy goals. This phenomenon is well-documented in the literature on fragmented authoritarianism Lieberthal and Lampton (1992); Mertha (2009). For instance, in China, the central government tends to favor softer approaches to protest management, while local officials—who directly command frontline agents—often prefer more repressive tactics to quickly restore order.

[2]There are numerous examples of shirking behaviors among police and secret police in their surveillance duties. For the GDR, see Gieseke (2014, p. 87); for the USSR and modern Russia, see Beissinger and Kotkin

(2014, pp. 144-146); for Czechoslovakia, see (Pucci, 2020, p. 240).

[3]`https://www.xndjw.gov.cn/website/contents/54/80255.html`; `https://www.lg.gov.cn/bmzz/lcjdb/csjs/zdxm/content/post_12119524.html`

[4]Grid workers sometimes are synonymous with or encompass pre-existing street-level bureaucrats, such as *Chengguan* (urban management personnel) and workers in *Ju Wei Hui* (residents' committee). They may also include people who have connections with the party-state, such as retired party cadres, veterans, and senior members of the clan (in rural areas).

[5]We further cleaned the data by removing events that have more than two issues, resulting in 130,153 events. There are also questions related to why there were still discussions about collective action on Chinese social media, given that studies by King et al. (2013) found that censorship is prevalent. Zhang and Pan (2019) found that the reason is that censorship is porous; large protests almost always have uncensored discussions on Weibo, whereas small protests are less likely to be deleted (Roberts, 2018). Zhang and Pan (2019) also provided evidence that Weibo censorship only causes a 2.7% underestimation of the overall event count.

[6]Human validation confirms its high accuracy, with an F1 score of 0.84, a false positive rate of 0.15, and a false negative rate of 0.1 (Zhang and Pan, 2019).

[7]CASM-China captures 10 to 100 times more collective action events than newspaper-based datasets and covers a broader range of issues.

[8]`https://www.ccgp.gov.cn`.

[9]See Pei (2024) for estimates of various street-level agents in China.

[10]Notably, social media surveillance can also be used to identify mobilizations that occur online and potentially lead to offline protests, but we do not have measures for its strength. Fortunately, social media surveillance is unlikely to vary drastically across cities because the central government directly supervises the censorship apparatus over major platforms like Weibo. To the best of the authors' knowledge, previous research has found no evidence that city governments can influence content moderation by national platforms (King et al., 2013; Roberts, 2018). Therefore, while social media surveillance's effect on protests is itself important, it does not bias our estimates due to the fixed effects.

[11]As evidenced by the "skewness" column in Table F.1.

[12]When we add each surveillance measure (on citizens or on grid workers) as an independent variable separately, each shows a significant reduction in the number of protests.

[13]Note that Beraja et al. (2023) used different outcome measures from GDELT.

[14]The measurement details are discussed in Section A.3.

[15]Chen (2020) is 15 times smaller than CASM-China, so we did not perform city-quarter level regression.

[16] We created measures for the target for each protest based on supervised machine learning. More information on these variables can be found in Appendix E.

[17] `https://www.epochtimes.com/gb/24/8/26/n14317658.htm`

References

Abraham, Sarah and Liyang Sun. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects". *Journal of Econometrics* 225 (2): 175–199.

Allison, Paul D. and Richard P. Waterman. 2002. "Fixed-Effects Negative Binomial Regression Models". *Sociological Methodology* 32 (1): 247–265.

Arellano, Manuel and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations". *The Review of Economic Studies* 58 (2): 277–297.

Ariel, Barak, William A. Farrar, and Alex Sutherland. 2015. "The Effect of Police Body-Worn Cameras on Use of Force and Citizens' Complaints Against the Police: A Randomized Controlled Trial". *Journal of Quantitative Criminology* 31 (3): 509–535.

Ball, Kirstie. 2010. "Workplace surveillance: An overview". *Labor History* 51 (1): 87–106.

Beissinger, Mark and Stephen Kotkin (Eds.)2014. *Historical Legacies of Communism in Russia and Eastern Europe*. New York: Cambridge University Press.

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

Brayne, Sarah. 2020. *Predict and Surveil: Data, Discretion, and the Future of Policing* (1st ed.). New York, NY: Oxford University Press.

Brehm, John O. and Scott Gates. 1999. *Working, Shirking, and Sabotage: Bureaucratic Response to a Democratic Public*. Ann Arbor: University of Michigan Press.

Cai, Yongshun. 2010. *Collective Resistance in China: Why Popular Protests Succeed or Fail.* Stanford University Press.

Chen, Chih-Jou Jay. 2020. "A Protest Society Evaluated: Popular Protests in China, 2000–2019". *Mobilization: An International Quarterly* 25 (SI): 641–660.

Chen, Huirong and Sheena Chestnut Greitens. 2022. "Information capacity and social order: The local politics of information integration in China". *Governance* 35 (2): 497–523.

Chen, Xi. 2012. *Social Protest and Contentious Authoritarianism in China.* Cambridge University Press.

Chin, Josh and Liza Lin. 2022. *Surveillance State: Inside China's Quest to Launch a New Era of Social Control.* New York: St. Martin's Press.

Conover, Emily, Daniel Kraynak, and Prakarsh Singh. 2023. "The effect of traffic cameras on police effort: Evidence from India". *Journal of Development Economics* 160 : 102953.

Danneman, Nathan and Emily Hencken Ritter. 2014. "Contagious Rebellion and Preemptive Repression". *Journal of Conflict Resolution* 58 (2): 254–279.

Davenport, Christian. 2005. "Understanding Covert Repressive Action: The Case of the U.S. Government against the Republic of New Africa". *The Journal of Conflict Resolution* 49 (1): 120–140.

de Brito, Charlotte and Barak Ariel. 2017. "Does Tracking and Feedback Boost Patrol Time in Hot Spots? Two Tests". *Cambridge Journal of Evidence-Based Policing* 1 (4): 244–262.

De Jaegher, Kris and Britta Hoyer. 2019. "Preemptive Repression: Deterrence, Backfiring, Iron Fists, and Velvet Gloves". *Journal of Conflict Resolution* 63 (2): 502–527.

Deng, Yanhua and Kevin J. O'Brien. 2013. "Relational Repression in China: Using Social Ties to Demobilize Protesters". *The China Quarterly* 215 : 533–552.

Dragu, Tiberiu and Adam Przeworski. 2019. “Preventive Repression: Two Types of Moral Hazard”. *American Political Science Review* 113 (1): 77–87.

Earl, Jennifer. 2003. “Tanks, Tear Gas, and Taxes: Toward a Theory of Movement Repression”. *Sociological Theory* 21 (1): 44–68.

Earl, Jennifer, Andrew Martin, John D. McCarthy, and Sarah A. Soule. 2004. “The Use of Newspaper Data in the Study of Collective Action”. *Annual Review of Sociology* 30 : 65–80.

Feldstein, Steven. 2021. *The Rise of Digital Repression: How Technology Is Reshaping Power, Politics, and Resistance* (1st edition ed.). New York, NY: Oxford University Press.

Foucault, Michel. 1977. *Discipline and Punish: The Birth of the Prison*. Vintage Books.

Frantz, Erica, Andrea Kendall-Taylor, and Joseph Wright. 2020. “Digital Repression in Autocracies”.

Gandhi, Jennifer and Adam Przeworski. 2006. “Cooperation, Cooptation, and Rebellion Under Dictatorships”. *Economics & Politics* 18 (1): 1–26.

Gieseke, Jens. 2014. *The History of the Stasi: East Germany’s Secret Police, 1945-1990*. New York: Berghahn Books.

Goebel, Christian and Lynette H. Ong. 2012. “Social Unrest in China”. SSRN Scholarly Paper ID 2173073, Social Science Research Network, Rochester, NY.

Goebel, Christian and H. Steinhardt. 2019. “Better coverage, less bias: Using social media to measure protest in authoritarian regimes”.

Gohdes, Anita R. 2020. “Repression Technology: Internet Accessibility and State Violence”. *American Journal of Political Science* 64 (3): 488–503.

Goold, Benjamin J. 2003. “Public Area Surveillance and Police Work: The impact of CCTV on police behaviour and autonomy”. *Surveillance & Society* 1 (2): 191–203.

Greitens, Sheena Chestnut. 2016. *Dictators and Their Secret Police: Coercive Institutions and State Violence.* Cambridge Studies in Contentious Politics. Cambridge: Cambridge University Press.

Hopkins, Daniel J. and Gary King. 2010. "A Method of Automated Nonparametric Content Analysis for Social Science". *American Journal of Political Science* 54 (1): 229–247.

Huang, Kaihui, Weijie Luo, Weiwei Zhang, and Jinhai Li. 2021. "Characteristics and Problems of Smart City Development in China". *Smart Cities* 4 (4): 1403–1419.

Imai, Kosuke, Luke Keele, and Dustin Tingley. 2010. "A general approach to causal mediation analysis". *Psychological Methods* 15 (4): 309–334.

Kendall-Taylor, Andrea, Erica Frantz, and Joseph Wright. 2021. "The Digital Dictators".

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 (02): 326–343.

Lee, Ching Kwan and Yonghong Zhang. 2013. "The Power of Instability: Unraveling the Microfoundations of Bargained Authoritarianism in China". *American Journal of Sociology* 118 (6): 1475–1508.

Levy, Karen. 2022. *Data Driven: Truckers, Technology, and the New Workplace Surveillance.* Princeton University Press.

Li, Hongbin and Li-An Zhou. 2005. "Political turnover and economic performance: The incentive role of personnel control in China". *Journal of Public Economics* 89 (9): 1743–1762.

Lieberthal, Kenneth and Michel Oksenberg. 1988. *Policy Making in China: Leaders, Structures, and Processes.* Princeton University Press.

Lieberthal, Kenneth G. and David M. Lampton (Eds.)1992. *Bureaucracy, Politics, and Decision Making in Post-Mao China* (1 ed.), Volume 14. University of California Press.

Liu, Howard and Christopher M. Sullivan. 2021. "And the Heat Goes On: Police Repression and the Modalities of Power". *Journal of Conflict Resolution* 65 (10): 1657–1679.

Mattingly, Daniel C. 2019. *The Art of Political Control in China.* Cambridge University Press.

Mellon, Jonathan. 2021. "Rain, Rain, Go Away: 176 Potential Exclusion-Restriction Violations for Studies Using Weather as an Instrumental Variable". SSRN Scholarly Paper 3715610, Social Science Research Network, Rochester, NY.

Mertha, Andrew. 2009. ""Fragmented Authoritarianism 2.0": Political Pluralization in the Chinese Policy Process". *The China Quarterly* 200 : 995–1012.

Miao, Michelle. 2024. "Coded Social Control: China's Normalization of Biometric Surveillance in the Post COVID-19 Era". *Washington Journal of Law, Technology & Arts* 19 : 53.

Mittelstaedt, Jean Christopher. 2022. "The grid management system in contemporary China: Grass-roots governance in social surveillance and service provision". *China Information* 36 (1): 3–22.

Nickell, Stephen. 1981. "Biases in Dynamic Models with Fixed Effects". *Econometrica* 49 (6): 1417–1426.

O'Brien, Kevin J. 2017. "China's Disaffected Insiders". *Journal of Democracy* 28 (3): 5–13.

O'Brien, Kevin J. and Lianjiang Li. 2006. *Rightful Resistance in Rural China.* Cambridge University Press Cambridge.

Oi, Jean C. 1999. *Rural China Takes Off: Institutional Foundations of Economic Reform.*

Ong, Lynette H. and Kevin Wei Luo. 2024. "Stability Maintenance, Preventive Repression, and Contentious Politics in China under Xi Jinping's Rule". *Issues & Studies*: 2450015.

Pan, Jennifer. 2020. *Welfare for Autocrats: How Social Assistance in China Cares for Its Rulers*. Oxford, New York: Oxford University Press.

Pei, Minxin. 2024. *The Sentinel State: Surveillance and the Survival of Dictatorship in China*. Cambridge, Massachusetts London, England: Harvard University Press.

Pucci, Molly. 2020. *Security Empire: The Secret Police in Communist Eastern Europe*. New Haven ; London: Yale University Press.

Ritter, Emily Hencken and Courtenay R. Conrad. 2016. "Preventing and Responding to Dissent: The Observational Challenges of Explaining Strategic Repression". *American Political Science Review* 110 (1): 85–99.

Roberts, Margaret E. 2018. *Censored: Distraction And Diversion Inside China's Great Firewall*. Princeton University Press.

Scoggins, Suzanne E. 2021. *Policing China: Street-Level Cops in the Shadow of Protest*. Ithaca New York: Cornell University Press.

Scoggins, Suzanne E. and Kevin J. O'Brien. 2016. "China's Unhappy Police". *Asian Survey* 56 (2): 225–242.

Sewell, Graham and James R. Barker. 2006. "Coercion versus Care: Using Irony to Make Sense of Organizational Surveillance". *The Academy of Management Review* 31 (4): 934–961.

Sherman, Lawrence W. 2013. "The Rise of Evidence-Based Policing: Targeting, Testing, and Tracking". *Crime and Justice* 42 (1): 377–451.

Shirk, Susan L. 1993. *The Political Logic of Economic Reform in China*. Berkeley: University of California Press.

Tirole, Jean. 2021. “Digital Dystopia”. *American Economic Review* 111 (6): 2007–2048.

Truex, Rory. 2019. “Focal Points, Dissident Calendars, and Preemptive Repression”. *Journal of Conflict Resolution* 63 (4): 1032–1052.

Waterman, Richard W. and Kenneth J. Meier. 1998. “Principal-Agent Models: An Expansion?”. *Journal of Public Administration Research and Theory* 8 (2): 173–202.

Williams, Monica. 2021. “Explaining public support for body-worn cameras in law enforcement”. *Police Practice and Research* 22 (6): 1603–1617.

Wilson, James Q. 1989. *Bureaucracy* (First Edition ed.). New York: Basic Books.

Wright, Teresa. 2018. *Popular Protest in China*. Cambridge, UK ; Medford, MA: Polity.

Xu, Xu. 2021. “To Repress or to Co-opt? Authoritarian Control in the Age of Digital Surveillance”. *American Journal of Political Science* 65 (2): 309–325.

Zhang, Han and Jennifer Pan. 2019. “CASM: A Deep-Learning Approach for Identifying Collective Action Events with Text and Image Data from Social Media”. *Sociological Methodology* 49 (1): 1–57.

Zhou, Xueguang. 2010. “The Institutional Logic of Collusion among Local Governments in China”. *Modern China* 36 (1): 47–78.

Zhou, Xueguang, Hong Lian, Leonard Ortolano, and Yinyu Ye. 2013. “A Behavioral Model of “Muddling Through” in the Chinese Bureaucracy: The Case of Environmental Protection”. *The China Journal* (70): 120–147.

Appendices

A Data on grid management

A.1 Data cleaning

We downloaded contracts containing one of the following keywords: 网格化社会管理, 网格管理, 社会网格, 网格员, 网格治理. These terms are alternative ways to express the same concept of social grid management systems. We obtained 51,611 contracts from 2010 to 2017. We then filtered out contracts that included the purchase of surveillance software to monitor grid workers (and other street-level agents who are often employed as part-time grid workers). Further selection was made for contracts containing at least one keyword in the below quoted list and dated between 2013 and 2016, resulting in 3,384 contracts. Research assistants read every procurement contract in detail to determine whether they involved surveillance on grid workers and other street-level agents. A contract was considered eligible and retained if it exhibited the following procurement contents: 1) Software aimed at collecting real-time GPS locations of grid workers and allowing the uploading of photos/videos/reports from incident scenes. 2) Software and hardware required to build "intelligence command centers" that aim to analyze digital traces of grid workers, from whether they deviate from their routes to how frequently they do so. After manual assessment, 330 contracts were retained.

> 轨迹 ("*trajectory*"); 网格 ("*grid*"); 跟踪, 追踪 ("*track*"); 巡逻, 巡检, 巡查 ("*patrol, inspection*"); 路线 ("*route*"); 痕迹 ("*traces*"); 行程 ("*schedule*"); 大屏, 屏幕 ("*screen, large screen*"); 留痕 ("*log retention*"); 实时位置 ("*real time positioning*"); 寻访记录 ("*community service recording*"); 台帐 ("*ledger*"); 手机 ("*mobile phone*"); *App (–)*

A.2 Auxiliary data on time of each city's implementation of grid management

We first reviewed each city's contracts chronologically to identify their earliest digital surveillance purchase. However, the Chinese Ministry of Finance's procurement website only started in 2013. A city might have purchased grid management equipment in 2012, but our contract data would only show purchases from 2013 onward. To address this limitation, we used two additional sources:

- Official local CCP and government newspaper articles about grid management, which frequently report on the system's implementation
- Academic articles authored by government officials, researchers with government appointments (*guazhi*), or scholars who conducted field research on local grid management.

For cities without procurement documents in 2013, our research assistants searched these additional sources for evidence of pre-2013 implementation. Finally, we created city-year dummy variables equal to one for all years after a city's earliest documented implementation of grid management.

A.3 Auxiliary data for placebo tests

We used contracts that were not among the final 330 identified contracts as irrelevant contracts. These were contracts that initially appeared to be related to surveilling grid workers (through keyword searches) but were actually not, as determined through human verification. We used this data in Table F.5.

B Mapping China's digital surveillance systems and study variables

Table B.1: This table reports the 10 cities with the highest total number of protests during the study period. It also shows the total number of government procurement contracts related to workplace surveillance and facial recognition technologies. The full dataset is available at `https://osf.io/wq5aj?view_only=537af1a5b56c4c9e97fe64980ac41b84`.

City (in Chinese)	Total Protests	Total Contracts for Workplace Surveillance	Total Contracts for Facial Recognition
Beijing	1907	50	622
GuangZhou	1809	11	592
Chongqing	1782	2	0
ZhengZhou	1700	4	101
Shanghai	1675	38	717
ShenZhen	1517	3	157
ChengDou	1503	0	2
XiAn	1491	0	46
NanJing	1277	1	16
WuHan	1256	1	21

B.1 China's digital surveillance systems

China's digital surveillance infrastructure comprises several overlapping national initiatives, each with distinct objectives and implications for social control. This appendix clarifies the relationship between three major programs—Skynet (*Tianwang*, 天网), Sharp Eyes (*Xueliang*, 雪亮工程), and smart-city grid management systems—and the two categories of procurement contracts analyzed in our study: Facial Recognition Contracts and Workplace Surveillance Contracts.

Skynet, launched in the early 2010s by the Ministry of Public Security, expanded visual surveillance in urban areas through widespread installation of CCTV cameras and real-time facial recognition systems. Our Facial Recognition Contracts variable captures the software components of this initiative—specifically, identity-matching systems embedded within broader camera infrastructures.

Sharp Eyes, initiated in 2016 and rolled out nationally between 2018 and 2020, sought to extend surveillance to rural and residential areas. It emphasized public-facing surveillance, enabling residents to view local camera feeds via set-top boxes or mobile apps and report anomalies to authorities. While Sharp Eyes expanded the state's visibility into neighborhoods, its design centers on integrating community participation-not on supervising government personnel. Thus, it is analytically distinct from workplace surveillance systems.

In contrast to citizen-facing systems like Sharp Eyes, workplace surveillance systems target state agents rather than the general public. These systems have often emerged through local implementations of smart-city grid-management reforms aimed at digitizing frontline governance. Moreover, some cities were selected as part of a broader national Smart City pilot program, which supported a wide range of digital governance initiatives—including infrastructure integration, data platforms, and task automation—that often embedded monitoring tools for grid workers. The resulting surveillance architecture reflects a logic of internal discipline: curbing shirking, enforcing compliance, and enhancing bureaucratic responsiveness. This agent-focused logic forms the core of our theoretical framework on surveillance as a tool for intra-bureaucratic control.

B.2 Why Workplace Surveillance is Not Part of Sharp Eyes

Workplace surveillance contracts are analytically distinct from the SkyNet (天网工程) or Sharp Eyes (雪亮工程) initiative. Skynet and Sharp Eyes were designed to expand state visibility in community spaces by encouraging residents to participate in local monitoring. While grid workers may use Sharp Eyes platforms to observe residents, these systems are not configured to monitor the workers themselves. In contrast, the workplace surveillance contracts in our analysis involve software designed explicitly to track street-level agents, using GPS logs, real-time photo uploads, digital attendance records, and task checklists. These two types of systems may coexist in the same locality, but they operate on fundamentally different logics. Sharp Eyes is citizen-facing, aimed at enhancing spatial surveillance and participatory governance; workplace surveillance systems are agent-facing, focused on internal accountability and bureaucratic discipline.

C Using Smart City Pilots as Instrument for Workplace Surveillance

first-stage relevance More broadly, China's Smart City initiative sought to modernize urban governance through the application of digital technologies, with stated goals including improving administrative efficiency, enhancing public security, and enabling data-driven service provision. Local governments were encouraged to develop interconnected platforms combining urban infrastructure, public service delivery, and social governance. These platforms often relied on extensive sensor networks, digital records of personnel movements, and AI-driven analytics—technologies that naturally lent themselves to surveillance applications. As such, in addition to transportation and environmental monitoring, many cities used Smart City projects to build out real-time personnel tracking systems and workplace monitoring tools, particularly for street-level bureaucrats such as grid workers.

The Smart City Pilot Implementation Guidelines, jointly issued by the Ministry of Housing and Urban-Rural Development (MOHURD) and the Ministry of Industry and Information Technology (MIIT), list a "social-governance module" as one of the core components that participating cities were required to build. This module aimed to integrate frontline personnel across multiple departments—including grid workers, community police, urban management personnel, private security staff, and volunteer patrol teams—into a unified digital platform to enable real-time coordination and supervision. These features help establish the first-stage relevance of the Smart City program: it increased the likelihood that participating municipalities would adopt workplace surveillance systems focused on grid management.

Exogeneity A key identifying assumption is that Smart City pilot selection was independent to short-term protest dynamics. According to official statements and contemporaneous policy documents, the selection of pilot cities was based primarily on administrative capacity, pre-existing digital infrastructure, geographic balance, and willingness to experiment with digital governance innovations. There is no evidence that cities with unusually high or rising protest levels were systematically favored or excluded from the pilot waves. In fact, over two-thirds of prefecture-level cities were never selected throughout the three national waves. While the first batch of smart city pilots was heavily concentrated in eastern coastal regions, subsequent rounds were more evenly distributed across eastern, central, and western China and within each province (Huang et al., 2021). This supports the assumption that Smart City entry was plausibly exogenous to local contentious dynamics.

Exclusion restriction The exclusion restriction requires that Smart City pilot status affects protest only through its impact on surveillance expansion—particularly the rollout of digital monitoring systems targeting grassroots governance staff. This assumption is credible for three reasons. First, the Smart City initiative was framed as a broad governance modernization effort, with emphasis on infrastructure integration, information sharing, and predictive policy capacity, rather than on direct social control or repression. By contrast, the other two surveillance projects—Tianwang and Xueliang—are more citizen-facing. Second, while the program included many components, most had limited direct connection to protest dynamics (e.g., smart transportation, e-government platforms). By contrast, the core security-related component—construction of a "social governance module"-focused on digitizing the daily activities of local-level actors such as grid workers, rather than deterring protest directly. Finally, procurement contracts and platform designs linked to Smart City often specify performance-tracking features for routine administrative workers rather than tactical protest response. Together, these features suggest that the main channel through which Smart City affects protest is via the surveillance of frontline agents, not other policy channels.

Facial Recognition Systems and Exclusion Restriction In our Table 4, we reported two sets of IV estimates, with or without facial recognition contracts. It is because a potential threat to the exclusion restriction is the adoption of facial recognition systems. Smart City pilots were not limited to workplace surveillance: in some cities, they also accelerated the deployment of facial recognition technologies, which could plausibly affect protest behavior through mechanisms unrelated to routine personnel monitoring (e.g., by deterring public anonymity in demonstrations). Including facial recognition as a control variable, however, introduces a post-treatment covariate that may itself be affected by the instrument. Controlling for it directly would therefore require an additional instrument to separately identify its effect. Since there is no perfect solution, we estimated both versions. The results remain consistent, suggesting that our main findings are not driven by this potential omitted variable.

D Constructing conducive weather measures

Our design to construct conducive weather measures follows Beraja et al. (2023). The difference is that Beraja et al. (2023) examined purchases of facial recognition technology on reducing protests, and we focus on software purchases for monitoring grid workers on reducing protests. We also used different protest measures since we believe CASM-China offers a more fine-grained dataset that is specifically designed for protests in China. We first ran a two-stage least square (2SLS) regression to access what weather variables are predictive of protests. Specifically we used a LASSO regression in the first stage to select the most predictive predictors among a collection of variables (ivreghdfe package in Stata):

- The instrument is a collection of weather variables. Our weather data comes from weather stations maintained by the World Meteorological Organization (WMO) and made available through the National Oceanic and Atmospheric Administration (NOAA). The stations collect daily measurements including temperature means, precipitation levels, various weather conditions (fog, rain, hail, thunder), maximum wind speeds, and visibility. Each prefecture-level city is matched with its nearest weather station. The final collection of instrumental variables includes 18 weather variables in each city-quarter, and all 153 of their pair-wise interactions.

- The treatment variable is the number of protests in CASM in each prefecture-level city and quarter that we used in the main text.

- The outcome variable is the purchase of surveillance technology to monitor grid workers in each city-quarter that we used in the main text.

Table F.1: Summary Statistics of Key Variables

Variable	**Mean**	**SD**	**Skewness**
No. Protests in a city-quarter	18 (22)	22	3
Facial Recognition Contracts	1.79 (6.04)	6.04	8.95
Workplace Surveillance Contracts	0.0636 (0.3969)	0.3969	11.5919
count_times_number	3.1291 (32.0029)	32.0029	20.7234
count_times_number_sqrt	0.2805 (1.7467)	1.7467	8.9293

This approach enabled us to select 9 weather variables most predictive of protests from all available weather measures. We then estimated a city-quarter fixed effect regression with:

- The outcome variable is the number of protests in each city-quarter
- The independent variables are:
  - the 9 weather variables
  - the number of protests occurring in the other cities in the same province.

This allowed us to predict protest numbers based on exogenous variations in both weather conditions and protests in neighboring areas. The predicted values are defined as "conducive weather" in Beraja et al. (2023). This is what we used in Table 3 in the main text.

E Data on protest features

Targets are coded into two categories (with one category has two subcategories). A sample of 10,000 events was randomly selected from the entire CASM-China dataset, and research assistants were instructed to code the tactics, targets, and police presence of each event. We used Support Vector Machine algorithm to train machine learning classifiers on these 10,000 events and applied the trained model to generate the target variable for the rest events. The machine predictions are highly accurate. The area under the ROC curve (i.e., the AUC score) is higher than 0.92 for most outcome categories, suggesting an very good performance.

- the Chinese party-state, such as the party or government, their offices, properties, and officials.
- non-government actors, including public or private institutions or firms; we further divide them by:

F Additional results

Panel A: Workplace Surveillance Contracts

	No contract	Has contract	**p-value**
No. Protests	64 (70)	132 (139)	<0.001
Log(Population)	0.79 (0.94)	1.10 (0.82)	0.002
Log(GDP)	4.27 (1.05)	4.95 (1.18)	<0.001
Log(Tax Revenue)	1.53 (1.13)	2.30 (1.53)	<0.001

Panel B: Facial Recognition Contracts

Characteristic	No contract	Has contract	**p-value**
No. Protests	64 (70)	132 (139)	<0.001
Log(Population)	0.79 (0.94)	1.10 (0.82)	0.002
Log(GDP)	4.27 (1.05)	4.95 (1.18)	<0.001
Log(Tax Revenue)	1.53 (1.13)	2.30 (1.53)	<0.001

[1] Mean (SD)

[2] Wilcoxon rank sum test

Table F.2: Comparison of City-Quarters With and Without Surveillance Contracts. For each variable, we report the mean and standard deviation (in parentheses), and p-values from Wilcoxon rank-sum tests.

	(1) OLS-Fixed effects	(2) Arellano-Bond
Coefficients		
No. Protests	0.366^{***}	-0.127
	(0.032)	(0.066)
Facial Recognition Contracts	-0.395^{***}	-0.613^{***}
	(0.063)	(0.146)
Workplace Surveillance Contracts	-1.482^{*}	-2.759^{*}
	(0.679)	(1.162)
Diagnostics		
Observations	5,312	4,648
Instruments (lags)	–	$t-2$ to $t-15$
# of instruments	–	208
Sargan Test: dfχ^2 (p)	–	228 (0.17)
Sargan Test: p-value		0.168
Wald Test Coefficients: chisq		21.478
Wald Test Coefficients: df		3
Wald Test Coefficients: p-value		0.000
Wald Test Time Dummies: chisq		241.860
Wald Test Time Dummies: df		14
Wald Test Time Dummies: p-value		0.000
AR(1) z (p)	–	-4.11 (0.000)
AR(2) z (p)	–	-1.35 (0.18)

Notes: Column 1 is a two-way fixed-effects model with a lagged dependent variable, susceptible to Nickell bias. Column 2 reports two-step difference-GMM (Arellano–Bond) estimates using lags 2–15 of each endogenous variable as instruments. Robust Windmeijer-corrected standard errors are shown; significance: $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$.

Table F.3: Lagged dependent variable linear regression model (Column 1), and Arellano-Bond estimator (two-step GMM; Column 2). Both models include two-way fixed effects. The dependent variables are shifted forward by one period.

Panel A: City-Quarter Fixed Effects		
Dependent Variables:	Facial Recognition Contracts	Workplace Surveillance Contracts
Lagged dependent variable	-0.079 (0.045)	-0.002 (0.001)
Fit statistics		
Observations	4,980	4,980

Panel B: City-Year Fixed Effects		
Dependent Variables:	Facial Recognition Contracts	Workplace Surveillance Contracts
Lagged dependent variable	-0.079 (0.045)	-0.002 (0.001)
Observations	4,980	4,980

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table F.4: Effects of lagged protests on two types of surveillance, using two-way fixed effects regression (OLS). The outcome is lagged for one quarter (Panel A) or one year (Panel B). Standard errors are clustered in two way.

	OLS	Poi	NB(unconditional)	NB(conditional)
Facial Recognition Contracts	-0.73^{***}	-0.00^{***}	-0.00^{**}	-0.00^{**}
	(0.07)	(0.00)	(0.00)	(0.00)
Contracts unrelated to workplace surveillance	0.00	-0.00	-0.00^{*}	-0.00^{*}
	(0.01)	(0.00)	(0.00)	(0.00)
Population (M, log)	-9.90	-0.37	-0.19	-0.19
	(12.49)	(0.21)	(0.32)	(0.26)
GDP (B¥, log)	-9.47	-0.09	-0.14	-0.14
	(9.35)	(0.12)	(0.15)	(0.13)
Tax revenue (B¥, log)	-2.02	0.04	0.06	0.06
	(4.56)	(0.05)	(0.08)	(0.05)
Num. obs.	1325	1325	1325	1325

The outcome variable is the number of protests from CASM-China for each city-year. The main variable of interest in this placebo test is the number of contracts that are mentions 'grid management' in Chinese, but are not directly relevant to monitoring grid workers, due to word polysemy. It is not expected to impact protests. Column 1 presents results using OLS; column 2 uses quasi-Poisson regression, which accounts for overdispersion in event counts; column 3 employs unconditional negative binomial regression; and column 4 applies conditional negative binomial regression. Standard errors are clustered at the city level to account for potential heteroscedasticity and correlation within cities over time.

Table F.5: Effect of purchase unrelated to monitoring grid workers on protests, using two-way (city-year) fixed effects regression

	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts	$-1.757^{\cdot}$	-0.036^{*}	-0.036^{*}	-0.036^{*}
	(0.962)	(0.015)	(0.015)	(0.015)
Facial Recognition Contracts	-1.628^{*}	-0.046^{**}	-0.046^{**}	-0.046^{**}
	(0.657)	(0.017)	(0.017)	(0.017)
Population (M, log)	-1.999	-0.869	-0.867	-0.867
	(2.235)	(0.627)	(0.638)	(0.627)
GDP (B¥, log)	-0.576	-0.550	-0.551	-0.551
	(1.279)	(0.446)	(0.454)	(0.446)
Tax revenue (B¥, log)	-1.771^{*}	-0.296^{*}	-0.296^{*}	-0.296^{*}
	(0.857)	(0.136)	(0.138)	(0.136)
Num. obs.	1224	1224	1224	1224
Num. groups: city_code	306	306	306	306
Num. groups: year	4	4	4	4

The outcome variable is the number of protests from Chen (2019) for each city-year. The main variable of interest is the interaction between the number of workplace surveillance contracts designed to monitor grid workers and the number of facial recognition contracts designed to monitor citizens. The product of these two variables is skewed, so we took the square root to reduce skewness. Column 1 uses OLS; column 2 uses quasi-Poisson regression, which accounts for overdispersion in event counts; column 3 uses unconditional negative binomial regression; and column 4 uses conditional negative binomial regression. Standard errors are clustered at the city level.

Table F.6: Interaction effect of surveillance on protests from from Chen (2019), using two-way (city-year) fixed effects regression.

	Panel A: City-Quarter Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts (sqrt)	-2.762^{**}	-0.080^{***}	-0.091^{**}	-0.091^{***}
	(0.894)	(0.023)	(0.032)	(0.026)
Facial Recognition Contracts (sqrt)	-0.909^{*}	$-0.021^{\cdot}$	-0.017	-0.017
	(0.446)	(0.012)	(0.013)	(0.011)
Num. obs.	5280	5280	5280	5280

	Panel B: City-Year Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts	$-3.685^{\cdot}$	-0.024^{***}	-0.017^{*}	-0.017^{**}
	(2.067)	(0.004)	(0.008)	(0.006)
Facial Recognition Contracts	-0.359^{*}	-0.002^{**}	-0.002^{*}	-0.002^{*}
	(0.145)	(0.001)	(0.001)	(0.001)
Num. obs.	1321	1321	1321	1321

The outcome variable is the number of protests from CASM-China for each city-quarter, restricted to the events that were identified from both text and images. The main variable of interest is the number of workplace surveillance contracts designed to monitor grid workers. Standard errors are clustered at the city and quarter level.

Table F.7: Effects of surveillance on protests using two-way fixed effects regression, restricted to the events that were identified from both text and images.

	Panel A: City-Quarter Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts (sqrt)	2.320^{***}	-0.048	-0.081	-0.081
	(0.671)	(0.037)	(0.057)	(0.047)
Facial Recognition Contracts (sqrt)	-0.496^{***}	-0.017^{*}	-0.016	-0.016
	(0.138)	(0.008)	(0.013)	(0.010)
Facial Recognition (sqrt) * Workplace Surveillance (sqrt)	-1.969^{***}	-0.009	-0.003	-0.003
	(0.187)	(0.008)	(0.012)	(0.012)
Num. obs.	5280	5280	5280	5280

	Panel B: City-Year Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts (sqrt)	10.704^{*}	0.016	0.030	0.030
	(4.206)	(0.043)	(0.043)	(0.037)
Facial Recognition Contracts (sqrt)	$1.435^{\cdot}$	0.006	0.003	0.003
	(0.801)	(0.012)	(0.013)	(0.009)
Facial Recognition (sqrt) * Workplace Surveillance (sqrt)	-4.197^{***}	-0.016^{*}	-0.017^{*}	-0.017^{**}
	(0.996)	(0.006)	(0.007)	(0.006)
Num. obs.	1321	1321	1321	1321

The outcome variable is the number of protests from CASM-China for each city-quarter, restricted to the events that were identified from both text and images. The main variable of interest is the interaction between the number of workplace surveillance contracts designed to monitor grid workers and the number of facial recognition contracts designed to monitor citizens. Standard errors are clustered at the city and quarter level.

Table F.8: Effects of surveillance on protests using two-way fixed effects regression, interaction effects, restricted to the events that were identified from both text and images.

	Panel A: City-Quarter Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts	-2.124^{*}	-0.023^{**}	-0.027^{*}	-0.027^{**}
	(0.843)	(0.007)	(0.012)	(0.010)
Facial Recognition Contracts	-0.518^{***}	-0.004^{**}	-0.004^{**}	-0.004^{**}
	(0.143)	(0.002)	(0.002)	(0.002)
Num. obs.	5296	5296	5296	5296

	Panel B: City-Year Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts	$-8.947^{\cdot}$	-0.021^{***}	-0.014^{*}	-0.014^{*}
	(5.382)	(0.005)	(0.007)	(0.006)
Facial Recognition Contracts	-0.860^{**}	-0.002^{**}	$-0.002^{\cdot}$	-0.002^{*}
	(0.295)	(0.001)	(0.001)	(0.001)
Num. obs.	1325	1325	1325	1325

The outcome variable is the number of protests from CASM-China for each city-quarter, with statistical correction after King and Hopkins (2010). The main variable of interest is the number of workplace surveillance contracts designed to monitor grid workers. Standard errors are clustered at the city and quarter level.

Table F.9: Effects of surveillance on protests using two-way fixed effects regression, with statistical correction after King and Hopkins (2010).

	Panel A: City-Quarter Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Facial Recognition (sqrt) * Workplace Surveillance (sqrt)	-3.163^{***}	-0.011	-0.005	-0.005
	(0.249)	(0.007)	(0.010)	(0.010)
Num. obs.	5296	5296	5296	5296

	Panel B: City-Year Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Facial Recognition (sqrt) * Workplace Surveillance (sqrt)	-9.775^{***}	-0.013^{**}	-0.013^{*}	-0.013^{*}
	(1.961)	(0.005)	(0.005)	(0.006)
Num. obs.	1325	1325	1325	1325

The outcome variable is the number of protests from CASM-China for each city-quarter, with statistical correction after King and Hopkins (2010). The main variable of interest is the interaction between the number of workplace surveillance contracts designed to monitor grid workers and the number of facial recognition contracts designed to monitor citizens. Standard errors are clustered at the city and quarter level.

Table F.10: Effects of surveillance on protests using two-way fixed effects regression, interaction effects, with statistical correction after King and Hopkins (2010). **Owing to space constraints, we only report the interaction term here; the corresponding main effects and city-year control variables were included in the regression and are available upon request.**

	Model 1
Facial Recognition Contracts	0.036
	(0.026)
Facial Recognition Contracts	0.036
	(0.026)
Workplace Surveillance Contracts	−0.032
	(0.320)
R^2	0.639
Adj. R^2	0.613
Num. obs.	5312

Table F.11: Effect of surveillance on difference between protests against government or not, using two-way (city-quarter) fixed effects regression

	OLS	Poisson	NB(unconditional)	NB(conditional)
	Panel A: City-Quarter Fixed Effects			
Workplace Surveillance Contracts	−3.100	−0.002	−0.004	−0.004
	(3.215)	(0.019)	(0.020)	(0.027)
Facial Recognition Contracts	-0.638^{***}	-0.004^{***}	-0.006^{***}	-0.006^{**}
	(0.106)	(0.001)	(0.002)	(0.002)
Num. obs.	5312	5296	5296	5312

	Panel B: City-Year Fixed Effects			
	OLS	Poisson	NB(unconditional)	NB(conditional)
Workplace Surveillance Contracts	1.773	−0.005	0.004	0.004
	(3.688)	(0.005)	(0.011)	(0.011)
Facial Recognition Contracts	0.360	−0.000	−0.000	−0.000
	(0.353)	(0.001)	(0.001)	(0.001)
Num. obs.	1324	1324	1324	1324

The outcome variable is the number of grievances posts from CASM-China for each city-quarter. The main variable of interest is the number of workplace surveillance contracts designed to monitor grid workers. Standard errors are clustered at the city and quarter level.

Table F.12: Effects of surveillance on online grievances using two-way fixed effects regression.