

Concealing Corruption: How Chinese Officials Distort Upward Reporting of Online Grievances

JENNIFER PAN *Stanford University*

KAIPING CHEN *Stanford University*

A prerequisite for the durability of authoritarian regimes as well as their effective governance is the regime's ability to gather reliable information about the actions of lower-tier officials. Allowing public participation in the form of online complaints is one approach authoritarian regimes have taken to improve monitoring of lower-tier officials. In this paper, we gain rare access to internal communications between a monitoring agency and upper-level officials in China. We show that citizen grievances posted publicly online that contain complaints of corruption are systematically concealed from upper-level authorities when they implicate lower-tier officials or associates connected to lower-tier officials through patronage ties. Information manipulation occurs primarily through omission of wrongdoing rather than censorship or falsification, suggesting that even in the digital age, in a highly determined and capable regime where reports of corruption are actively and publicly voiced, monitoring the behavior of regime agents remains a challenge.

INTRODUCTION

A prerequisite for the durability of authoritarian regimes as well as their effective governance is the ability to gather information about the actions of government officials (Wintrobe 1998). However, most mechanisms for monitoring government officials—secret police, oversight committees—fail to provide reliable information because officials have incentives to collude with monitors to suppress information that would jeopardize their access to positions of power and associated rents. To sidestep unreliable monitoring agencies, authoritarian regimes have borrowed the strategy of bottom-up citizen participation used in democratic contexts to increase accountability (O'Donnell 1999; Olken 2007). By adopting channels where the public can lodge complaints and provide oversight over local officials, authoritarian regimes can in theory gather more reliable information because citizens have little incentive to collude with corrupt officials (Cai 2013; Dimitrov 2014a, 2014b). In the digital age, these citizen complaints are often posted online where they remain uncensored and publicly accessible (King, Pan, and Roberts 2013). Some argue that autocrats can effectively glean information on officials' misconduct from posts shared on social media and punish officials based on this information (Dimitrov 2014a; Nathan 2003; Qin, Stromberg, and

Wu 2017a), and others suggests that these online channels may increase authoritarian accountability (Gunitsky 2015; Yong 2005; Noesselt 2014; Qiang 2011). Authoritarian regimes as diverse as China, Egypt, Pakistan, Russia, Saudi Arabia, Singapore, and Vietnam have adopted online platforms where citizens can express their grievances.¹

In this paper, we gain rare access to internal communications between a monitoring agency and upper-level authorities in China, and we show how citizen reports of malfeasance implicating lower-level officials are systematically concealed from upper-level authorities. Our evidence comes from an analysis of email archives leaked from the Chinese Communist Party (CCP) Propaganda Department of a prefecture in south central China, which we call J. Prefecture. Although the leaked emails have received press coverage and remain publicly available, the archive is complex and large, and it has not been systematically analyzed. Using large-scale hand-coding and machine learning methods, we identified 643 Online Sentiment Monitoring Reports (輿情报告) produced by the J. Prefecture Propaganda Department, which contain the details of 3,423 online complaints generating public anger and discontent from 2012 to 2014. Most importantly, we can differentiate between the Online Sentiment Monitoring Reports sent to upper-level provincial authorities and those that are kept for internal circulation within the prefecture,

Jennifer Pan, Assistant Professor, Department of Communication, Building 120, Room 110 450 Serra Mall, Stanford University, Stanford, CA 94305-2050 (jenpan.com).

Kaiping Chen, Ph.D. Candidate, Department of Communication, Building 120, Room 110450 Serra Mall, Stanford University, Stanford, CA 94305-2050

Our thanks to Jianghong An, Rita Lu, Yanchen Song, Feiya Suo, Zhiheng Xu for excellent research assistance; to Steven Balla, Peter Lorentzen, Guillermo Rosas, Lily Tsai, the SMAPP Global meeting, and Bay Area China Social Science Workshop participants for their extremely helpful comments and suggestions; and to the Stanford Asia-Pacific Scholars Fund, Stanford China Fund, and Stanford IRiSS Faculty Fellows Program for research support.

Received: December 23, 2016; revised: September 18, 2017; accepted: April 12, 2018.

¹ Citizen complaints in China will be discussed in the next Section. In Egypt, citizens can complain through ministry and other government websites (OECD 2013). Khyber Pakhtunkhwa province in Pakistan created the KP Citizens Portal (<http://smart.pmru.gkp.pk/>) to gather complaints. In Russia, ministries have federal and regional websites where citizens can submit complaints, for example, 36.mvd.ru/ for the police department and 36.mchs.gov.ru/ for the ministry of emergency services in the Voronezh region. The Saudi Arabian website shakwa.net works with officials to gather citizen complaints. In Singapore, citizen complaints are gathered through various online channels (Rodan and Jayasuriya 2007). Various cities in Vietnam have set up online complaint sites (e.g., <http://egov.danang.gov.vn/gop-y>). Note that online complaint forums may also serve other functions such as giving citizens an outlet to vent their anger or providing a way to address grievances.

allowing us to compare online citizen complaints reported upward with complaints identified by the monitoring agency as important but not reported upward.

This data shows that complaints implicating the prefecture government in wrongdoing—including embezzlement, graft, and violence—are less likely to be reported upward to provincial superiors compared to complaints related to governance issues such as education quality and pollution. This data also shows that wrongdoing by subordinate counties with workplace and birthplace ties to prefecture leaders are less likely to be reported to provincial superiors, whereas citizen complaints of wrongdoing implicating counties without political connections to prefecture-level officials are more likely to be reported up to the province. Finally, we find that citizen complaints of government wrongdoing are censored by J. Prefecture when complaints are posted on websites managed by J. Prefecture; however, the large majority of complaints remain publicly available because they are posted to websites where J. Prefecture has no censorship authority.

The information manipulation we identify is distinct from the known problem of falsification of economic and fiscal data in China (Cai 2000; Park and Wang 2001; Tsai 2008; Wallace 2016). While falsification of economic statistics entails manipulation through falsification of performance targets, the information manipulation we identify is rooted in concealment. Some public complaints are reported upward, and these complaints reflect real grievances and public concerns; however, information distortion occurs because what is reported upward simply does not reveal the entire picture, potentially increasing the difficulty of penalizing officials for engaging in this form of information manipulation. Furthermore, while economic falsification is focused on distorting performance-based targets, manipulation of public opinion directly distorts understandings of public satisfaction, contention, and regime support. Our results show that even in a highly determined and capable authoritarian regime where grievances are actively and publicly voiced, lower-level officials continue to distort and manipulate information.

These findings have implications for our understanding of political communication and information conflicts in authoritarian regimes. Authoritarian regimes face two types of conflicts over information. The first conflict occurs between the regime and the public over what information is publicly available. The second conflict occurs among regime elites, specifically, between upper-level authorities who want to monitor lower-level agents and lower-level agents who want to hide malfeasance. Research in political communication has emphasized the first conflict, with an emerging consensus that authoritarian regimes continue to control publicly available information in the age of social media (Kalathil and Boas 2010; MacKinnon 2012; Morozov 2012). Our results reveal a different picture by focusing on the second conflict, joining a new strand of research that shows the constraints of autocrats' ability to gather reliable information about regime agents in the digital age (Pan 2016). Even though authoritarian regimes may learn much by monitoring social media

and digital data as scholars have suggested,² access to large-scale data by no means guarantees autocrats omniscience over regime agents.

These results enrich our understanding of the relationship between nonelectoral forms of citizen participation and accountability in authoritarian contexts. Recent research shows how local governments in authoritarian regimes acknowledge and respond to citizen complaints at relatively high rates (Chen, Pan, and Xu 2016; Distelhorst and Hou 2017; Meng, Pan, and Yang 2017). There is a great deal of optimism surrounding the benefits of citizen participation in China, in nurturing an informed citizenship, in informing the regime of the public's preferences, and even in facilitating government accountability (Fishkin et al. 2010; He and Warren 2011; Ma 2012). However, government officials can only be held accountable if information is made available to those that have power to sanction these officials, and if those with sanctioning power have the incentives and resources to punish offenders. In authoritarian contexts such as China, this sanctioning power resides with upper-level authorities that control the promotion prospects of lower-tier officials, not the public. Therefore, our results, showing how the public's grievances and complaints do not always reliably make their way to upper-level authorities, reveal obstacles to bottom-up accountability in authoritarian regimes even when citizen participation increases. Our results do not contradict findings that online and offline citizen complaints can uncover government corruption and lead to the punishment of government officials; they simply suggest that many citizen reports of corruption may not lead to sanctions, and citizen participation is likely insufficient to root out corruption and to create the conditions for full accountability.³

INFORMATION GATHERING AND INFORMATION MANIPULATION

Authoritarian regimes are often administered through multiple layers of government, and central autocrats delegate responsibilities for governance to multiple subnational levels of government.⁴ In this section, we describe how dual incentives for information gathering and information manipulation exist at each

² See Qin, Stromberg, and Wu (2017a), Qin, Stromberg, and Wu (2017b), and MacKinnon (2012) for discussion of China; Pearce and Kendzior (2012) of Azerbaijan; Gunitsky (2015) of Bahrain; House (2015) for Uzbekistan; and MacKinnon (2012), Morozov (2012), and Gunitsky (2015) of Russia.

³ As of December 2016, none of the prefecture-level officials of J. Prefecture, including those accused of corruption in citizen complaints, have been implicated in Xi Jinping's extensive anticorruption campaign.

⁴ For example, Russia consists of federal subjects, which are divided into administrative districts. Administrative districts consist of cities, towns, and rural settlements. Cities are further subdivided into city districts. China is divided into 31 provincial-level administrative divisions, which are further subdivided into prefectures; prefectures are subdivided into counties (also called districts), and counties are subdivided into townships (also call subdistricts). Similarly, Iran is divided into 31 provinces, which are further subdivided into counties; counties are subdivided into districts, and districts are subdivided into cities and rural districts.

subnational level, and we show how variation in incentives and resources between subnational levels facilitate or hinder surveillance capacities.

Theory of Information Gathering and Manipulation

We assume the interest of political leaders at all subnational levels in an authoritarian regime is to maximize rents, and the primary way to maximize rents is to stay or advance in political office.⁵ The logic we describe can be applied to any superior-subordinate relationship where the superior controls the career prospects of the subordinate, and where the superior is also subordinate to some higher-level authority. For example, prefecture party secretaries in China (the superior) control the career advancement of county party secretaries (the subordinate), but prefecture party secretaries are themselves subordinate to provincial leaders. For our set-up, there are two actors: (1) the superior who we refer to as the upper-level official and (2) the subordinate who we refer to as the lower-level official.⁶

To remain or advance in political office, upper-level officials delegate tasks to lower-level subordinates—for instance, attract investment, control crime, improve infrastructure. However, in the course of carrying out these tasks, lower-level officials engage in corrupt practices to enrich themselves.⁷ Malfeasance by lower-level officials, and the by-products of corruption such as public dissatisfaction and unrest, could lead to scandals that impede or ruin the political careers of upper-level officials. Upper-level officials are interested in avoiding such scandals, and to do so, upper-level officials may want to strengthen monitoring of lower-level officials to determine whether the activities of subordinates will impede their own political careers and to remove from office subordinates whose actions jeopardize their own careers. At the same time, since the rents of lower-tier officials are associated with political office, lower-tier officials have strong incentives to hide actions that

would lead to scandals of their own and their removal from political office. This means that to survive or advance in political office, political leaders at each subnational level of an authoritarian regime have dual incentives to (a) hide their own wrongdoing from superiors who control their career prospects and (b) obtain enough information about the subordinates they control so the activities of subordinates do not jeopardize their own career prospects. How these dual incentives play out in terms of the effectiveness of information gathering vs. information manipulation depends on the interaction between the incentives and resources of lower-level and upper-level officials. Below, we first examine each component, and then examine their interactions.

Variation in Incentives. While these dual incentives exist among all subnational officials, the strength of both the incentive to hide wrongdoing from superiors and the incentive to obtain accurate information about subordinates varies among officials depending on two main factors: (1) the career ambition of the official and (2) the conditions of the region the official governs.

Officials' levels of concern with scandal varies with their career ambitions—an extremely ambitious official actively seeking career advancement is likely more concerned with scandal than an official who is only interested in preserving his or her current political position. As a result, more ambitious officials have greater incentives to hide wrongdoing from superiors and greater incentives to monitor their subordinates than officials who only want to remain in their current position. Lower-level officials who simply want to remain in office may only be interested in hiding wrongdoing that would result in major scandals, while lower-level officials who are actively pursuing advancement may have incentives to suppress all forms of wrongdoing. This logic is similar to the finding of Kung and Chen (2011), that provincial leaders with stronger incentives for career advancement were more likely to falsify agricultural procurement data, which led to higher rates of mortality during China's Great Leap Famine. Along the same lines, upper-level leaders who want to remain in office may be only interested in monitoring to detect major scandals, while upper-level leaders actively pursuing advancement may have incentives to detect and head off all potential scandals.

Second, in terms of regional conditions, officials who govern regions whose social, economic, or political situation is more precarious—for example, a region with a history of protest or corruption—may have stronger incentives to hide wrongdoing from superiors as well as incentives to monitor the activities of their subordinates than officials in regions on more solid footing. In regions with more precarious conditions, scandals are more likely to jeopardize officials' chances of advancement.

Given career ambition and regional conditions, we expect some officials to be more determined to hide their wrongdoing than others, and some officials to be more determined to monitor their subordinates than others. Those more determined to hide wrongdoing

⁵ Research has established links between political office and monetary rewards (Truex 2014; Zhang, Giles, and Rozelle 2012). While it would not be a stretch to imagine that the interest of central-level leaders is also to maximize rents, we focus on subnational levels because the way to maximize rents for central-level leaders differs from subnational leaders. For example, central leaders might maximize rents by pursuing the survival of the regime since they can no longer advance to higher political positions.

⁶ We do not conceptualize monitors or monitoring agencies as independent actors. We assume that monitors are either aligned with the lower-level official they are supposed to monitor or aligned with the upper-level official they are gathering information for. Monitors can also be conceptualized as independent actors—secret police and other repressive apparatus can be powerful and can act independently. However, in our empirical context, monitoring agencies are separate from the repressive apparatus and have limited resources relative to lower-level and upper-level officials.

⁷ Although officials have a choice in whether or not to engage in corruption, we assume that corruption is endemic to officials in authoritarian regimes. Corruption is often defined as the abuse of public office for private gain, but in many autocracies, access to the spoils of political office is implicitly promised to political officeholders as a way of managing intra-elite conflict and coopting elites (Magaloni 2008; Magaloni and Kricheli 2010).

from superiors have greater incentives to collude with monitoring agencies in suppressing information, and those more determined to monitor subordinates are more likely to adopt mechanisms to counteract information manipulation by subordinates.

Variation in Resources. The resources available to officials for hiding wrongdoing and obtaining accurate information also varies. By resources, we mean assets at the disposal of officials that can be used to exert control over information and monitoring agencies. These assets include the authority to hire and promote, social relationship, and fiscal resources. Some officials will be better equipped to control monitoring agencies (and prevent monitors from reporting malfeasance) because institutional structures give officials control over the career prospects of monitors. Other officials can exert control because social relationships such as kinship ties allow officials to wield influence over monitors. Finally, officials with plentiful fiscal resources are more likely to overcome information manipulation by lower levels—for example, they can pay monitors more so monitors are less likely to be bribed or they can pay to have more monitors / monitoring systems.

Interaction between Incentives and Resources in Monitoring. Officials' incentives to monitor interact with the availability of resources for monitoring. For example, an ambitious upper-level official in a region with plentiful resources will likely put stronger monitoring mechanisms in place to overcome information manipulation by lower-level officials than an upper-level official in a region with plentiful resources who simply wants to remain in office. In general, if plentiful resources are available, monitoring may vary as a function of incentives, with stronger monitoring when there are stronger incentives.

What happens if resources are limited or unavailable? If an upper-level official does not have the resources to strengthen monitoring, we might expect upper-level officials to take one of two strategies. First, the upper-level official may choose not to change or bolster information collection procedures, knowing that lower-level officials are likely hiding information but not knowing what is being hidden. If scandal breaks out, the upper-level official can plead innocence, deflect blame onto the lower level for concealing information, and hope that these excuses can overcome the fallout of the scandal.⁸ Second, the upper-level official might seek to collude with lower-level officials in corrupt, rent-seeking activity, and collude with lower-level officials to suppress information.

Interaction between Superiors and Subordinates. The degree to which upper-level officials can monitor and the degree to which lower-level officials can hide wrongdoing results from the interaction between the incentives and resources of upper-level officials and the

incentives and resources of lower-level officials. When upper-level officials have plentiful resources for monitoring, whether we observe lower-level officials concealing wrongdoing depends on the relative incentives and resources between superiors and subordinates. For instance, if lower-level officials have relatively stronger incentives and more resources than their superiors, we are more likely to see information manipulation and concealment of wrongdoing by lower-level officials. If upper-level officials have relatively stronger incentives and more resources than their subordinates, we are less likely to see concealment of wrongdoing by lower-level officials. When upper-level officials have limited resources, we will likely observe lower-level officials concealing wrongdoing. In some of these cases, upper-level officials may suspect information manipulation is taking place but be unable to stop it, and in other cases, upper-level officials may be complicit in the information manipulation. Altogether, this suggests that unless upper-level officials have stronger incentives and greater resources than lower-level officials, information manipulation by lower levels will take place.

Information Gathering and Manipulation in China

In this section, we describe how the Chinese regime aims to use public participation via citizen complaints to monitor lower-tier officials. We then show how the dynamics of obfuscation and monitoring we have described manifest in China and why information manipulation can persist despite publicly voiced grievances.

How China's Propaganda Department Monitors Officials with Public Complaints. China has devoted substantial resources to monitoring the performance of lower-tier officials. The CCP has adopted limited versions of elections and free media, while embracing citizen complaints as an essential part of the CCP's strategy of monitoring lower-tier officials.⁹ There are many channels for individuals to share their grievances and complaints: the Bureau of Letters and Visits (信访局) where citizens can complain in person (Chen 2009; Dimitrov 2014a), telephone hotlines (Economist 2017), government-managed websites where citizens can complain online (Chen, Pan, and Xu 2016; Distelhorst and Hou 2017), as well as web and mobile apps designed for individuals to complain to the government (China Daily 2015). The volume of citizen complaints, especially those submitted online, has skyrocketed (Jiang, Meng, and Zhang 2016).

The CCP propaganda department (宣传部) is the primary agency responsible for monitoring public opinion and sentiment through citizen complaints (Central Propaganda Department 2009; National

⁸ The upper-level official runs the risk that claims of innocence are insufficient if the scandal is severe, which is why, if resources are available, the official would choose to strengthen monitoring and avoid being blindsided.

⁹ The CCP has implemented elections, but only at the village and neighborhood levels (O'Brien and Li 2000). China has legislative institutions, but delegates are not freely selected (Manion 2016; Truex 2016). Chinese media is commercialized, and some forms of investigative journalism are allowed (Lorentzen 2013; Stockmann 2013), but media outlets remain tightly state-controlled (Qin, Strömberg, and Wu 2016).

Grassroots Party Work Key Textbook 2013).¹⁰ Propaganda departments exist at the central, provincial, prefecture, and county levels. Each level is responsible for monitoring public sentiment in its geographic jurisdiction, including public reports of corruption.¹¹ The propaganda department is also responsible for reporting this information upward to the propaganda departments at the next level up in what is called “level-by-level” reporting so that the CCP and Chinese government can use this information to help guide policy and political decisions (Cai 2000; Central Committee of the Communist Party of China 2016; Huang 1995; Oi 1995). Evaluation metrics contained in internal documents such as the *Propaganda and Thought Work Management Evaluation Form* (宣传思想工作指导管理评分表) reveal the importance of monitoring work for the propaganda department. Approximately half of the areas on which CCP propaganda departments are evaluated relate to monitoring public sentiment and reporting this information to upper-level superiors.¹² Specific evaluation criteria include the number of citizen complaints collected and analyzed,¹³ the number of bursts of positive and negative online discussions identified, and whether there is regular reporting of information to upper levels of the government and party. In contrast, it is not within the purview of the propaganda department to validate the veracity of citizen complaints, nor is it within its jurisdiction to discipline individuals or organizations based on public complaints. The task of the propaganda department is to determine what key issues and/or events generate interest and attention among the public, and to report this information upward (Central Propaganda Department 2009).

¹⁰ Government bureaus and offices (e.g., bureau of letters and visits, public security bureau), and China’s Cyberspace Administration (网信办) also deal with citizen complaints. However, instead of monitoring, government bureaus are tasked with collecting, investigating, and ultimately resolving complaints related to their areas of work in their geographic jurisdictions. China’s Cyberspace Administration (CAC), also known at the central level as the Office of the Central Leading Group for Cyberspace Affairs, is tasked with governing the internet. CAC, located at central and provincial levels of government, regulates the activities of internet content providers, from news portals to social media platforms, and penalizes companies and individuals that fail to comply with government regulations (Cyberspace Administration of China 2017).

¹¹ The central role of the propaganda department in monitoring public opinion has been echoed by central as well as provincial leaders. Provincial party secretaries in provinces ranging from Henan to Guangdong have emphasized the central role of the propaganda department in monitoring public opinion (see <http://bit.ly/2eJUPlw>, <http://bit.ly/2wM8W1U>, accessed Sept. 5, 2017). Li Xi, former deputy secretary of Shanghai and current party secretary of Liaoning Province, wrote an editorial in *People’s Daily*, the official newspaper of the Chinese Communist Party, that provincial decisions should take into account public opinion information collected and analyzed by the propaganda department (see <http://bit.ly/2eK1M6f>, accessed Sept. 5, 2017).

¹² The other half of the functions relate to cultural development and ideological thought work.

¹³ Documents in the leaked emails show that the provincial propaganda department explicitly tasks the prefecture propaganda department with monitoring complaints related to their jurisdiction across social media sites and internet platforms, including national-level platforms the prefecture does not manage.

Institutional Incentives for Concealing Corruption. China is a single-party authoritarian regime with five levels of state administration: central, provincial, prefectural, county, and township. We use the term “top leaders” to refer to leaders at *each level* of administration who hold senior CCP positions—specifically that of party secretary, vice party secretary, and politburo standing committee member.¹⁴ Individuals holding these top party positions also typically occupy key government positions, so that crucial positions of power are all under CCP control.

At each level of administration, top leaders determine the advancement of top leaders at the next level down, in what is called one-level-down management (下管一级) (O’Brien and Li 1999), by evaluating them against performance targets (Edin 2003; Whiting 2004). Points are assigned to various dimensions of performance, but failure to achieve “veto” targets, such as failure to ensure social stability—by preventing collective action or rooting out corruption—jeopardizes advancement prospects even if the performance is strong in other areas.¹⁵ This means that top leaders at each subnational level in China have incentives to hide their own wrongdoing from superiors one level up who control their career prospects. At the same time, top leaders at each level have incentives to obtain enough information about their subordinates one level down so the behavior or malfeasance of these subordinates does not jeopardize their own career prospects.

Opportunities for Information Manipulation. Top leaders have opportunities to conceal malfeasance from superiors one level up because they can exert control over the propaganda department that monitors them and because of China’s level-by-level reporting structure. Top leaders in China can exert control through institutional incentives as well as social ties to motivate collusion with the propaganda department. In terms of institutional incentives, the promotion of heads of various bureaucracies at each level (e.g., the prefecture propaganda department) is determined by the top leaders of that same level (e.g., prefecture party secretary) rather than the bureau one level up (e.g., provincial propaganda department) (Edin 2003; Kou

¹⁴ Note that when we refer to “top leaders,” we are referring to individuals, at each level of administration, who hold these key party positions. For the case of China, the upper-level and lower-level officials we described in the theory section are all top leaders.

¹⁵ Based on internal documents of criteria for political advancement from several central and eastern provinces, we find that, although the overall weight given to economic development (e.g., 100 out of 350 points) and social stability (75 out of 350 points) are highest, these broad categories are broken down into smaller subactivities—for example, social stability includes managing petitions (5 points) and strengthening public security (18 points) while economic development includes GDP growth (10 points) and service sector growth (10 points)—the subcategory with the highest point value is “construction of clean government” (党风廉政建设) worth 30 points. This subcategory refers to party discipline and the need for clean, non-corrupt government, which has been documented to be a veto target (see <http://bit.ly/2xMvIXs> and <http://bit.ly/2j1IH19> (accessed Sept. 3, 2017).

and Tsai 2014; Whiting 2004; Zhan and Yu 2011).¹⁶ As a result, although Chinese bureaucracies are part of a matrix organization (条块关系) where bureaucracies report to top leaders at the same level and also report to their functional equivalents one level up, satisfying the interests of top leaders at the same level takes priority over the interests of functional superiors at the next level up (Li 2012; Lieberthal and Lampton 1992).¹⁷ In addition, social relations between top leaders and propaganda officials at the same level of administration can motivate monitors to collude with top leaders. For example, the prefectural head of propaganda and the prefecture politburo members are all elites within the prefecture and may know each other through work, school, or even kinship ties.

Information manipulation occurs when local propaganda departments collude with top leaders because the local propaganda department can withhold information from superiors. For example, if a prefecture propaganda department identifies online complaints of corruption pertaining to prefecture officials, in order for provincial authorities to obtain this information, the information must be transmitted by the prefecture propaganda department to the provincial propaganda department, and from the provincial propaganda department to provincial top leaders. In other words, since the propaganda department system for information gathering system relies on level-by-level reporting, there are opportunities for each subnational level to conceal information.

Propaganda departments can be penalized for failing in their task of monitoring by CCP discipline inspection commissions, which are the party organization tasked with enforcing internal party rules. We can think of these discipline inspection commissions as a third party operating outside the normal procedures for advancement. However, collusion between top leaders and monitors may persist because the risk of penalty is relatively low when information manipulation of citizen complaints occurs through *concealment* rather than *falsification*. If propaganda departments regularly report public opinion upward but do not reveal the full picture, propaganda departments are fulfilling the letter, but not spirit, of monitoring regulations, and punishment may be less likely relative to other forms of information manipulation such as falsification of economic statistics.

Opportunities to Strengthen Monitoring. Since the shortcomings of level-by-level reporting are well

known, China's central authorities have worked to increase the number of channels for citizen complaints and by directly gathering information in specific situations—for example, making surprise visits to lower-level governments and creating mobile apps for citizens to report corruption to the central level (Gao 2016; Xinhua 2012). National-level social media platforms are a particularly fruitful channel for obtaining information because subnational levels of government do not have the ability to censor these platforms directly. As long as officials have access to human and/or computational resources, large-scale social media text data can be mined for information on the activities of lower-level officials.

Indeed, in recent years, China's central authorities have dramatically increased monitoring of the internet. Qin, Stromberg, and Wu (2017a) find that the posts on Sina Weibo mentioning the names of officials and corruption predict charges of corruption by the central regime. However, this does not mean the central regime, which has the most human and computational resources of any level of government, is able to uncover all complaints of lower-level corruption posted to social media. Indeed, China's central government is still in the process of building a national-level surveillance system that can extract information from various social media platforms, and this system is only intended to reach down to the prefecture level (leaving out county, township, and lower-level government offices and party organizations).¹⁸

The capacity for monitoring social media using computational methods is weaker among provincial and lower-level governments (Mai and Liber 2015). While subnational governments can devote human resources to monitoring social media,¹⁹ whether they do so likely depends on their incentives and availability of resources. At present, subnational governments remain primarily reliant on internal, level-by-level reporting to identify lower-level corruption (Hsu, Zhao, and Wu 2011). This suggests that while the Chinese regime has invested in public channels for expressing grievances to improve monitoring of lower-tier officials, information problems persist, especially between subnational levels of government.

In sum, China's central regime has embraced public complaints as a core component of information gathering. At subnational levels, top leaders have opportunities to manipulate information through control of local propaganda departments and they also have opportunities to strengthen monitoring. In the following section, we describe our unique dataset of communications between the J. Prefecture propaganda department and its provincial superiors, which allows us to test some aspects of our theory. An observable implication of the theory is that unless upper-level officials have stronger incentives and greater resources

¹⁶ This differs from the promotion of top leaders at each level, which is controlled by top leaders one level up.

¹⁷ The only situation in which bureaucratic relationships to top leaders is subservient to the relationship with functional superiors one level up is when the bureaucracy is under vertical management (垂直管理). This arrangement applies to a handful of bureaucracies that deal with cross-regional issues, including the railway administration, civil aviation administration, customs administration, state security, land and resources administration, taxation, finance, and bureau of statistics. The propaganda department is not a bureaucracy under vertical management. The career advancement prospects of propaganda department officials are determined by top leaders at the same level, not by propaganda officials one level up.

¹⁸ The central government aims to complete this system by 2020; information based on interviews with individuals working with the government on computational surveillance.

¹⁹ Qin, Stromberg, and Wu (2017b) estimate that it would take 2,080 person-hours to identify the strikes over 3 years from the related posts on Sina Weibo.

than lower-level officials, information manipulation by lower levels will take place. As we will see, in the case of J. Prefecture, provincial and prefectural levels have similar levels of resources, but prefecture officials likely have stronger incentives for information manipulation than provincial officials have for information gathering. We would thus expect to see information manipulation by J. Prefecture, and our data allows us to ascertain whether this is the case.

LEAKED PROPAGANDA DEPARTMENT COMMUNICATIONS

A problem with existing studies of the effectiveness of gathering information through citizen complaints is that they rely on data found at specific levels of government, for example, central government archives or city government reports. As a result, we never see what information is concealed by lower levels of government and not transmitted upward.

This paper changes the situation. In December 2014, an email archive from the Propaganda Department (宣传部) of J. Prefecture was publicly leaked (Henochoicz 2014; Sonnad 2014). The email archive contains Online Sentiment Monitoring Reports produced by the J. Prefecture Propaganda Department, as well as many other communications to and from the department. The leak was reported and the archive of emails remains publicly available (Henochoicz 2014; Sonnad 2014). The email archive is large and complicated by multiple email storage formats, diverse document types, numerous attachments, and many links to outside information. Because of the complexity of this data, no systematic analysis has been conducted. To systematize this rich data source, we developed and applied a variety of methods and procedures, from large scale hand coding, to specially tuned and adapted methods of named entity recognition, and methods of automated text analysis and extraction. Because of the considerable effort entailed, the extracted data, as well as replication materials for the analyses in this paper are available, see Pan and Chen (2018).

We identified 2,768 emails sent to and from the J. Prefecture Propaganda Department between May 14, 2012 and December 6, 2014. Figure 1 shows the number of emails in the archive by month between 2012 and 2014. There are substantially more emails in the archive in 2014 than in the previous two years.²⁰ We do not know if this is because there are missing emails from the previous years, or if it is because the email account was used more actively in 2014. We focus our analysis on the time period when the data is most plentiful, between January 1, 2014 and November 30, 2014.²¹

FIGURE 1. Number of emails in archive by month.

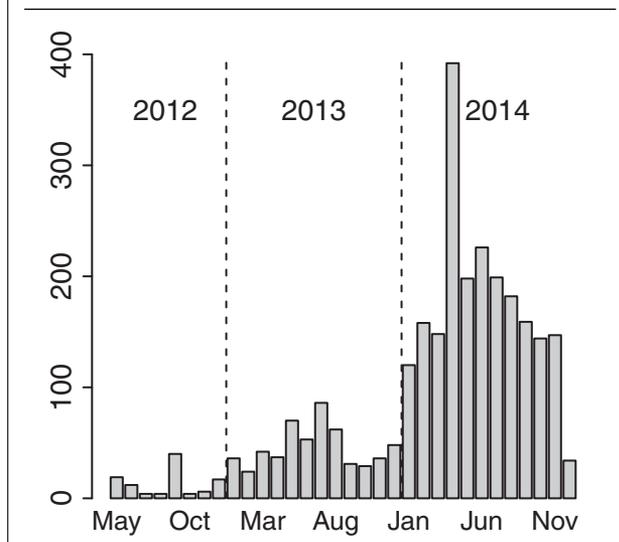
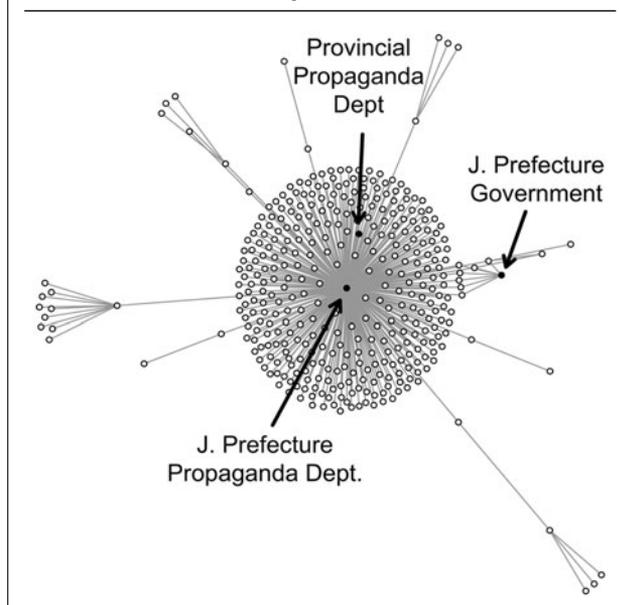


FIGURE 2. Network structure of emails. Circles are email correspondents, lines indicate email correspondence.



The overall structure of communications captured by these emails is shown with the network graph in Figure 2. Each circle is a specific email account and each line denotes where one or more emails was sent from and to. The most central node is the J. Prefecture Propaganda Department, and the accounts it corresponds most frequently with include the J. Prefecture government (市政府) and the provincial propaganda department (省宣传部).

²⁰ The large number of emails in April 2014 is largely due to duplication, which may be related to how the data was obtained. For this reason, we remove exact duplicates in our analysis.

²¹ We do not include December since the archive was leaked in early December. When the entire dataset is used, none of our substantive results change.

TABLE 1. Three Types of J. Prefecture Online Sentiment Monitoring Reports

	2012 to 2014	Jan. to Nov. 2014
Daily reports	514	249
Weekly reports	114	55
Bi-monthly reports	15	15
Total	643	319

Among all emails, 643 contain Online Sentiment Monitoring Reports produced by J. Prefecture.²² Between January and November 2014, 319 emails contain these Online Sentiment Monitoring Reports. There are three types of Online Sentiment Monitoring Reports: Daily, Weekly, and Bi-monthly, with Bi-monthly report only appearing in 2014 (see Table 1). All reports contain three main sections: (a) review of current public sentiment, (b) trending news items, and (c) national online trends. Most relevant to our analysis is the first section, which details issues that have generated local online discontent. Section (b) on trending news items includes keywords of trending topics, often overlapping with the more detailed analysis of section (a). Section (c) is a brief summary of national-level top news items, primarily from traditional news sources.

Section (a) is further divided into four subsections: (a1) the time period covered by the monitoring report (daily reports can cover a range of 1–4 days), (a2) the number of online posts analyzed for the report (e.g., 185 posts), (a3) an evaluation of the overall tendency of public sentiment: negative (负面), neutral (中性), positive (正面), and (a4) a summary of key issues including descriptions of the topics that have generated negative sentiment. The overall assessment of sentiment is almost always positive.²³ Thus, we focus on the actual posts classified as negative from section (a4) of the report. Typically, the issue is described in a sentence or two, and a link to the original post and discussion is sometimes included.

In total, we extract 3,423 negative sentiment issues from the Online Sentiment Monitoring Reports; after removing exact duplicates, 2,879 remain.²⁴ Between January and November 2014, there are 1,925 negative sentiment issues, and after removing exact duplicates, 1,412 issues remain. We examine the 1,412 nonduplicate issues more closely and identify 1,038 unique complaints.²⁵ In addition to using the Propaganda Department's terminology of "negative sentiment issues," we

also refer to these posts as citizen complaints.

In the process of extracting data from the email archives, we noticed that a subset of Online Sentiment Monitoring Reports was addressed to higher-level officials, while others were kept for internal circulation within the prefecture propaganda department. We realized that we had data on all of the negative sentiment issues the J. Prefecture Propaganda Department deemed to be important, and the subset which they reported to upper-level officials.

Reports sent to provincial superiors meet three conditions, which are always met simultaneously. First, using email metadata, we find that the report is emailed to the provincial propaganda department email address. Second, the subject of the email requests the attention of provincial officials, e.g., "Provincial Leaders Please Inspect" (请省领导查收). Third, the report itself contains the line "cc: City Party Secretary, City Vice Party Secretary, City People's Congress Chairman, City People's Political Consultative Conference Chairman, City Politburo Standing Committee, City Vice Mayor, County Propaganda Departments, other relevant city departments." Among the 1,412 unique negative sentiment issues identified by the J. Prefecture Propaganda Department from January to November 2014, 590 are sent to upper-level leaders.

Observable Implications of Theory in J. Prefecture

This leaked data comes from one prefecture in China. Is this prefecture more or less likely to manipulate information than other prefectures? Is the province where J. Prefecture is located more or less likely to uncover manipulation by the prefecture? As discussed in the previous Section, information manipulation results from the strategic interaction between superior and subordinate levels of government, which in turn depends on the incentives and resources available to each level of government. Here, we discuss the observable implications of the theory for J. Prefecture by describing the incentives and resources of J. Prefecture leaders and their provincial superiors. In doing so, we set the scope conditions of our empirical analysis.

Officials in J. Prefecture may have stronger incentives for career advancement than their provincial counterparts, which would result in stronger incentives for information manipulation at the prefecture level than incentives for information gathering at the provincial level. There are two aspects that influence incentives—regional conditions and career ambition. In terms of regional conditions, neither J. Prefecture nor its superior province has high levels of protest or a

²² There is a total of 653 Online Sentiment Monitoring Reports, ten of these reports are not produced by J. Prefecture, but instead produced by subordinate counties.

²³ Only 20 sentiment monitoring reports among all reports classify the overall tendency of public sentiment as negative, and in all of these cases, the negative assessment is qualified and described as "tending toward the negative" (较为负面) or "neutral to negative" (中性偏负面).

²⁴ Some issues are highlighted in multiple reports because they persist over time. We eliminate exact duplicates for our main analysis. If duplicates are retained, the substantive results remain unchanged.

²⁵ The 1,412 issues do not contain exact duplicates but contain some complaints that discuss the same issue, perhaps using different words

or by citing different sources. The 1,038 unique complaints are complaints that refer to different issues, in different locations, at different points in time. We use 1,412 as the unit of analysis in the paper, but substantive results remain unchanged if we use the 1,038 unique complaints (see the Supplemental Appendix).

recent history of corruption scandals relative to other prefectures and provinces in China.²⁶

In terms of career ambition, we examine the party position of J. Prefecture and provincial leaders following Kung and Chen (2011). Provincial leaders hold more senior party positions than prefecture leaders so we cannot directly compare their positions. Rather, we compare the actual party position of the leaders relative to positions leaders at comparable levels could hold in the CCP. For provincial leaders, we examine whether the party secretary and governor are alternate or full members of the CCP Central Committee—if both are full members, we label the province as “less ambitious” because leaders have already attained a CCP position that is senior for their level. If either is an alternate Central Committee member, we label the province as “more ambitious” because one of its top leaders could attain a party position of greater seniority. For prefecture leaders, we examine whether the prefecture party secretary and mayor are in the provincial party committee or provincial politburo—if either are in the provincial politburo, we label the prefecture as “less ambitious” because a prefecture leader has already attained a party position that is senior for officials at the prefecture level. If either the party secretary or mayor are in the provincial party committee but neither have attained the provincial politburo, we label the prefecture as “more ambitious” because one of its top leaders could attain a party position with greater seniority. Both the provincial party secretary and governor of J. Prefecture’s superior province were full members of the CCP Central Committee as of 2014. In contrast, the prefecture party secretary and mayor of J. Prefecture were only members of the provincial party committee as of 2014. This suggests that J. Prefecture top leaders faced stronger career ambitions than their provincial superiors. In 2016, the J. Prefecture party secretary was elevated to the provincial politburo, which further suggests that, as of 2014, activities in J. Prefecture were motivated by incentives for political advancement.

How does the arrangement of career ambition we observe in J. Prefecture compare to other Chinese prefectures? We randomly sampled two prefectures from every province in China,²⁷ and we collected the biographies of the provincial party secretary, provincial governor, prefecture party secretary, and prefecture mayor from these prefectures in 2014. Among 54 prefectures, 45 (83 percent) had the same arrangement as

J. Prefecture—provincial top leaders were full Central Committee members, and prefecture leaders were only provincial party committee members. As of early 2018, in 11 prefectures (20 percent), at least one top prefectural leader has attained the provincial politburo. We do not know how many other prefectural leaders tried but failed to reach the provincial politburo, but these estimates suggest that somewhere between 20 to 83 percent of Chinese prefectures have a similar arrangement of career incentives as J. Prefecture and its province.

We do not observe substantial difference between J. Prefecture and its superior province in resources. Top leaders in J. Prefecture, as in all other prefectures in China, exert institutional control over the prefecture propaganda department. In addition, J. Prefecture is one where resources are relatively plentiful. Compared with the other 10 prefectures in its province, J. Prefecture was third in terms of economic production in 2013 and was in the middle of the pack in terms of fiscal revenue. Compared to all other Chinese prefectures, J. Prefecture appears in the middle of the pack across a number of indicators, including GDP, foreign trade, and consumer demand (for details see Supplemental Appendix). The province to which J. Prefecture belongs also has resources to strengthen monitoring. The province has been consistently ranked in the middle of China’s provinces in terms of fiscal revenue.

In sum, when we compare J. Prefecture and its province, while there are no clear differences in terms of resource capacity or regional conditions, the incentives of J. Prefecture officials for career advancements are likely stronger than those of provincial superiors. As a result, we would expect to observe information manipulation by J. Prefecture based on our theory. It is important to note that neither the province nor J. Prefecture represents an extreme case in terms of incentives or resources. Many other prefectures exhibit the same pattern of career incentives as J. Prefecture, and J. Prefecture’s control of the prefecture propaganda department is a feature of China’s administrative structure.

Limitations

This dataset is unique because it provides a view of the transfer of information between different levels of government; however, it faces two main limitations. The first limitation relates to the veracity of the data. The inference we make in this paper depends on the veracity of the leaked email archive we analyzed. While we cannot know for certain whether the data are genuine, the size and extraordinary complexity of this archive make it highly unlikely to be fake. In addition, there are no signs the archive was generated by automated means.

The second limitation of this dataset is that it does not represent the full set of communications between J. Prefecture and its superiors. Although a wide range of issues, including topics marked for internal consumption, appear in this email archive, we do not have

²⁶ For information on protest, we used data from <http://bit.ly/1MOQzdz> and <http://bit.ly/1nk2Azf> to calculate protests at the prefectural and provincial levels (accessed Dec 15, 2016). For information on corruption, we used data from *Procuratorial Daily* (检察日报) between January 1, 2013 and October 11, 2014. Among 825 reports of officials convicted of corruption, only one official from J. Prefecture was mentioned, and only 43 officials from the province were named, among them none were top provincial leaders. Relative to other prefectures and provinces, the rate of conviction for corruption is low in J. Prefecture and the upper-level province.

²⁷ We do not include the municipalities of Beijing, Shanghai, Tianjin, and Chongqing since their levels of administration differ from other provinces. Special administrative regions of China such as Hong Kong are also excluded.

records of in-person conversations, phone calls, text messages, or the transmission of information in other bureaucracies between the prefecture and upper levels of government. This raises the concern that, even if information transmitted by email between the prefecture propaganda department and upper levels of government is biased, upper-level superiors could have an accurate picture of the activities and performance of J. Prefecture.

Upon more careful consideration, however, the incompleteness of our data presents a hard test for finding systematic concealment of wrongdoing. If upper-level authorities have a full and accurate picture of the performance of J. Prefecture, then it would be riskier for the propaganda department of J. Prefecture to manipulate and suppress information because the threat of discovery increases. Because there are other channels of communication between the prefecture and upper levels of government, we should be less likely to find information manipulation in one channel. If we do observe information manipulation in these propaganda department emails, it suggests provincial authorities lack a full picture of the activities of lower-level officials.

Legal and Ethical Considerations

The ethics of conducting research using leaked data has been a subject of intense discussion, especially as leaked data has become ever more important and prevalent in the digital era (Michael 2014). Here, we discuss the legal and ethical implications of working with this data.

From a legal perspective, we do not know how the data was originally obtained. However, the method by which we are obtaining the data—downloading from a publicly available website after learning of the data from journalists—is legal.

Ethically, the phenomenon we are studying with this data pertains to government institutions. The emails that we analyze belong to government offices and bureaucracies where more than one individual may be sending and receiving correspondence. We do not know the identities of any individual involved, and we do not collect any personal, identifying information from the leaked email archive. Nevertheless, we obtained approval from our university Institutional Review Board for analyzing this data.

Finally, although the names of localities are available in the publicly available email archive, we do not include any place names in this paper since names are not germane to our central arguments. Our paper is not intended as a report on the activities of any particular local government in China but rather a study to improve our understanding of information controls under authoritarian rule.

BIASED UPWARD REPORTING

In this section, we show how citizen complaints reported upward by J. Prefecture is biased to con-

ceal complaints that implicate the prefecture of wrongdoing—of breaking the law and engaging in corrupt practices—as well as to conceal complaints of wrongdoing, implicating counties where prefecture leaders have birthplace or workplace ties. We first use an unsupervised method of text analysis and close reading of the text to explore the citizen complaints described in the J. Prefecture Online Sentiment Monitoring Reports. Second, we hand code all unique complaints identified in the Online Sentiment Monitoring Reports to determine whether complaints of wrongdoing by the government are less likely to be reported to upper-level superiors, while accounting for alternative explanations. Finally, we explore the role of censorship in concealing complaints from upper-level authorities.

Content of Citizen Complaints

We use a structural topic model (STM) to gain a better understanding of the types of citizen complaints highlighted by the J. Prefecture Propaganda Department in its Online Sentiment Monitoring Reports, and whether there are differences among topics in terms of upward reporting (Lucas et al. 2015; Roberts et al. 2013, 2014). We are interested in examining whether a complaint was reported to provincial superiors, and we want to understand the relationship between this covariate and the content of complaints. By using an STM, we allow whether a complaint is reported upward to affect the proportion of issues focused on a topic and the distribution of words that characterize a topic.

To determine the number of topics, we compare the held-out likelihood, residual, semantic coherence of models with 25 to 80 topics. The model with 40 topics yielded the most intuitive results and forms the focus of subsequent analysis (Chang et al. 2009). We hand labeled each topic by reading documents associated with the topic, by examining the words that appear with highest probability in that topic, and by examining the words that are frequent and exclusive to that topic. We were able to hand label 30 of the topics.²⁸

Among the 30 topics we could label, half deal with complaints of government corruption and malfeasance.²⁹ Complaints accuse government of embezzlement and misuse of public funds, for example:

²⁸ A coherent topic was not easily discernible for the remaining 10 topics.

²⁹ These topics include Y. county gov't misconduct, P. county gov't malfeasance in land development, D. county gov't malfeasance in land development, Complaints about J. Prefecture governance, Prefecture gov't malfeasance in public welfare, J. county gov't malfeasance in land seizures, Complaints about gov't misconduct and corruption, Embezzlement by civil servants, Gov't malfeasance in public goods provision, Collection of illegal fees, D. county gov't malfeasance, County gov't violence and misconduct, County gov't embezzlement in land, D. county gov't corruption and malfeasance, X. county gov't illegal land development and violence. The estimated topic proportion of all labeled corruption-related topics totals 39.05 percent.

“Complaint from village doctors: relevant laws allow 50 to 60 year old village doctors to participate in skill-based selection examination free of charge. However, the X. county health department took hard earned money from 98 rural doctors by charging them 250 yuan per person for taking the exam.” (X. county 乡村医生反映, 原本五十周岁到六十周岁乡村医生可以免费参与相关法律规定的一技之长选拔考试, 可 X. county 卫生局却变相借机收取九十八名乡村医生的血汗钱, 每人收费250元)

Several topics deal with rent seeking in the form of illegal land development and land taking, for example:

“Online complaints claim that the P. county H. village government spent four years using various methods such as deception and force to expropriate tens of acres of forest and farm land from villagers. And for four years, villagers saw that these forcibly taken lands were unused. Believing the land was greatly wasted, villagers planted crops on them. However, the seedlings were secretly destroyed by village officials based on instructions from the township government officials.” (网称, P. county H. village 政府花四年时间, 利用哄、骗、逼等各种手段, 把繁荣村计家弄村庄几十亩山林和土地强行征去。村民们看到被强行征去的这几亩土地整整抛荒了四年, 感到很可惜, 便种上庄稼, 却被乡政府官员指使村干部偷偷地把青苗全部毁掉)

The remaining half of topics we could label deal with governance issues, such as complaints of pollution, complaints about inadequate public goods provision (roads, schools, health care access, utilities), complaints about student safety and welfare issues, as well as economic disputes between the public and businesses.³⁰ An example of a governance-related complaint writes:

“The garbage transfer station in P. county YM road is extremely smelly and nearby residents are miserable. Online complaints reported that although P. county is building a health-friendly city, the garbage transfer station near YM lake park is noxious all day long...internet users have complained repeatedly to the relevant departments, but the smell persists...hope that the responsible agency can deal with the problem soon and move the garbage transfer station elsewhere.” (P. county YM road 垃圾中转站臭味熏天, 附近居民苦不堪言。网民反映, P. county 正在创建省级卫生城, 可是 YM 湖公园旁边的有个垃圾中转站...这从早到晚一股臭味飘浮在空中...网友先后多次投诉, 相关部门也多次处理, 但臭味依然...希望有关部门能尽快处理, 将垃圾中转站搬走)

Figure 3 shows the effect of upward reporting on topic prevalence for each of the 30 labeled topics. The point estimate is the mean effect of upward reporting, and the lines are 95% confidence intervals. If the estimate and its confidence interval cross the vertical zero line, then whether or not the topic was reported upward does not affect the proportion of complaints focused on this particular topic. Estimates above zero are topics that are more likely to be reported upward. Estimates below zero are topics that are less likely to be reported upward.

³⁰ The estimated topic proportion of all labeled governance-related topics totals 33.78 percent.

From Figure 3, we can see that seven topics, mostly related to corruption, are less likely to be reported upward: Y. county gov't misconduct, P. county gov't malfeasance in land development, “X. county taxi driver strike, D. county gov't malfeasance in land development, Complaints about J. Prefecture governance 1, Prefecture gov't malfeasance in public welfare, and J. county gov't malfeasance in land seizures. There are four topics that are more likely to be reported upward, and one relates to corruption: Complaints about taxi service, Student safety and welfare issues, X. county gov't illegal land development and violence, and Complaints about public goods provision 4. For the remaining labeled topics, which deal with a mix of governance and corruption issues, upward reporting does not have a statistically significant effect on the prevalence of the topic.

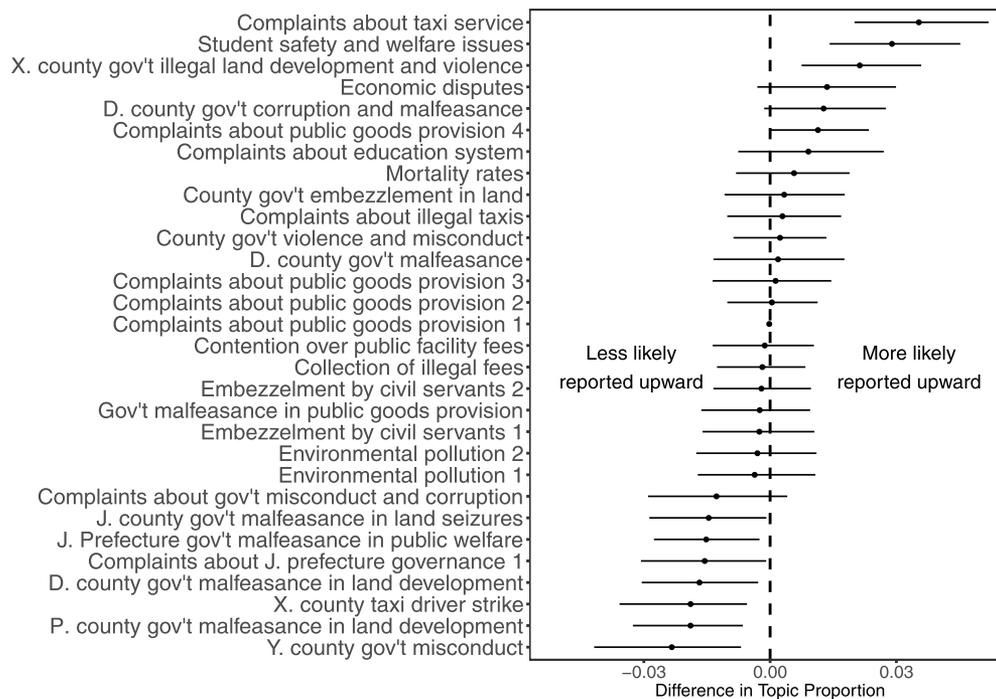
These descriptive results suggest that upward reporting may be motivated by a desire to protect the interests of prefecture-level officials. Two topics criticizing prefecture government performance—Complaints about J. Prefecture governance 1 and Prefecture government malfeasance in public welfare—are less likely to be reported upward, while no topics critical of the prefecture government are more likely to be reported upward.

A hazier picture emerges regarding the corruption of county officials. Figure 3 shows that some topics pertaining to county-level corruption are more likely reported upward (e.g., X. county gov't illegal land development and violence) while other county-related topics with similar content are less likely reported upward (e.g., P. county gov't malfeasance in land development and J. county gov't malfeasance in land seizures). When we examine grievances related to county corruption more closely, and we find that complaints of wrongdoing in some counties are sent up, while complaints of wrongdoing in others are not. Below are two grievances related to X. county, which are reported upward, and two grievances related to Y. county, which are not reported upward. In each pair of complaints, one relates to financial corruption and the other to violence perpetuated by county officials. These two posts about X. county are reported upward:

“Villagers of X. county, S. village reported that the former party secretary HCY of X. county, C. township illegally sold acres of reservoir land in the village.” (X. county S. village 委会谢园片区全体村民报该乡原党委书记现 X. county, C. township 党委书记 HCY 违规贱卖片区千亩水库)

“X. county's police station chief ZMH was reported in online complaints to have led station staff and police in using torture to extort confessions; victims were not given water for around 30 hours and were coerced into confessing the crimes they were accused of.” (网友举报 X. county 派出所所长 ZMH 带领该所教导员和民警刑讯逼供, 受害人被连续审讯用刑近30小时滴水未进, 不得不认罪)

These two posts about Y. county are not reported upward:

FIGURE 3. Effect of upward reporting on topic prevalence, with mean and 95% confidence intervals.

“Online complaints question whether the health insurance bureau in Y. county misused public funds. When paying for health insurance at the beginning of 2012, online complaints report that the bureau forced them to pay for two additional years’ fees, totaling more than 2,700 yuan [per person]. But, the additional fees were not put into the people’s health spending accounts nor given to employers as subsidies, so health accounts had no funds, and the bureau provided no receipts for the fees.” (网友质疑 Y. county 医保局敛财, 2012年初办理医保时, Y. county 医保局强迫必须多缴两年无编制期间的医保统筹金, 共计2700多元, 否则不予办理医保。统筹金不打入个人医保帐户, 单位也没有任何补助, 卡里没有一分钱, 等于是白交, 还不开任何发票凭证)

“Online complaints say that around ten city management police in Y. county severely beat citizens. On Feb 27th, while the county development and reform committee was holding a provincial conference on petitions and maintaining social stability, the city management police were beating citizens for their amusement outside the building.” (网称 Y. county 城管十几人疯狂殴打老百姓。2月27日, 县发改委在楼上开全省信访维稳电视电话会, 楼下城管就在和老百姓玩群殴)

Accusations targeting Y. county are similar to those leveled at X. county, yet only the complaints related to X. county are reported by J. Prefecture to the province.

Our examination of the content of citizen complaints reveals that complaints of government corruption appear frequently among the complaints collected by J. Prefecture’s Propaganda Department. However, citizen complaints reported up to provincial authorities

do not reflect the distribution of topics that J. Prefecture has identified by monitoring citizen complaints. Certain topics are reported upward with greater frequency, and others are reported upward with lesser frequency. This means the subset of information reported to provincial officials does not simply reduce the scale of information, and the complaints shared with the upper level do not reflect the overall prevalence of topics gathered by the J. Prefecture Propaganda Department.

Concealing Complaints of Wrongdoing

To systematically test whether upward reporting is skewed to protect the interests of prefecture-level officials, we use logistic regression to estimate the effect of wrongdoing by J. Prefecture and the effect of wrongdoing by counties with patronage ties to J. Prefecture officials on upward reporting while controlling for alternative explanations.

Following the theory described in the theory Section, we assume that top leaders in J. Prefecture want to maximize rents, which may include engaging in corrupt practices, and want to minimize the risk of being disciplined and removed from office by their provincial superiors. Prefecture officials can access rents and improve their hold on political power by protecting the interests of those within their patronage networks. As Hillman (2010) notes in his study of local-level factions in China, patronage networks between local political leaders and their subordinates help patrons advance

TABLE 2. Upward Reporting for Prefecture and County Wrongdoing

	Reported upward	Not reported upward	Total
Prefecture wrongdoing	17	61	78
Wrongdoing in patronage counties	71	116	187
Wrongdoing in non-patronage counties	68	68	136

politically and increase rent-seeking opportunities for patrons. In terms of observable implications, we expect complaints that could lead to the punishment of J. Prefecture officials to be less likely reported upward, and complaints that incriminate those in the patronage network of J. Prefecture leaders to be less likely reported upward.³¹

Key Variables. Our unit of analysis is each citizen complaint aggregated by the J. Prefecture Propaganda Department. Our dependent variable is whether or not that post is reported upward to provincial-level officials. Our main independent variable is whether the post accuses J. Prefecture government officials or government agencies of wrongdoing (*Prefecture Wrongdoing*), which includes accusations of corruption and violence, as well as violations of laws and regulations.³² We also create a variable that determines whether a post accuses any of J. Prefecture's subordinate counties of wrongdoing (*County Wrongdoing*).³³

Among the 1,412 unique complaints identified in the Online Sentiment Monitoring Reports, 28 percent relate to government wrongdoing: 78 complaints target the prefecture level, and 323 the county level (see Table 2). Among the 78 reports of prefecture-level wrongdoing, 31% (17) are reported upward to provincial superiors. Among the 323 reports of county-level wrongdoing, 43% (139) are reported upward.

To differentiate between counties that are in or out of the patronage network of J. Prefecture officials, we examine the biographies of J. Prefecture politburo members to identify county birthplace and workplace

³¹ Note that patronage may explain information manipulation by the prefecture on behalf of its client counties, but patronage is less likely to explain information manipulation between the prefecture and province. Even if J. Prefecture officials were in the patronage network of provincial leaders, we would expect J. Prefecture to report its wrongdoing to provincial superiors, but we would expect the province not to report J. Prefecture wrongdoing to the center.

³² We hand code posts because it offers us the most precise estimate of our quantity of interest. We do not use the results of the STM because STM, as with topic models more generally, allows any particular document (complaint) to contain multiple topics, and interpreting topics generated by the model is more ambiguous than using explicit coding rules.

³³ If a complaint accuses both the prefecture and county government of wrongdoing, the complaint is coded as 1 for *Prefecture Wrongdoing* and for *County Wrongdoing*.

ties. Among the 14 members of the J. Prefecture politburo, three were born outside of the province, six were born in another prefecture of the same province, and five were born in J. Prefecture or its subsidiary counties. All of the five politburo members born in J. Prefecture come from one of two counties. If we also consider counties where current J. Prefecture politburo members have worked, there are six counties³⁴ out of twelve where two or more politburo members have worked or were born. We code the *Connection* variable as 1 if the complaint pertains to any one of these six county-level units, and 0 otherwise.³⁵ Table 2 shows upward reports of county-level wrongdoing for counties with and without birthplace and workplace ties to J. Prefecture officials. For counties that have connections, 38% (71) complaints of wrongdoing are reported upward, and for counties that do not have connections to the prefecture, 50% (68) complaints of wrongdoing are reported upward. The difference between upward reports of wrongdoing in counties with and without connection is statistically significant with a p -value <0.01 .³⁶

Alternative Explanations. We account for four alternative explanations that could explain upward reporting of complaints. The first relates to the importance of the issue; the second to the reliability of the complaint; the third to alternative sources of information; and the fourth to the division of responsibilities between prefecture and provincial levels of government.

It may be that the subset of information reported upward represents the content provincial officials have deemed to be most important. In other words, prefecture propaganda departments remove content related to less important issues so provincial superiors can focus on the topics of greatest concern. Since the CCP's goals for monitoring citizen complaints is in large part related to identifying poorly performing government officials, if importance motivated upward reporting, we should see more upward reporting of government wrongdoing.

However, other characteristics of citizen complaints might make complaints more or less important to upper-level officials. We include three variables to proxy different dimensions of importance. The first variable (*Prevalence*) denotes whether the complaint relates to an issue area that has garnered the greatest public attention in J. Prefecture. If there is a recurring theme to citizen complaints, perhaps this is what prefecture official would highlight as important for their superiors.³⁷ *Prevalence* is a binary variable that takes on the value of 1 if the topic with the highest topic

³⁴ Specifically, there are five counties and one county-level district.

³⁵ Approximately 40 percent of complaints related to counties within the patronage network of prefecture leaders (see Supplemental Appendix for summary statistics).

³⁶ The difference between upward reports of prefecture wrongdoing and reports of wrongdoing in counties without connection is also statistically significant with a p -value <0.01 .

³⁷ We also examine the correlation between topic proportion from the STM model and upward reporting. Topics that the J. Prefecture Propaganda Department has highlighted more frequently are not those that the department reports up to provincial-level officials. The

proportion from the STM model for that post is one of four topics most prevalent across all posts, and 0 otherwise.³⁸ The second variable (*Group Issue*) measures the number of people impacted by the complaints, since complaints that affect more people may be more important than complaints affecting one person or one household. For each complaint, *Group Issue* takes on the value of 1 if the complaint is shared by more than one individual or household, and 0 otherwise. For example, a complaint that a drunken police officer crashed his vehicle into a pedestrian would be coded as an individual issue (0), while a complaint that village cadres illegally seized land from villagers would be coded as an issue pertaining to multiple households (1).³⁹ The third variable, the (*Sentiment*) of complaints may be another dimension of importance to measure the intensity of the complaint. We determine the sentiment of posts using dictionary-based and probabilistic methods, which produce similar results.⁴⁰ Larger values denote more positive sentiment while smaller values denote more negative sentiment.

The second alternative explanation relates to whether the complaint is based on personal, direct experience, or whether the complaint is based on indirect observations of incidents or events that the person writing the complaints has not experienced directly. When complaints are not based in personal experience, they may be regarded as more speculative and less worthy of upward reporting. We create a variable (*Personal Experience*), which takes on the value of 1 when the post is based on direct, personal experience, and 0 otherwise.⁴¹

The third alternative explanation relates to alternative sources of information. Lower-tier officials may be more likely to report information to superiors if their superiors receive the same type of information through other sources, such that information manipulation is easier to detect. We included two variables—*Collective Action* and *Petition*—to denote information provincial-level superiors may also obtain through other channels. Upper levels may obtain information about collective action events from the public security bureau and information about petitions from the Bureau of Letters and Visits. For *Collective Action*, posts discussing real-world collective action events are coded as 1. For *Petition*, posts related to real-world petitions are coded as 1.⁴² Note that like the prefecture

propaganda department, public security and letters and visits offices at the prefecture level are primarily accountable to the prefecture government. As a result, if J. Prefecture suppresses information across different information gathering channels, complaints related to *Collective Action* and *Petition* would not be more likely to be reported upward.

The last alternative explanation relates to the division of responsibilities between prefecture and provincial governments. Perhaps certain citizen complaints are not reported upward because the responsibility for dealing with these complaints falls to the prefecture government rather than the provincial government. In other words, the prefecture government only reports upward complaints that require provincial governments to take action, and do not report upward complaints that pertain to issues the prefecture should manage itself. We include *Provincial Jurisdiction* to denote complaints that likely necessitate provincial involvement, which are issues that require intervention outside of the prefecture.⁴³ For example, citizens in J. Prefecture complain about water pollution in a local river; however, the contaminants come from an upstream lake bordering several prefectures. This issue requires provincial intervention so the relevant complaints would be coded as 1. Note that the division of labor should not, in theory, prevent upward reporting because monitoring is concerned with providing upper-level superiors an understanding of public opinion. For example, if a prefecture government engages in corrupt practices in an area under its jurisdiction, upper-level superiors should be interested in learning this information since they are responsible for evaluating the performance of lower-tier officials.

Regression Results. Table 3 shows coefficient estimates and standard errors of logistic regression with three specifications.⁴⁴ In column (1), we estimate the effect of *Prefecture Wrongdoing* alone on upward reporting. The result indicates that complaints related to J. Prefecture wrongdoing are less likely to be reported upward, and this result is statistically significant. In column (2), we examine the effect of *Prefecture Wrongdoing* as well as the interaction between *County Wrongdoing* and patronage connections on upward reporting. Complaints related to J. Prefecture wrongdoing remain less likely to be reported upward, but in addition, complaints related to wrongdoing by counties where prefecture-level leaders have birthplace or workplace ties are less likely to be reported upward while complaints related to wrongdoing in the remaining, politically unconnected counties are more likely to be reported upward.

correlation between expected topic proportion and upward reporting is 0.12.

³⁸ Approximately 18 percent of complaints take on the value of 1 for prevalence (see summary statistics in Supplemental Appendix).

³⁹ The majority of complaints, 95 percent, deal with group-based issues (see summary statistics in Supplemental Appendix).

⁴⁰ We measure sentiment using the National Taiwan University Sentiment Dictionary as well as a multinomial logistic regression model trained on the sentiment of Chinese language movie reviews. Results presented in the paper are based on the second measure.

⁴¹ Slightly less than 20 percent of complaints are based on personal or direct experience (see summary statistics in Supplemental Appendix).

⁴² Only 4 percent of complaints relate to collective action, and 1 percent of complaints relate to petitions (see summary statistics in Supplemental Appendix).

⁴³ Only 1 percent of complaints pertain fall under provincial jurisdiction (see summary statistics in Supplemental Appendix).

⁴⁴ We do not cluster standard errors because the vast majority of complaints identified in the monitoring reports are about different issues, individuals, and incidents. Out of 1,412 complaints, there are 1,038 unique issues. We conduct an additional analysis at the issue level, and our results remain unchanged (see Supplemental Appendix for regression results at the issue level).

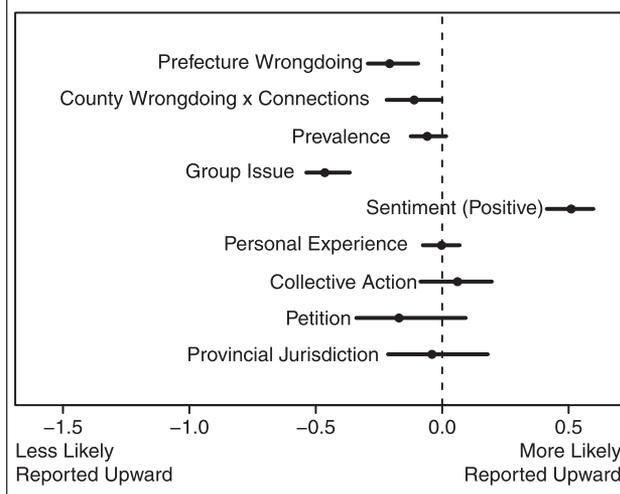
TABLE 3. Predictors of Upward Reporting

	Upward Reporting		
	(1)	(2)	(3)
Prefecture Wrongdoing	-0.994*** (0.280)	-0.997*** (0.286)	-0.990*** (0.297)
County Wrongdoing		0.289 (0.190)	0.273 (0.199)
Connections		-0.012 (0.130)	0.002 (0.138)
County Wrongdoing × Connections		-0.482* (0.263)	-0.461* (0.275)
Prevalence			-0.244 (0.153)
Group Issue			-2.328*** (0.368)
Sentiment			2.304*** (0.291)
Personal Experience			-0.016 (0.153)
Collective Action			0.245 (0.286)
Petitions			-0.800 (0.594)
Provincial Jurisdiction			-0.199 (0.466)
Intercept	-0.284*** (0.055)	-0.282*** (0.081)	1.330*** (0.381)
Observations	1,412	1,412	1,412

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Column (3) includes controls for alternative explanations. The main results remain robust. Wrongdoing by J. Prefecture is less likely to be reported upward and wrongdoing by politically connected counties is less likely to be reported upward. From column (3), we can also see that upward reporting is not related to the prevalence of the issue, and issues which pertain to groups rather than individuals are less likely to be sent up. Instead of extremely negative complaints being more likely to be reported upward, we find that complaints with more positive sentiment are more likely to be reported upward. Complaints based on personal experience are not more likely to be reported upward. There are no statistically significant relationships between collective action and upward reporting or petitions and upward reporting. Finally, issues that are likely to require regional coordination and fall under the jurisdiction of the provincial government are not more likely to be sent up.

Figure 4 plots first differences from the logistic regression specified in column (3) of Table 3 so we can better interpret these results. We can see from Figure 4 that complaints related to prefecture wrongdoing are 21% less likely to be reported upward than complaints that do not mention prefecture wrongdoing. For complaints related to county-level wrongdoing, content pertaining to counties that have ties to prefecture politburo members are 11% less likely than content implicating other counties to be reported upward. There

FIGURE 4. First differences.

is little difference in upward reporting (6%) between posts that pertain to prevalent topics and those that do not. Issues that encompass greater numbers of people are 46% less likely to be reported upward than issues related to individuals. Comparing posts with the most positive sentiment and most negative sentiment, content with the most positive sentiment is 51% more

likely to be reported upward. There is virtually no difference ($<1\%$) in posts based on personal, direct experience and those based on indirect information. Complaints that discuss real-world collection action are 6% more likely to be reported upward, and complaints pertaining to petitions are 17% less likely to be reported upward, but neither result is statistically significant. Finally, when a complaint requires the intervention of provincial authorities, it is 4% less likely to be reported upward, but again, the difference is not statistically significant.

These results clearly show that upward reporting is driven by the incentives of prefecture-level officials to protect their interests, to conceal online complaints accusing the prefecture government of wrongdoing as well as complaints of wrongdoing by counties where prefecture officials have birthplace and workplace ties.

Wrongdoing Not Censored

Finally, we are interested in whether J. Prefecture conceals wrongdoing solely by not reporting the information upward, or whether they also censor the content they do not report to upper levels. Examining whether the content that is not reported upward is censored provides us with information about the local governments' tactics for information manipulation, but perhaps more importantly, sheds light on whether upper levels of government are also monitoring online complaints directed at subordinate prefectures and counties. Presumably, if upper levels of government are also monitoring complaints, upper levels are more likely to detect information manipulation if the concealed information were still available online.

We create two new variables to facilitate this analysis: the dependent variable *Censorship*, based on whether the complaint has been censored or if it remains publicly viewable, and an independent variable *Prefecture Censorship Authority* denoting whether the prefecture has the authority to censor a complaint.

To create the *Censorship* variable, we use a three-step process to check whether the content is still publicly available online. First, we use the Google custom search API to search the text of each complaint and collect the first ten URLs of search results.⁴⁵ Second, we load the content of each returned URL, and we use a variety of automated text matching methods—e.g., cosine similarity, sub-string matching—to determine whether the content of the complaint matches the content of the returned URL. Third, we use extensive human validation to ensure the validity of our text matching, hence the censorship measure. If the complaint remains publicly viewable on any platform, we code *Censorship* as 0, and if it is no longer available, then this variable takes on the value of 1.

⁴⁵ If we are not able to find the complaints using the Google custom search API, we conduct an additional manual search using a variety of search engines such as Baidu.cn.

Based on interviews, we find local-level officials have censorship authority only for websites they manage. For example, J. Prefecture operates a prefecture government website with a public forum, and the J. Prefecture government can censor content on this forum. However, for most other websites (e.g., national-level platforms such as Sina Weibo or Tianya), the J. Prefecture government has no censorship authority. To create the *Prefecture Censorship Authority* variable, we code a complaint as 1 if it was posted to websites owned or operated by J. Prefecture, and 0 otherwise. We find that 16% of complaints are posted to websites run by J. Prefecture, which the prefecture can censor, while the remaining bulk of complaints are posted to sites where J. Prefecture has no censorship authority.

Table 4 shows the result of logistic regressions where the dependent variable is *Censorship*. Column (1) shows a model where prefecture wrongdoing is not interacted with censorship authority, while column (2) contains this interaction. In column (1), we see that posts related to prefecture wrongdoing are more likely to be censored and this effect is exacerbated when the prefecture has censorship authority. In column (2), when these variables are interacted, we see that posts related to J. Prefecture wrongdoing are more likely to be censored, but only when the prefecture has censorship authority over the site where the complaint has been made. Since 84 percent of complaints are not posted to sites where J. Prefecture has censorship authority, this does not apply to the bulk of complaints. We also find that complaints that contain more positive sentiment are less likely to be censored, and complaints based on personal experience are less likely to be censored.⁴⁶

These results show that, when possible, local officials use all tools at their disposal to manipulate information—by distorting what is reported upward and by censoring complaints of prefecture wrongdoing on websites they control. However, since most complaints are posted to websites that J. Prefecture cannot censor, information manipulation likely occurs primarily through the upward reporting process rather than censorship.

Incomplete censorship of J. Prefecture wrongdoing leaves open the possibility that provincial and central leaders could find and punish malefactors based on complaints posted to sites where J. Prefecture cannot censor. However, our finding that prefecture officials persist in concealing wrongdoing despite this possibility strengthens the conclusion that upper-level leaders are not investing the time or resources to evaluate all information pertaining to lower-level governments. As a result, lower-level officials can engage in information manipulation despite incomplete censorship. If provincial officials were able to detect prefecture-

⁴⁶ Since most complaints are posted to websites that J. Prefecture cannot censor, we do not think the negative effect of personal experience on censorship means that more reliable content is less likely to be censored. A more likely explanation is that posts based on personal experience may get more attention or are less likely to be flagged as spam content, and hence less likely to be deleted.

TABLE 4. Predictors of Censorship

	Censorship	
	(1)	(2)
Prefecture Wrongdoing	0.435 (0.272)	0.160 (0.313)
Prefecture Censorship Authority	0.392** (0.177)	0.290 (0.186)
Prefecture Wrongdoing × Prefecture Censorship Authority		1.601** (0.761)
Prevalence	−0.878*** (0.220)	−0.878*** (0.220)
Group Issue	−0.270 (0.332)	−0.269 (0.332)
Sentiment	−0.653* (0.359)	−0.654* (0.360)
Personal experience	−1.648*** (0.278)	−1.654*** (0.279)
Collective Action	0.242 (0.343)	0.245 (0.343)
Petition	−0.813 (0.785)	−0.794 (0.785)
Provincial Jurisdiction	−0.936 (0.754)	−0.956 (0.754)
Intercept	−0.687** (0.347)	−0.668* (0.348)
Observations	1,337	1,337

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

level corruption (from the complaints posted to websites the prefecture cannot censor), the province would more likely be able to identify discrepancies between prefecture propaganda reports and provincial information. Then, if the province so desires, it could sanction, or punish, prefecture propaganda departments for manipulating information. The strong empirical result that the prefecture propaganda department systematically manipulates information along with our existing knowledge of China's information-monitoring system suggests provincial superiors in this region are not directly expending resources to mine and analyze data from online complaints for identifying lower-level corruption.

CONCLUSION

China is often characterized as a regime that has prevailed against the information problems that plague autocrats, in large part because it has built many channels to gather information through citizen complaints and grievances. Our evidence indicates the opposite—that there are systematic shortcomings in China's ability to gather reliable and accurate information about the actions of regime agents through citizen participation. China remains reliant on monitoring agencies to gather, distill, and verify the large quantities of information voiced by the public.

Our data, based on rarely seen internal communications between a monitoring agency and upper-level authorities in J. Prefecture, reveals information manipulation by the prefecture government. These results correspond with our theoretical expectations given the relative incentives and resources of top leaders in J. Prefecture to hide information relative to the incentives and resources of top provincial leaders to overcome information manipulation. However, in other localities where the relative incentives and resources of upper-level and lower-level officials differ from those in J. Prefecture, the strategic interaction between upper-level and lower-level officials over information may yield different outcomes. We hope future research will examine this dynamic in other regions of China, and in other authoritarian contexts.

We might also expect deviations from the findings in this analysis when it comes to monitoring by the central government. China's central authorities have more resources and stronger capabilities to directly monitor regime agents than subnational levels of government. However, the center has made clear that it is interested primarily in monitoring national- and regional-level activities.⁴⁷ The purpose of China's hierarchical administrative structure is to delegate, and likewise,

⁴⁷ As discussed previously, we know that central authorities are currently developing systems that would directly gather and analyze citizen complaint data from social media platforms down to the prefecture level.

responsibility for regularly monitoring low-level officials has been delegated to subnational leaders. Officials at these lower levels—the county and below—are those primarily responsible for policy implementation and those who most frequently interact with the public. Thus, information manipulation and concealment of corruption at county and lower levels will likely persist despite increased central monitoring. This is consequential for the regime since public trust and satisfaction with county and lower levels of government in China is low, hindering the country's ability to govern, to carry out policies, and to prevent mass incidents (Li 2004; Whyte 2010; Yan and Peng 2010).

These results demonstrate an alternative way in which information manipulation occurs—not primarily through censorship or the deletion of undesirable information and not through falsification or fabrication, but through partial concealment. Because upper-level authorities rely on monitoring agencies to synthesize large quantities of online data, monitoring agencies can satisfy upper-level demands for information with truthful but incomplete information that systematically hides corruption.

Our findings bring additional nuance to our understanding of the relationship between nonelectoral forms of citizen participation and accountability in authoritarian regimes. Although individuals living under authoritarian rule are in some ways free to publicly express their grievances, whether this information can lead to accountability depends on whether this information can make its way to those with sanctioning power, and whether those who can sanction actually do so. Our results, showing how the public's complaints of lower-level malfeasance do not reliably make their way to upper-level authorities, reveal one way in which public participation is insufficient for accountability in nondemocratic contexts.

Finally, these results show that conflicts over information remain a challenge for authoritarian regimes in the digital age. Even in a highly determined and capable authoritarian regime where grievances are actively and publicly voiced, information manipulation by lower-level officials persists. Even in the era of large-scale data and increasingly sophisticated methods for analyzing large quantities of data, political incentives continue to motivate information manipulation. Even though authoritarian regimes across the world are adopting online systems for public complaints, and autocrats may learn a great deal from monitoring online content and social media, this proliferation of recorded information by no means guarantees autocrats omniscience. We hope our results will generate greater focus on political communication and information conflicts among elites and regime insiders to complement existing work on information conflicts between the regime and society.

SUPPLEMENTARY MATERIAL

To view supplementary material for this article, please visit <https://doi.org/10.1017/S0003055418000205>.

REFERENCES

- Cai, Yongshun. 2000. "Between State and Peasant: Local Cadres and Statistical Reporting in Rural China." *The China Quarterly* 163: 783–805.
- Cai, Yongshun. 2013. "Information Collection and Anticorruption in China." In *Elites and Governance in China*, eds. Xiaowei Zang and Chien wen Kou. New York: Routledge, 136–52.
- Central Committee of the Communist Party of China. 2016. "Zhongguo gongchandang dangnei jiandu tiaoli [China Communist Party Internal Supervision Regulations]."
- Central Propaganda Department. 2009. *Wangluo Yuqing Xinxi Gongzuo Lilun yu Shiwu* [Theories and Guidance on Online Public Sentiment]. Xuexi chuban she.
- Chang, Jonathan, Jordan L. Boyd-Graber, Sean Gerrish, Chong Wang, and David M. Blei. 2009. "Reading Tea Leaves: How Humans Interpret Topic Models." In *Nips*, Vol. 31, pp. 1–9.
- Chen, Jidong, Jennifer Pan, and Yiqing Xu. 2016. "Sources of Authoritarian Responsiveness: A Field Experiment in China." *American Journal of Political Science* 60 (2): 383–400.
- Chen, Xi. 2009. "Power of Troublemaking: Chinese Petitioners' Tactics and Their Efficacy." *Comparative Politics* 41 (4): 451–71.
- China Daily. 2015. "New Anti-Corruption App Sparks Instant Flood of Public Reports." Accessed December 13, 2016. <http://bit.ly/1GhcZQL>.
- Cyberspace Administration of China. 2017. "Hulianwang Xinwen Xinxi Fuwu Guanli Guiding [Internet News and Information Service Management Regulations]." Beijing, China.
- Dimitrov, Martin K. 2014a. "Internal Government Assessments of the Quality of Governance in China." *Studies in Comparative International Development*, 1–23.
- Dimitrov, Martin K. 2014b. "What the Party Wanted to Know Citizen Complaints as a "Barometer of Public Opinion" in Communist Bulgaria." *East European Politics & Societies* 28 (2): 271–95.
- Distelhorst, Greg, and Yue Hou. 2017. "Constituency Service Under Nondemocratic Rule: Evidence from China." *Journal of Politics* 79 (3).
- Economist. 2017. "Chinese Officials Use Hotlines to Take the Public's Pulse." Accessed August 20, 2017. <http://whr.tn/1imqq1x>.
- Edin, Maria. 2003. "State Capacity and Local agent Control in China: CPP Cadre Management from a Township Perspective." *China Quarterly* 173: 35–52.
- Fishkin, James S., Baogang He, Robert C. Luskin, and Alice Siu. 2010. "Deliberative Democracy in an Unlikely Place: Deliberative Polling in China." *British Journal of Political Science* 40 (2): 435–48.
- Freedom House. 2015. *Freedom in the World 2015: The Annual Survey of Political Rights and Civil Liberties*. Lanham, MD: Rowman & Littlefield.
- Gao, Jie. 2016. "Bypass the Lying Mouths': How Does the CCP Tackle Information Distortion at Local Levels?" *The China Quarterly*. p. doi:10.1017/S0305741016001491.
- Gunitsky, Seva. 2015. "Corrupting the Cyber-Commons: Social Media as a Tool of Autocratic Stability." *Perspectives on Politics* 13 (01): 42–54.
- He, Baogang, and Mark Warren. 2011. "Authoritarian Deliberation: The Deliberative Turn in Chinese Political Development." *Perspectives on Politics* 9 (2): 269–89.
- Henochowicz, Anne. 2014. "Thousands of Local Internet Propaganda Emails Leaked." *China Digital Times*. <http://j.mp/leakedEmails>.
- Hillman, Ben. 2010. "Factions and Spoils: Examining Political Behavior within the Local State in China." *The China Journal* 64: 1–18.
- Hsu, S. Philip, Suisheng Zhao, and Yushan Wu. 2011. *In Search of China's Development Model: Beyond the Beijing Consensus*. Abingdon, UK: Routledge.
- Huang, Yasheng. 1995. "Administrative Monitoring in China." *The China Quarterly* 143: 828–43.
- Jiang, Junyan, Tianguang Meng, and Qing Zhang. 2016. "Closing the Influence Gap: The Policy Consequences of Online Participation in China." *APSA Annual Meeting*.
- Kalathil, Shanthi, and Taylor C. Boas. 2010. *Open Networks, Closed Regimes: The Impact of the Internet on Authoritarian Rule*. Washington D.C.: Carnegie Endowment.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences

- Collective Expression." *American Political Science Review* 107 (2): 1–18.
- Kou, Chien-wen, and Wen-Hsuan Tsai. 2014. "Sprinting with Small Steps' Towards Promotion: Solutions for the Age Dilemma in the CCP Cadre Appointment System." *China Journal* 71: 153–71.
- Kung, James Kai-Sing, and Shou Chen. 2011. "The Tragedy of the Nomenclatura: Career Incentives and Political Radicalism during China's Great Leap Famine." *American Political Science Review* 105 (1): 27–45.
- Li, Lianjiang. 2004. "Political Trust in Rural China." *Modern China* 30 (2): 228–58.
- Li, Ruichang. 2012. *Zhengfujian Wangluo Zhili: Chuizhi Guanli Bumen Yu Difang Zhengfujian Guanxi Yanjiu* [Network Governance: Top-down Management of Departments and Inter-Governmental Relationships]. Shanghai, China: Fudan University Press.
- Lieberthal, Kenneth, and David Lampton. 1992. *Bureaucracy, Politics, and Decision Making in Post-Mao China*. Berkeley, CA: University of California Press.
- Lorentzen, Peter. 2013. "Regularizing Rioting: Permitting Protest in an Authoritarian Regime." *Quarterly Journal of Political Science* 8 (2): 127–58.
- Lucas, Christopher, Richard A. Nielsen, Margaret E. Roberts, Brandon M. Stewart, Alex Storer, and Dustin Tingley. 2015. "Computer-Assisted Text Analysis for Comparative Politics." *Political Analysis* p. mpu019.
- Ma, Jun. 2012. "The Rise of Social Accountability in China." *Australian Journal of Public Administration* 71 (2): 111–21.
- MacKinnon, Rebecca. 2012. *Consent of the Networked: The Worldwide Struggle For Internet Freedom*. New York, NY: Basic Books.
- Magaloni, Beatriz. 2008. "Credible Power-Sharing and the Longevity of Authoritarian Rule." *Comparative Political Studies* 41 (4/5): 715–41.
- Magaloni, Beatriz, and Ruth Kricheli. 2010. "Political Order and One-Party Rule." *Annual Review of Political Science* 13: 123–43.
- Mai, Bo, and Tim Liber. 2015. "Web Tracking with Chinese Characteristics." 65th Annual International Communication Association Conference.
- Manion, Melanie. 2016. *Information for Autocrats: Representation in Chinese Local Congresses*. Cambridge, UK: Cambridge University Press.
- Meng, Tianguang, Jennifer Pan, and Ping Yang. 2017. "Conditional Receptivity to Citizen Participation: Evidence from a Survey Experiment in China." *Comparative Political Studies* 50 (4): 399–433.
- Michael, Gabriel J. Forthcoming. "Who's Afraid of Wikileaks? Missed Opportunities in Political Science Research." *Review of Policy Research*.
- Morozov, Evgeny. 2012. *The Net Delusion: the Dark Side of Internet Freedom*. New York, NY: Public Affairs.
- Nathan, Andrew. 2003. "Authoritarian Resilience." *Journal of Democracy* 14 (1): 6–17.
- National Grassroots Party Work Key Textbook. 2013. *Zuixin Dang de Xuanchuan Gongzuo Caozuo Fangfa yu Chuangxin Shiwu* [Latest Guidelines and Procedures on How to Perform Propaganda Work]. Beijing, China: Hongqi Press.
- Noesselt, Nele. 2014. "Microblogs and the Adaptation of the Chinese Party-State's Governance Strategy." *Governance* 27 (3): 449–68.
- O'Brien, Kevin, and Lianjiang Li. 1999. "Selective Policy Implementation in Rural China." *Comparative Politics* 31 (2): 167–86.
- O'Brien, Kevin J., and Lianjiang Li. 2000. "Accommodating 'democracy' in a one-party state: Introducing village elections in China." *The China Quarterly* 162: 465–89.
- O'Donnell, Guillermo. 1999. Horizontal Accountability in New Democracies. In *The Self-Restraining State: Power and Accountability in New Democracies*, eds. Larry Diamond, Andreas Schedler, and Marc F. Plattner. Boulder, CO: Lynne Rienner Publishers, Inc., pp. 29–51.
- OECD. 2013. *OECD e-Government Studies: Egypt 2013*. Paris, France: OECD Publishing.
- Oi, Jean C. 1995. "The Role of the Local State in China's Transitional Economy." *The China Quarterly* 144: 1132–49.
- Olken, Benjamin A. 2007. "Monitoring Corruption: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 115 (2): 200–49.
- Pan, Jennifer. Forthcoming. "How Chinese Officials Use the Internet to Construct their Public Image." *Political Science Research and Methods*.
- Pan, Jennifer, and Kaiping Chen. 2018. "Replication Data for Concealing Corruption: How Chinese Officials Distort Upward Reporting of Online Grievances." <https://doi.org/10.7910/DVN/2HRRCT>.
- Park, Albert, and Sangui Wang. 2001. "China's Poverty Statistics." *China Economic Review* 12 (4): 384–98.
- Pearce, Katy E., and Sarah Kendzior. 2012. "Networked Authoritarianism and Social Media in Azerbaijan." *Journal of Communication* 62 (2): 283–98.
- Qiang, Xiao. 2011. "The Battle for the Chinese Internet." *Journal of Democracy* 22 (2): 47–61.
- Qin, Bei, David Strömberg, and Yanhui Wu. 2016. "Media Bias in China." Working paper.
- Qin, Bei, David Stromberg, and Yanhui Wu. 2017a. "Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda." CEPR working paper DP11778.
- Qin, Bei, David Stromberg, and Yanhui Wu. 2017b. "Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda." *Journal of Economic Perspectives* 31 (1): 117–40.
- Roberts, Margaret E., Brandon M. Stewart, Dustin Tingley, Chris Lucas, Jetson Leder-Luis, Bethany Albertson, Shana Gadarian, and David Rand. 2014. "Topic Models for Open Ended Survey Responses with Applications to Experiments." *American Journal of Political Science* 58: 1064–82.
- Roberts, Margaret E., Brandon M. Stewart, Dustin Tingley, and Edoardo M. Airolidi. 2013. "The Structural Topic Model and Applied Social Science." In *Advances in Neural Information Processing Systems Workshop on Topic Models: Computation, Application, and Evaluation*.
- Rodan, Garry, and Kanishka Jayasuriya. 2007. "The Technocratic Politics of Administrative Participation: Case Studies of Singapore and Vietnam." *Democratization* 14 (5): 795–815.
- Sonnad, Nikhil. 2014. "Hacked emails reveal China's elaborate and absurd internet propaganda machine." *Quartz*. <http://j.mp/Sonnad>.
- Stockmann, Daniela. 2013. *Media Commercialization and Authoritarian Rule in China*. New York: Cambridge University Press.
- Truex, Rory. 2014. "The Returns to Office in a 'Rubber Stamp' Parliament." *American Political Science Review* 108 (02): 235–51.
- Truex, Rory. 2016. *Making Autocracy Work*. Cambridge, UK: Cambridge University Press.
- Tsai, Lily. 2008. "Understanding the Falsification of Village Income Statistics." *The China Quarterly* 196: 805–26.
- Wallace, Jeremy L. 2016. "Juking the Stats? Authoritarian Information Problems in China." *British Journal of Political Science* 46 (01): 11–29.
- Whiting, Susan. 2004. "The Cadre Evaluation System at the Grass Roots: The Paradox of Party Rule." In *Holding China Together: Diversity and National Integration in the post-Deng Era*, eds. Dali Yang and Barry Naughton. Cambridge, UK: Cambridge University Press, 101–19.
- Whyte, Martin. 2010. *Myth of the Social Volcano: Perceptions of Inequality and Distributive Injustice in Contemporary China*. Stanford, CA: Stanford University Press.
- Wintrobe, Ronald. 1998. *The Political Economy of Dictatorship*. Cambridge, UK: Cambridge University Press.
- Xinhua. 2012. "China's craze for online anti-corruption." Accessed September 10, 2016. <http://bit.ly/2blgIYz>.
- Yan, Ye, and Min Peng. 2010. "The Unbalance Political Trust Distribution between Central Government and Local Government in China." *The Journal of Gansu Administration Institute* 3 (008).
- Yong, Hu. 2005. *Blogs in China: China Media Project Case Study*. Technical report, Peking University.
- Zhan, Weiping, and Yongda Yu. 2011. "Study of the Relationship Between Cadre Promotion and Career Path." *Human Resources Development of China* 7: 75–8.
- Zhang, Jian, John Giles, and Scott Rozelle. 2012. "Does it Pay to be a Cadre? Estimating the Returns to being a Local Official in Rural China." *Journal of Comparative Economics* 40 (3): 337–56.