

Social Media, Information Networks, and Protests in China

Bei Qin, David Strömberg, and Yanhui Wu*

Preliminary draft

July 2019

Abstract

This paper studies whether the explosive growth of social media in China affects the spread and incidence of protests. We combine a unique dataset of 13.2 billion microblog posts published on Sina Weibo—the Chinese equivalent to Twitter—during 2009-2013 with detailed information on thousands of protests and strikes during 2006-2016. Using retweets by users in one city of blogposts from users in other cities, we construct a network of social media information flow across cities. The rapid network expansion generates time-series variation in these information flows that we use to identify causal effects. In particular, we develop a novel difference-in-difference methodology to estimate the impact of network interactions. We find that information diffusion over the social media network has a sizeable and significant effect on the spread of both protests and strikes. The spread of events caused by information flow over social media is fast and predominantly local – between events within the same category (e.g., cause and industry); event spread across categories is still significant, albeit weaker. Furthermore, we find that social media networks increase the incidence of protests and strikes. These findings shed light on the recent debate regarding the political role of social media in autocracies by showing that in the attempt of using social media for surveillance and preempting anti-regime collective action, autocratic governments face the risk of triggering widespread protests.

*Qin: The University of Hong Kong, beiqin@hku.hk; Stromberg: University of Stockholm, david.stromberg@ne.su.se; Wu: University of Southern California, yanhuiwu@marshall.usc.edu.

1 Introduction

Whether social media promote political participation and improve accountability is heatedly debated in recent years. On one hand, social media facilitate the free flow of public information, which may help citizens organize political action that constrains politicians' misbehavior. On the other hand, this information may be covertly used or manipulated by politicians against citizens' interest. This debate is even more prominent in autocracies with few alternative free media sources (e.g. Shirky, 2011; Morozov, 2012). However, rigorous empirical study that sheds light on this subject is scant, with a notable exception being Enikolopov et al. (2018). In this paper, we provide some of the first large-scale evidence regarding the effects of social media on the incidence and spread of protests and the ensuing government reaction in China.

Authoritarian regimes are characterized by the lack of political opportunities for collective action and strict control of the media. However, this does not mean that popular resistance and public communication of political information are absent in such regimes. In China, protests against unpopular local government policies or corruption take place from time to time. Strikes in which people seek redress for economic issues have increased drastically in the last two decades (China Labor Bulletin 2015, 2018). These costly events provide credible ways for citizens to express their attitude towards the government and their concerns about public policies. To the extent that these events help identify social problems, monitor officials, and avoid policy oversight, the Chinese government allows a certain level of freedom in media coverage of them.¹

Despite the informational value of protests, it is crucial for an authoritarian regime to keep protests contained within regions and social groups because widespread protests can have severe political consequences. When protests spread across regions, initially small events can exhibit a scale shift, leading to social unrest that may threaten the existing regime. A cautionary example is Solidarity in Poland, which originated from limited demand for workers' rights and better economic conditions but soon developed into pervasive resistance that was fatal to the regime. Examining how social media affect an authoritarian regime's balancing of the benefit of learning local information from isolated protests against the threat of widespread protests is the focus of our study. In particular, we investigate whether and how information flow in the social media network causes protests and strikes to break the geographical boundaries and generate far-reaching social movement.

Since 2009, social media have played an important role in the Chinese public sphere. Information circulating in social media can affect protests in various ways. It may be used by protesters to organize events and coordinate action (e.g., Acemoglu et al. 2018; Enikolopov et al., 2018). Moreover, the emergence of social media provides unprecedented opportunities

¹This is the so-called "mass line" in the political dogma of the Chinese Communist Party. Its essence is to learn about public opinion and policy outcomes from people (the mass) through observing their activities, particularly collective action (Zhao 1998).

for both citizens and governments to learn about policy consequences, corruption, public opinion, and generally the state of the world. From the perspectives of citizens, discussion about the cause, process, and consequences of protests on social media can inform them of bad policy outcomes, reinforce their belief about local conditions and the type of the regime, and even inspire them to learn effective protest tactics (e.g. Little 2016; Chan and Suen 2016). From the perspectives of the government, social media can be used to surveil potential protesters and preempt unfavorable collective action. A priori, it is unclear whether social media will increase the incidence of protests. The effect of social media becomes even more complex when the concern is about the spread of protests. On the one hand, simultaneous protests across regions are likely to receive favorable government response (e.g., concession), and thus information about protests in one region may spur additional protests. On the other hand, spreading protests are more regime-threatening than isolated protests and will inspire governments to take more aggressive preemption action.² The overall effect of social media on protests depends on citizens' use of social media information for coordination and learning, the regime's use of social media for surveillance, and governments' response to protests (e.g., preemption and concession). Our empirical investigation aims to assess the relative importance of these different mechanisms so as to draw implications beyond the current setting.

Given the Chinese context in which censorship is pervasive, the first empirical question is whether information about protests and strikes circulates in Chinese social media at all. To this end, we study the content of a unique social media dataset consisting of 13.2 billion posts published on Sina Weibo—the most prominent microblogging platform in China—during 2009-2013. We find that there are approximately 4 million microblog posts mentioning protest or strikes, and these posts predict real-world protests and strikes (see Qin et al. 2017 for more details). Moreover, we find approximately 40 million forwards of the protest and strike posts. In a subset of 3 million forwards for which we have precise timing and location information, around 30% of the forwards occur within one hour after the posting of the original messages and 80% within one day; after one hour, the mean distance between the user who posts a message and the user who forwards it is more than 800 km. Evidently, Chinese social media are abundant in information about protests and strikes, and the rapid diffusion of this information suggests that social media have the potential to generate sweeping waves of protest activities. Combining this social media dataset with detailed information on thousands of protests and strikes from 2006 to 2016, we examine how information diffusion over social media affects the spread of these events.

Estimating the effect of information flow over social media on the spread of real-world events is intricate. Cities with strong informational ties through social media are likely to have other channels for communication between them (e.g., phones, meetings). Moreover,

²Another argument for a negative effect of social media on protests is the so-called slacktivism view in which people protest online to release their anger and would not protest offline.

social media users tend to communicate more extensively with users from cities that are more similar to theirs. Events taking place in socially well-connected and similar cities are naturally more closely correlated, regardless of the impact of social media. Therefore, it is challenging to identify the informational effect of social media on events. This type of identification challenge has been widely discussed in the studies of social networks, pioneered by Manski (1993). Existing solutions include exploiting the network structure to identify instruments (e.g. Bramoulle et al. 2009), using instruments that are uncorrelated with the error terms and the network (e.g. Acemoglu et al., 2015; Konig et al., 2017), matching (Aral et al., 2009), and explicit randomization (e.g. Bakshy et al., 2012).

We propose a novel identification approach taking advantage of the network dynamics. In particular, we exploit the time-series variation from the rapid network expansion to identify causal effects. We approximate the information flow between cities over social media by the number of forwards in one city of the posts originated from another city— posts about all topics except protests and strikes. We allow the average spread of events between cities to be correlated with the average network strength while assuming that the exact timing of the network expansion is unrelated to other changes in the spread of events. Practically, we employ two research designs. First, we use a differences-in-differences estimator that allows for arbitrary time-invariant heterogeneity in the interactions between cities. Second, we treat the average connectedness between cities as fixed and then investigate how the average spread of events evolves across cities that are eventually closely connected through social media over time. Monte Carlo simulations justify the validity of our methods. We also show that we can consistently test the null hypothesis that the social media network does not affect event spread even if social media enable researchers to observe more events. The solution to such an observability problem lies in the justifiable assumption that increased observability per se does not affect the temporal clustering of events across cities. In contrast, one cannot accurately estimate the effects of social media penetration on the incidence of events if social media affect observability.

The main findings of this study are summarized as follows. First, our two research designs consistently provide strong evidence that protests and strikes spread rapidly (within 2 days) through the social media network. We estimate that, because of information flows through social media, a strike within one prefecture in the last two days increases the probability of a strike in all other prefectures by .14. This corresponds to an increase of 33 percent in terms of probability of event, relative to the average strike probability.³ The corresponding number for protests is a 55 percent increase in probability relative to the average protest probability. We find that the spread of events is predominantly local— events spread more within the same event category (for the same social cause or in the same industry). For example, strikes occurring in the manufacturing industry in a province spread more to other strikes in

³The average strike probability for a given prefecture and day is .0015 and there are 282 other prefectures in our sample: $.14/(282*.0015)=.33$. The corresponding number for protests is $.23/(271*.0015)=0.55$.

the same industry within the same province than non-manufacturing strikes outside of the province. This strong within-category spread does not mean that events do not spread across categories at all. Instead, we find that social media also induce weaker but still significant spread of events across provinces and event categories. Since the total number of events across all categories is ten times larger than within the same category, the aggregate effect of across-category spread is larger than within-category spread. In this sense, social media help protesters break geographical, industrial, and societal barriers and generate widespread social movement.

Second, we find that after a peak in 2013-2014, spread of strikes through the social media network slowed down and even ceased after 2016 and that protests continued to spread but at a substantially lower level. This is likely to be caused by the Chinese central government's more stringent control over the social media and its crackdown on collective action after 2013.⁴

Third, based on a difference-in-differences approach, we find that the penetration of Sina Weibo measured by the number of posts per capita is associated with an increase in the numbers of both protests and strikes. The effects are statistically significant and sizeable: they imply that the increases in the number of protests and strikes attributed to Sina Weibo are 13 and 11 per month respectively, which account for one third of the total increase in these events over our sample period. Admittedly, part of this increase may be driven by increased observability due to the use of social media as mentioned above. However, the heterogenous effect across regions with different access to information before the launch of Sina Weibo suggests that the penetration of Sina Weibo still plays an important role in increasing event incidence even after accounting for the observability problem. Moreover, we find that greater Weibo penetration is more strongly associated with an increase in strikes met with concession than those met with repression.

This paper primarily contributes to the emerging literature on the effect of social media and more generally information and communication technology (ICT) on collective political action in non-democratic countries. Inspired by the "colour revolution", particularly the Arab Spring, some scholars believe that social media can provide an important means for citizens to organize collective action opposing and even overthrowing the ruler, and that this power released by social media helps hold authoritarian governments accountable (e.g., Shirky 2011; Tucker et al. 2016; Acemoglu et al. 2018). Recently, Manacorda and Tesei (2018) show that mobile phones are instrumental to mass mobilization during economic downturns in Africa. Enikolopov et al. (2018) show that the penetration of a dominant Russian online social network led to more protest activity in Russia. Our study advances this line of research in two important aspects. First, our research focuses on a large number of local protests and

⁴The Chinese government's stricter control of social media after 2013 is reflected in the declining media freedom index constructed by the Freedom House. Moreover, the Chinese government's procurement of media-based surveillance equipment increased massively from 2013 to 2016. Source: China Government Procurement Website.

strikes, which are not intended to overthrow the regime but instead can provide substantial value for the regime to collect information and monitor local officials. Without a scale-shift effect through event spread, the regime will be tolerant to such events. Therefore, the effect of social media on grass-root political action and the degree of media control depend crucially on the effect on the spread of events. We precisely study this effect and the underlying mechanisms. Second, we also study how social media affect government reaction to protests. This has important implications for the effect of social media on accountability.

In a broader context, our study is related to the literature that studies the effect of the media so as to understand the role of government-controlled media in autocracies. The media has long been regarded as a powerful propaganda instrument for autocrats to exert political control. This traditional view has gained support in a number of studies (eg., Enikolopov et al. 2011; Yanagizawa-Drott 2014; Adena et al. 2015). Another strand of literature stresses the role of the media in autocracies as an effective tool to monitor local officials (e.g., Egorov et al. 2009; Lorentzen 2014). In the era of social media, some researchers push this strategic-media-control view further by arguing that an authoritarian government can even increase regime stability and enhance power by exploiting social media for surveillance and propaganda (e.g., Morozov 2012; Edmond 2013; Lorentzen 2017). This is partially supported by recent evidence on the Chinese government’s strategic censorship on the social media (e.g., King et al. 2013, 2014) and the use of social media for surveillance and propaganda (e.g., Qin et al. 2017). In the current paper, by demonstrating nuanced effects of social media on the spread of protests in China, we show that while authoritarian regimes may control social media to curtail some outcomes, they may not learn information as effectively as citizens and thus cannot turn social media completely in their favor. In the Chinese setting, one important factor that cripples the government’s ability to learn from bottom-up information is rooted in the decentralized political system, in which local governments lack incentives to communicate true information.

The remainder of this paper proceeds as follows. Section 2 describes the background of this study, based on which we draw the major empirical implications that we will test. Section 3 describes the data, and Section 4 presents the main results. Section 5 provides further analysis to examine some subtle mechanisms. Section 6 concludes.

2 Background and discussion

2.1 Social Media in China

Historical development. We primarily study the sample period of 2006-2013, during which the use of social media surged in China. In 2006, some Chinese people became aware of Twitter; several Chinese counterparts were launched the next year. With the Chinese government blocking Twitter and Facebook and strictly controlling domestic microblogging services, the use of social media was limited until Sina Weibo appeared in August 2009.

Sina Weibo is a hybrid of Twitter and Facebook, allowing users to tweet and retweet short messages with embedded pictures or videos and send private messages and write comments. By 2010, Sina Weibo had 50 million registered users, and this number doubled in 2011, reaching a peak of over 500 million at the end of 2012 (China Internet Network Information Center 2014). Over our main sample period (2009-2013), Sina Weibo was the dominant microblogging platform in China. Since 2013, it has lost ground to WeChat, a cellphone-based social networking service, but has remained an influential platform for public communication.

Political communication over Sina Weibo. Users of Sina Weibo and other blogging platforms actively posted their opinions on public events, government policies, and social and economic problems. They also frequently expressed their sentiment towards government officials and social events. Occasionally, some users revealed their attitude towards the leadership and ideology of the Chinese Communist Party (CCP). In Qin et al. (2017), we document the existence of tens of millions of posts published on Sina Weibo discussing sensitive issues such as corruption, protests, strikes, and petitions during 2009-2013. The enormous volume of these posts and the ensuing forwards and comments over social media created unprecedented public communication of political information in China.

Government control and censorship. Certainly, the information flow on Chinese social media is subject to the control of the Chinese government. The first control tool is policing—to punish users who post sensitive content, which is carried out by tens of thousands of information officers and internet monitors (Chen and Ang 2011). Another control tool is censoring. Censorship on Chinese social media is regulated by the national Propaganda Department of the CCP and is implemented largely by private service providers who are registered in Beijing. The estimated extent of censorship of Sina Weibo ranges from 0.01 percent of posts published by a sample of prioritized users (Fu, Chan, and Chau 2013) to 13 percent of posts on selected sensitive topics (King, Pan, and Roberts 2013). Zhu et al. (2013) find that among the censored posts, 30 percent of deletions occur within the first half hour and 90 percent within 24 hours.

2.2 Protests and Strikes

Authoritarian regimes are often characterized by the lack of opportunities for political collective action such as protests despite the presence of popular resistance. However, In China, there have been numerous instances of collective resistance in the last two decades. While many of these instances are small-scale protests and strikes confined within certain localities, some of them are large scaled, spreading across regions, and disruptive. For instance, an event that was widely reported by Western media is the Wukan Event in 2011 when thousands of farmers in a city in Guangdong province protested against corruption of local officials. The event led to direct confrontation between farmers and officials, violent conflicts between protesters and police, and demonstrations in multiple cities in support to the farmers. The Chinese government appears more tolerant of collective action than the typical authoritarian

regime, for example, by requiring local officials to handle collective action events strategically rather than simply suppressing all events.

The Chinese government's tolerance to political collective action may give rise to multiple benefits. First, absolute suppression of collective action may generate distrust and undermine the legitimacy of the ruling party. Top CCP leaders have made public statements urging restraint in police handling of protests (Steinhardt 2017). In many anti-corruption protests such as the Wukan event, high-rank CCP officials were eventually sent to converse with the protesters so as to re-establish the public's trust in the government. Second, in a large and diverse country like China, most political and economic decisions are decentralized to local governments. Allowing for collective action from bottom up provides information for the central government to correct policy oversights, gauge public sentiment, and monitoring local officials (Lorentzen, 2017). Obvious cases are the protests for environmental protection or those for unpaid pensions. Finally, some collective action events such as strikes may be welfare improving and even productivity enhancing if they result in fairer competition, redistribution and better working conditions (Cai, 2010; O'Brien and Li, 2006; Su and He, 2010; Lee and Zhang, 2013).

Certainly, allowing for collective action exposes the Chinese government to the risk of regime instability. First, widely publicized protests about serious issues may undermine the legitimacy (or trust) of the CCP. Second, large-scaled collective action events may disturb social orders, disrupt policy implementation, and even cause rebellion. Finally, protests may exhibit a scale shift that leads to revolts intending to overthrow the government. When wide spread, initially small protests can evolve into better-organized political action and social movement that are fatal to the regime. Therefore, whether to repress collective action or not is a tricky task faced by the Chinese government, which must balance the above benefits and costs according to the cause, scale, and contagious nature of an event.

From the perspective of an individual citizen, the cost of organizing political collective action is very high. Protesters are likely to be arrested and punished when the government represses the protest. Even if the government chooses concession, organizers and active participants can face the risk of losing jobs or being under close watch by the government. Another cost is the vulnerability of coordination because of the lack of public information to reveal preferences, communicate beliefs, and coordinate actions.

2.3 Social Media and Protests

The ability to organize collective action relies crucially on information. This is an important reason why authoritarian governments tend to strictly control the media and suppress public political communication. The popularity of social media in China greatly facilitates information flow and generates new information.

Social media have the potential to spur collective action events such as protests and strikes in a number of ways. First, social media help inform citizens of corruption, bad policies,

human rights, and other social problems. When more people are informed, collective action is more likely to occur. Second, the massive communication feature of social media aides in generating common knowledge that allows citizens to coordinate on certain focal points and mobilizes people to political action. Third, visibility on social media may attract attention from the central government who cares more about people's trust of the regime and discourage local governments' use of repression, which in turn increases peoples' incentives to participate in collection action.

Equally important, social media are likely to affect the spread and dynamics of protest events. Communication in social media is speedy and long distanced. These features are particularly relevant to the spread of initially isolated and small collective action events. When a protest or strike in a locality succeeds, people in other areas, even far away, may be inspired by its success and learn effective tactics to organize similar events (e.g. Little, 2016; Chan and Suen, 2016). Furthermore, with speedy communication, a group of people initiate a protest, and other groups facing similar problems have incentives to follow suit immediately as simultaneous events tend to receive more visibility, intervention from higher-level governments, and concession. For example, Chang and Hess (2018) document a wave of school teacher strikes and demonstrations spanning from the fall of 2014 to the spring of 2015. Starting from two inland provinces (Henan and Hubei), earlier participants heavily used social media— primarily Sina Weibo— to post images of protests and circulate tactics. According to interviews with teachers in this study, the location, timing, and format of demonstrations were carefully planned to attract the attention of provincial and national authorities, who care more about social unrest and would press local administrators to solve problems. This tactic was quickly adopted by teachers in other places, and the teacher strikes soon spread to the entire country. In our own interviews, China Labor Bulletin— a non-governmental organization that promotes workers' rights in China— shared a similar view that social media help break the geographic bounds of localized events and generate widespread, cross-regional protest (strike) waves that place greater pressure on subnational authorities to give in to protesters' demand.

However, information on social media may also be used by the authorities to identify and quench protests at an early stage. This is the surveillance function of social media in autocracies, as documented in Qin et al. (2017). Government officials can even disseminate propaganda through their accounts on social media to derail protest activities. One example is in May 2013, when citizens in Chengdu blogged on Sina Weibo, calling for a demonstration against the local government's decision to launch a large-scale plant that produces the potentially hazardous chemical compound P-Xylene (PX). These posts were picked up by the Chengdu government, which immediately took measures including making Saturday and Sunday of that week into working days and requiring students to be in school on those days. Consequently, the demonstration was limited to a small group of people.

3 Theoretical Framework

In this section, we present a regime-official-citizen theoretical framework to guide the empirical analysis, leaving a formal model to the online appendix. Fitting the empirical setting, our theoretical analysis focuses on small protests that aim to solve local problems such as reversing unfair policies, demanding higher compensation for lost lands, and reallocating factories that cause pollution. Although these problems are likely to affect the public trust of the government, they are not intended to overthrow the regime.

Central in this framework are two features: (1) information asymmetries and preference misalignment between the regime and local officials and (2) protests serving as a credible means for political communication between citizens and the regime. We assume that the regime cares about economic growth and political stability (legitimacy). The regime observes economic growth more precisely than it observes how much citizens are hurt by its negative side-effects, such as pollution or poor working conditions. Although it may not care about citizens' utility per se, the regime wishes to maintain regime legitimacy (citizen trust) and thus wants to reverse unpopular pro-growth policies with large negative effects.

Since China is a very large autocracy, the implementation of policies relies on local officials who care about their own private benefits. The regime faces the formidable task of monitoring local officials's implementation of unpopular policies that have a substantial and uncertain negative impact on regime legitimacy. To properly gauge this impact, the regime needs to acquire information about local conditions and citizens' attitude towards specific policies and their beliefs about the regime. However, local officials have strong incentives to silence the media and distort negative information reported upwards to the central leaders. The resulting skewed observability creates strong incentives for officials to implement pro-growth policies but not to internalize the social cost of these policies imposed on citizens.

One important means for citizens to transmit political information is through protests and demonstrations, which incur high costs and thus convey informative signals to the public and the central leaders. Using protests to gauge public sentiment and collect bottom-up information has a long tradition in the Chinese history (citations) and has been particularly emphasized by the CCP as part of the so-called "Mass Line" policy (citations). This has the additional benefit that the regime can encourage local officials to implement ex ante more risky pro-growth policies because citizen protests will aid the correct policy oversight ex post.

However, there is a trade-off because the protests may become widespread, shifting its focus from local to central leaders and threatening regime legitimacy. Revolving around these two features, our theory not only delivers testable predictions regarding how social media affect the incidence and spread of protests, but also generates implications for the effect of social media on political accountability and the regime's media control strategy in autocracies.

In our theory, how social media affect protests hinges critically on their effect on event

visibility. Without observing the occurrence of a protest, the regime is not informed of the social costs of unpopular policies and will not reverse them to avoid citizens' welfare loss. This, in turn, discourages citizens from protesting. Moreover, the lack of information about protests in other regions (or industries), citizens are less inspired to participate in protests because an isolated protest is less likely to trigger intervention from the regime than simultaneous protests. Unlike traditional media (e.g., newspapers and broadcasters) whose production of information is relatively easy to control, social media have millions of information producers and are thus difficult for the local officials to silence. Unlike other communication tools (e.g., telephone), social media enable instant and wide reach of a message and public discussion of an issue. Therefore, social media substantially increase the public visibility of social events, increasing both the speed and reach of protest information.

The increased protest visibility makes it more likely that the regime will reverse unpopular policies, which in turn increases citizens' incentives to protest. Furthermore, social media generates common knowledge among a vast population, which enables citizens from different regions to coordinate and synchronize their protests so as to generate more consequential outcomes that lead to concessions.

Social media raises the stakes in the regime's trade-off between using information from local protests and stopping them from becoming widespread. On one hand, visibility of protests on social media helps the regime to identify particularly harmful policies and reverse them. Therefore, social media increases the information value of protests. On the other hand, information flows over social media are fast and wide, increasing the risk of widespread protests. If the informational value of protests outweigh the threat to regime legitimacy due to spread of protests or the regime cannot censor social media effectively, our theory produces the following main hypotheses that we test empirically.

Hypothesis 1 *(1a) The use of social media increases the incidence of protests.*

(1b) Information flow over social media increases the spread of protests across regions.

Next, we enrich the model to allow officials to take preemptive action. Local officials have incentives to preempt protests in their administration since protests that are visible to the regime reduce their chance of being promoted. Given that preempting protests incurs high costs, local officials' preemptive action tends to target large-scale protests. In this case, the observed protests are truncated from above. Social media increases the official's incentives to preempt protests, by increasing both the visibility of protests, and hence their costs to the official, and the ability to surveil and identify planned protests. In consequence, the use of social media leads to the following result.

Hypothesis 2 *In the presence of surveillance and preemption, the use of social media will increase the incidence of small protests, while the effects on the incidence of large protests and the total number of protest are ambiguous.*

We now discuss different mechanisms by which social media may affect how protests spread from one city to another. So far, we have stressed the mechanism of policy complementarity—information flow over social media increases public visibility of events so that citizens in one region can organize protests immediately following protests in other regions mechanism in order to generate more impactful policy implications for the regime. This mechanism delivers several testable predictions. First, it implies that spread tends to take place among small protests because large protests themselves would have triggered the response of the regime and the value of action complementarity between protests is small. Second, it predicts that the spread of events occurs mostly within the same event category that has similar policy implications because the complementarity is specific to a certain policy. Finally, it is unlikely to generate a long-lasting effect on the spread of protests because the policy effect due to the complementarity between two protests is weak if the two events are temporally distant.

There are other mechanisms by which social media may affect the spread of protests. For example, it may allow citizens to learn effective tactics used in previous protests. This mechanism tends to produce a persistent effect on the incidence of protests following visible protests in other regions because an effective protest tactic is unlikely to lose its effectiveness quickly and social media posts about protest tactics (if not censored at all) would remain accessible for citizens to learn for a long period. Moreover, the effect is likely to be widespread, because the widely discussed protest tactics are general and not confined within a geographical boundary and within an event category. Another mechanism is by learning the regime’s preference through the regime’s response to existing protests (e.g., repression vs. concession). In this case, social media information on protests met with concessions is likely to induce spread while information on protests met with repression is likely to have the opposite effect. The effects are likely to be persistent given that the regime’s preference is stable. Finally, a more straightforward mechanism is that protests are driven by citizens’ emotional reaction, which is intensified when citizens observe protests (and particularly repression of protests) from social media. This mechanism predicts that information on protests met with repression is more likely to trigger spreading protests, and that the effect tends to be short-lived but diffusive across regions and event categories.

To discriminate between the above mechanisms, we will test the following auxiliary hypotheses.

Hypothesis 3 *Social media will increase the spread of protests through four mechanisms: policy complementarity, learning effective protest tactics, learning about regime preferences, and emotional reaction. Their implications can be summarized as follows:*

(3a) Protest complementarity: short-run, within category, medium sized protests.

(3b) Learning about effective tactics: long run, across categories.

(3c) Learning about regime preferences: long-run, within category, concession spreads more.

(3d) Anger and emotional reactions: short-run, repression spreads more

Regardless of the above mechanisms, the popularity of social media imposes a challenge to the regime. On one hand, the regime benefits from better vertical communication with citizens who can inform the regime of the social costs of policies more effectively. On the other hand, social media facilitate horizontal communication between citizens in different regions and have the potential to spur widespread protests that threaten regime legitimacy. Ideally, the regime wishes to restrict communication within each isolated region by shutting down information flow across regions. This information-segregation strategy would allow the regime to better achieve its dual politico-economic goal while keeping legitimacy unchecked. However, the cost of implementing such strategy is probably prohibitively large when open platforms such as social media become the tool of communication. The regime has to use more aggressive strategy to censor social media information about protests and even shut down social media when it perceives a substantial negative effect of free information flow on regime legitimacy.

4 Data

We assemble a unique dataset combining detailed information about thousands of collective action events from 2006 to 2017 together with posts published on Sina Weibo over the period of 2009-2013. We now describe how this data was collected and describe it briefly.

4.1 Protest and strike data

Data about protests and strikes in China are not available from any official sources and media coverage of these kinds of events in mainland China is rather limited. Hence, our collection of data has to rely on sources outside of mainland China. Below, we explain how we collect these event data, and then provide a description of the data.

The data on protests are collected manually from the website of Radio Free Asia (RFA), which is a private nonprofit international broadcasting corporation based in Washington DC. We focus on the news content published in Chinese. The news reported on the RFA website comes from news collected and written by RFA’s special correspondents, as well as media outlets in mainland China, Hong Kong, Taiwan and Western media outlets such as the New York Times and the BBC. News produced by the RFA’s special correspondents are widely used by Chinese news portals outside of mainland China. To the best of our knowledge, the RFA website is the most comprehensive and well-structured data source for protest events in mainland China. We searched key words related to “protest” and “demonstration” (in Chinese) on the RFA website and downloaded the relevant news reports. Several research assistants were hired to first verify that the news is indeed about events of interest and to purge duplicate information. Then, they extracted relevant information from each news report and coded the date, location, cause, and scale (number of participants) of each event.

During the period of July 2006 and December 2013, there are in total 1172 protest events

in our dataset. As shown in Table 1, most protests concern governments (policy, police and court, and housing and land reforms) and livelihood issues (employment, environment, and health). In terms of scale, more than 70% of the events involve 100-1000s people. Geographically, these protest events span 272 prefectures, slightly over 80% of all prefectures in China. The right panel of Figure 1 plots the density distribution of protests across prefectures. It shows while the distribution is right-skewed, many locations experience over ten events.

We collect data on strikes mostly from the China Labor Bulletin (CLB), a non-governmental organization based in Hong Kong. Since its start in 1994, CLB has evolved into an influential organization that supports the development of trade unions in China and the enforcement of China’s labor laws. It has endeavored to collect data about strikes that take place in Mainland China since 2007. The CLB website releases its complete collection of strikes since 2011, with detailed information on the timing, location, employers involved, industry, scale, worker action, and government responses for each event. For strikes before 2011, we extracted relevant data from the annual reports published by CLB and supplemented them with data collected from Boxun, a widely-read Chinese website that is based in the US and specializes in political news.

According to interviews with CLB, their data collection before 2011 relies mostly on overseas Chinese media outlets and occasionally on information provided by labor movement activists in China. Starting from 2011, CLB also collected information on strikes by searching over social media and following accounts on Sina Weibo that specialized in labor dispute settlements and trade union activities in Mainland China. During the period of 2013-2016, CLB also collected strike information from Wickedonna, a mass-event-tracing blog that gathered information on mass demonstrations in China. After the arrest of the founders of Wickedonna by the Chinese government in June 2016 this source stopped updating its data. The CLB instead extended its data collection to other Chinese social media such as WeChat. While admitting that social media help data collection, especially for small-scale events that are less likely to be covered by traditional media, CLB expressed little concern about data quality being affected by censorship. One director of CLB said that strikes in China are typically not regarded as sensitive by the Chinese government unless a strike evolves into a violent event that threatens public security. This suggests that a strike of significant size, once reported, will become known on social media.

During the period between January 2007 and December 2013, there are in total 1551 strikes in our dataset. As shown in Table 2, strikes occur in a wide range of economic sectors, with concentration in manufacturing and transportation (including taxi) industries. More than 60% of the strikes involve more than 100 people. The most common cause of strikes is to demand payment of wage arrears. Geographically, these strikes cover 283 prefectures, approximately 85% of all prefectures in China. Figure 1 (the left panel) shows that the distribution of strikes across prefectures is less skewed than the distribution of protests.

4.2 Social media data

Our social media data was extracted from a database including 13.2 billion posts published on Sina Weibo from 2009 to 2013. The database was collected by Weibook Corp, which crawled the posts published by more than 200 million authentic active users. According to our estimates, the Weibook dataset contains approximately 95% of the total posts published on Sina Weibo before 2012. From this Weibook database, we obtain two datasets.

The first one is the Weibo penetration data, which counts the aggregate number of posts per prefecture and month, based on all 13.2 billion posts in the Weibook data. The penetration data will be used to measure how the use of Sina Weibo was spreading across time and prefectures.

The second data set contains individual microblog posts extracted from the 13.2 billion posts in the Weibook data. These posts are of two types: original posts and forwarding posts. The original posts are the 202 million blog posts that mention any of approximately 5,000 key words related to various social and political topics. The forwarding posts are the 133 million retweets of the original posts. We only have the direct forwards of the original post, not the forwards of these forwards. For each original and forwarding post, we obtain the textual content as well as the posting time, how many times it has been forwarded, and the location of its author.⁵ This data will be used to measure how information about protests and strikes spread, and more generally to measure the expansion of communication through social media across Chinese cities.

4.3 Basic statistics and correlations

We now describe the temporal and spatial variations of events (strikes and protests) and social media penetration, followed by a description of the blog post data. We then explain how we use forwards (retweets) of posts to measure information flows over social media, based on which we depict the pattern of social media communication across cities.

4.3.1 Spatial and temporal distribution

Figure 2 shows the geographical distribution of events and social media penetration. The upper panel shows the total number of protests and strikes by prefecture. These events are spread widely across China. For strikes, the developed coastal areas are over-represented, notably in Guangdong, Shanghai and Jiangsu, but a significant number of strikes occurred in some inland areas such as Chengdu and Chongqing. For protests, with Beijing being an outlier with 95 protests, coastal areas in Shanghai and Guangdong and inland areas in Hebei, Shaanxi, Chongqing and Sichuan witnessed frequent occurrences of protests. The bottom panel of Figure 2 shows the total number of blog posts across prefectures.

⁵Details about Weibook’s data collection strategy and our selection of keywords and extraction of posts can be found in Qin et al. (2017b).

Over time, there is a positive correlation between the incidence of protests/strikes and the Weibo penetration. Figure 3 shows the total number of protests and strikes per month along with the Weibo use per month. There were around five strikes per month in the period 2007-2010. The number of strikes rapidly rose to over 20 per month in 2013. The pattern for protests is similar, with around 8 protests per month until 2010, followed by a rapid increase to around 15 protests per month in 2013. The green dots show the number of Weibo posts per capita, which increased drastically after early 2010. The figure shows that the increase in protests coincides with the increase in Sina Weibo penetration while the increase in strikes takes place with a lag. This trend of increasing protests and strikes has not gone unnoticed. It has been commented on in numerous news sources including the BBC, CNN, the New York Times and the Washington Post.⁶ Of course, this relationship between social media and events may not be causal. For example, social media may only have affected the observability of the events rather than the underlying event frequency.

4.3.2 Posting and forwarding on protests and strikes

A first empirical question is whether social media posts about protests and strikes are allowed to exist and spread at all, given the pervasive censorship. Inspecting the content of individual posts, we find extensive coverage of protests and strikes on Sina Weibo from 2019 to 2013. From the original post dataset, we find 2.5 million posts mentioning keywords related to protest and 1.3 million posts mentioning keywords related strikes. We drew a random sample of 1000 posts in each category and read them manually. Around 30 percent of the posts are indeed about protests and strikes. Among them, 58 percent of the strike posts cover forthcoming or ongoing events, and 95% of the protest posts are about past events and general comments. As shown in Qin et al. (2017), these posts predict and identify the real-world protests and strikes. The post data contain a variable measuring the total number of times each post was forwarded. Based on this variable, we observe that the 3.8 million protest and strike posts are forwarded in 37 million retweets, which implies that on average, there are around ten forwards per original post. This average masks a huge dispersion: some posts are forwarded millions of times but others none. Conditional on being forwarded at least once, there are on average 50 forwards per strike post and 100 forwards per protest post.

We use the data on individual retweets of protests and strike posts to examine how quickly and widely news about these events spread, exploiting the time and user location of the original post and the retweet. Since we only have the direct forwards of the original post, our dataset contains around 3 million of the 37 million forwards. For protests and strikes, around 30% of these forwards occur within 1 hour of the original post and 80% of

⁶"Can China keep its workers happy as strikes and protests rise?", Mukul Devichand, BBC, 15 December 2011; "China on strike", James Griffiths, CNN, 30 March, 2016; "Strikes by Taxi Drivers Spread Across China", Andrew Jacobs, Jan. 14, 2015; "Strikes and workers" protests multiply in China, testing party authority", Simon Denyer, Washington Post, February 25, 2016.

the forwards within one day of the original post. Forwards within the first hour are more likely to be geographically close. After that, distance plays no role: the average distance between the user who posted the original post and the user who forwarded it is the same as the unconditional distance between users. In general, the forwarding data shows that information about protests and strikes on social media disperses rapidly and widely across China.

4.3.3 Using forwards to proxy information flows in social media

To study how social media affects the spread of protests and strikes across prefectures, we want to measure how information flows through social media across these prefectures. We will use forwards across all topics to measure how information flows across prefectures change as the use of Sina Weibo expands. A post being forwarded implies that someone has read the post and decided to forward it. Of course, many others may read the post without forwarding it, so our forwarding measure is a conservative measure of information spread.⁷ We will proxy the information flow from city i to city j by the number of posts by users in city i that are forwarded by users in city j . We use the 133 million forwards across all topics except the 3 million forwards of posts discussing protests or strikes.

We now describe the above measure of social media connectedness between Chinese cities, with a particular focus on cities with many strikes and protests. People in areas with many Weibo users are likely to write and read more microblog posts. However, connections also depend on how much people communicate, controlling for the average frequency of posting and forwarding. To investigate this, we plot the log number of forwards of posts from a location i by users in another locations j , relative to the average number of forwards of posts from location i and the average number of forwards by users in location j .⁸ We call this relative forwards. Figure 4 plots this for four of the prefectures with most strikes. The area in black is the source of the original posts. For instance, in the upper left panel of Figure 4, the black area is Shenzhen with 144 strikes. Geographical proximity clearly matters for forwarding. For example, areas close to Shenzhen are forwarding more of their posts than the average level. However, there is also many areas far from Shenzhen that are closely connected to it through the social media network.

One reason that geographically distant cities are closely connected is that they are close in other dimensions. Figure 5 shows that cities that are more strongly connected through the social media network are similar in several characteristics measured in 2008 (geographically, population size, population density, tertiary share, FDI, landline and internet penetration). People tend to form links with people who are similar to them, a network property called

⁷A much less conservative measure is to use followers, since many followers will not read each blog post. One of the most cited Twitter studies (Kwak et al., 2010) asks whether the number of followers or the number of retweets is a better measure of influence and settles on retweets.

⁸We regress the log number of forwards of post from i by users in j on fixed effects for locations i and j and plot the residual.

homophily. For this reason, information from other sources is likely correlated with information through the social media network. In addition, errors are likely to be correlated across cities that are more connected through the social media network. For example, because they share similar industrial characteristics, strikes are likely to erupt at similar points in time in areas which are tightly connected through the social media network, even if there was no such network. Therefore, one challenge in our empirical analysis is to overcome this homophily effect, which is one focus of our empirical design as will be discussed in detail.

5 Empirical analysis

We first describe our baseline specifications and then discuss potential identification problems and how to modify the basic specifications to address them.

5.1 Baseline specification

We analyze whether information on Sina Weibo affects the spread of protests and strikes across Chinese prefectures (cities) using a panel of prefecture at a daily frequency. The panel consists of 272 prefectures with at least one protest and 283 prefectures with at least one strike, for these respective outcomes.

Consider how information on social media affects the spread of protests across N cities. Let G_t be a $N \times N$ nonnegative matrix representing information flows with G_{it} being the i th row vector. Each element of the G_{it} vector, G_{ijt} , measures the number of people in city i at day t who are informed about any event y_{jt-1} in city j two days before t . We focus on effects within two days because social media are more likely than other informational mechanisms to affect the spread of events over long distances within a short time span. We assume a linear probability model:

$$\Pr(y_{it}) = \alpha y_{it-1} + G_{it} y_{t-1} + \gamma' x_{it} + \varepsilon_{it},$$

where the outcome y_{it} is affected by the lagged outcomes in other cities, weighted by the information flows between city i and the other cities, G_{it} , as well as an auto-regressive component, y_{it-1} , and some other variables, collected in x_{it} .

We care mainly about how information flows are affected by the reading of posts on Sina Weibo. To proxy for these information flows, we will use f_{ijt} , the number of forwards between cities i and j at time t , defined more precisely below:

$$G_{ijt} = \gamma c_{ij} + \beta * f_{ijt},$$

where c_{ij} is the number informed from other sources, depending on factors such as geograph-

ical distance. We will estimate the following baseline specification:

$$\Pr(y_{it}) = \alpha h(y_{it-1}) + \beta h\left(\sum_{i \neq j} f_{ijt-1} y_{jt-1}\right) + \gamma h\left(\sum_{i \neq j} d_{ij} y_{jt-1}\right) + \beta_0 w_{it} + \theta' x_{it} + \delta_t + \delta_i + \varepsilon_{it}. \quad (1)$$

Here, y_{it} is a binary event dummy, y_{jt-1} is the number of events in location j within two days prior to t , and the $h(\cdot)$ functions are transformation of y_{jt-1} weighted by various matrices as will be explained later. To control for the potential direct effect of Sina Weibo on events, the above specification includes a measure of Sina Weibo penetration, w_{it} , that equals the log number of Sina Weibo posts per capita +1. Furthermore, x_{it} is a set of controls including population, GDP, tertiary share, industrial share, and the number of cell phone users and landline users. Finally, δ_t and δ_i are time and prefecture fixed effects.

The key variables in (1) are the measure of information flow between cities and its functional transformation. For the information flow variable, to avoid the endogeneity, we use the cumulative forwards by users in i by posts from users in j on all subjects other than protests and strikes up until one week before day t (f_{ijt-1}) as the proxy. We exclude forwards of posts about protests and strikes and forwards less than one week before day t to avoid reverse causality that may arise from the possibility that people who are planning protests and strikes are more likely to forward posts about protests and strikes. As a first proxy for c_{ij} , we use d_{ij} , the inverse geographic distance between cities i and j . The weighting matrices F with the forwarding element f_{ijt-1} and D with the distance element d_{ij} are normalized to increase interpretability. F is arranged as T stacked $N \times N$ matrices F_t , and hence has $N \times T$ rows and N columns. We normalize these matrices so that the average sum of all elements in a row of a weighting matrix equals one.⁹ To address the issue of stationarity, we use the functions $h(x) = \ln(5x + 1)$ and $h(x) = x$, respectively. The identity function is simpler and more straightforward to interpret, but the process generates too much feedback to be stable whereas the log-transformation is sufficiently concave to ensure stationarity of the process.

There remain a number of econometric issues that we need to deal with. First, logistic and probit models are biased in rare events data (King and Zeng, 2001) and do not work well in panel data with a large set of fixed effects. Thus, we estimate a linear probability model that is immune from these problems. Second, our model includes location-fixed effects and lagged dependent variables. In general, the estimates in this type of model are inconsistent with T fixed (e.g. Hsiao 2014, Arrelano 2003, Baltagi 2005). In the current model, T is large and the bias is likely to be small. We will show that the bias is indeed small using Mont

⁹A common form of normalization is row-normalization by which all coefficients in a row sum to one. However, in our case this normalization would imply that all locations would be equally affected (on average) by protests in other locations, which is clearly incorrect (e.g. Elhorst, 2001, argues against row-normalization of distance matrices for analogous reasons). Our weighting maintains the relative magnitude between all elements in the weighting matrix. In a constant $N \times N$ matrix, such as the distance matrix, the total sum of all elements is N and in a $(N \times T) \times N$ matrix, such as the forwarding matrix, this sum is $N \times T$.

Carlo simulations.¹⁰ Third, we need to check whether the estimated process is stationary.¹¹ Fourth, consistency requires no serial autocorrelation in the errors. Serial correlated in the error term implies that ε_{it} will be correlated with y_{it-1} . We will test for serial autocorrelation in the first-differenced residuals. Finally, errors may be correlated across both time and spatial units. We use two-way clustering of errors in the time and spatial dimensions.

5.1.1 Baseline results

Columns I-IV of Table 3 show the results from estimating how protests and strikes spread across Chinese cities, using the specification (1). The first two columns use the identity function $h(x) = x$, while the following columns use $h(x) = \ln(5x + 1)$. Both protests and strikes are autoregressive processes. Having an event during the last two days significantly increases the probability of an event occurring in the same location. Strikes spread to nearby locations as the distance-weighted effect of lagged strikes is positive. However, this is not true for protests. The incidence of protests and strikes are both increasing in Weibo penetration.

Our main interest is in the coefficient on the variable with lagged events, weighted by the amount of forwarding of posts from city j by users from city i . This is significant and positive across specifications. The coefficient estimate is not much affected by the inclusion of controls; see columns II and IV.

To assess the bias in these regressions we run a number of Monte Carlo simulations. We first estimate the parameters $\alpha, \beta, \gamma, \delta_t$ and δ_i from a regression specified as in Equation 1 but for simplicity without Weibo penetration and controls. We then generate data using the estimated parameters, adjusted so that $\delta_t + \delta_i \geq 0$ and estimate the model on this data. We repeat this procedure 100 times. The bias is very small, the difference between the true and the mean estimated β is in the third value digit for both protests and strikes. Figure 6 shows the distributions of the t-statistics against the standard normal density.

Marginal effects To assess the magnitude of the estimated effects, consider first the case $h(x) = x$. The marginal effect of a strike at $y_{j,t-1}$ on the strike probability $\Pr(y_{it})$, through the forwards-weighted term, equals βf_{ijt-1} . Our weighting matrices, D , and F , are normalized so that the average row-sum equals one. Hence, β measures the average increase in strike probability if there was a strike in all locations on the previous day. Since the average column-sum of F_t is one, β also measures the expected total increase in strike probability across all locations at date t due to one strike at a random locality the previous day. For example, the estimate $\hat{\beta} = 0.14$ in column 2 of Panel A implies that a strike in a given location the

¹⁰If the bias was larger, one could address this issue by using the GMM estimators of Arrelano and Bond (1991) or Blundell and Bond (1998). However, instrumenting rare events, such as our protests and strikes, with lagged differences and levels is not likely to perform well.

¹¹Sufficient analytical conditions for stationarity are that the errors are not auto-correlated and that $|\beta\omega_{F_t,\max}| + |\gamma\omega_{D,\max}| < 1$, where $\omega_{D,\max}$ and $\omega_{F_t,\max}$ are the largest eigenvalue of the matrices D and F respectively (largest negative if $\gamma < 0$ or $\delta < 0$). We will evaluate this criterion and check whether the process is stationary in Monte Carlo simulations.

previous two days increased the expected number of locations with strikes this day by 0.14. Relative to the mean strike probability, this is an increase of 33 percent.¹² For protests, the estimate $\widehat{\beta} = 0.23$ implies an increase in the protest probability of 55 percent, relative to the sample average probability.

For the case $h(x) = \ln(5x + 1)$, the marginal effect equals $\beta s_t f_{ijt-1}$, where the scaling factor s_t equals 4.63 at the sample mean.¹³ So the coefficient in column IV of Panel A should be multiplied by 4.63 to be comparable to the estimate in column II. For example, the estimated coefficient of $\widehat{\beta} = 0.038$ in column IV of Panel A implies that a strike in a given location during the previous two days increased the expected number of locations with strikes on the current day by 0.18. Because $h(x) = \ln(5x + 1)$ is a concave function, β , the marginal impact of additional previous strikes is falling in the number of such strikes. Hence the average increase in strike probability if there was a strike in all locations on the previous day is considerably lower than 0.18.¹⁴

Stationarity and widespread protests. In our setting, stationarity of the protest and strike processes is not simply an econometric issue. Whether these processes are stable or exploding is likely to be a core concern of the top political leadership. Stationarity in a dynamic spatial panel data model depends on the parameters of the model as well as on the spatial weights matrix that determine the amount of feedback in the process. For the location's own autoregressive term and the distance weighted term, this feedback is constant over time (since α , γ and the distance matrix D are constant). However, increased use of social media increases the feedback because each individual row in the forwarding matrix, F , does not sum up to one. This implies that marginal effect of a change in y_{t-1} on the probability of a protest or strike differs across localities and time. In the linear model, the average effect at a particular date t equals $\beta \bar{f}_{t-1}$, where \bar{f}_{t-1} is the average row sum across locations that day. The maximum such row sum equals 10.5, an order of magnitude larger than the average row sum. In the linear model, this implies so much feedback that the process

¹²The mean of the strike incidence variable is .0015 and there are 282 other prefectures than the one where the first strike took place, $.14/(282*.0015)=0.33$. For protests, the corresponding number is $.23/(271*.0015465)=.55$.

¹³The scaling factor

$$s_t = \frac{5}{5 \sum_{i \neq j} f_{ijt-1} y_{jt-1} + 1}.$$

The sample mean value of

$$\sum_{i \neq j} f_{ijt-1} y_{jt-1} = 0.0134.$$

(spread_analysis_sumstat_strike_1.csv.)

¹⁴In this case, the scaling factor

$$s_t = \frac{5}{5 \sum_{i \neq j} f_{ijt-1} + 1} = \frac{5}{5 + 1} = 0.83$$

since the average row sum of f_{ijt-1} is one. Hence the increase is $\widehat{\beta} s_t = 0.038 * 0.3 = 0.032$.

is unstable.¹⁵

However, the actual data do not exhibit the explosive path suggested by the estimated linear model. This could be because the government observed that the process was exploding and stepped in and struck down on protests and strikes, but we do not have direct evidence of this. A more likely reason is that the model is mis-specified in that the marginal effects of protests and strikes are assumed linear in the number of cross-sectional events, in other words, in how wide spread protests are. The linearity assumption may be incorrect for several reasons. If the mechanism is through information about protest methods, then it is likely that the marginal value of information from an additional event is falling. If the mechanism is through information about the number of protests about an issue, then the incentive to organize an additional protest so as to increase government awareness is likely to be decreasing. It may also be the case that there is a limited number of areas where people are upset enough about a particular issue to potentially protest or strike if they see that other people protesting for this issue. Hence there is a cap to the total number of protests and the process will eventually be concave. We model this by using the concave function $h(x) = \ln(5x + 1)$. This function was chosen to be sufficiently concave to make the process stable.

Observability Obviously, Weibo-induced observability can create a spurious correlation between Weibo penetration and observed events. We will provide some evidence on the size of this bias by comparing effects for events which are likely to be more or less affected by observability, for example, by comparing effects for large vs. small events or events in inland vs. coastal regions. However, the measured effect of social media on the spread of events across prefectures is less likely to be affected by observability. The reason is that while observability increases the number of observed events across cities, it is unlikely to increase the probability that one observed event occurs just the day after an another, relative to other days.

We use Monte Carlo simulations to verify that we can consistently test our null hypothesis of no effects of social media even if event observability is increasing in Weibo penetration. In particular, we generate data using the estimated parameters from Equation 1 but without any spread of events through the social media network, that is we set $\beta = 0$. We assume that the probability of observing a simulated event increases linearly in log Weibo penetration, setting the size of the effect by assuming that the correlation between Weibo penetration and event incidence arises entirely because of increased observability.¹⁶ We then estimate

¹⁵Formally, sufficient conditions for stationarity are $|\beta\omega_{F_t, \max}| + |\gamma\omega_{D, \max}| < 1$, where $\omega_{D, \max}$ is the largest real characteristic root of the matrix D (largest negative if $\gamma < 0$) and $\omega_{F_t, \max}$ correspondingly for the F_t matrix (from Elhorst, 2014). The greatest characteristic roots z of a irreducible non-negative matrix A with maximum row-sum $R(A)$ satisfy $|z| < R(A)$. (Brauer and Gentry, 1970). Hence, the second criterion is fulfilled if $|\gamma R(D)| + |R(\beta F_t)| < 1$. It is clear that this condition is not fulfilled in our case with $h(x) = x$. The estimated process is explosive. We also verify this by simulating the process.

¹⁶The correlation between Weibo penetration and event incidence is estimated in section 6.3 and the results

the model in Equation 1. Figure 7 shows that no spurious correlation in spread is generated because social media increase observability. Hence, we can consistently test the hypothesis that $\beta = 0$ even if social media affect observability.

5.2 Extended specification

Another identification concern is that it is difficult to identify the role of social media in spreading events if information from other sources or shocks are correlated with the social media network. We showed in the previous section that the social media network exhibits homophily: cities that are similar in different ways are more connected through the social media network. Consequently, strikes and protests may spread more between cities that are closely connected in the social media network even though social media play no role, as these cities may have strong information flows through other channels as well as correlated shocks. Such shocks would make ε_{it} correlated with the lagged outcome in other cities y_{jt-1} . A large literature, pioneered by Manski (1993), has discussed this issue and related identification problems. Solutions include exploiting the network structure to identify instruments (e.g. Bramouille et al. 2009), of identifying instruments that are uncorrelated with the error terms and the network (e.g. Acemoglu et al., 2015, and König et al., 2017), matching (Aral et al., 2009) and explicit randomization (e.g. Bakshy et al., 2012).

We instead use time-series variation in the rapidly expanding network to identify effects. We will allow strikes and protest to spread through correlated channels and that errors are correlated with the network, but assume that these effects and correlations are constant during our short sample period. Although the average spread of strikes between cities may be related to the average network strength, the exact timing of the network expansion is unrelated to other changes in the spread of strikes between cities.

We explore three different ways of using time-series variation in the network. Our first solution is to use a model with interaction-fixed effects.

$$\Pr(y_{it}) = \alpha h(y_{it-1}) + \beta h\left(\sum_{i \neq j} f_{ijt-1} y_{jt-1}\right) + \sum_{i \neq j} c_{ij} y_{jt-1} + \beta_0 w_{it} + \theta' x_{it} + \delta_t + \delta_i + \varepsilon_{it}, \quad (2)$$

where the coefficient c_{ij} captures any time-invariant heterogeneity in event spread from location j to i . Our second solution is to use a weighting matrix of average forwards, allowing the coefficient on this term β_b to vary by half-year period, b ,

$$\Pr(y_{it}) = \alpha y_{it-1} + \beta_b h\left(\sum_{i \neq j} \bar{f}_{ij} y_{jt-1}\right) + \beta_0 w_{it} + \theta' x_{it} + \delta_t + \delta_i + \varepsilon_{it}. \quad (3)$$

Here, the element \bar{f}_{ij} is proportional to the average of forwards f_{ijt-1} over time and normalized so that the sum of all elements \bar{f}_{ij} equals N . Finally, we estimated a model of dyads of

are reported in Table 5.

locations with time-constant dyad-fixed coefficients $\beta_{i,j}$. However, Monte Carlo simulations showed that this third method produced strongly (Nickell) biased estimates and hence this approach was abandoned.

Note that in the above two specifications, β can be consistently estimated even though the errors are correlated with the network, as long as this correlation is time constant. For example, for the error process $\varepsilon_{it} = \kappa w y_{t-1} + v_{it}$, equation (2) is correctly specified and the coefficient $c_{ij} = \kappa w_{ij}$ is estimated consistently as T goes to infinity.

In Equation 3, any spurious correlation between the variable $\sum_{i \neq j} \bar{f}_{ij} y_{jt-1}$ and the error term can be estimated during the period before Sina Weibo was introduced. The effect of Sina Weibo is the difference between β_b estimated in the pre-period and the post-period. Note that using the average forwards, \bar{f}_{ij} , introduces a weighting error because $\beta_b \bar{f}_{ij}$ is different than the βf_{ijt-1} in the data generating process.

5.2.1 Results

We first estimate the specification that allows for interaction-fixed effects and identifies the coefficients through the time-series variation in the interactions, as described in Equation (2). The results are shown in columns V and VI of Table 3. The pure auto-regressive and the distance-weighted effects are not identified in the specification with interaction-fixed effects. The estimates of β remain highly significant and are, in the case of strikes, slightly larger than those in columns III and IV.

However, this slight increase in magnitude is likely to be caused by the so-called Nickell bias. Figure 8 shows the distribution of β -estimates from Monte Carlo simulations. The graphs to the left show results from estimations without interaction-fixed effects (Equation 1), the graphs to the right include interaction-fixed effects. As previously mentioned, the bias is very small in the graphs to the left. The graphs to the right, show a clear, albeit small, bias. The blue line shows the coefficients from the estimations without interaction-fixed effects (Equation 1) while the red line shows the mean coefficients from the simulated data. The green line shows the interactions-fixed effects estimate in the original (not simulated) data.¹⁷ The slightly larger estimates from the interaction-fixed effects model are explained by the Nickell bias. The bias-corrected estimates from the interaction-fixed effects model would be almost identical to those from the model without these effects.

We next use the specification that uses a weighting matrix of average forwards, allowing the coefficient on this term β_m to vary by half-year period, as described in Equation (3). Figure 9 shows the estimated coefficients. In the period before the start of Sina Weibo, protests and strikes were not spreading across cities that subsequently became tightly connected through the social media network. This provides reassuring evidence that the shocks other than Sina Weibo (e.g., information shocks through other channels) are unlikely to

¹⁷These estimates are slightly different from those in Column V of Table 3 because the simulation does not include the control for Weibo posts.

spread the events disproportionately across cities. However, starting from the second half of 2010, these events began to significantly spread across cities connected through social media.

We evaluate this model using Monte Carlo simulations. Specifically, we use the baseline model in Equation (1) with its estimated parameters to generate data. We then estimate the model with time-varying coefficients on the simulated data. Since the data generating process is different from the estimated model, we compare the implied increase in strike probability from the data generating process with the equivalent parameter based on the estimation.¹⁸ The top two graphs of Figure 10 show that this statistic is similar for the data generating process and the estimates. The bottom two graphs of Figure 10 show that the coefficient is zero for the period before 2010. There is no bias in the coefficient estimate for this period. The estimated coefficient levels off after 2012. The reason is that the marginal effect of the DGP diminishes because the scaling factor s_t drops as to the increasingly many strikes and protests reduce the marginal impact of additional events.

6 Additional results

6.1 Spread of protests and strikes after 2013

After 2013, there are clear signs that the regime in China has taken a tougher stance on protest, strikes and media freedom. During our sample period, 2006 to 2013, media freedom was at the average level during the period 1990-2016 for which Freedom House provides data. However, post 2013, media freedom in China has dropped to an historical low, comparable to the levels just after the Tianamen Square protests. We also see a clear shift in the response to strikes in our data. Before 2013, around one strike in three is met by concessions and one in ten with repression. After 2013, one in twenty strikes are met with concessions and one in six with repression. There are also signs that China is increasing surveillance of social media. Government procurement of media-based surveillance equipment has increased from around 150 units per year 2011-2013 to 413 units in 2016.¹⁹ In Qin et al. (2017) we show that social media surveillance of protests and strikes is very effective. However, it may still have taken some years for the regime to fully develop this capacity.

A key question is whether protests and strikes still spreading through the social media network in the new environment after 2013. We can answer this question using the model with time-varying coefficients on the average ties from the social media network under the

¹⁸The implied increase in probability from the DGP is

$$\beta h\left(\sum_{i \neq j} f_{ijt-1} y_{jt-1}\right) \quad (4)$$

with the equivalent parameter based on the estimation is

$$\hat{\beta}_m h\left(\sum_{i \neq j} \bar{f}_{ij} y_{jt-1}\right).$$

¹⁹Source: China Government Procurement Website.

assumption that the average ties in 2013 is a good proxy for average ties post 2013. This is likely the case. By 2013, the social media network was already extensive. In addition, the closeness in terms of geography, size and development that explain social media ties between cities do not change much over a few years. We can evaluate this assumption by investigating how much our measure of social media connectedness change in the last two years of our sample. A regression of $\log(\text{forwards})$ the last day of 2013 on the same variable the last day of 2012 yields a very high R2 of 0.96 (and a coefficient greater than one capturing the proportional increase in forwards).

Figure 11 shows that the spreading of strikes through social media peaked in 2013-2014. After that, this spreading was reduced and by 2016 it had completely stopped. A similar pattern is found for protests, although the peak was earlier and the spread was positive even in 2016. The reason for the reduced spread of protests could, for example, be increased censoring or self-censoring on social media, the lower share of successful examples met with concessions rather than repression. We do not have the data coverage of Weibo posts after 2013, so we cannot evaluate these channels. Perhaps as a consequence of the reduced spread and tougher stance of the regime, the number of protests and strikes have fallen, from a peak at around 40 per month in 2014 and 2015 to around 20 per months in 2017.

6.2 Mechanisms and heterogeneity

We will now investigate the mechanisms and assumptions of our model by studying heterogeneous effects. We have discussed three different mechanisms whereby protests and strikes may spread: learning about the likely government response, learning about efficient protest tactics and policy complementarity. These mechanisms have different empirical implications. Learning about the likelihood of concessions implies that events met with concessions will spread whereas those not met with concessions will reduce the probability of future events. The policy complementarity mechanism implies that events of intermediate size are most likely to be affected. These are not in themselves large enough to induce a response but who will do so in the presence of other simultaneous protests for the same cause. This mechanism will induce protests with short time span and within cause, since they are targeting one policy action. Both learning mechanisms are likely to produce a more persistent increase in event probability since it is unlikely that the preference of the regime or the efficient protest tactic will change every day or even week. Hence this type of learning can affect protest probability immediately, but also in the medium to long run.

To investigate these predictions, we study heterogeneous responses. Our events can be partitioned into a number of categories, depending on the government response, but also size, strike industry, protest cause, etc. Let y_{it}^c is an indicator variable for whether an event of category c occurred, for example a strike in manufacturing. We will analyze whether events of a certain type spread more than other events using the specification

$$\Pr(y_{it}) = \alpha h(y_{it-1}) + \beta^c h\left(\sum_{i \neq j} f_{ijt-1} y_{jt-1}^c\right) + \gamma h\left(\sum_{i \neq j} d_{ij} y_{jt-1}\right) + \theta' x_{it} + \delta_t + \delta_i + \varepsilon_{it},$$

and report the coefficients β^c that capture how much an event in category c increase the probability of future events.

6.2.1 Do social media break industry and cause bounds?

Strikes are likely to spread within industries and protests within causes. For example, the wave of school teacher strikes in 2014 and 2015 documented by Chang and Hess spread within the same industry, and the spreading protests among farmers against corruption in Wukan were for the same cause. On the other hand, strikes and protests may also spread across industries and causes, for example, if people learn about general effective protest and strike tactics. We investigate this by splitting the protests into the categories listed in Table 1 (government policy and corruption, housing and land, etc.) and strikes by the industries listed in Table 2.

Table 4 shows the results from regressions where we have investigated separately the spread of events through social media within and across categories. For strikes we split events by industry and for protests by cause. The table shows that the spread through social media is 6-7 times higher within strike industry and protest cause than across, although the spread both within and across categories is statistically significant. These results consistent with the policy complementarity mechanism. Workers in the same industry tend to strike for similar appeals and will be encouraged to do so when they become aware of other strikes in the industry, because these simultaneous strikes create more pressure to respond. Similarly, if protests against a common cause, such as local corruption, becomes more widespread, then policy action to correct it becomes more likely.

While the effect of an individual event is smaller across that within categories, the total effect of spread across categories is larger. The reason is that there are many more events across all categories. The mean of the social media weighted events within and across categories is reported in the two last rows of the table. These are 8 and 10 times higher for strikes and industries in the first two columns. Hence, the total effect of protests spreading across protests causes is 80% higher than within. Similarly, for the total effect across strike industries is 10% larger than the total effect within industries. Social media are breaking the bounds of strikes industries and protest causes.

6.2.2 Do strikes met with concessions spread more?

Suppose that strikes spread because people update their beliefs about the likelihood of a concession or that they are inspired by successful strikes. Then strikes met with concessions should increase the probability of a strike in other locations more than strikes met with

repression. This is the prediction from the mechanism involving learning about the likely government response. On the hand, if people strike because they are angered by images of repression or because they learn about what tactics are more effective, then this will not be the case.

In the strike data, the CLB have recorded information about the response to collective action for around 1/3 of the observations. We code the government response as being repression if the description of the response contains the words *police*, *arrest*, *beaten*, *threatened* or *wounded*. We code the government response as concession if the description contains the words *mediation* or *negotiation*.

Figure 12 shows the results. It shows that strikes met with concession do not spread more through social media than strikes met with repression. This indicates that learning about the response is not the main reason for the spread of strikes through social media.

6.2.3 How long-lived are the effects?

This paper focuses on short run effects, since social media are more likely to spread events over long distances within a short time span. However, we now expand the effects window and study how effects persist over time.

Figure 13 shows the estimated spread of events occurring within 1-2 days, 3-7 days, 8-30 days 31-90 days and 91-180 days. Although the effect is largest just after the strike, strikes continue spread through social media (beta) for around one month. The effects for both protests and strikes die out after a month. The effect seems larger in the short run (within a week).

Hence, learning effects are less likely to be driving the effect, unless knowledge depreciates very quickly. The large effects in the very short run are consistent with a mechanism whereby protests are complements, for example, because simultaneous protests increase the probability of success or lower the probability of repression. They are also consistent with protests being driven by emotions aroused by observing other protests.

6.3 Social media penetration and event incidence

Table 5 shows that incidence of protests and strikes is increasing in Weibo posts in our prefecture-by-day panel. To explore this relationship more in detail, we collapse the data at the monthly level at which we have Weibo penetration data. We now estimate a linear probability model of the form

$$y_{im} = \alpha_i + \alpha_m + \beta_0 w_{im} + \beta' x_{im} + \varepsilon_{im},$$

where i indicates prefecture, m indicates month, y_{im} is an indicator variable for an event taking place in prefecture i at month m , w_{im} is Sina Weibo penetration, specifically, $\log(\#\text{weibo posts}/\text{capita} + 1)$, and x_{im} is a set of controls including the log of population, GDP, share of GDP from

agricultural-, industrial-, tertiary sectors, real foreign direct investment, university students, landline phone users, cell phone users and internet users. To study heterogenous responses, we use the difference-in-difference specification

$$y_{im}^c = \alpha_i + \alpha_m + \beta_0^c w_{im} + \beta' x_{im} + \varepsilon_{im},$$

where y_{im}^c is an indicator variable for whether an event of category c occurred, to test whether events of a certain type are more affected by social media.

Table 5 shows the results from the difference-in-differences regressions. The first three columns use an indicator variable for a strike as the dependent variable, whereas the last three columns use an indicator variable for protest as dependent variable. The number of protests and strikes are positively associated with the number of Weibo posts. Adding controls does not much affect the estimates. The magnitudes of the estimated correlations are large. In 2012, the variable Weibo posts has an average of 0.3. The estimated coefficients imply that an increase in Weibo posts by 0.3 is associated with an increase in the total number of protests by 13 per month ($.3 \cdot .157 \cdot 272$) and in the total number of strikes by 11 per month ($.3 \cdot .131 \cdot 283$).

If observability drives these correlations, then we would expect to see larger correlations where observability prior to social media was lowest. The RFA has better information in certain coastal provinces (Guangdong, Fujian, Zhejiang, Jiangsu) than inland. Table 6, columns I and VI, include an interaction term between our Weibo penetration variable and a dummy variable for areas other than these four coastal provinces. The effects on strike incidence is not significantly different in coastal and inland areas. However, the effect on protests is significantly larger in inland areas, consistent with the RFA having fewer sources there.

It is also likely that prior observability was worse for small events. In the remaining columns of Table 6, we study whether effects differ by the number of event participants (less than 100, in the hundreds, thousands or tens of thousands). We have too few observations for protests and strikes with more than ten thousand participants and for protests with less than a hundred participants. For the other event magnitudes, the estimated effects are significant and essentially proportional to the mean incidence. To conclude, we find some evidence that observability drive effects in inland areas, but little evidence that small events were affected more. Combining the results from above, we have Yes answer for both questions: the incidence of protests and strikes is increased by Weibo penetration and the observability is also increased by Weibo penetration.

6.3.1 What type of events increases more?

Finally, we investigate what types of events are more affected by Sina Weibo penetration. The graphs at the top of Figure 14 show the results from regressions where strikes are partitioned

by industry and protests by cause. The figure to the left shows that a high Weibo penetration is more strongly positively associated with the incidence of strikes in manufacturing than in other industries, although the differences are not statistically significant. The figure to the right shows that protest about government policy and corruption and housing and land are more associated with Weibo penetration than other causes, although the difference is only statistically significant for housing and land.

Previously, we argued that social media is likely to increase the share of protests that are met with concessions, because social media use increases the visibility of the protests, and protest visibility increases pressure of local governments to act and the awareness of upper-level governments to the problems. We will test whether strikes met with concessions increase more with Weibo penetration than those met with repression.

The graph at the bottom of Figure 14 shows the results. Higher Weibo penetration is more strongly associated with an increase in strikes met with concessions than with repression. This is consistent with Sina Weibo shifting the government response towards concessions. However, some caveats are in place. For example, it could be that the strikes met with repression are more observable absent social media.

7 Conclusion

Our study provides several insights, shedding new light on the debate regarding the political role of social media in autocracies. We find that the explosive increase in social media in China 2009-2013 is associated with a large and significant increase the incidence of protests and strikes. However, it is hard to completely rule out that at least part of this increase is driven by increased observability. We show that increased observability is less of a problem when estimating the effect of social media on the spread of protests and strikes. Using a new estimation method, based on time-series variation in the network, we find that information diffusion over the social media network has a sizeable and significant effect on the spread of both protests and strikes. The spread of events is fast and predominantly local – between events within the same category (e.g., province, cause, and industry), but spread across categories is still significant, albeit weaker. However, after 2016 the spreading of protests and strikes through social media is limited.

References

- [1] Acemoglu, Daron, Camilo Garcia-Jimeno, and James A. Robinson. "State capacity and economic development: A network approach." *American Economic Review* 105, no. 8 (2015): 2364-2409.

- [2] Acemoglu, Daron, Tarek A. Hassan, and Ahmed Tahoun. 2018. "The Power of the Street: Evidence from Egypt's Arab Spring." *The Review of Financial Studies* 31, no. 1(January 2018): 1-42
- [3] Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa, and Ekaterina Zhuravskaya. 2015. "Radio and the Rise of the Nazis in Prewar Germany." *The Quarterly Journal of Economics*, 130(4): 1885-1939, <https://doi.org/10.1093/qje/qjv030>
- [4] Aral, Sinan, Lev Muchnik, and Arun Sundararajan. "Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks." *Proceedings of the National Academy of Sciences* 106.51 (2009): 21544-21549.
- [5] Arellano, Manuel. *Panel data econometrics*. Oxford university press, 2003.
- [6] Arellano, Manuel, and Stephen Bond. "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations." *The review of economic studies* 58.2 (1991): 277-297.
- [7] Bakshy, Eytan, et al. "The role of social networks in information diffusion." *Proceedings of the 21st international conference on World Wide Web*. ACM, 2012.
- [8] Baltagi, Badi. *Econometric analysis of panel data*. John Wiley & Sons, 2008.
- [9] Bamman, David, Brendan O'Connor, and Noah A. Smith. 2012. "Censorship and Deletion Practices in Chinese Social Media." *First Monday* 17(3).
- [10] Barberà, Salvador, and Matthew O. Jackson. "A model of protests, revolution, and information." Matthew O., *A Model of Protests, Revolution, and Information* (March 1, 2018) (2018).
- [11] Battaglini, Marco. "Public protests and policy making." *The Quarterly Journal of Economics* 132.1 (2017): 485-549.
- [12] Blundell, Richard, and Stephen Bond. "Initial conditions and moment restrictions in dynamic panel data models." *Journal of econometrics* 87.1 (1998): 115-143.
- [13] Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin. "Identification of peer effects through social networks." *Journal of econometrics* 150.1 (2009): 41-55.
- [14] Brauer, Alfred, and Ivey C. Gentry. "Bounds for the greatest characteristic root of an irreducible nonnegative matrix." *Linear Algebra and its Applications* 8.2 (1974): 105-107.
- [15] Cai, Yongshun. *Collective resistance in China: Why popular protests succeed or fail*. Stanford University Press, 2010.

- [16] Cantoni, D., Yang, D. Y., Yuchtman, N., & Zhang, Y. J. (2017). Are protests games of strategic complements or substitutes? Experimental evidence from Hong Kong's democracy movement (No. w23110). National Bureau of Economic Research.
- [17] Chang, Shengping, and Steve Hess. "The Diffusion of Contention in Contemporary China: An investigation of the 2014–15 wave of teacher strikes." *Modern Asian Studies* 52.4 (2018): 1172-1193.
- [18] Chen, Heng, and Wing Suen. 2016. "Falling Dominoes: A Theory of Rare Events and Crisis Contagion." *American Economic Journal: Microeconomics*, 8(1):228-255
- [19] Chen, Xiaoyan and Peng Hwa Ang. 2011. "The Internet Police in China: Regulation, Scope and Myths." In *Online Society in China: Creating, Celebrating, and Instrumentalising the Online Carnival*, edited by David Kurt Herold and Peter Marolt, 40–52. New York: Routledge.
- [20] China Internet Network Information Center. 2011. "The 28nd Statistical Report on Internet Development in China." July 2011, Beijing.
- [21] China Internet Network Information Center. 2013. "The 32nd Statistical Report on Internet Development in China." January 2013, Beijing.
- [22] China Internet Network Information Center. 2014. "The 34th Statistical Report on Internet Development in China." January 2014, Beijing.
- [23] De Mesquita, Ethan Bueno. "Regime change and revolutionary entrepreneurs." *American Political Science Review* 104.3 (2010): 446-466.
- [24] Edmond, Chris. 2013. "Information Manipulation, Coordination, and Regime Change." *Review of Economic Studies* 80(4): 1422–1458.
- [25] Egorov, Georgy, Sergei Guriev, and Konstantin Sonin. 2009. "Why Resource-Poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data." *American Political Science Review* 103(4): 645–68.
- [26] Elhorst, J. Paul. "Dynamic models in space and time." *Geographical Analysis* 33.2 (2001): 119-140.
- [27] Elhorst, J. Paul. "Spatial panel data models." *Spatial econometrics*. Springer, Berlin, Heidelberg, 2014. 37-93.
- [28] Enikolopov, Ruben, Alexey Makarin, and Maria Petrova. 2018. "Social Media and Protest Participation: Evidence from Russia." Available at SSRN: <https://ssrn.com/abstract=2696236>.

- [29] Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya. 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review*, 111(7): 3253-85
- [30] Epstein, Gady. 2013. "China's Internet: A Giant Cage." *The Economist*, April 6. <http://www.economist.com/news/special-report/21574628-internet-was-expected-help-democratise-china-instead-it-has-enabled>.
- [31] Fu, King-wa, Chung-hong Chan, and Marie Chau. 2013. "Assessing Censorship on Microblogs in China: Discriminatory Keyword Analysis and the Real-Name Registration Policy." *IEEE Internet Computing* 17(3): 42-50.
- [32] Hsiao, Cheng. *Analysis of panel data*. No. 54. Cambridge university press, 2014.
- [33] International Telecommunication Union. 2013. "The World in 2013: ICT Facts and Figures." Geneva. <http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2013-e.pdf>.
- [34] 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
- [35] King, Gary, Jennifer Pan, and Margaret E Roberts. 2014. "Reverse-Engineering Censorship in China: Randomized Experimentation and Participant Observation." *Science* 345(6199): 1-10.
- [36] King, Gary, and Langche Zeng. "Logistic regression in rare events data." *Political analysis* 9.2 (2001): 137-163.
- [37] Kleinberg, Jon. 2006. "Complex Networks and Decentralized Search Algorithms." *Proceedings of the International Congress of Mathematicians (ICM)*.
- [38] König, Michael D., Dominic Rohner, Mathias Thoenig, and Fabrizio Zilibotti. "Networks in conflict: Theory and evidence from the great war of africa." *Econometrica* 85, no. 4 (2017): 1093-1132.
- [39] Kwak, Haewoon, et al. "What is Twitter, a social network or a news media?." *Proceedings of the 19th international conference on World wide web*. AcM, 2010.
- [40] Lee, Ching Kwan, and Yonghong Zhang. "The power of instability: Unraveling the microfoundations of bargained authoritarianism in China." *American Journal of Sociology* 118.6 (2013): 1475-1508.
- [41] Little, Andrew T. "Communication technology and protest." *The Journal of Politics* 78.1 (2016): 152-166.

- [42] Lorentzen, Peter. 2014. "China's Strategic Censorship." *American Journal of Political Science* 58(2): 402–414.
- [43] Lorentzen, Peter. "Designing contentious politics in post-1989 China." *Modern China* 43.5 (2017): 459-493.
- [44] Manski, C., 1993. Identification of endogenous social effects: The reflection problem. *Review of Economic Studies* 60 (3), 531–542.
- [45] Morozov, Evgeny. 2012. *The Net Delusion: The Dark Side of Internet Freedom*. Reprint edition. Public Affairs.
- [46] O'Brien, Kevin J., and Lianjiang Li. *Rightful resistance in rural China*. Cambridge University Press, 2006.
- [47] Passarelli, Francesco, and Guido Tabellini. "Emotions and political unrest." *Journal of Political Economy* 125.3 (2017): 903-946.
- [48] Reporters Without Borders. 2013. "2013 World Press Freedom Index: Dashed Hopes after Spring." <https://rsf.org/en/news/2013-world-press-freedom-index-dashed-hopes-after-spring>.
- [49] Shirky, Clay. 2011. "The Political Power of Social Media: Technology, the Public Sphere, and Political Change." *Foreign Affairs*, January/February.
- [50] Steinhardt, H. Christoph. 2017. "Discursive Accommodation: Popular Protest and Strategic Elite Communication in China." *European Political Science Review*, 9(4):539-560 <https://doi.org/10.1017/S1755773916000102>
- [51] Su, Yang, and Xin He. "Street as courtroom: state accommodation of labor protest in South China." *Law & Society Review* 44.1 (2010): 157-184.
- [52] Yanagizawa-Drott, David. 2014. "Propaganda and Conflict: Evidence from the Rwandan Genocide." *The Quarterly Journal of Economics*, 129(4): 1947-1994 <https://doi.org/10.1093/qje/qju020>.
- [53] Zhu, Tao, David Phipps, Adam Pridgen, Jedidiah R. Crandall, and Dan S. Wallach. 2013. "The Velocity of Censorship: High-Fidelity Detection of Microblog Post Deletions." arXiv preprint, arXiv:1303.0597.

Tables and Figures

Table 1: Protests by cause 2006-2013

cause	size					Total
	<100	100s	1,000s	10,000s	unknown	
government policy or	58	106	94	6	52	316
housing and land	23	113	80	5	62	283
labor	7	54	48	1	15	125
police and court	12	45	80	7	19	163
environment and health	3	36	51	6	26	122
firm and finance	0	20	12	3	5	40
ethnic and religion	9	21	8	3	9	50
school and education	5	17	12	0	5	39
social conflicts	0	5	9	2	1	17
international relations	0	0	1	0	1	2
transportation	0	0	0	0	1	1
human rights	2	2	2	0	5	11
unknown	0	0	0	1	2	3
Total	119	419	397	34	203	1,172

Table 2. Strikes by industry 2007-2013

industry	size					unknown	Total
	<100	100s	1,000s	>=10,000s			
manufacturing	98	400	216	5	72	791	
taxi	0	25	12	2	4	43	
transportation	163	116	15	0	38	332	
services	63	59	4	0	11	137	
education	36	32	10	0	6	84	
construction	49	29	12	1	3	94	
retail	15	13	2	0	5	35	
mining	1	8	3	0	3	15	
other	3	8	4	0	2	17	
unknown	0	0	3	0	0	3	
Total	428	690	281	8	144	1,551	

Table 3. Event spread across locations

	I	II	III	IV	V	VI
<i>Panel A: Dependent variable - Strike</i>						
y_{it-1}	0.029*** (0.007)	0.028*** (0.007)	0.016*** (0.004)	0.015*** (0.004)		
$h(\sum d_{ij} y_{jt-1})$	0.167*** (0.058)	0.159*** (0.053)	0.036*** (0.013)	0.034*** (0.012)		
$h(\sum f_{ijt-1} y_{jt-1})$	0.162*** (0.042)	0.140*** (0.038)	0.045*** (0.012)	0.038*** (0.011)	0.051*** (0.014)	0.045*** (0.013)
<i>Weibo posts</i>	0.008*** (0.003)	0.006*** (0.002)	0.008*** (0.003)	0.006*** (0.002)	0.007** (0.003)	0.005** (0.002)
Observations	665,630	665,630	665,630	665,630	665,630	665,630
R-squared	0.023	0.023	0.022	0.023	0.229	0.229
QPtest	0.13	0.15	0.10	0.16	0.00	0.00
<i>Panel B: Dependent variable – Protest</i>						
y_{it-1}	0.016** (0.007)	0.016** (0.007)	0.009** (0.004)	0.009** (0.004)		
$h(\sum d_{ij} y_{jt-1})$	-0.028 (0.036)	-0.027 (0.035)	-0.006 (0.008)	-0.006 (0.007)		
$h(\sum f_{ijt-1} y_{jt-1})$	0.242*** (0.059)	0.232*** (0.055)	0.073*** (0.016)	0.069*** (0.015)	0.072*** (0.019)	0.066*** (0.018)
<i>Weibo posts</i>	0.007*** (0.002)	0.006*** (0.002)	0.006*** (0.002)	0.005*** (0.002)	0.006** (0.002)	0.004** (0.002)
Observations	671,362	671,362	671,362	671,362	671,362	671,362
R-squared	0.017	0.017	0.017	0.017	0.222	0.223
QPtest	0.03	0.16	0.00	0.02	0.00	0.00
Controls	No	Yes	No	Yes	No	Yes
h function	x	x	ln(5x+1)	ln(5x+1)	ln(5x+1)	ln(5x+1)

Unit of observation: prefecture by day.

Standard errors two-way clustered by prefecture and month in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

Table 4. Event spread, within and across categories

	I strike	II protest
<i>Within</i>		
Number events 1-2 days prior, cumulative forwards weighted	0.0133*** (0.0032)	0.0174*** (0.0034)
Number events 1-2 days prior, distance weighted	0.0246*** (0.0069)	-0.0016 (0.0023)
<i>Across</i>		
Number events 1-2 days prior, cumulative forwards weighted	0.0018** (0.0007)	0.0030*** (0.0008)
Number events 1-2 days prior, distance weighted categories	0.0010 (0.0006)	-0.0004 (0.0005)
Observations	6,656,300	8,692,138
R-squared	0.0143	0.0077
Controls	Yes	Yes
Category	Industry	Cause
Mean within	0.0061	0.0056
Mean across	0.0502	0.0588

Unit of observation: prefecture by category by day.

Standard errors two-way clustered by prefecture and month in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Dependent variable event dummy

	I	II Strike	III	IV	V Protest	VI
Weibo posts	0.163*** (0.041)	0.131*** (0.042)	-0.123* (0.069)	0.199*** (0.043)	0.157*** (0.041)	0.131** (0.060)
F3 Weibo posts			0.060 (0.086)			-0.036 (0.057)
F6 Weibo posts			0.008 (0.050)			-0.021 (0.041)
L3 Weibo posts			-0.022 (0.091)			-0.028 (0.061)
L6 Weibo posts			0.291*** (0.066)			0.151*** (0.050)
Observations	22,142	21,938	18,800	22,083	21,622	18,825
R-squared	0.169	0.173	0.183	0.144	0.152	0.135
Controls	No	Yes	Yes	No	Yes	Yes
Leads (pvalue)			0.552			0.260
Lags (pvalue)			0.000			0.001
Lags+current (pvalue)			0.000			0.000

Unit of observation :prefecture by month. Standard errors two-way clustered by prefecture and month in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

Table 6: Heterogenous effects: dependent variable event dummy

	I	II	III	IV	V	VI	VII	VIII	IX	X
	Strikes					Protests				
VARIABLES										
Weibo posts	0.134*** (0.044)	0.075*** (0.023)	0.095** (0.038)	0.034** (0.017)	-0.001 (0.001)	0.110*** (0.029)	0.005 (0.010)	0.093*** (0.028)	0.069*** (0.021)	0.009 (0.008)
Weibo posts, inland	-0.012 (0.050)					0.229*** (0.045)				
Observations	21,938	21,938	21,938	21,938	21,938	21,622	21,622	21,622	21,622	21,622
R-squared	0.173	0.085	0.159	0.100	0.016	0.158	0.064	0.116	0.068	0.027
Mean	0.033	0.011	0.016	0.007	0.000	0.037	0.004	0.015	0.015	0.001
Mean, inland	0.024					0.031				
Size		<100	100s	1 000s	>=10 000s		<100	100s	1 000s	10 000s

All regressions include prefecture and running month fixed effects, as well as controls. Unit of observation :prefecture by month. Standard errors two-way clustered by prefecture and month in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

Figure 1. Distribution of collective action events across prefectuers





Figure 2. Strike count by prefecture 2007-2013 and protest count by prefecture 2006-2013




Figure 3. Number of events and Weibo posts per capita per month




Figure 4. Forward connections from prefectures with many strikes




Figure 5: Predictors of relative forwarding between cities




Figure 6. Monte Carlo Simulations: distribution of t-statistic of estimated parameter = true parameter




Figure 7. Monte Carlo simulations with observability driven by Weibo and no network spread effect




Figure 8. Monte Carlo Simulations: distribution of coefficients around DGP parameter values.




Figure 9. Time-varying coefficients and constant forwarding matrix




Figure 10. Monte Carlo simulations: Comparing average marginal effect from DGP and estimation




Figure 11. Time-varying coefficients and constant forwarding matrix post 2013




Figure 12. Strike spread by response




Figure 13. Spread effect, number of days prior to event




Figure 14. Effect of Weibo on events, by industry, cause, and government response

