The Political Economy of Social Media in China

Bei Qin, David Stromberg, and Yanhui Wu*

Abstract

This paper examines the role of Chinese social media in three areas: organizing collective action, surveillance of government officials, and propaganda. Our study is based on a data set of 13.2 billion blog posts to Sina Weibo – the most prominent Chinese microblogging platform – over the 2009-2013 period. We find millions of posts discussing explicit corruption allegations and collective action events, such as, protests, strikes, and demonstrations. Higher Sina Weibo use is associated with higher incidence of strikes and protests but not with large-scale conflicts and government-sanctioned demonstrations (e.g., anti-Japan protests). Social media are effective tools for surveillance: Sina Weibo content predicts collective action events the day before their occurrence, and corruption charges a year in advance. We estimate that our data contains 600,000 government-affiliated accounts that contribute 4 percent of all posts about politics and economics on Sina Weibo. The share of government microblog accounts is larger in areas with a higher level of internet censorship and where newspapers have a stronger pro-government bias. Overall, our findings suggest that the Chinese government regulates social media to balance threats to regime stability against the benefits of utilizing bottom-up information.

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

1 Introduction

In China, nearly half of the population has access to the internet, and two of every ten Chinese actively use Weibo (microblog in Chinese).[1] Every day, millions of blog posts are produced, exchanged, and commented upon. Many of these posts reach thousands or even millions of readers. However, Chinese microbloggers are subject to one of the greatest control efforts the world has seen, combining policing and punishing users, censoring websites and content, and posting propaganda. Indeed, Freedom House ranked China 186th of 199 countries on a scale of press freedom (Freedom House 2015). While Chinese microbloggers have set the social agenda in a handful of well-publicized cases,[2] it is unclear whether social media in China plays an important role in social and economic outcomes.

In this paper, we examine the political role of social media in China. We analyze a data set of 13.2 billion blog posts from Sina Weibo – the most prominent Chinese microblogging platform – over the 2009-2013 period. To the best of our knowledge, this is the largest sample of textual content obtained from Chinese social media. We combine statistical description, econometric analysis, and machine learning to document basic facts about these microblog posts and study whether they predict and affect real political and social outcomes.

We primarily study outcomes in three areas: organizing collective action, government surveillance of the public and local politicians, and propaganda. These areas are central to understanding the advantages and disadvantages of social media for an authoritarian regime (Egorov et al. 2009, Lorentzen, 2014). As starkly illustrated by the role played by social media in the Arab Spring uprisings against the regimes in Tunisia, Egypt, and Libya, social media can facilitate collective action, which may pose a threat to authoritarian regimes. However, the bottom-up information flows produced by the media allow a regime to identify and solve social problems before they become threatening to the regime. This surveillance role of the media – termed "the Mass Line" in Chinese politics – has been an important part of the propaganda policy of the Communist Party of China (hereafter, CPC). Government agencies across the country have invested heavily in software to track and analyze online activities, to gauge public opinion, and to contain threats before they spread.[3] Another important surveillance function of social media from the perspective of a central government is monitoring local governments and officials. In a large country such as China, many political and economic decisions are delegated to local governments. The implementation of policies and the evaluation of local politicians both require substantial efforts to collect information other than internal reports, which are likely to be distorted. A prominent recent example is the "online anti-corruption" campaign, which has become an important complement to

[1] According to the China Internet Network Information Center (2014), by the end of 2013, there were 618 million Chinese internet users, 281 million of whom were active microbloggers.

[2] Examples include the “My father is Li Gang” or “Guo Meimei” incidents that were widely covered in the Western media.

[3] The Economist, (2013) "China’s Internet: A Giant Cage," April 6th, 2013.

the "offline anti-corruption" campaigns launched by the Chinese central government.

The role of social media in authoritarian regimes is heatedly debated (e.g. Shirky, 2011, Morozov, 2012). Some believe that social media will play a positive role, increasing the public's ability to coordinate massive and rapid responses, their access to information and ability to engage in public speech. This will constrain authoritarian governments' ability to act without oversight. Others are more skeptical, either because they believe that low-cost activities on social media will act as a replacement for real-world action ("slacktivism"), or because they believe that authoritarian governments are becoming better at using these tools to suppress dissent. Empirical work on the subject is scant, and attempts to outline the effects of social media on political action are too often reduced to dueling anecdotes.

We attempt to address the role of social media in authoritarian regimes by providing large-scale evidence from China. We measure the amount of public speech on Weibo concerning sensitive topics such as collective action events, grievances about public policies, and explicit criticism of political leaders. We find literally millions of posts on these topics, for instance, 2.5 million posts about protests and 5.3 million on corruption. To characterize these posts, we list words describing hot topics. For example, the top three topics concerning protests are "demonstration," "sit-in," and "self-immolation"; the top three corruption-related topics are "embezzlement," "corrupt," and "government money."

We examine how informative the posts on Sina Weibo are in predicting real-world events. We analyze approximately 600 large collective action events that took place in Mainland China between 2010 and 2012, and 200 corruption cases involving high-ranking Chinese government or party leaders. We find that collective action events can be predicted based on posts the day before the event and that the corruption charges can be predicted based on posts published one year before the charge. In contrast, newspaper coverage of these topics does not predict the events. This result indicates that social media provide the public with information that is unavailable in traditional media. It also shows that social media represent potentially useful tools for government surveillance of protests and corruption. For instance, a government could use social media to identify collective action events before they take place. We explore machine learning techniques to further investigate this role of social media.

We use the staggered introduction of Sina Weibo across prefectures to analyze the potential effect of the platform on collective action events. We find that increased use of Sina Weibo is associated with a higher incidence of protests and strikes, but not with large-scale conflicts or anti-Japan events. The effect on protests and strikes refutes the slacktivism argument. However, effects on large-scale conflicts appear to be contained by the central government. Effects on anti-Japan events are unlikely because these events are supported by the Chinese government and can be organized through other channels.

Finally, we investigate government propaganda on Sina Weibo and its goals. We estimate the overall presence of government content on Sina Weibo by manu-

ally coding user's affiliation with the Chinese government in a random sample of users. We also use machine learning techniques to identify government-affiliated accounts based on the word patterns used in the posts. Both approaches reveal a tremendous amount of government content on Sina Weibo. We estimate that there are 600,000 government-affiliated users, whose posts accounts for 4% of all posts on political and economic issues in our database. Consistent with the view that propaganda is the primary goal of government-affiliated users on social media, we find that the predicted share of government-affiliated accounts is highly correlated with both internet censorship and pro-government bias in Chinese newspapers. Furthermore, we find that the share of government-affiliated accounts decreases in the GDP of a region, increases in a region's proximity to Beijing, and is larger in CPC strongholds.

Section 2 provides background information about the development and regulation of microblogging in China. Section 3 presents the data. In Sections 4 and 5, we examine the role of microblog coverage in organizing collective action and monitoring local politicians, respectively. Section 6 analyzes propaganda content published on Chinese social media. Section 7 concludes.

2 Background

Internet use in China has increased steadily since it became commercially available in the mid 1990s. By 2013, there were 618 million Chinese internet users, accounting for approximately 46% of the Chinese population. This rate is slightly higher than the global average of 39%.[4] Of these internet users, 281 million (45%) actively participated in microblogging.

Among all types of social media, microblogs provide the most extensive and vivid discussions of and debates on public issues. In 2012, for example, the two most popular Facebook-type social media platforms in China – Renren and Kaixin – covered the top 20 public events listed by the Public Opinion Monitoring Agency in 20 million posts, and Sina Weibo – the leading microblog site – covered the same events in more than 230 million posts.[5]

The popularity of microblogs is a recent phenomenon. In 2006, Chinese people became aware of Twitter; the next year, major Chinese counterparts - Fanfou, Digu, and Jiwai - were launched. However, the number of microbloggers grew slowly. After the Urumqi riots in July 2009, the Chinese government not only blocked Twitter and Facebook but also shut down most domestic microblogging services. The microblog market was essentially vacant until Sina Weibo appeared on the scene in August 2009, and NetEase, Sohu and Tencent followed suit in 2010. The number of microblog users surged from 63 million at the end of 2010 to 195 million by mid-2011.[6]

[4] China Internet Network Information Center (2014) and International Telecommunication Union (2013).

[5] Reports on the Online Public Opinion (2010-2013), published by the People's Daily Public Opinion Monitoring Agency.

[6] CNNIC (2011). The 28th Statistical Report on Internet Development in China.

Sina Weibo led the development of the Chinese microblogging market. It is a hybrid of Facebook and Twitter: up to 140 Chinese characters are allowed per tweet, and users can send private messages, embed pictures and videos, comment, and re-post. Because of its easy access and use, Sina Weibo soon became the most popular microblogging platform in China. By 2010, it had 50 million registered users, and this number doubled in 2011, reaching a peak of over 500 million at the end of 2012. Since 2013, Sina Weibo has lost ground to WeChat, a cellphone-based social networking service, but has remained an influential platform for shaping public opinion.[7]

The Chinese government strictly controls microblogs through censorship and policing. Censorship is regulated by the central CPC Propaganda Department and a number of national media control offices but is implemented largely by private service providers. With the exception of Tencent, all companies operating social media are registered in Beijing and are closely watched by the Beijing Network Information Office. Service providers are aware of the danger of deviating from the government's censorship policy. However, the Chinese social media market is highly competitive, and consumers' demand is elastic to censorship. Facing this dilemma, Chinese social media firms acknowledge that although they attempt to monitor the content posted by users on their platforms, they are unable to effectively control all content.[8]

Several recent studies examine censorship in China. The estimated extent of censorship of Sina Weibo ranges from 0.01% (Fu, 2013) to 12% (King et al, 2013) of all posts. King et al. (2013) find that the Chinese government allows criticism of officials and bureaucrats but prohibits information about collective action. Bamman (2012) and Fu (2013) find that internet censorship in China focuses on political and minority group issues. Zhu et al. (2013) find that the implementation of censorship is speedy: 30% of deletions occur within the first half hour and 90% within 24 hours.

In this paper, we study the content that is available on microblogs, rather than what is removed. This information crucially depends on self-censorship, that is, what people dare to blog about. Self-censorship, in turn, depends on government policing and punishment of users.

Tens of thousands of internet police officers and internet monitors operate at all levels of government.[9] While censorship is nationally regulated and implemented, local politicians may use their own internet police to pursue their own goals. For example, driven by career concern, they may suppress negative information about the regions under their administration, even if blogging about this information is tolerable or encouraged by the central government. Users who post undesired content may receive warnings or have their accounts shut down; they may be sent to "reeducation through labor" camps without a trial and even be imprisoned.[10] Although these targeted individuals represent

[7] The Microblog users dropped by 27.83 million and the utilization ratio dropped by 9.2 percentage points in 2013 (China Internet Network Information Center, 2014).

[8] See the 2015 annual report of Sina to the U.S. Security and Exchange Commission.

[9] Chen and Ang (2011).

[10] Freedom House (2012) reported that prison sentences often had a minimum of three years

a tiny fraction of the user population, harsh punishment discourages potential dissenters and encourages self-censorship. Personal punishments can occur only if the users are identified.

The Chinese government initially allowed users on Sina Weibo to post anonymously.[11] In March 2012, the media control authority launched a real-name reform, requiring users to reveal their identities to social media providers. However, three years later, service providers have yet to implement this reform in its entirety.

The CPC and the Chinese government maintain a large presence in the microblog sphere. Chinese governments at all levels have opened their own microblog accounts to participate in blog debates and discussions in an effort to steer public opinion. In 2012, Sina Weibo reported that approximately 50,000 accounts were operated by government offices or individual officials. It is widely believed that Chinese governments hired armies of professional internet commentators, nicknamed "the 50-cent party" because some are paid at a piece-rate of 50 cents per post.

On Sina Weibo, users provide content and the central government censors it. Consequently, we expect Sina Weibo to affect outcomes where these actors have a common interests. This might include addressing local problems. For example, local users have an incentive to write about local corruption because they know that the central government may then address it. A caveat is that these effects may be muted by local internet policing. Critiques of central government policies that are not key to the CPC may also be allowed. We would not expect social media posts to contain critical discussions of the CPC's fundamental ideology, major policies, or national leadership. Nor do we expect social media to have a large impact on large-scale collective action events, such as riots or conflicts between governments and the public.

3 Data

Our primary data, Sina Weibo posts, were collected by Weibook Corp. Since September 2009, this company has executed a massive data collection strategy to download the posts of active users. First, they identify users as authentic active persons based on the individual's information and interaction with other users. In total, they identify 200-300 active million users. Second, they categorize users into 6 tiers based on the number of followers. They download the microblogs of the top-tier users at least daily, and the download frequency decreases by tier, with the second and third tiers every 2-3 days and the lowest tier downloaded on a weekly basis.[12] For each post, we obtain the content, posting time, and

and sometimes as long as life imprisonment.

[11] Even if users do not provide their real names, they may be identified via their IP-addresses. However, IP-addresses can be hidden using services such as Tor (The Onion Router) and VPN (Virtual Private Network) services

[12] For users whose posts are downloaded once per day, some posts would be downloaded immediately after posting while other posts would be downloaded 24 hours after posting. This implies that, except for data that are immediately censored, the data include posts that

user information (including self-reported location).

In total, the data set contains 13.2 billion posts published from 2009 to 2013. For the 2009-2011 sub-period, we construct an independent measure of the volume of posts on Sina Weibo.[13] According to our estimation, the Weibook data contains approximately 95% of the posts published on Sina Weibo. As illustrated in Figure 1, the blue line indicates the number of posts per month included in the Weibook data, and the red line is our estimate of the total number of posts published on Sina Weibo.

From this Weibook database, we extract microblogs mentioning any of approximately 5,000 key words that are related to important social and political topics. The key words consist of two groups. The first group refers to categories of issues, including major political positions from the central to the village levels, names of top political leaders, social and economic issues (such as corruption, pollution, food and drug problems, disasters and accidents, and crimes), and collective action events (such as strikes, protests, petitions, and mass conflicts). Some words occur at a very high frequency. We collect only a 10-percent random sample of the posts mentioning these words. The second group of key words refers to specific events that we have recorded, including those in censorship directives issued by the Chinese media control authorities and a large number of massive events from 2009 to 2013. In total, our extracted data contain 202 million posts from 30.6 million different users.

To analyze word frequencies in the Chinese text, we use the Stanford Word Segmenter to segment the words in each microblog post. We remove stopwords, punctuations, URLs, usernames and non-Chinese characters except meaningful English abbreviations) from the text. We exclude words with more than 30 characters and words occurring fewer than 5 times. We obtain 3.2 million distinct words and 6.0 billion tokens (word occurrences).

4 Collective Action

Social media can be used to coordinate and organize collective action events. However, because of extensive policing and censorship, it is an open question whether they play such a role in China. To do so, the relevant content must first be present in social media. Second, social media coverage must precede, or at least be concurrent with, the events. Finally, social media entry is likely to be associated with more collective action events. We investigate these three conditions in turn.

We analyze approximately 600 large collective action events that took place in Mainland China between 2010 and 2012. The events were recorded by Radio Free Asia, a non-profit radio station based in Washington D.C. For most

are later censored.

[13] Using the Sina Weibo public API, we downloaded all posts containing the neutral words "ya" or "hei" during four 5-minutes intervals each day and then divide by the average share of posts that contain these words and the average share posts contained in these 5-minutes intervals in a day. We could not do this for later years because the public timeline API denied access.

events, we are able to identify the start date from news reports. We classify the collective action events into four categories, ranked by their sensitivity – the expected probability of triggering censorship. The first category contains the most sensitive events. These are social conflicts, including riots and massive conflicts between governments and the public. The second category contains protests, including street demonstrations and mass protests. The third category contains strikes, including strikes in factories and schools and among taxi drivers. The last category includes anti-Japan demonstrations, which have occurred periodically in China over the last two decades.

We select key words that identify posts about each event type and extract all posts that mention these key words from the entire Weibook data set. These key words are described in the appendix.

4.1 Content and Users

We first analyze whether there exists relevant coverage of collective action events on social media. There are reasons to believe that this coverage would be limited. It is well documented that Chinese internet users have been punished after posting information about protests and other collective action events (e.g. Freedom House 2012). In an ambitious study of social media censorship in China, King et al. (2014) examine four collective action events and conclude that posts on such events are censored: "Chinese people can write the most vitriolic blog posts about even the top Chinese leaders without fear of censorship, but if they write in support of or opposition to an ongoing protest—or even about a rally in favor of a popular policy or leader—they will be censored." (p. 891).

However, we find a large number of posts covering even the most sensitive collective action events based on our classification. In our data, we identify 382,000 posts in the conflict category and over 2.5 million posts in the protest category. To describe the content, we characterize the "hot topics" in posts about collective action. These topics are identified by words that are used more often in collective action posts than in the entire sample of posts.[14]

Table 1 presents the results. The words are ordered by statistical significance. For example, in the conflict category, "suppression" has the most abnormally high use. It is used 322,797 times in 382,232 posts. Note that the topic word ranking is not simply increasing in the frequency of the words. For example, "tear-gas bomb" is ranked above "government" because the latter word is more commonly used in general. Other topic words in this category include "police community," "violence," "revolt," and "opening fire."

To characterize these data further, we investigate a random sample of 1,000 posts for each category. We manually code whether and how the posts cover a particular type of event (see Table 2). The share of posts that actually cover the events ranges from 50.4 percent for the anti-Japan category to 31.2 percent for the strike category. To estimate the total number of posts in a given category,

[14] More precisely, we compare the frequency of each word in a given category with the overall frequency of this word in our data set, assuming that the words are independently drawn from a multinomial distribution, as in Kleinberg (2006).

one should multiply the shares by the total number of posts in that category. For example, the number of posts about ongoing protests in our data is approximately 48,000 (2.526 million times .019). Examples are listed below to convey a sense of our coding:

- "I saw hundreds of policemen armed with weapons. Fire was everywhere, after some gas containers were bombed." [Conflict, ongoing]
- "A big crowd is gathering in front of the government building, holding 'No Forced Demolition of House' signs." [Protest, ongoing]
- "The money from selling lands all went into the pockets of officials. They are nothing but gangsters. We have no choice but to rebel." [Protest, general]
- "Seriously? Taxi-drivers strike again!" [Strike, ongoing.]
- "Low wages, cheap labor. We make tons of Made-in-China, but receive little in return. Migrant workers, strike!" [Strike, general]
- "We will march towards the Japanese Embassy today. Gathering at the People's Square at 10 am. Anyone wanna join?" [Anti-Japan, forthcoming]

To investigate whether a regular user who posts this type of content is punished, we examine the subsequent posts of the users who posted about collective action events. We find that users who posted on these topics continued to post to a similar extent as other users, indicating that their accounts were not closed.[15]

Users also do not seem afraid of posting this type of content. Users who anticipate punishment may post about sensitive events through a separate Sina Weibo account, the IP address of which is hidden and that contains few posts that can be used to identify the user. Thus, we expect that the posts on sensitive topics come from user accounts with few posts in our data. However, the average number of posts from users who post on sensitive topics is not significantly lower than that from a randomly drawn comparison sample of users (drawn using the number of posts by each user as sampling weights).

We find some patterns in the data that indicate censorship or self-censorship. The scale of collective action events in our data is substantial enough to capture media attention, and bloggers would write extensively about these events if they were not concerned about being punished. For a significant number of events, we find no blogs posts. Panel A in Table 3 tabulates two binary indicators for the presence of an event (the column variables) and any post about the event (the row variable) at the prefecture-day level. For example, all but two of 44 anti-Japan events are covered presumably because these events are sanctioned by the government. For a prefecture-day wherein no anti-Japan event is recorded, there

[15] In our data, 16 percent of the posts are the last post published by a user. In the conflict and protest categories, the corresponding rates are 17 and 23 percent, respectively. The share of users who exit from our data within five or ten more posts is slightly higher in the full data (38 and 49 percent) than in the conflict and protest categories (33-34 and 41-42 percent).

are 112,486 (out of 398,900) posts discussing anti-Japan events. In general, the number of events that is not covered by any blog posts is increasing in sensitivity: strikes (8), protests (37), and conflict (60). This result is consistent with the view that the more sensitive an event, the more likely it is to be censored and self-censored. Because we do not find much evidence that people self-censor when posting about these topics, censorship is a more likely explanation.

The last two columns of Table 3 examine coal-mine accidents. We add these as a placebo used for checking our specification. These accidents are likely to be discussed on Sina Weibo, but not likely to be caused or predicted by Weibo posts. We obtain data on the locations and days of 253 coal mine accidents during the 2010-2012 period from the State Administration of Coal Mine Safety.[16] We search for ten word strings related to coal mine accidents in our dataset. On the day of a coal-mine accident and in the affected prefecture, there are on average three posts about the accident. On the day before an accident or on days without accidents, there is approximately one post about an accident.

4.2 Prediction and surveillance

A key question is whether collective action events are caused or amplified by social media posts. Such effects could arise, for example, because blog posts explicitly call for participation or because information about ongoing events can spur and coordinate participation. A necessary condition for blog posts to affect an event is that they precede, or are at least concurrent with, the event. This is also important for whether social media are effective for surveillance. If the government is going to use information from social media to take action to prevent a collective action event, it needs to know about the event in advance or at least concurrently.

We examine whether blog posts cover upcoming or ongoing events in a random sample of 1,000 posts. Compared with less sensitive events, more sensitive events (e.g., conflict and protests) receive more coverage in the form of general and retrospective comments (see Table 2). In the conflict category, 1 in 1,000 posts covers an upcoming event, and 11 cover ongoing events. For the strike and protest categories, these figures nearly double; for the anti-Japan category these numbers are even higher.

Next, we investigate whether social media activity predicts collective action events. The last two rows of Table 3 report the average number of posts about each event type published by users in the prefecture in which the event took place on the day of the event and on the day before. For example, large-scale conflicts are covered in only 6 posts on the event days, compared with 1897 posts per day for government-sanctioned anti-Japan events. The average number of posts is much higher on the day of and the day before a collective action events than on other days. This pattern holds for all types of collective action. Although only 3.3 posts discuss conflicts on the day before these events, this number is still much higher than on other days (.6). In contrast, while coal mine accidents

[16] http://media.chinasafety.gov.cn:8090/iSystem/shigumain.jsp

are discussed much more the day of the accident, they are not more discussed the day before the accident than other days.

Consequently, the number of microblog posts mentioning a collective action event is highly significant in identifying and predicting that event. Table 4 presents the results of a regression wherein the dependent variable is an indicator for whether an event of a particular type took place in a particular prefecture on a given day. The key independent variable is the number of Weibo posts from users in a prefecture that mention the event on the event day (panel A), or on the previous day (panel B). The regression includes the number of newspaper articles mentioning the event type at the prefecture-day level[17] and controls for the total number of Sina Weibo posts in the entire Weibook data set at the prefecture-month level. Standard errors are clustered by prefecture.

Table 4 indicates that microblogs are highly significant in predicting where and when collective action events take place. In contrast, newspaper coverage is uninformative. The fact that people begin discussing events before they happen indicates that Sina Weibo may be used to organize or at least to coordinate collective action events, and that the government may employ the platform to conduct surveillance of these events. This is the case even for the most sensitive "conflict" events, which pose a potential threat to the regime. One may suspect that the predictive power of the number of posts about an event one day before the event is spurious. To assess this possibility, we analyze coal-mine accidents, which are covered in the media only after an accident occurs (see the last column of Table 4). The effect of the posts on the same day of an accident is as strong as the effects on the collective action events; however, the effect of posts one day before an accident is statistically insignificant (and even negative).

We now discuss how useful social media information is for conducting surveillance of collective action events. The Chinese government presumably wants an indicator of the prefectures in which and the days on which collective action events are likely to occur. These instances can then be monitored more closely. In the terminology of information retrieval, such an indicator is called a classifier – the fraction of collective action events that are indicated by the classifier is called "recall," and the fraction of collective action events among the indicated events is called "precision." A risk-averse authoritarian government primarily seeks a classifier with high recall to avoid the cost of missing regime-threatening events, but it also desires high precision to reduce the cost of investigating a large number of false positives.

However, a classifier's dual goals - high recall and high precision – conflict with one another. To illustrate this conflict, we use anti-Japan events as an example because these events are not censored, and thus, our data provide the same information that the Chinese government would use. A simple classifier with high recall can be obtained by simply investigating all locations and days with any posts mentioning an event. The first two rows of Table 3 show that anti-Japan posts were published for 42 out of 44 anti-Japan events. On the

[17] Newspaper coverage is based on 62 general-interest newspapers that ever covered these events during the 2010-2012 period.

one hand, only two anti-Japan events fly under the social media radar, and this simple classifier has nearly perfect recall. On the other hand, such a classifier has poor precision, 42/(112,486+42). To identify all of these 42 events, the government would have to investigate posts for more than 112,000 prefecture-days.

Putting ourselves in the Chinese government's shoes, we explore machine learning methods to improve the classifier. In particular, we use a method that combines a Support Vector Machine (SVM) with regression analysis, similar to that developed by Sasaki et al. (2010) to detect earthquakes based on Twitter flows. This approach dramatically increases precision (by 97%) but at the cost of missing 2 anti-Japan events. Facing such a trade-off, a risk-averse government may be forced to use very costly methods to reduce the risk of missing regime-threatening events. Although it is possible that the Chinese government could develop better classification methods, improvement is challenging given that machine learning techniques have not been very successful in identifying other large-scale events, such as flu epidemics (Lazer et al., 2014).

An additional advantage of the above machine learning approach concerns data collection. For example, based on the predictions of the SVM method, we can identify collective action events that do not belong to the events in our original collection. Specifically, we investigate the 100 prefecture-days in which the predicted probability of a strike is the highest based on the social media data. By manual inspection of posts published for these prefecture-days, we are able to identify 23 strikes that are not included in our original data and 14 strikes that are.

4.3 Effects

Much has been written on the role of social media in triggering and facilitating the coordination of protests, for instance, during the Arab Spring. Nevertheless, there has been little systematic analysis of the impact of social media on protest activity.[18] Thus far, we have provided evidence that Chinese social media contain extensive and vivid discussions of collective action events and that blog posts about the events both precede and predict these events. While this suggests that social media posts have helped coordinate events, the evidence is far from conclusive. Note that collective action events may also be affected by blog posts that are entirely different from those discussed above. For example, collective action events may be sparked by information about injustices, corruption, ethnic tensions or collective action events in other areas.

We explore these potential types of effects using the staggered diffusion of Sina Weibo across prefectures. After Sina Weibo was introduced in September 2009, its use increased more rapidly in some areas, notably, those with a high number of mobile phone users, greater educational expenditures, and larger tertiary sector shares of GDP (Qin, 2013). Using a difference-in-difference design,

[18] An exception is Acemoglu et al. (2015) who test whether activity on Twitter predicts protests in Tahrir Square.

we can test whether the propensity for events such as conflicts and strikes increases with the use of Sina Weibo.

Table 5 reports the results of regressing a dummy for each type of event on Weibo penetration (defined as $log\ (\#posts/population + 1)$). The specifications with column heading (a) control for month and prefecture fixed effects, while those with column heading (b) further control for prefecture-by-year fixed effects. We find that strikes are consistently associated with higher Weibo penetration. Protests are positively correlated with Weibo penetration; this correlation is not significant after controlling for prefecture-by-year fixed effects. The coefficient changed little, but the standard errors doubled, indicating that the fixed effects remove much of the variation.

The above estimates suggest that the use of Sina Weibo increases strikes and perhaps protests but not conflicts or anti-Japan events. These results are sensible. People do not need social media to organize anti-Japan events, which are sanctioned by the government and can be organized via other information channels, such as newspapers. The strikes and protests are usually on a small scale and confined to a small region, and thus, they do not threaten the regime. By contrast, massive conflicts pose a potential threat to the regime, and thus, their coverage is suppressed.

Our findings suggest that Chinese social media platforms are unlikely to be threatening to the regime, which has contained the effect of Sina Weibo to protests and strikes, leaving massive-conflict events unaffected. The lack of an effect on anti-Japan demonstrations suggests that Sina Weibo is not effective in enhancing collective action that is supported by the government.

5 Monitoring Local Politicians

In this section, we investigate whether social media provide information relevant to holding local politicians accountable to higher-up politicians. Such an investigation requires that the relevant content be available and informative in identifying corrupt officials. We will first describe the content on Sina Weibo related to corruption and petitions. We then analyze 186 corruption cases involving high-ranking Chinese government or CPC leaders.[19] Finally, we examine whether the corruption posts were published before government action and whether they are informative in identifying politicians who were subsequently charged with corruption.

5.1 Content and Users

5.1.1 Petitions

Petitioning is an old procedure whereby people can air their grievances to local officials or even travel to Beijing to bring their issues to the attention of central

[19] The sources of the corruption data are the CPC Central Disciplinary Committee and Ministry of Supervision, news reports published by Xinhua News.

government leaders. Local politicians often dislike petitioning because they fear direct reprisals from the central government or because it negatively affects their careers. Some officials have been accused of intercepting petitioners and forcing them to return home or even imprisoning them in illicit detention facilities. For example, a report by Human Rights Watch from 2009 charged that numerous petitioners, including children, were detained; the report also documented several alleged cases of torture and mistreatment.[20] With the increased availability of social media, petitioners do not have to travel to Beijing; they can simply post about their issues on Sina Weibo, although they still face the risk of being identified and punished by local officials.

To document this phenomenon, we retrieve over 1 million posts that contain two key words that are widely used in petitioning. We conduct a topic analysis of the posts in this category. Table 6 shows that the topic words are mainly related to the practice of petitioning, such as "appealing for help" and "petitioners" or related to the punishment of petitioners, such as "reeducation through labor," "imprisoning," "police," and "specialized hospital."[21]

We randomly select 1,000 posts on petitions and manually inspect them. Only 93 posts are irrelevant to petitioning in China; 35 are about ongoing or upcoming petitions, 469 are about past petitions, and 406 are about petitioning in general. There is little evidence that users who post content about petitioning are punished. Users who post on this topic continue to be as active as other users, implying that their accounts are not closed. Petition posts are not generated from special accounts with few posts. Rather, they are generated by users who publish more posts than an average user.

In summary, there is a vivid microblog discussion of petitions, mostly concerning petitioning in general or past cases. This practice explains why the hot topic words in this category are about the petitioning process or punishment. However, we estimate that approximately 40,000 posts in our data concern upcoming or ongoing petitions.

5.1.2 Corruption

How much discussion of corruption involving political leaders should appear in Chinese microblogs is unclear a priori. On the one hand, the central government and social media users have a common interest in reducing corruption. On the other hand, corruption among local politicians reflects poorly on the CPC and may undermine its legitimacy. For this reason, the Chinese central government may wish to silence discussions of corruption. In addition, local politicians have strong incentives to suppress corruption coverage in their jurisdictions to avoid jeopardizing their careers. Consequently, users may self-censor to avoid being identified and punished by powerful local politicians.

[20] Human Rights Watch, "An Alleyway in Hell", Nov 12 2009
http://www.hrw.org/en/reports/2009/11/12/alleyway-hell-0

[21] A topic word that does not appear to belong to these categories is "Tang Hui." However, this is the name of a Chinese mother who was imprisoned in a labor camp for demanding justice for the rape, kidnapping and prostitution of her 11-year-old daughter.

To examine the coverage of corruption, we combine two types of microblog posts: those mentioning politicians or political positions and those mentioning corrupt behavior. For the first category, we retrieve posts that mention any major political positions at the central, provincial, prefectural, county, and village levels. We obtain over 11 million total posts in this category. Column I of Table 7 shows the number of posts covering each position. The table is sorted by Column II – the number of posts per office (e.g., 33 offices for provincial-level positions). Xi Jinping, the current president of China and the general secretary of the CPC, is the most discussed leader, with over 1.3 million posts mentioning his name, followed by Wen Jiabao, the former prime minister of China. In general, officials at higher levels are more extensively discussed, and executive positions are more covered than are the party secretaries.

For the second category, we search for words that are widely used to describe corrupt behavior, wrongdoing, and punishment of officials (e.g., bribery, cronyism, and misuse of power). We identify over 5.3 million posts in this category. The hot topic words in this category are "embezzlement," "corrupt," "government money," and "bribes" (see Table 6).

To characterize the corruption posts, we manually inspect 1,000 randomly selected posts. Most of these posts make general comments on corruption. Of the 419 posts that discuss specific corruption cases, 293 were written after the government had taken action. However, 126 posts discuss particular instances of corruption before government action. These 126 posts consist of two types of posts. One type targets specific government officials and are usually very short, as showed by the following examples.

- "XXX, the Party secretary of XXX village, misused the monetary transferred from the central government for low-income villagers to pay his family members and relatives."
- "XXX, the chief officer of XXX county, embezzled public money by awarding all major government project contracts to his brother's company. Even worse, he hired gangsters to stab people who reported his corruption to the upper-level government."

The other type conveys resentment of and anger toward certain corrupt officials. In most cases, these posts discuss positions and government divisions without specifying the names of the officials. Several examples are documented as follows.

- "The black market for government positions in XXX prefecture is rampant. The price is getting higher and higher, the top officials in this prefecture are becoming richer and richer, and corruption will be more and more severe because the buyers need to make sufficient money to cover their costs."
- "Without support from the prefecture party secretary and the vice governor, how dare these prefecture officials sell government positions? Crack down on the tigers!"

- "Billions of money went into the pockets of local officials and their business partners! President Xi, Premier Li, and Secretary Wang in the Central Discipline Inspection Department, do you read our microblogs? Can you hear our voice? Please eradicate these corrupt officials! Right now!"

Based on the share in the random sample, we estimate that our data contains approximately 668,000 posts that discuss specific instances of corruption before government action. This provides a wealth of information to upper-level governments that seek to hold local politicians accountable. It is clear that posts of this type are not censored by the central government. It also seems that people are not afraid of posting concrete corruption allegations implicating powerful local politicians. Further, we find that users who post about corruption do not disappear from our data and that corruption posts are not generated from special accounts with few posts. We even find that some posts explicitly criticize top national leaders, although these posts do not contain explicit corruption allegations. Such posts claim, for example, that democracy and social stability decreased under Hu Jintao's reign, that the campaign against Bo Xilai was initiated by Xi Jinping as part of a political fight, and that Wen Jiabao shifted capital to Wenzhou to help the children of some top leaders.

We now describe the distribution of corruption content in Sina Weibo across major political positions. We first identify the posts that actually discuss particular instances of corruption, rather than, for example, discussing government anti-corruption policy. Adopting a widely-used SVM classification algorithm,[22] we estimate the probability that each post discusses a specific case of corruption. In the sample of 1,000 coded posts, we use SVM to classify the 419 posts that discuss specific cases of corruption based on the frequencies of words used.[23] We classify each of the 1,000 posts using leave-one-out predictions. Precision is .82 (306 of the 374 posts about corruption are correct) and recall is .73 (306 of 419 corruption posts are correctly classified). To obtain the estimated classification probability, we estimate a probit regression of an indicator variable for the post being about corruption on the SVM classification parameter (Platt scaling[24]). We will use these probit-based estimates of the probability that a post is about a specific corruption case.

We then compute these probabilities for each of the 5.3 million posts in our corruption category. Column IV of Table 7 show the average probability of a post being in our corruption category and mentioning a political position. Using only the raw number of posts mentioning corruption would clearly be misleading. For example, we estimate that only a small fraction (approximately 10 percent) of the corruption posts that mention national leaders actually discuss specific corruption cases. Most of these posts simply report the government stance on

[22] Based on performance in other classification tasks, SVMs have been identified as one of the most efficient classification methods (Dumais et al., 1998, Joachims, 1998, Sebastiani, 2002).

[23] The word frequencies of in each post are computed after the pre-processing described at the end of Section 3. As input to the SVM, we use term-frequency inverse document frequencies. We use the software SVM-light by Joachims (1999).

[24] Platt (1999)

corruption. The share of posts that discuss specific cases of corruption is much higher for lower offices, ranging from 0.40% for village chiefs to 0.63% for village party secretaries. Column III shows the percentage of posts mentioning a leader position that also discuss specific corruption cases.

To obtain a broader measure of how people feel about their leaders, we also conduct a sentiment analysis of all posts mentioning these leaders. We adopt a simple sentiment word count approach, using the National Taiwan University Sentiment Dictionary (NTUSD), which contains a list of 2810 positive words and 8276 negative words. For each post, we subtract the number of negative words from the number of positive words. Column V of Table 7 provides the average sentiment of posts mentioning each political position.

Figure 2 plots the percentage of corruption posts in Column IV against the average sentiment in Column V. Sentiment is most positive for national politicians, who are also infrequently mentioned in connection with corruption. The result might be driven by censorship and self-censoring on issues involving top leaders. County and village party secretaries have the most negative sentiment and the largest share of corruption posts. These two types of officials are usually viewed as the most powerful low-ranked politicians who have good chance of being corrupt. Another view is that they are the most vulnerable officials in anti-corruption campaigns because they are at the bottom of the Chinese government hierarchy.

5.2 Prediction and Surveillance

We relate the social media coverage to 185 corruption cases involving high-ranking Chinese government or party leaders. Of these cases, 39 occur at the provincial level, 114 at the prefecture level and 32 at the county level. To investigate whether social media posts predict corruption charges, we regress a dummy variable for whether a politician is charged with corruption in a location and month on the number of posts in our corruption category in the prefecture 2-7 months before the charge. We control for the total number of Weibo posts in the location and month and include year fixed effects. Column II of Table 8 also controls for location fixed effects. The regression indicates show that the volume of corruption posts in an area predicts future corruption charges in that area.

We also examine corruption charges of individual officials, by inspecting posts that mention their names. We start with a larger sample of 200 corruption charges, including charges against 15 national politicians. For comparison, we construct a matched control sample of 480 politicians who were not charged with corruption. The control politicians hold similar political positions to and are located in geographically nearby areas to the charged politicians.

We count the number of posts mentioning each of these 680 politicians and the number of posts that mention both the politician and any word in our corruption category. We calculate the number of posts 2-7 months (as well as 12-23 months) before a corruption charge. Table 9 shows that corrupt officials are mentioned, on average, in 49 posts within 2-7 months before a corruption

charge, while non-corrupt officials are mentioned in 44.4 posts. However, corrupt officials appear much more frequently in posts that mention our corruption words (3.9 compared to .4). A similar pattern is found in posts published 12-23 months before a charge. Given the substantial difference in the number of posts, it is not surprising that corruption posts are highly predictive of corruption charges.

Table 10 presents the results of a regression of the corruption-charge indicator variable on the number of posts mentioning an official's name and corruption.[25] Columns II, IV and V include dummy variables for case indicators, which are assigned the same value for an official charged with corruption and his or her matched officials. Future corruption charges are strongly predicted by the number of posts mentioning corruption 2-7 and 12-23 months before the first legal action.

This result indicates that social media information could be useful for surveilling corruption. However, a significant number of corrupt officials fly under the social media radar. In particular, 133 corrupt officials were never mentioned in a corruption post in the two months before the first government action against them. From the perspective of the Chinese government, aiming to crack down on corruption, a simple rule would be to investigate all officials with at least one corruption post (two months before the charge). As shown in Table 11, in our sample, this rule would lead to the investigation of 192 officials, 67 of whom were later charged with corruption. This simple classifier has a recall of 67/200 and a precision of 67/192. To increase recall, one would have to examine features of the microblog posts that fall outside our corruption category.

In summary, we find a massive volume of posts discussing corruption in Chinese microblogs. These posts help identify the political positions, regions, times, and individuals involved in instances of corruption. The results show that for the Chinese government, improved monitoring of lower-level officials outweighs the negative publicity of corruption coverage. They also suggest that local politicians are unable to impose self-censorship on users.

6 Propaganda

In this section, we examine propaganda appearing on Sina Weibo by studying the content of posts generated from government-affiliated accounts. We use two approaches to estimate the government presence on Sina Weibo. On a small scale, we manually code posts published by randomly selected users; on a large scale, we use machine-learning techniques to learn the language patterns used by well-identified government users and thus predict which accounts are affiliated with the Chinese government. We then use the predicted share of government-affiliated accounts at the regional level to investigate the goals of these government-affiliated users.

[25] The regression also includes the number of posts mentioning only the official's name, but this variable is never significant and is not shown.

6.1 Volume

Most propaganda content on Sina Weibo is generated by government-affiliated users, including government departments, mass organizations such as schools and hospitals, and media, which are state-owned. To accurately measure the share of such government-affiliated users, we manually code a sample of 1,000 users randomly selected from our entire database of 30 million users. A research assistant read all posts published by a selected users. A user is classified as a government user if the posts explicitly reveal the user's identity or are mostly related to the activities of a government function; mass-organizations users are analogously coded. An account is classified as a media account if the posts reveal that the user is a media outlet or a division.

A first observation is that the names of many government-affiliated users do not reveal user identities. While some user names explicitly mention a government function or division (e.g., police bureau or anti-corruption office), some names do not (e.g., "Falling Flowers", "I am A Super Fan", or merely some combination of letters and numbers).

We estimate a substantial government presence on Sina Weibo. Table 12 displays the result based on the random sample of 1,000 users. There are 5 (or .5%) government users, implying that there are approximately 150,000 (with a standard deviation of 67,000) government users in our entire data set. The state-owned media and mass-organization users contribute an even larger share. In total, these three types of government-affiliated users comprise 2% – or 600,000 – users. In 2012, Sina Weibo reported that approximately 50,000 accounts on Sina Weibo were operated by government offices or individual officials. Our estimation shows that even when limited to the most restrictive definition of government user (excluding mass-organization and media users), this reported number substantially under-estimates the government presence on Sina Weibo.

Knowing the number of posts by each user, we can estimate the share of posts published by government-affiliated users. In total, we estimate that the government-affiliated accounts contribute 3.6% of all posts in our database (with a bootstrapped standard errors of 1.6 percent). This percentage is greater than the 2% share of government-affiliated users, because these users publish more posts than others. Note that the above estimates are restricted to our sample of posts that mention words related to political and economic issues. Because we do not include users who write on other topics, the total number of government-affiliated accounts on Sina Weibo is likely to be higher than our estimates. However, the share of government posts may be substantially lower on topics outside political and economic issues.

6.2 Identifying government-affiliation by language

The above approach is limited in its ability to estimate the share of government-affiliated users at the disaggregated level. To overcome this limitation, we use a linguistics approach to predict the probability that a user is affiliated with the government. We now restrict attention to the 5.6 million users with more

than 5 posts in our data. These users contribute more than three quarters of all posts. In essence, we construct a "training set" consisting of a large number of government-affiliated users identified by user names. We then estimate a SVM to identify these users based on word frequencies in their posts. Finally, we use these estimates to predict the probability that each of the 5.6 million users is government-affiliated.

The first step of this procedure is to assemble a training data set. For this purpose, we search all user names for characters that are usually part of government or office names. We then manually inspect the retrieved user names by reviewing their Sina Weibo web-pages. Similarly, we inspect the official accounts of the main parts of Chinese legal system – the People's Court and Procurate at all administrative levels. In this way, we identify 1,043 government accounts in our database. To obtain media accounts, we search for newspaper names from the comprehensive newspaper directory compiled by Qin et al. (2013).[26] Together with the 538 identified newspaper accounts, we obtain a sample of government and media users who published 72,691 posts. Our final "training" set consists of these users and a random sample of 29,234 other users.

We use the training set to estimate a model to identify government-affiliated accounts. In particular, we estimate an SVM to classify the government-affiliated users in the training set based on word frequencies.[27] Such an estimated model indicates what combination of word frequencies predict government-affiliated accounts. We use the estimated model to predict SVM classification parameters based on the word frequencies used by the 5.6 million users who published at least 5 posts in our data set. In a new random sample of 500 users, we then estimate a probit model of the probability of being a government account conditional on the SVM parameter,(Platt scaling).

Finally, we use this to compute the probability that each of the 5.6 million users in the current data is government-affiliated, based on the word frequencies in their posts. We compute the average of this probability in total, by province, and by prefecture. This provides us with a measure of the share of government-affiliated users across geographic regions.

At the national level, we estimate that 3.1 percent of the 5.6 million users are government-affiliated (with a standard error of .8 percent). This is higher than the 2 percent in the overall sample because government users contribute more posts and are hence more strongly represented in the sample of users with more than five posts. The estimated share posts by government affiliated users in this sample is 3.9 percent (with a standard deviation of 1.0 percent).

[26] These identified newspaper accounts include the official accounts representing newspapers, related accounts such as supplements, non-regular editions, articles from different offices and in different formats, and accounts created by editors or reporters who are associated with the newspapers.

[27] As in the corruption case above, the word frequencies of in each post are computed after the pre-processing described at the end of Section 3. As input to the SVM, we use term-frequency inverse document frequencies. We use the software SVM-light Joachims (1999).

6.3 Goals of Government Users

We use the geographically disaggregated estimates to test several hypotheses with regard to the goals of government users on social media. Government users may provide neutral information or propaganda. Propaganda is intended to influence peoples' beliefs and actions. This goal can be achieved by both removing and adding microblog content. In areas where the government perceives that the need for influence is high, we should observe more of both censorship and propaganda. Consequently, if government users largely deliver propaganda, we should observe a strong positive correlation between censorship and posts from government users. We should also observe a positive correlation between posts from government users and pro-government bias in traditional media, which are subject to greater government control. Conversely, these correlations should be absent if government users mainly provide neutral information. We test these hypotheses using a measure of censorship developed by Bamman et al. (2012) and a measure of bias in Chinese newspapers developed by Qin et al. (2013). The latter measure is based on nine content categories, including leader mentions, cites of the official CPC news agency, and coverage of stories that criticize the regime.

Informative propaganda may be more effective on audiences that share the message sender's view, while the effect of propaganda may be negative when the audience holds opposing views. This argument follows from a rational model of Bayesian persuasion (e.g., DellaVigna and Gentzkow, 2010) or from psychological theories of cognitive dissonance. Empirically, Adena et al. (2014) find that Nazi radio in the 1930s was most effective in places where anti-Semitism was historically high and had a negative effect on support for Nazi policies in places with historically low levels of anti-Semitism. Similarly, in a laboratory experiment, DellaVigna et al. (2014) identify that Serbian radio exposure causes anti-Serbian sentiment among Croats. If the Chinese regime believes in this argument, then we would expect to find more official accounts in CPC strongholds.

Propaganda is likely to reduce consumers' valuation of social media. To the extent that service providers can affect the amount of propaganda, we should see few official accounts in areas where the advertisement market is valuable and where competition for consumers is high. We lack direct measures of these factors, but they are positively related to local incomes and we will use GDP per capita as a proxy.

We test the above hypotheses using the geographical distribution of the share of government-affiliated users on Sina Weibo. Figures 3 and 4 plot the estimated share of government users against the share of deleted posts (from Bamman et al.) and the measure of media bias in the daily newspapers strictly controlled by the CPC (from Qin et al. 2013). Guangdong has the lowest share of government users (2.5 percent) whereas Ningxia and Gansu have the highest share (6 percent). The graph remains virtually unchanged up to a scale factor if we use the share of posts published by government users instead of the share of government users.

The estimated share of government users is strongly correlated with both the share of deleted posts and newspaper bias (the correlation coefficient is 0.7 in both cases). This evidence strongly indicates that censorship, newspaper bias, and official accounts on Sina Weibo are used for the same purpose, namely, for propaganda. Note that Tibet has more deleted posts than expected based on the estimated share of government users on Sina Weibo. Perhaps this is an indication that propaganda is not viewed as particularly effective in Tibet because of weaker underlying support for the Chinese central government.

We investigate how the use of government users in Sina Weibo covaries with CPC support and economic development at the prefecture level. We use GDP as a measure of economic development. We define a variable, CPC stronghold, which equals the share of counties in a prefecture through which CPC armies passed during the Long March (1933-1935) or that were part of a CPC Soviet before 1949.[28] Other areas have a history of Western influence, notably, the areas that were part of a Treaty Port controlled by Western powers during the 1840-1910 period (Jia 2014). We also include the distance to Beijing, latitude, longitude and population in the regression.

Table 13 displays the results. The estimated share of official accounts is significantly lower in areas with a high level of GDP and is higher in CPC strongholds. The latter result is consistent with the view that propaganda is more effective in areas where the audience shares the ideology of the senders. The estimated share of government users also appears higher in areas that are closer to Beijing and in areas that are more populous.

To sum up, we find a very large government presence on Sina Weibo. We estimate that 600,000 accounts in our data are government-affiliated and that they contribute 4 percent of all posts. Language is highly predictive of these government-affiliated accounts. We find a strong positive correlation between the share of government-affiliated users on Sina Weibo and both censoring and pro-government bias in newspapers. This is consistent with propaganda being the main goal of this content. Finally, we find that the share of government-affiliated accounts is declining in GDP, increasing in closeness to Beijing and higher in CPC strongholds.

7 Conclusions

We analyze a large data set of blog posts from the most prominent Chinese microblogging platform over the 2009-2013 period. Our main findings are as follows. We measure the amount of discussion Weibo concerning sensitive topics such as collective action events, grievances about public policies, and explicit criticism of political leaders. We find literally millions of posts on these topics, for instance, 2.5 million posts about protests and 5.3 million on corruption. To characterize these posts, we list words describing hot topics. For example, the

[28] If affected by the Long March or the CCP Soviet was located in the metropolitan center of the prefecture, CCPStronghold is coded as one.

top three topics concerning protests are "demonstration," "sit-in," and "self-immolation"; the top three corruption-related topics are "embezzlement," "corrupt," and "government money."

We examine how informative the posts on Sina Weibo are in predicting real-world events. We analyze approximately 600 large collective action events that took place in Mainland China between 2010 and 2012, and 200 corruption cases involving high-ranking Chinese government or party leaders. We find that collective action events can be predicted based on posts the day before the event and that the corruption charges can be predicted based on posts published one year before the charge. In contrast, newspaper coverage of these topics does not predict the events. This result indicates that social media provide the public with information that is unavailable in traditional media. It also shows that social media represent potentially useful tools for government surveillance of protests and corruption. For instance, a government could use social media to identify collective action events before they take place.

We use the staggered introduction of Sina Weibo across prefectures to analyze the potential effect of the platform on collective action events. We find that increased use of Sina Weibo is associated with a higher incidence of protests and strikes, but not with large-scale conflicts or anti-Japan events. The effect on protests and strikes refutes the slacktivism argument. However, effects on large-scale conflicts appear to be contained by the central government. Effects on anti-Japan events are unlikely because these events are supported by the Chinese government and can be organized through other channels.

Finally, we investigate government propaganda on Sina Weibo and its goals. We estimate the overall presence of government content on Sina Weibo by manually coding user's affiliation with the Chinese government in a random sample of users. We also use machine learning techniques to identify government-affiliated accounts based on the word patterns used in the posts. Both approaches reveal a tremendous amount of government content on Sina Weibo. We estimate that there are 600,000 government-affiliated users, whose posts accounts for 4% of all posts on political and economic issues in our database. Consistent with the view that propaganda is the primary goal of government-affiliated users on social media, we find that the predicted share of government-affiliated accounts is highly correlated with both internet censorship and pro-government bias in Chinese newspapers. Furthermore, we find that the share of government-affiliated accounts decreases in the GDP of a region, increases in a region's proximity to Beijing, and is larger in CPC strongholds.

Our findings suggest that social media platforms in China are likely to improve the accountability of local governments and unlikely to affect regime change. We find an enormous number of posts on corruption and wrongdoings among local government officials and bureaucrats and posts concerning local strikes and protests. In contrast, we find few posts concerning large-scale collective action events that may threaten the regime. We find that social media affect the frequency of local collective action events but not that of large-scale events. This likely reflects the Chinese government's trade-off between maintaining regime stability and utilizing bottom-up information.

References

[1] Acemoglu, Daron, Tarek A. Hassan, and Ahmed Tahoun. 2015. "The Power of the Street: Evidence from Egypt's Arab Spring". working paper

[2] Adena, Maja, Ruben Enikolopov, Veronica Santarosa, and Katia Zhuravskaya. 2014. "Radio and the Rise of Nazis in Pre-War Germany", forthcoming in Quarterly Journal of Economics

[3] Bamman, David, Brendan O'Connor, and Noah Smith. 2012. "Censorship and Deletion Practices in Chinese Social Media", First Monday 17.3.

[4] China Internet Network Information Center. 2014. "The 34th Statistical Report on Internet Development in China" January 2014, Beijing.

[5] China Internet Network Information Center. 2013. "The 32nd Statistical Report on Internet Development in China" January 2013, Beijing.

[6] China Internet Network Information Center. 2011. "The 28nd Statistical Report on Internet Development in China" July 2011, Beijing.

[7] Chen, Xiaoyan and Peng Hwa Ang. 2011. "Internet Police in China: Regulation, Scope and Myths". In Online Society in China: Creating, Celebrating, and Instrumentalising the Online Carnival, ed. David Herold and Peter Marolt. New York: Routledge pp. 40–52.

[8] DellaVigna, Stefano and Matthew Gentzkow, 2010. "Persuasion: Empirical Evidence," Annual Review of Economics, Annual Reviews, vol. 2(1), pages 643-669.

[9] DellaVigna, Stefano, Ruben Enikolopov, Vera Mironova, Maria Petrova and Ekaterina Zhuravskaya. 2014. "Cross-border media and nationalism: Evidence from Serbian radio in Croatia", AEJ: Applied Economics, forthcoming.

[10] Dumais, S., Platt, J., Heckerman, D., & Sahami, M.. 1998. "Inductive learning algorithms and representations for text categorization". Proceedings of the 7th international conference on information and knowledge management, 48-155. Retrieved May 28, 2007, from ACM Digital Library.

[11] 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 103.04: 645-668.

[12] Festinger, Leon.1962. A theory of cognitive dissonance. Vol. 2. Stanford university press.

[13] Freedom House. 2015. "2015 Freedom of the Press Data"https://freedomhouse.org/report-types/freedom-press

[14] Freedom House. 2012. “Freedom on the Net: China”https://freedomhouse.org/report/freedom-net/2012/china

[15] 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." Internet Computing, IEEE 17.3: 42-50.

[16] Human Rights Watch. 2009. "An Alleyway in Hell", Nov 12 2009, http://www.hrw.org/en/reports/2009/11/12/alleyway-hell-0

[17] 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

[18] Joachims, T. 1998. "Text categorization with Support Vector Machines: learning with many relevant features". 10th European Conference on Machine Learning, volume 1398 of Lecture Notes in Computer Science, 137-142. Berlin: Springer Verlag.

[19] Joachims, T.. 1999. "Making large-Scale SVM Learning Practical". Advances in Kernel Methods - Support Vector Learning, B. Schölkopf and C. Burges and A. Smola (ed.), MIT-Press.

[20] Kleinberg, Jon. 2006. “Complex Networks and Decentralized Search Algorithms.”, Proceedings of the International Congress of Mathematicians (ICM).

[21] 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(May)): 1-18

[22] King, Gary, Jennifer Pan, and Margaret E Roberts. 2014. “Reverse-Engineering Censorship in China: Randomized Experimentation and Participant Observation.” Science 345 (6199): 1-10.

[23] Lazer, D.; Kennedy, R.; King, G.; and Vespignani, A. 2014. The parable of Google flu: Traps in big data analysis. Science 343(6176):1203?1205.

[24] Lorentzen, Peter. 2014. "China’s Strategic Censorship." American Journal of Political Science 58.2 : 402-414.

[25] Ng, Jason Q.. 2015. "Politics, Rumors, and Ambiguity: Tracking Censorship on WeChat’s Public Accounts Platform", mimeo, University of Toronto.

[26] Platt, John. "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods." Advances in large margin classifiers 10.3 (1999): 61-74.

[27] Qin, Bei. 2013. “Chinese Microblogs and Drug Quality”, working paper.

[28] Qin, Bei, David Stromberg and Yanhui Wu. 2013. “Media Bias in China”

[29] Jia, Ruixue. 2014. "The Legacies of Forced Freedom: China’s Treaty Ports", Review of Economics and Statistics, Vol.96(4): 596-608

[30] Sakaki, Takeshi, Makoto Okazaki, and Yutaka Matsuo. 2010. "Earthquake shakes Twitter users: real-time event detection by social sensors." Proceedings of the 19th international conference on World wide web. ACM.

[31] Sebastiani, F.. 2002. "Machine learning in automated text categorization", ACM Computing Surveys, 34(1), 1–47.

[32] The Economist. 2013. "China’s Internet: A Giant Cage," April 6th, 2013.

[33] The 2015 annual report of Weibo Coporation to the U.S. Security and Exchange Commission, Form 20-F, http://phx.corporate-ir.net/External.File?item=UGFyZW50SUQ9NTc4NDQ2fENoaWxkSUQ9MjgyNzU5fFR5cGU9MQ==&t

[34] Zhu, Tao, David Phipps, and Adam Pridgen. 2013. "The velocity of censorship: High-fidelity detection of microblog post deletions." arXiv preprint arXiv:1303.0597.

Table 1. Hot topics by collective action category

Conflict			Protest			Strike			Anti-Japan		
Sensitivity: Very High #posts: 382,232			High 2,526,325			Medium 1,348,964			Low 2,506,944		
Freq.	Word	Translation	Freq.	Word	Translation	Freq.	Word	Translation	Freq.	Word	Translation
322,797	镇压	Suppression	647,711	示威	Demonstration	1,361,854	罢工	Strike	1,358,585	抗日	Opposition to Japan
32,117	冲突	Conflict	534,784	静坐	Sit-in	69,068	罢课	Student strike	1,041,104	日货	Japanese goods
19,124	警民	Police and People	430,112	自焚	Self-immolation	101,887	工人	Workers	1,013,233	抵制	Resisting
17,460	催泪弹	Tear-gas bomb	260,574	讨薪	Ask for compensation	98,822	电脑	Computer	754,032	日本	Japan
31,161	矛盾	Contradictory	346,836	游行	Parade	65,557	出租车	Taxi	257,310	反日	Anti-Japanese
40,286	警察	Police	164,367	请愿	Petition	164,549	泪	Tears	181,809	抗日战争	Sino-Japanese War
14,271	官民	Officials and people	113,936	示威者	Demonstrators	46,219	工会	Trade union	253,308	爱国	Patriotic
31,935	暴力	Violence	109,339	堵路	Stops up the road	91,051	抓狂	Driven nuts	209,467	钓鱼岛	Diaoyu Island
130,036	被	By	166,600	抗议	Protest	55,687	司机	Drivers	178,758	战争	War
74,391	政府	Government	101,845	集会	Assembly	48,845	集体	Collective	126,250	抗战	Sino-Japanese War
12,002	宽恕	Forgiveness	118,262	农民工	Migrant workers	52,066	员工	Staff	137,774	日本人	Japanese
12,764	武力	Military force	103,975	思	Thinking	157,937	今天	Today	201,444	游行	Parade
18,951	军队	Army	80,481	静静	Static	24,477	的士	Taxi	99,910	鬼子	Devils
29,566	民众	Populace	60,237	闲谈	Chat	22,559	法国人	French	84,505	国军	National troops
14,701	叙利亚	Syria	58,318	人非	People are not	51,479	上班	Going to work	201,785	中国人	Chinese
20,170	抗议	Protest	72,753	民工	Laborers	16,290	罢市	Merchant strike	130,590	英雄	Heroes
60,068	人民	People	63,719	白宫	White House	40,827	抗议	Protest	680,886	中国	China
21,521	村民	Villagers	130,198	坐	Sitting	86,612	手机	Handset	99,135	剧	Drama/Play
10,264	起义	Revolt	60,957	己	Oneself	17,679	罢	Strike	113,488	同胞	Compatriots
10,150	开枪	Opening fire	37904	玩火自焚	Being made to pay for one's evil doings	41586	工资	Wages	104276	理性	Rationality

Table 2. Collective action posts

		Random 1000 post sample				
	Total	About event	Forthcoming events	On-going events	Past event	General comments
conflict	382,232	398	1	11	156	230
protest	2,526,325	317	2	19	172	124
strike	1,348,964	312	5	178	39	90
antijapan	2,506,944	504	9	188	42	265

Results from a regression of an event dummy on the number of Sina Weibo posts on each type of event (log(number posts+1)/1000). The unit of observation is prefecture by day. The regression controls for the log total number of Sina Weibo posts in the prefecture that month (log (total number posts+1)). Standard errors are clustered by prefecture.

Table 3. Posts and observed collective action events 2010-2012

		Conflict		Protest		Strike		Anti-Japan		Coal Mine Accident	
		Yes	No	Yes	No	Yes	No	Yes	No	Yes	No
post	Yes	45	52,058	220	146,284	122	105,873	42	112,486	92	84,475
	No	61	346,780	44	252,396	9	292,940	2	286,414	161	314,216
#posts (t)	mean	6.0	0.6	61.8	3.8	166.0	2.2	1987.2	3.9	2.9	1.0
#post (t-1)	mean	3.3	0.6	53.6	3.8	47.7	2.2	903.3	4.0	0.7	1.0

Table 4. Prediction and identification

VARIABLES	conflict	protest	strike	antijapan	coalmine accident
			Panel A		
Regression coefficient					
# Weibo posts	0.670***	0.981***	1.774***	1.054***	1.231***
	(0.202)	(0.162)	(0.305)	(0.205)	(0.283)
# newspaper articles	0.002*	0.002	0.001	-0.000	
	(0.001)	(0.001)	(0.002)	(0.001)	
Observations	398,944	398,944	398,944	398,944	398,944
R-squared	0.002	0.006	0.007	0.005	0.004
			Panel B		
Regression coefficient					
# Weibo posts day before event	0.393***	0.618***	0.792***	0.591***	-0.124
	(0.136)	(0.141)	(0.197)	(0.131)	(0.078)
# newspaper articles articles	-0.000	0.001	0.000	0.000	
day before event	(0.001)	(0.001)	(0.002)	(0.000)	
Observations	398,944	398,944	398,944	398,944	398,944
R-squared	0.001	0.006	0.005	0.003	0.004

Unit of observation is prefecture and day. The dependent variable is a dummy for the occurrence of an event. The key independent variables are the log of 1 + the number of Sina Weibo posts mentioning words related to the event and the log of 1 + the number of newspaper articles mentioning event words. Controls include prefecture and year fixed effects. Standard errors, clustered by prefecture, in parenthesis.

Table 5. Effect of Weibo on collective action events

VARIABLES	antijapan		strike		protest		conflict	
	(a)	(b)	(a)	(b)	(a)	(b)	(a)	(b)
Weibo penetration	-0.001	-0.004	0.015***	0.009**	0.016***	0.012	0.002	0.001
	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)	(0.009)	(0.004)	(0.005)
Observations	1,368	1,368	2,976	2,976	4,596	4,596	3,972	3,972
R-squared	0.526	0.565	0.103	0.165	0.101	0.171	0.029	0.089

Results from a regression of an event dummy on the total number Weibo posts. The unit of observation is prefecture and month. Only prefectures which had at least one occurrence of the event is included in the regression. The regression includes prefecture and month fixed effects, the specification in columns labeled (b) also contains prefecture-by-year fixed effects.

Table 6. Hot topics in corruption and petition posts

Corruption			Petition		
#posts	5,326,897			1,151,563	
1,455,878	贪污	Embezzlement	1,069,371	上访	Appealing for help
1,658,687	腐败	Corrupt	96,491	上访者	Petitioners
681,055	公款	Government money	110,757	访	Visit
674,503	受贿	Bribe	497,029	被	Quilt
556,609	贿赂	Bribe	71,508	劳教	Reeducation through labor
975,187	官员	Officials	38,766	上访户	Appealing for help household
393,125	廉政	Honest government	43,155	唐慧	Tang Hui
639,293	利益	Benefit	63,820	信访	Inquiry
1,002,491	政府	Government	72,680	民	People
245,606	挪用	Diverting	38,696	进京	Going to the capital
512,006	集团	Group	75,209	村民	Villagers
201,891	吃喝	Food and drink	28,313	访民	Visitors
153,731	职权	Authority	196,236	政府	Government
572,569	钱	Money	29,530	关押	Imprisoning
247,942	贪官	Corrupt officials	31,586	信访局	Bureau of Letters and Calls
156,363	滥用	Abusiveness	32,143	精神病	Neurosis
291,309	原	Original	67,136	警察	Police
288,287	干部	Cadres	24,279	病院	Specialized hospital
123,827	行贿	Bribery	69,074	解决	Solution
126,820	情妇	Lovers	23,482	丢了	Lost

Table 7. Coverage of politicians

	I	II	III	IV	V
position	# posts	# posts per position	Nontalk	Pct corruption	Sentiment
Xi Jinping	1,374,780	1,374,780	0.12	0.23	0.88
Wen Jiabao	1,338,882	1,338,882	0.09	0.15	0.51
Li Keqiang	401,451	401,451	0.08	0.14	0.81
Hu Jintao	347,158	347,158	0.13	0.11	1.16
Provincial Governor	728,386	22,072	0.61	1.88	-0.19
Provincial Party Secretary	403,074	12,214	0.51	1.91	0.52
City Mayor	3,541,029	10,634	0.52	1.36	0.17
City Party Secretary	718,856	2,159	0.60	2.81	0.28
County Governor	719,634	251	0.49	1.21	-0.70
County Party Secretary	324,522	113	0.65	4.40	-0.88
Village Chief	1,053,346	25	0.40	0.65	-0.51
Village Party Secretary	144,742	3	0.63	4.26	-1.40

Table 8. Dependent variable corruption case dummy

VARIABLES	I	II
# posts about corruption 2-7 months prior	3.857***	2.880***
	(0.862)	(0.870)
Observations	19,494	19,494
R-squared	0.006	0.032
Fixed Effects	Year	Year, Prefecture

The regression controls for the log total number of posts +1. The key independent variable is ln(1+ # posts about corruption 2-7 months before the corruption case)/1000. The division by 1000 is just scaling. Unit of observation is prefecture by month. Standard errors clustered by prefecture.

Table 9. Mean number of posts, by corruption charge

	2-7 month lag		12-23 month lag	
	name	corruption	name	corruption
Corrupt official	49.0	3.9	148.3	4.7
Non-corrupt official	44.4	0.4	121.1	1.8

Table 10. Dependent variable is corruption case dummy

VARIABLES	(1)	(2)	(3)	(4)	(5)
# posts mentioning name and corruption	0.0042***	0.0065***			0.0038***
(2-7 months before first action)	(0.0010)	(0.0015)			(0.0009)
# posts mentioning name and corruption			0.0035**	0.0050**	0.0029
(12-23 months before first action)			(0.0014)	(0.0024)	(0.0019)
Observations	680	680	680	680	680
R-squared	0.014	0.053	0.009	0.044	0.052
Fixed Effects	No	Case Id	No	Case Id	Case Id

Unit of observation is official. The regression also includes the number of posts mentioning the official's name. This variable is always insignificant. Standard errors in parenthesis, clustered by case id (charged leader and matched control leaders).

Table 11. A simple corruption classifier.

		Corrupt		Total
		0	1	Total
Any posts	0	355	133	488
	1	125	67	192
Total	Total	480	200	680

Table 12 Government presence on Sina Weibo

	Users			Posts	
Type	Percent	Est. Number	St. Dev	Percent	St. Dev
Government	.5	149,746	66,801	.2	.1
Media	.5	149,746	66,801	2.3	1.6
Public	1.0	299,491	94,233	1.1	.5
Government-affiliated	2.0	598,982	132,590	3.6	1.6
Others	98.0	29,350,118	132,590		

Table 13. Dependent variable: share government users

	I
GDP	-0.849***
	(0.103)
CPC stronghold	0.533**
	(0.236)
Treaty port	-0.079
	(0.166)
Distance to Beijing	-0.464***
	(0.165)
Population	0.366***
	(0.129)
Latitude	0.052***
	(0.016)
Longitude	-0.037***
	(0.014)
Observations	259
R-squared	0.358

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Figure 1

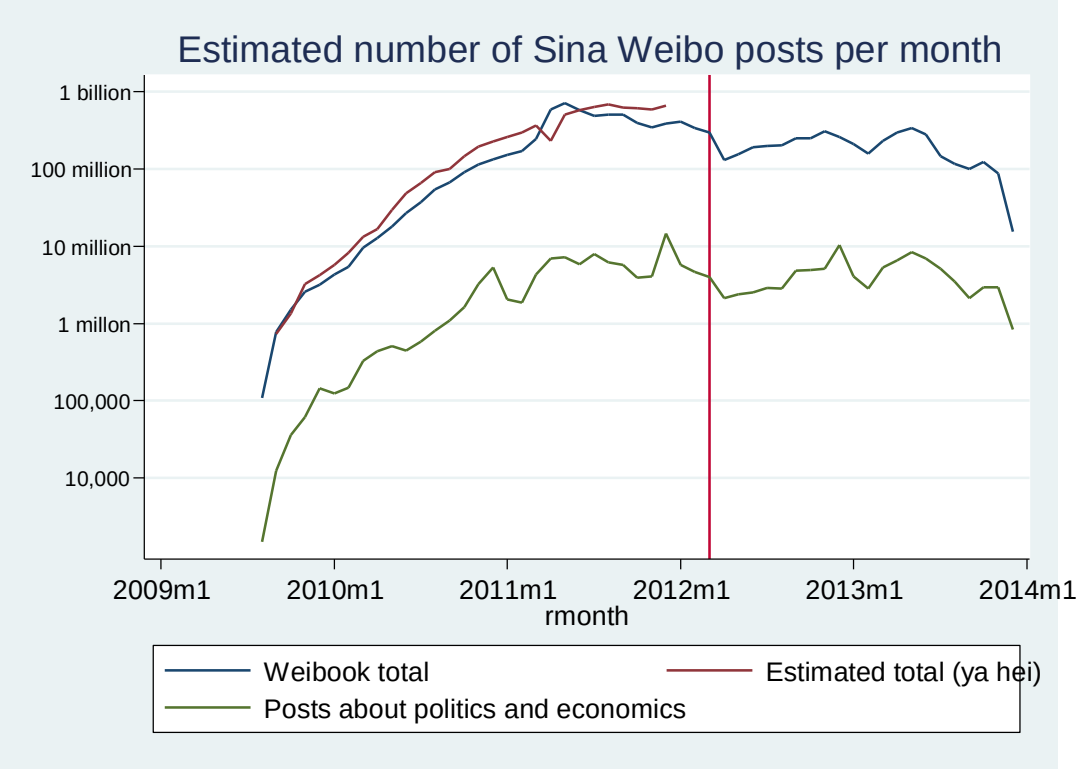
Estimated number of Sina Weibo posts per month
1 billion
100 million
10 million
1 millon
100,000
10,000
2009m1
2010m1
2011m1
2012m1
2013m1
2014m1
rmonth
Weibook total
Estimated total (ya hei)
Posts about politics and economics

Figure 2

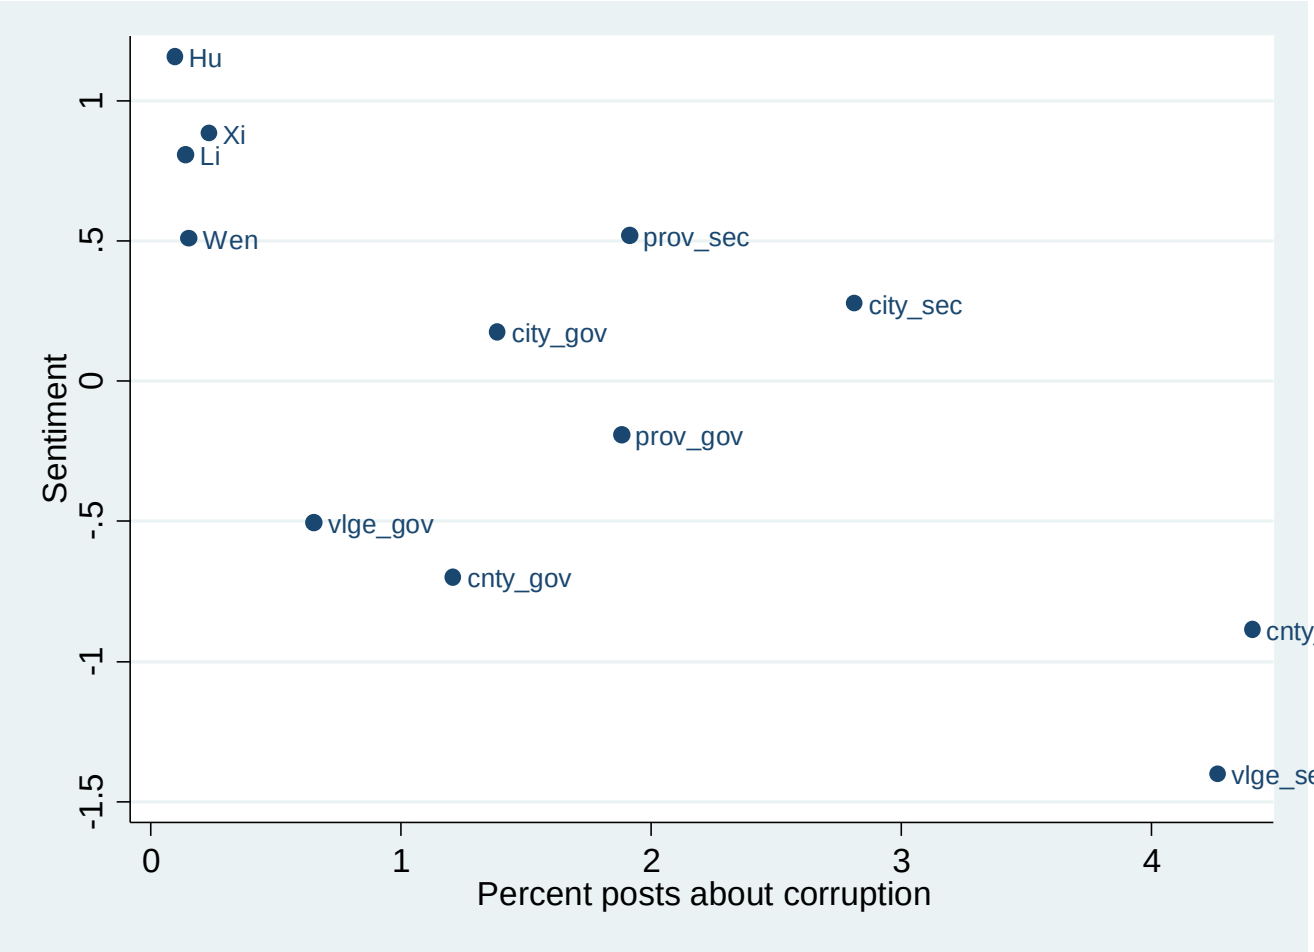
Hu
Xi
Li
Wen
prov_sec
city_sec
city_gov
prov_gov
vlge_gov
cnty_gov
cnty_sec
vlge_se
Sentiment
1
.5
0
-.5
-1
-1.5
0
1
2
3
4
Percent posts about corruption

Figure 3

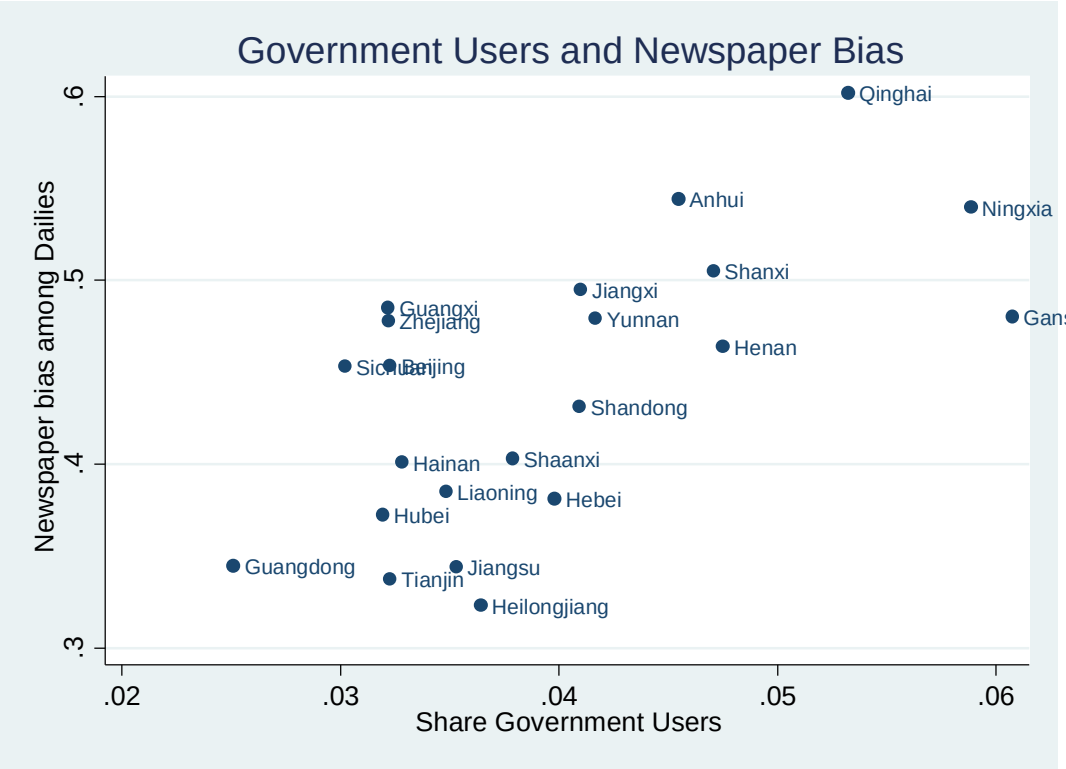
Government Users and Newspaper Bias
Newspaper bias among Dailies
.3
.4
.5
.6
.02
.03
.04
.05
.06
Share Government Users
Qinghai
Anhui
Ningxia
Shanxi
Jiangxi
Guangxi
Zhejiang
Yunnan
Gansu
Henan
Sichuan
Beijing
Shandong
Hainan
Shaanxi
Liaoning
Hebei
Hubei
Guangdong
Tianjin
Jiangsu
Heilongjiang

Figure 4

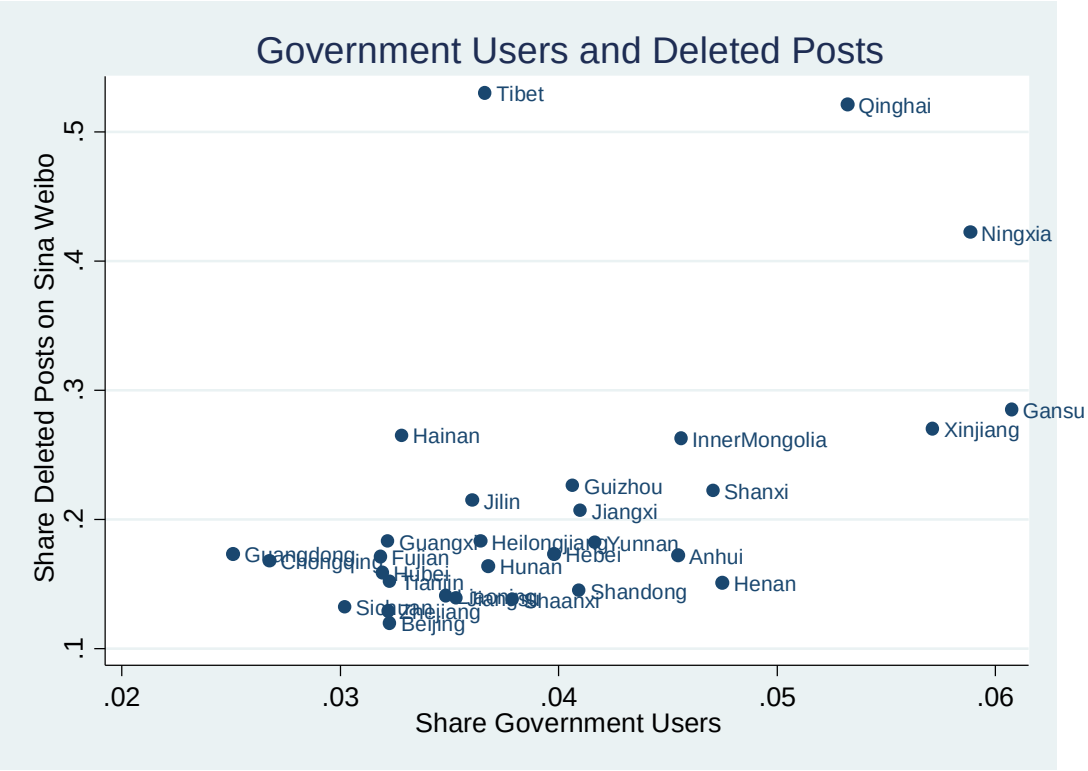
Government Users and Deleted Posts
Share Deleted Posts on Sina Weibo
.1
.2
.3
.4
.5
Share Government Users
.02
.03
.04
.05
.06
Tibet
Qinghai
Ningxia
Gansu
Xinjiang
Hainan
InnerMongolia
Guizhou
Shanxi
Jilin
Jiangxi
Guangxi
Yunnan
Hebei
Anhui
Fujian
Hunan
Hubei
Henan
Shandong
Shaanxi
Beijing