The Political Economy of Social Media in China

Bei Qin, David Stromberg, and Yanhui Wu*

Abstract

We analyze a data set of 13.2 billion blog posts from Sina Weibo – the most prominent Chinese microblogging platform – over the period of 2009-2013. We find not millions of posts covering corruption of government officials and top political leaders and millions of posts discussing collective action events such as protests, strikes, and demonstration. This contrasts with previous studies who conclude that these events are censored.

Second, we find that microblog posts are highly informative in predicting real-world outcomes, such as collective action events and corruption charges. We analyze a data set of 600 large collective action events during 2009-2013 and find that postings on Sina Weibo predict these events taking place the next day. Matching 200 politicians who were charged with corruption with a sample of politicians who held similar positions but were not charged, we find that postings in microblogs 12-23 months before the events strongly predict charges.

Third, using the gradual entry of Sina Weibo across different regions of China, we find that microblog postings have a significant and positive effect on the incidence of strikes and protests, but not on the incidence of large-scale massive conflicts or anti-Japan demonstrations. This suggests that the effects of microblogs are contained, they do not affect events that could pose a threat to the regime.

Finally, we use machine learning tools to estimate the number of government microblog accounts. This is greater in areas with a higher level of censoring and where newspapers have a stronger pro-governmetn bias. This suggests that the government accounts are used for propaganda. We also find that number of government accounts is higher in the strongholds of the Communist Party of China and lower in rich areas.

*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

Out of ten Chinese, nearly half have access to the internet and two actively use Weibo (microblog in Chinese).[1] Every day, millions of blog posts are produced, exchanged, and commented. Many of these posts reach thousands and even millions of readers. However, the Chinese microbloggers are also subject to one of the greatest control efforts the world has seen, combining policing and punishment of users, censoring of websites and content, and posting of propaganda. 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 well-publicized cases,[2] it is unclear whether social media in China play any systematic or important role for social and economic outcomes.

In this paper, we provide an overview of 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 period of 2009-2013. To the best of our knowledge, this is the largest sample of textual content from Chinese social media. We combine statistical description, econometric analysis, and machine learning to document basic facts about the microblog postings, and to study whether they predict and affect real political and social outcomes.

We focus on outcomes in three areas, collective action, monitoring of politicians, and the gauging of public sentiment. These areas are central to understanding the pros and cons of social media to the regime (Egorov et al. 2009, Lorentzen, 2014). On one hand, social media's role in facilitating collective action represents a potential threat to the regime. This is starkly illustrated by the role of social media in the Arab Spring uprisings against regimes in Tunisia, Egypt, Libya.

One the other hand, the bottom-up information flows in media allow the central government to identify and solve social problems before they become regime threatening. This role of media has been an important part of the CPC's propaganda policy, under the name of "the Mass Line". 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 function of social media from the perspective of the central regime is to monitor local governments and officials. China is a large country, and many political and economic decisions are decentralized to local governments. In recent years, "online anti-corruption" has become an important complement to the "offline anti-corruption" campaigns launched by the Chinese central government.

We examine four questions related to effects in these areas. First, how much do the microblog posts cover collective action events, people's grievances

[1] According to the China Internet Network Information Center (2014), by the end of 2013, the number of Chinese internet users was 618 million, among whom 281 million are 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.

with public policies and explicit criticism of political leaders? This will help us understand if this coverage is censored and if users are punished after posting. Second, how informative are social media posts in predicting outcomes such as corruption charges, strikes, or protests? Third, how are social media used to organize collective action events? Finally, how do Chinese governments use social media for propaganda, and is propaganda a substitute for or complement to censorship?

Section 2 provides background information about the development and regulation of microblogging in China. Section 3 presents our data. The microblogs role for collective action and monitoring of local politicians are discussed in sections 4 and 5, respectively. Section 6 analyzes government content and section 7 concludes.

2 Background

Internet use in China increased steadily after becoming commercially available in the mid 1990s. By 2013, the number of Chinese internet users rose to 618 million, approximately 46% of the Chinese population and slightly higher than the global average level of 39%.[4] Among these internet users, 281 million (45%) actively used microblog services.

Among social media, microblogs have the most extensive and vivid discussion and debate on public issues. In 2012, for example, the two most popular Facebook-type social media in China – Renren and Kaixin – covered the top 20 public events listed by the Public Opinion Monitoring Agency in 20 million posts, whereas 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 the presence of Twitter; in the next year, its major Chinese counterparts - Fanfou,Digu, and Jiwai - launched their services. 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 of the domestic microblogging services. The microblog market was almost vacant until Sina Weibo came on the scene in August 2009. 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 in mid 2011.[6]

Sina Weibo led the development of the Chinese microblogging market. It is a hybrid of Facebook and Twitter: in one tweet, up to 140 Chinese characters are allowed, but users can send private messages and embed pictures and videos in addition to adding comments and re-posting. Because of its easy access and use, Sina Weibo soon became the most popular microblog in China. By 2010,

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

[5] Reports on the Online Public Opinion (2010-2013), provided by the Public Opinion Monitoring Agency.

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

it had 50 million registered users, and the number doubled in 2011 and reached the 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 remained an influential platform for public opinion.[7]

The Chinese government strictly controls the microblogs through censoring and policing. Censoring is regulated by the central authority - the CPC Propaganda Department and a number of national media control offices - and is largely implemented by private service providers. Except for Tencent, companies running social media are all registered in Beijing and are closely watched by the Beijing Network Information Office. The service providers are aware of the danger of deviating from the government's censoring policy. On the other hand, the Chinese social media market is highly competitive and consumers' demand is elastic to censoring. Facing such a 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]

A number of papers study censorship in China. The estimated scope of censorship on Sina Weibo ranges from 0.01% (Fu, 2013) to 12% (King et al, 2013) of all posts. Studies have found that censoring allows criticism but not collective action (King et al, 2013), and that censoring focus on political and minority groups issues (Bamman, 2012, Fu, 2013, Zhu, 2013). Censorship is fast, 30% of deletions happens in the first half hour and 90% within 24 hours (Zhu et al., 2013). In this paper, we focus on the content that is available on microblogs, rather that what is taken out.

Equally important to this question is policing. Tens of thousands of internet police and internet monitors operating at all levels of government.[9] They keep track of users who post sensitive content. These users may receive warnings, have their accounts shut down, be sent to "reeducation through labor" camps without trial, and even be sentenced to prison.[10] Although these targeted individuals represent a tiny fraction of the user population, the harsh punishment chills potential dissenters and encourages self-censorship.

While censoring is nationally regulated and implemented, the local leaders have their own internet police. Local leaders may use this to pursue their own goals, for example suppressing negative information about the regions under their administration, even if this information may be tolerable or encouraged by the central government. However, they can only punish users that they can identify.

Initially, as a way to commit not to punish users and thus encourage information flows, the Chinese government allowed users on Sina Weibo to post anonymously.[11] However, in March 2012, the media control authority launched

[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 and sometimes as long as life imprisonment.

[11] Even if users do not provide their real names, they may be identified via their IP-adresses.

a Real Name Reform, requiring users to reveal their true identity to the social media providers. The tension between the economic and political goals faced by Chinese social media is strikingly reflected in the implementation of this reform. In 2012, Sina, the provider of Sina Weibo and a listed company in the U.S. stock market, reported to the U.S. Security and Exchange Commission that a full compliance to the Real-name Reform would dramatically reduce its traffic and usage, and that its ongoing failure to conform to the rules exposed the company to potentially "severe punishment" by Beijing, ranging from the disabling certain services to the "termination of Weibo operations." Three years later, this reform has not yet been entirely implemented by service providers.

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

Against this background, we would expect social media to most strongly affect most social outcomes that bloggers and the central regime have a common interest. This could include addressing local problems such as corruption, and small strikes and protests. A caveat is that these effects could be muted by local internet police. Critique of central government policies that are not key to the CPC may also be allowed. We would not expect social media to contain critical discussions of the CPC fundamental ideology, major policies, and national leadership. Neither do we expect social media to have a large impact on large-scale collective action events, such as riots and massive conflicts between governments and the public.

The regime also has a common interest with the service providers, because the Chinese IT sector is one important driver of economic growth in China. Thus, it is likely to see less-strict implementation of content regulation when economic costs are higher, for example, in areas where demand is more elastic and the advertising market is larger. Similarly, stricter enforcement of regulation may be observed during times of higher political tensions.

3 Data

Our data was collected by Weibook Corp. Since September 2009, they have executed a massive data collection strategy, downloading the posts from the active users that they can identify. First, by crawling, they identify users as authentic active persons according to the individual information and the interaction with friends. They locate the 200-300 active million users. Second, according to the

However, IP-addresses can be hidden using services such as Tor (The Onion Router) and VPN (Virtual Private Network) services

number of followers, they categorize users into 6 tiers. They download the microblogs of the top tier at least daily, the second to third every 2-3 days, and the frequencies decreasing with tiers till a week. For each post, we have the content, the posting time and user information (including self-reported location).

In total, the data set contains 13.2 billion posts 2009 to 2013. For the subperiod 2009 to 2011, we construct an independent measure of the volume of posts on Sina Weibo.[12] The Weibook data contains around 95% of the total amount of posts on Sina Weibo by this measure. This is illustrated in Figure 1, where the blue line contains the number of posts per month in the Weibook data and the red line is our estimate of the total number of posts on Sina Weibo.

From this data, we pull out microblogs mentioning any of around 5,000 words related to politics or economics. This includes two groups of keywords. The first group refer to categories of issues: names of all listed companies and important business leaders; major political positions from central to village level; names of top political leaders; social and economic issues such as corruption, pollution, food and drug problems, disasters and accidents, crime; collective action events, such as strikes, protests, petitions, mass conflicts. For words occurring with high frequency, we only collect a 10 percent random sample of posts mentioning these. The second group of keywords refer to specific events that we have recorded, including events in censoring directives and massive events from 2009 to 2013. In total, we end up with 202 million posts from 30.6 million users.

Part of the analysis uses word frequencies. Computing these requires some pre-processing. We used the Stanford Word Segmenter to segment the words in each microblog post. We then removed stopwords, punctuations, URLs, usernames and non-Chinese characters (other than pm and px) from the text. We drop words with more than 30 characters and words occurring less than 5 times. We end up with 3.2 million distinct words and 6.0 billion tokens (word occurrences).

4 Collective Action

Social media potentially play a large role in coordinating and organizing collective action events. However, because of extensive policing and censoring, it is an open question wether they can do so in China. To play a role, the relevant content must at a minimum be present in social media. Social media coverage must also proceed, or at least be concurrent with, the event. Finally, social media entry should be associated with more collective action events. We investigate these three conditions in turn.

We analyze around 600 large collective action events that took place in mainland China between 2010 and 2012. The events were recorded by Radio

[12] 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 periods since we lost access to the public timeline API.

Free Asia, a non-profitable radio station based in Washington D.C. For most events, we are able to identify the start date from news reports. We classify the collective action events into four categories, ranked by their expected sensitivity. The first category contains the most sensitive events. These are social conflicts, including riots and conflicts between governments and the public. The second category is protest, including street demonstrations and mass protests. The third category is strike, including strikes in factories and schools and those among taxi drivers. The last category is anti-Japan demonstrations.

We select keywords that identify posts about each event type and extract all posts that mention these keywords in the full Weibook data. These keywords are described in the appendix.

4.1 Content and Users

We first analyze whether there is relevant coverage of collective action on social media. There are reasons to believe that this coverage would be limited. There are many documented cases of people being punished after posting information about a protest and other collective action events (see e.g. Freedom House 2012). Perhaps the most ambitious study of social media censoring in China, King et al. (2014), study four collective action events in China. They conclude that collective action 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)

To our surprise, we find a large amount of posts covering even the most sensitive collective action events. We identify 382 thousand posts in the conflict category and over 2.5 million posts in the protest category; see Table 1. To describe the content, we first characterize "hot topics" in posts about collective action. These topics are identified by words used more often in collective action posts than in our overall sample.[13]

The results are presented in Table 1. The words are ordered by statistical significance. For example, for the conflict category, "suppression" is the word that has the most abnormally high use. It is used 322,797 times in the 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 more highly ranked than "government" because the latter word is more commonly used in general. Other topic words in this category are "police community", "violence", "revolt" and "opening fire".

To characterize this data more in detail, we investigate a random sample of 1,000 posts for each category (that we will post online). We manually code whether and how the posts actually cover this type of event; see Table 2. The share posts which are actually about the event ranges from 50.4 percent for

[13] 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) p. 11.

the anti-japan events to 31.2 percent for strikes. To get an estimate of total number of posts in a given category, the shares should be multiplied with the total number of posts in the category. For example, the number of posts about ongoing protests in our data is around 48,000 (2.526 million times .019). To give a sense of our coding, some examples are:

- "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 to 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 10am. Anyone wanna join?" [Anti-Japan, forthcoming]

We do not find much evidence that regular users who post this type of content are being punished. We can investigate this question by looking at future posts from the users who post about collective action events. We find that users who post on these topics continue to post to a similar extent as other users so their accounts are not closed.[14]

Users also do not seem afraid of posting this type of content. Users who anticipate punishment could post about sensitive events in a separate anonymous Sina Weibo account, where their IP addresses are hidden and with few other posts that can be used to identify them. We would see this in the data as the posts on sensitive topics coming from user accounts with few posts. However, the average number of posts from users who post on these topics is not significantly lower than a similarly sized randomly drawn comparison sample of users (drawn using the number of posts by each user as sampling weights).

4.2 Identifying and Predicting Events

We do find some patterns that indicate censoring. The collective action events in our data are large enough to be mentioned by the news media. It seems likely that at least some blogger would write about all of them, if they do not fear being punished. Still, we find no blogs for many events. Table 3 shows whether

[14] In our data, 16 percent of the posts are the last post of a user. In the conflict and protest categories, the corresponding numbers are 17 and 23 percent. The share 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 the conflict and protest categories (33-34 and 41-42 percent).

there was an event, and whether there was any post about this type of event in a prefecture and day. For example, we find posts covering all 46 cases of anti-Japan events. The likely reason is that these events are not censored. The number of events not covered by any blogpost is increasing in sensitivity: strikes (8), protests (37) and conflict (60). The likely reason is that the more sensitive events were censored.

A key question is whether collective action events are caused or amplified by posts in social media. This could either be because of explicit calls for participation, or because information about ongoing conflicts can spur and coordinate participation. A necessary condition for this is that the post are covering a forthcoming or ongoing event.

We can analyze this in detail in our random sample of 1,000 posts. Compared to the less sensitive events, the more sensitive events (conflict and protests) are covered more in general comments and in retrospect; see Table 2. However, some posts cover the sensitive types of events while they are ongoing, and some even discuss these events in advance. In the conflict category, 1 in 1,000 posts cover the event in advance and 11 cover ongoing events. For protests, the numbers are roughly doubled. In comparison, 178 and 188 posts cover ongoing strikes and anti-Japan events, respectively, and 5 and 9 posts cover forthcoming events.

Another less direct way to address this question is to see whether the number posts on topic predicts the events. Table 3 shows the average number of posts about each event type by users in the prefecture where the event took place, the day of the event and the day before. The coverage on Sina Weibo decreases in the sensitivity of events. The highly sensitive large-scale conflicts are only covered in 6 posts on the event days, compared to 2000 posts per day for the government sanctioned anti-Japan events. Although there are only 3.3 posts discussing conflicts the day before these events, this is still much higher than in other days (.6).

Consequently, the number of microblog posts mentioning an event type is highly significant in identifying and predicting the event. Table 4 shows the results from a regression where the dependent variable is an indicator for whether an event of this type took place in a particular prefecture on a given day. The key independent variables are the number of Weibo posts from users in that prefecture that mention the event type in that day (panel A), or the previous day (panel B). The regression also includes the number of newspaper articles mentioning the event type in that prefecture and day. Newspaper coverage is taken from the WiseNews archives. For the 2010-2012 sample, this includes 62 general-interest newspapers. The regression controls for the total number of Sina Weibo posts in the total Weibook data in that prefecture and month. Standard errors are clustered by prefecture.

Table 4 shows that microblogs are highly significant in predicting where and when these collective action events take place. In contrast, newspaper coverage is uninformative. That people start discussing the events before they happen is an indication that Sina Weibo may be used to organize or at least coordinate on the collective action events. This is true for even the most sensitive "conflict" category, and hence potentially a threat to the regime.

We now turn the tables and ask: how useful this data is to a regime that wants to identify potentially regime-threatening collective action events? The regime wants a classifier that indicates what prefectures and days are likely to have collective action events. Suppose that the central government is risk-averse since they have a high cost of missing regime threatening events. In information retrieval, the fraction of relevant events that are identified is called recall. A risk-averse authoritarian government thus wants a classifier with high recall. On the other hand, it is costly to investigate a large number of false positives. The share actual events of among the indicated possible events is called precision. To lower cost, the regime wants a classifier with high precision.

A simple classifier with high recall is to investigate all locations and days with any posts mentioning the event. We focus on anti-Japan events because these are not censored and hence our information is similar to that used by the government. Table 3 shows there were anti-Japan posts in all 46 cases where there was an anti-Japan event. No anti-Japan event is under the social media radar, and the simple classifier has perfect recall. On the other hand, the simple classifier has terrible precision, 46/(125,359+46) and hence is very costly.

We explored machine-learning methods to improve the classifier. More precisely, we use a method similar to that developed by Sasaki et al. (2010) to detect earthquakes based on twitter flows. This combines a Support Vector Machine (SVM) with regression analysis. This dramatically increases precision (by 97%), but at a cost of missing 4 anti-Japan events. We are not able to increase precision much without reducing recall. Hence, a risk-averse regime may be forced to use very costly methods. It is more than likely that the Chinese government can develop better methods; however, it is not trivial. Machine learning techniques have also not been very successful in identifying other large-scale events, such as flu epidemics. The bottom line is that social media data is very effective in identifying large-scale collective action events, but that doing this is likely to be costly.

Our event data in China is incomplete. The social media data can also be used to identify missing events. For example, to identify strikes that are not in our data, we look at the 100 prefecture-dates where the predicted probabilities of a strike are the largest based on the social media data. By manual inspection of posts in these prefecture-days, we are able to identify 23 strike events that is not in our original data (as well as 14 strikes that are in our data).

4.2.1 Social Media Effects

We find a large number of posts about collective action events, and that these posts precede and predict these events. A natural question to ask is whether Sina Weibo increased the prevalence these events. We can explore this by using the staggered diffusion of Sina Weibo across prefectures. After Sina Weibo was introduced in September 2009, its use picked up faster in some areas, notably those with a high number of mobile phone users, educational expenditures, and the tertiary sector share of GDP (Qin, 2013). Using a difference-in-difference design, we can ask whether the propensity for conflicts, strikes, etc., is increasing

in the use of Sina Weibo.

Table 5 shows the results from regressing a dummy for each type of event on the Weibo penetration (log(per capita # posts + 1)). The specifications with column heading (a) control for running month and prefecture fixed effect, those with column heading (b) control also for prefecture by year fixed effects. 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 estimate is not much affected, but the standard errors double, indicating that the fixed effects remove much variation. The estimates suggest that Weibo use has increased strikes and perhaps protests, but not anti-Japan events and conflicts.

The above results are sensible. People do not need social media to organize anti-Japan events, because these are supported by the government and organized e.g. by newspaper coverage. Sina Weibo does not affect the highly sensitive conflicts, because these events are censored.

The lack of effect of Sina Weibo on conflicts suggests that social media are unlikely to be regime threatening. Large-scale events in conflict with central government interests are not significantly affected. The lack of effect on anti-Japan demonstrations also suggests that Sina Weibo is not very effective in boosting collective action supported by the government. The positive effects for strikes and protests indicate that social media may affect collective action events that are of a small scale, confined within a small region, and do not threaten the central regime.

5 Monitoring Local Politicians

In this section, we investigate whether social media provide information relevant to hold local politicians accountable to higher-up politicians. This requires that the relevant content is 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-rank Chinese government or CPC leaders.[15] Finally, we examine if the corruption posts are posted before government action, and if they are informative in identifying politicians that are later charged with corruption.

5.1 Content and Users

5.1.1 Petitions

Petitioning is an old system by which people can bring their grievances to local officials, or even travel to Beijing to bring their issues to the attention of the central leaders. Local politicians often dislike petitioning because they fear direct reprisals from the center or because it negatively affects their careers. They

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

have been accused of intercepting petitioners and forcing them to return home and even imprisoning them in illicit detention facilities. For example, a report by Human Rights Watch in 2009 charged that large number of petitioners, including children, were detained, and documented several alleged cases of torture and mistreatment.[16] With the availability of social media, petitioners do not have to travel to Beijing; they simply post about their issues on Sina Weibo. However, they still risk being identified and punished by the local leaders.

To document this phenomenon, we retrieve over 1 million posts that contain two keywords closely related to petitioning. We do a topic analysis on the posts in this category. Table 6 shows that the topic words are mainly related to the practice of petitioning (e.g. "appealing of help" and "petitioners"), or to different punishments, such as "reeducation through labor", "imprisoning", "police" and "specialized hospital". A word that does not appear to belong to these categories is "Tang Hui". However, this is the name of a Chinese mother who was locked up in a labor camp for demanding justice for the rape, kidnap and prostitution of her 11-year-old daughter.

We randomly select 1,000 posts of this kind and manually read them. Only 93 are irrelevant to petitioning in China; 35 are about ongoing or forthcoming petitioning cases, 469 about past petitioning cases, and 406 about petitioning in general. We find little evidence that users who post this type of content are punished. Users who post on these topics continue to post to a similar extent as other users, implying that their accounts are not closed. They also do not petition from special accounts with few posts. Rather, petitions are generated from users who have more posts than an average user.

To sum up, there is a vivid microblog discussion of petitions, mostly about petition in general or past cases. This is likely why the hot topic words in this category are about the petitioning method or punishments. However, we estimate that around 40,000 posts in our data are about forthcoming or ongoing petitioning cases.

5.1.2 Corruption

On one hand, one might expect plenty of posts discussing local corruption because the central regime and regular users have a common interest in reducing this. On the other hand, corruption among local leaders reflects badly on the CPC and may reduce its legitimacy. Perhaps more importantly, people may self-censor in order to avoid being identified and punished by the powerful local leaders that they are accusing of corruption. The local leaders have strong incentives to suppress this coverage. Corruption charges in microblogs potentially carry a strong penalty. Even if there is no legal action, news of corruption in their jurisdiction may negatively affect the leaders' careers.

To study coverage of corruption, we combine two types of microblog posts: those mentioning politicians or political positions and those mentioning corrupt behavior. In the first category we retrieve posts that mention any top political

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

position, from central to village level. In total, we find over 11 million 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 is the most discussed leader, with over 1.3 million posts, followed by Wen Jiabao. Officials in higher offices are more discussed, and executive positions are more covered than Party Secretaries.

In the second category, we search for words that are widely used to describe corrupt behavior (e.g., bribery, cronyism, and the misuse of power) words that are used to describe general wrongdoings and punishment. We find 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 these posts more in detail, we manually inspect 1,000 randomly selected posts. Most of these just talk about corruption in general. Of the 419 posts that talk about specific cases of corruption, 293 are written after the government has taken action. However, 126 posts talk discuss particular instances of corruption before the government has taken action. This category includes two types of posts. One type targets specific government officials and are usually very short.

- "XXX, the Party secretary of XXX village, misused the monetary transfer from the central government to low-income villagers to pay his family members and relatives."

- "XXX, the chief officer of XXX county, embezzled public money by contracting all major government projects 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 shows resentment and anger about certain corrupt officials and in most cases only talking about the positions and government divisions without specifying the names of officials.

- "The black market of 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 dared these prefecture officials sell government positions? Crack down on the tigers!"

Based on the share in the random sample, we estimate that our data contains around 668,000 posts that discuss particular instances of corruption before the government has taken action. This is a wealth of information for superiors seeking to hold local politicians accountable. It is clear that posts of this type

are not censored from the center. Perhaps more surprisingly, it also seems that people do not fear posting concrete corruption allegations implicating powerful local leaders. Again, we find no evidence that users who post about corruption disappear from our data, or post from special accounts with few posts.

It is easy to find posts that explicitly criticize even the four very top leaders, although not with explicit corruption allegations. Such posts claim, for example, that democracy and social stability fell 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 the capital to Wenzhou to help some children of the top leaders.

We next describe how this discussion of particular corruption cases is distributed across top positions. To do this, we first identify the social media posts that are actually discussing particular instances of corruption. The top leaders often talk about corruption policy, and a naive count of posts would pick this up instead of corruption allegations.

We estimate the probability that each post discusses a specific case of corruption. To this end, we apply a widely used text classification algorithm - Support Vector Machines (SVM). 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). In the 1,000 sample of coded posts we use the SVM to classify the 419 posts that discuss specific cases of corruption, based on the frequencies of words used in the posts.[17] We classify each of the 1,000 posts using leave-one-out predictions. Precision is .82 (306 of the 374 posts classified as being about corruption are correct) and recall is .73 (306 of 419 corruption posts are correctly classified). The SVM output does not include the estimated classification probability. To get this, run a probit regression of an indicator variable for the post being about corruption on the SVM classification parameter. Figure 2 shows that the probit-estimated probabilities (blue line) are close to the non-parametric average share corruption posts (red line). The x-axis is the SVM classification parameter. We will use these probit-based estimates of the probability of a post being about a specific case of corruption.

We then compute these probabilities for each of the 5.3 million posts in our corruption category. Column III of Table 7 shows the average such probability for posts in our corruption category who also mention different political positions. Just using the raw number of posts mentioning corruption would clearly be misleading. For example, we estimate that only a small fraction (around 10 percent) of the corruption posts that mention national leaders actually discuss specific corruption cases. Most of these posts are just government talk about corruption. The share posts that talk about specific cases of corruption is much higher for lower offices, ranging from 40% for village chiefs to 63% for village party secretaries. Column IV shows the percent posts mentioning a leader position that also discuss specific corruption cases.

[17] 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).

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

Figure 3 plots the percent of corruption posts in column IV against the average sentiment in column V. Sentiment is most positive for the national politicians, who are also infrequently mentioned in connection with corruption.

County and village party secretaries are have the most negative sentiment and have the highest share corruption posts. In comparison, village chiefs have almost as positive sentiment as the national leaders and feature in fewer posts about corruption than any other non-national leaders. One possible explanation is that some village chiefs are elected. Another possible explanation is the practice of provincial governments to send highly educated officers to villages as village chiefs. These are less likely to be involved with corruption because they have strong career concerns and may have less effective control of resources and lack social networks (they typically only stay in villages for a few years).

5.2 Predicting Corruption Charges

We finally relate this coverage to 185 corruption cases involving high-ranking Chinese government or party leaders. Of these cases, 39 are at the provincial, and 114 are at the prefecture level and 32 are 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 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 year-fixed effects. Column II also controls for location fixed effects. The regression results show that the volume of corruption posts in an area predicts future corruption charges in that area.

We also study the corruption charges of individuals, by examining posts that mentions their names. Here we start with a larger sample of 200 corruption charges, including also national politicians. For comparison, we construct a matched control sample of 480 politicians that were not charged with corruption. The control politicians hold similar political positions 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 the politician and a word that is in our corruption category. We do this for posts 2 to 7 months before the corruption charge, and also for an earlier period 12-23 months before the charge. Table 9 shows that the corrupt officials are mentioned in 49 posts on average 2-7 months before the corruption charge, while the non-corrupt officials are mentioned in 44.4 posts. However, the corrupt officials are much more frequently mentioned in posts that mention our corruption words (3.9 compared to .4). A similar

pattern is found in for posts 12-23 months before the charge. Given the large differences, it is not surprising that corruption posts are highly predictive of corruption charges.

Table 10 shows the results from a regression of the corruption-charge indicator-variable on the number of posts mentioning the official's name and corruption. The regression also includes the number of posts just mentioning the official's name, but this variable is never significant and is not shown. Columns II, IV and V include dummy variables for case ids. The case id is the same for the 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.

Still, many of the corrupt officials are under the social media radar; 133 of the corrupt officials were never mentioned in a corruption post two months before the first government action against them. A simple rule for the government would be to investigates all officials with at least one corruption post (two months before the charge). In our sample, this would lead to the investigation of 192 officials, 67 of which 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 look features of the microblog posts, outside of our corruption category. Of course, it is possible that there are some false negatives in this group: corrupt officials that have not been charged.

6 Propaganda

In this section, we turn to government posting on microblogs. A government microblog post can be published either by disclosed government account or by a hidden account. In this paper, we focus on the official accounts. We will investigate how the use of official accounts covaries across regions with censoring, bias in Chinese newspapers, CPC support and economic development.

We first discuss some possible hypotheses. Official accounts could provide neutral information or propaganda. Propaganda aims at influencing peoples' beliefs and actions. This goal can be achieved both by removing and by adding microblog content. Consequently, if they official accounts largely deliver propaganda, then we should see a strong positive correlation between censoring and posting from official accounts. In this case, we should also see a positive correlation with pro-government bias in the traditional media. We will test this using the measure of Bamman et al. (2012) of deletions of social media posts across regions and the measure in Qin et al. (2013) of bias in Chinese newspapers. The latter is based on nine content categories, including leader mentions, cites of the central Xinhua news agency and coverage of regime critical stories.

Informative propaganda may be more effective on audiences that share the propagandist's view. This follows from a rational model of Bayesian persuasion (e.g. DellaVigna and Gentzkow, 2010), and from psychological theories of cognitive dissonance. There is some empirical support for this view, and even that the effect of propaganda is negative when the audience holds opposing views to

the sender. 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 the support for Nazi messages in places with historically low anti-Semitism. Similarly, DellaVigna et al. (2014) find in a laboratory experiment that Serbian radio exposure causes anti-Serbian sentiment among Croats. To the extent that official accounts carry propaganda, we should see more such accounts in CPC strongholds.

Finally, propaganda is likely to lower the consumers' valuation of social media. To the extent that the service providers can affect the amount of propaganda, we should see less in areas where the advertisement market is valuable and where competition for consumers is high.

6.1 Identifying Official Accounts

We first identify official government and newspaper accounts. For government accounts, we search all user names for characters that are usually part of government or office names. We then manually inspect the retrieved user names, e.g. by inspecting their Sina Weibo websites. We also similarly identify the official accounts for the main parts of Chinese legal system: People's Court and Procurate at all administrative levels. In this case, the second step is not needed as we find no individuals or non-relevant agencies that use the two legal functions as part of their usernames. Actually, no one dare to do so. This way, we identify a total of a set of 1,043 government accounts in our data.

For newspapers, we both identify the official accounts representing the newspaper and related accounts, for example, supplements, non-regular editions, articles from different offices and in different formats, and accounts created by editors or reporters who are associated with the newspaper. To obtain the newspaper accounts, we first, we search the user names for newspaper names from the comprehensive newspaper directory in Qin et al. (2013). To further confirm, we compare the newspaper location from the newspaper directory that from Weibo. If these match, they are coded as the official newspaper accounts. If they do not match, then we check manually by locating the microblog of the user in Sina Weibo by the user ID in our data. This way, we identify 538 newspaper accounts. In our data, these identified government and newspaper users contribute 72,691 posts.

We then investigate how well we can identify these users, based on the language in their posts. We first compute word frequencies in the microblog posts of all 30 million users in our data. We then create a "learning" data set, consisting of the government and newspaper users and a random sample of one percent (29,234) of other users. We apply a Support Vector Machine to classify the government and newspaper users based on word frequencies.[18] Word frequencies are very helpful in identifying government/newspaper users. Of the 786 users who were classified by the SVM to be government/newspaper users, 640 (81%)

[18] 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).

actually were. We compute the probability that each account is a government account using the methodology described in section 5.

We then use the word frequencies of all 30 million users in our data to compute the probability that each user is a government or newspaper. We then collapse these probabilities for all users by province. This gives us a measure of the amount of government content on Sina Weibo, by province.

6.2 The Correlates of Government Content

We now investigate the geographical distribution of government official content on Sina Weibo. Figures 4 and 5 plot our estimated share government users against the share deleted posts (from Bamman et al.) and media bias in Party Daily Newspapers (from Qin et al. 2013). There is a strong positive correlation (0.7) between the estimated share government users and both deleted posts and newspaper bias. This evidence strongly supports that censoring, newspaper bias and government accounts on Sina Weibo are used for the same purpose, namely propaganda. Note that Tibet and Hainan has more deleted posts than expected based on the estimated share government users on Sina Weibo. Perhaps this is an indication that propaganda is not viewed as being very effective in Tibet because of a weaker underlying support for the central Chinese regime.

We next investigate how the use of government content in microblogs covaries with CPC support and economic development at the prefecture level. We use GDP as a measure for economic development. We define CPC strongholds as the prefectures that were part of the CPC soviet in 1933 or were passed in the CPC Long March in 1933-1935. Other areas have a history of Western influence, notably the areas that were part of a Treaty Port, controlled by Western powers, during the period of 1840-1910 (Jia 2012). We also include distance to Beijing, latitude, longitude and population in the regression.

The results are shows in table 12. The estimated share of government users is significantly negatively correlated with GDP and positively higher in CPC strongholds. The latter is perhaps surprising given that the need for propaganda is smaller in these areas. However, it is consistent with the view that propaganda is most effective there. The estimated share government users is also higher for areas closer to Beijing, and more populous areas.

7 Conclusions

We analyze a data set of 13.2 billion blog posts from Sina Weibo – the most prominent Chinese microblogging platform – over the period of 2009-2013. Our main findings are as follows. First, we find not only an enormous number of posts covering corruption of government officials and corruption charges involving top political leaders, but also millions of posts related to collective action events such as protests, strikes, and demonstration, some of which even appeared in the censorship directives issued by the Chinese media control authority. This finding draws a sharp contrast with a number of well-documented cases in which

social media users were sentenced to long-term imprisoning because of posting similarly sensitive content. Using machine learning to characterize these sensitive posts, we find that posts mentioning local (county and village) political leaders contain a high frequency of corruption charges and a stronger negative sentiment.

Second, we find that microblog posts are highly informative in predicting collective action events and corruption charges. We analyze a data set of 600 large collective action events during 2009-2013 and find that postings on Sina Weibo predict these events taking place the next day. Matching 200 politicians who were charged with corruption with a sample of politicians who held similar positions but were not charged, we find that postings in microblogs 12-23 months before the events strongly predict charges. Third, using the gradual entry of Sina Weibo across different regions of China, we find that microblog postings have a significant and positive effect on the incidence of strikes and protests, but not on the incidence of large-scale massive conflicts or anti-Japan demonstrations.

Our last result concerns an important but yet understood aspect of social media in China – the government use of social media. Based on the content of posts generated from more than 1,000 government accounts on Sina Weibo, we identify words that are strong predictors of government accounts. Putting this "government speak" knowledge in machine learning, we compute the probability for each of the 30 million users in our data base to be a government account. We find that the predicted number of government accounts is greater in areas with a higher level of censoring. It is also larger in the strongholds of the Communist Party of China.

Our findings suggest that social media in China are likely to improve the accountability of local governments and unlikely to affect regime change. We find an enormous number of posts about corruption and wrongdoings of local governments and bureaucrats and posts about local strikes and protests. In contrast, we find few posts about 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 large-scale events. This perhaps reflects the Chinese government's trade-off between maintaining regime stability and utilizing bottom-up information.

References

[1] 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

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

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

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

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

[6] 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.

[7] DellaVigna, Stefano and Matthew Gentzkow, 2010. "Persuasion: Empirical Evidence," Annual Review of Economics, Annual Reviews, vol. 2(1), pages 643-669, 09.

[8] 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.

[9] 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.

[10] Egorov, Georgy, Sergei Guriev, and Konstantin Sonin. "Why resource-poor dictators allow freer media: A theory and evidence from panel data." American Political Science Review 103.04 (2009): 645-668.

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

[12] Freedom House (2015), “2015 Freedom of the Press Data”, 2015, https://freedomhouse.org/report-types/freedom-press

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

[14] Fu, King-wa, Chung-hong Chan, and Marie Chau. "Assessing censorship on microblogs in China: Discriminatory keyword analysis and the real-name registration policy." Internet Computing, IEEE 17.3 (2013): 42-50.

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

[16] International Telecommunication Union, (2013) "The World in 2013: ICT Facts and Figures," Geneva, 2013, http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2013-e.pdf

[17] 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.

[18] 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, 1999.

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

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

[21] Lorentzen, Peter. "China's Strategic Censorship." American Journal of Political Science 58.2 (2014): 402-414.

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

[23] Qin, Bei, 2013, "Chinese Microblogs and Drug Quality", working paper

[24] Qin, Bei, David Stromberg and Yanhui Wu, 2013, "Media Bias in China"

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

[26] 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

[27] (Yuanxi and Jing, 2009)

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

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

[30] Zhu, Tao, et al. "The velocity of censorship: High-fidelity detection of microblog post deletions." arXiv preprint arXiv:1303.0597 (2013).

Figure 1.

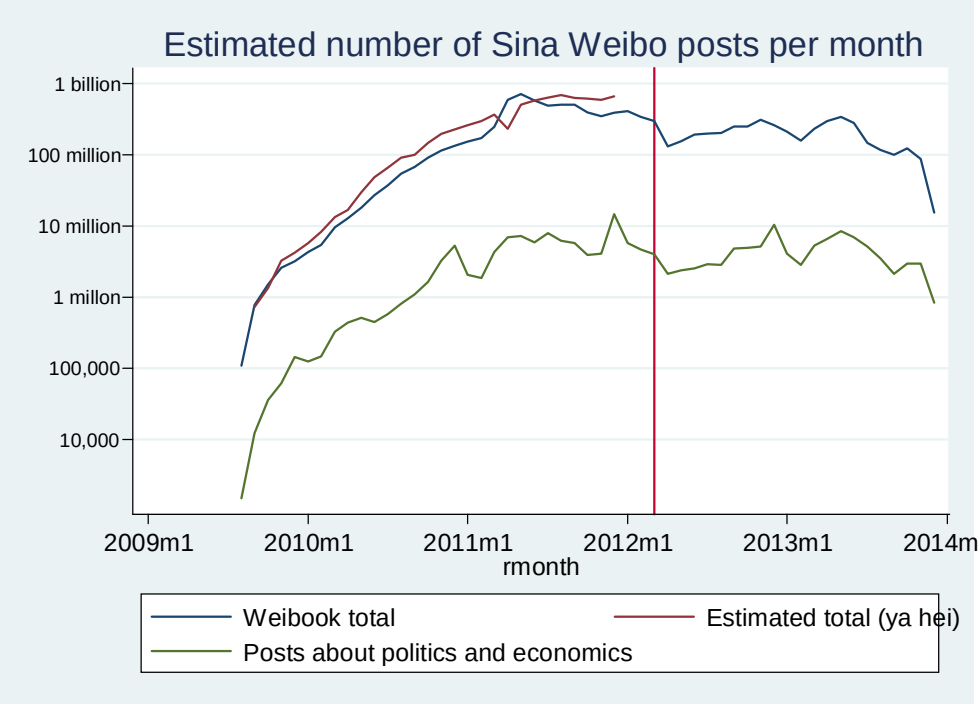


Figure 2.

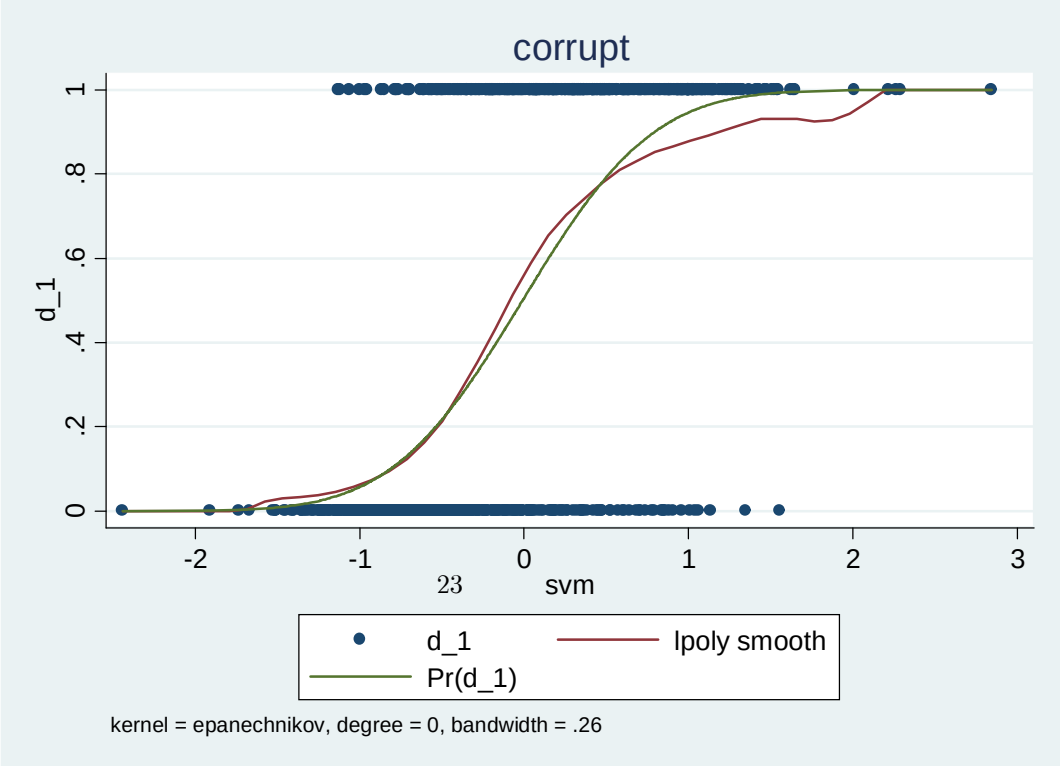


Figure 3.

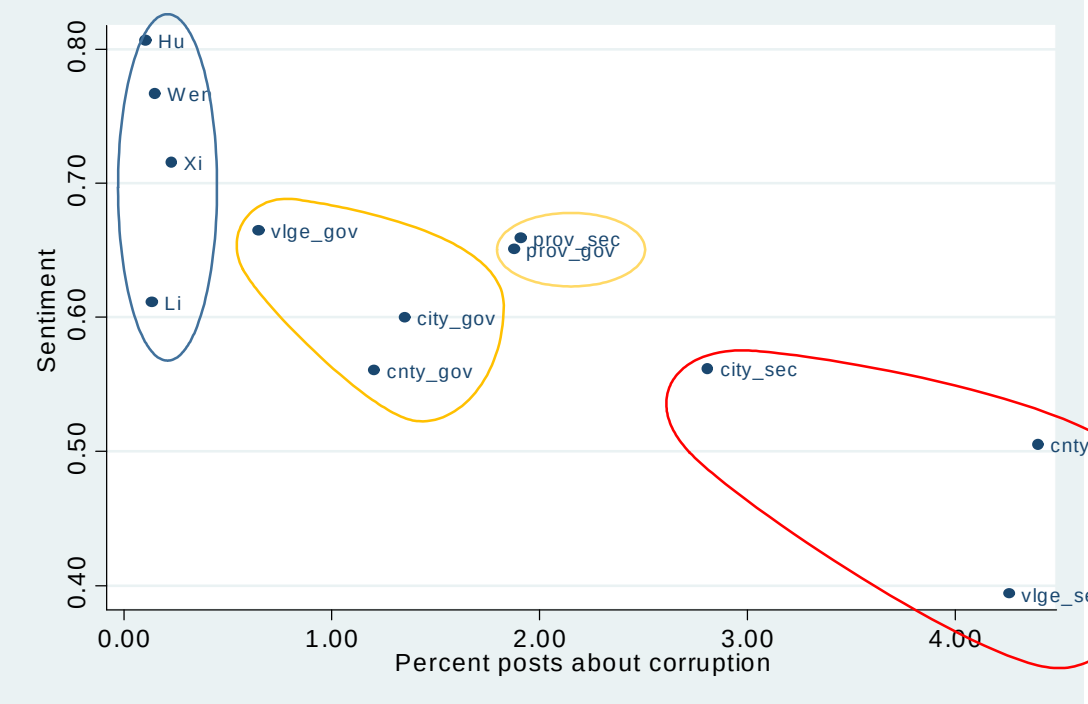


Figure 4.

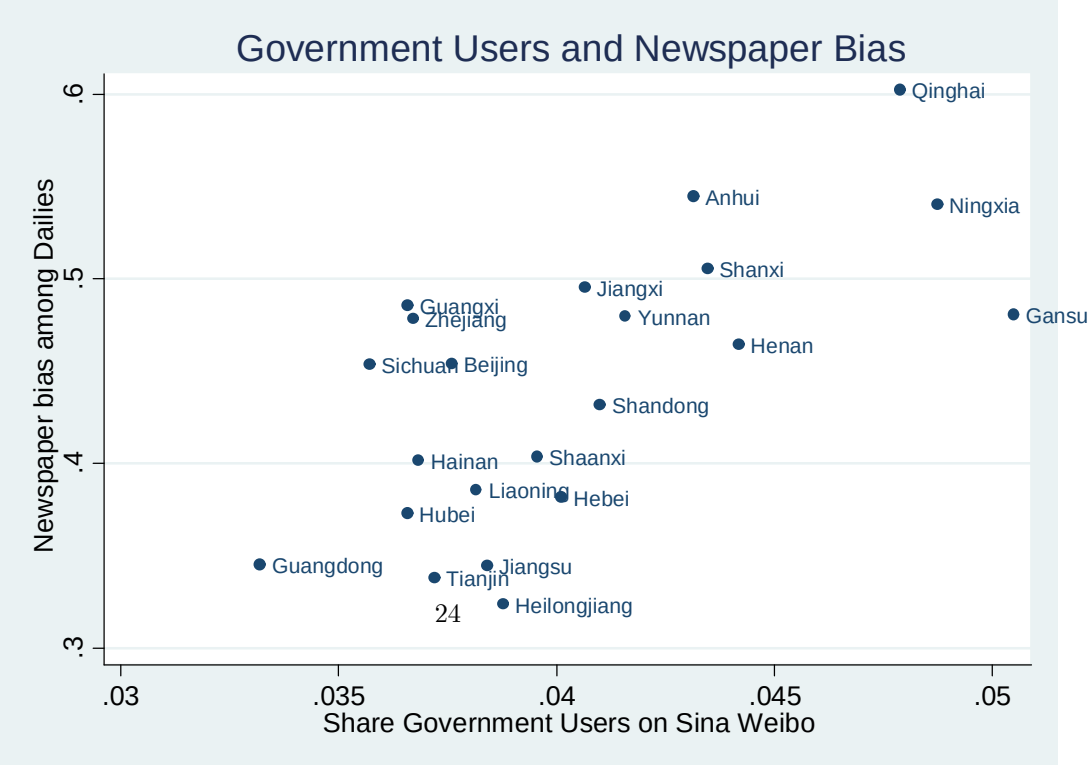


Figure 5.

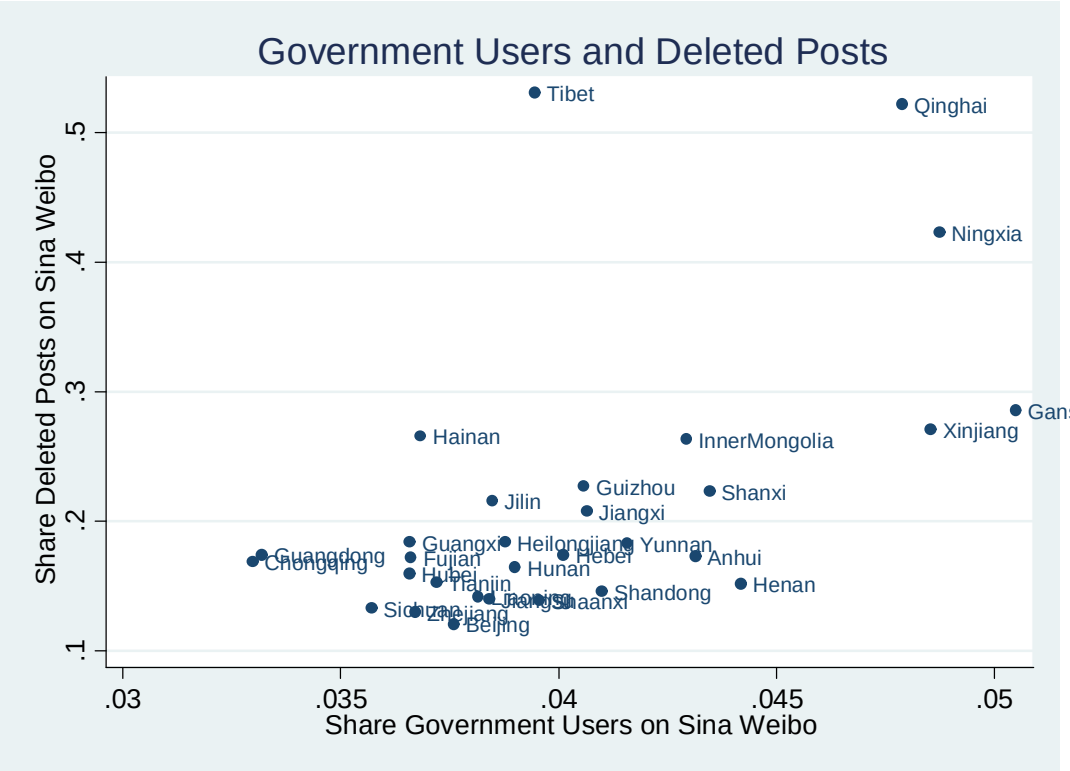

Government Users and Deleted Posts
Share Deleted Posts on Sina Weibo
.1
.2
.3
.4
.5
.03
.035
.04
.045
.05
Share Government Users on Sina Weibo
Tibet
Qinghai
Ningxia
Gansu
Xinjiang
Hainan
InnerMongolia
Guizhou
Shanxi
Jilin
Jiangxi
Guangxi
Heilongjiang
Yunnan
Guangdong
Fujian
Hebei
Anhui
Chongqing
Hubei
Hunan
Tianjin
Shandong
Henan
Shaanxi
Beijing

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	
		Yes	No	Yes	No	Yes	No	Yes	No
post	Yes	47	57,942	237	163,328	126	116,912	46	125,359
	No	60	346,251	37	40,834	8	286,286	0	278,977
#posts (t)	mean	6.0	0.6	65.5	4.1	220.9	3.3	2000	4.3
#post (t-1)	mean	3.3	0.6	57.4	4.1	104.4	3.3	916.0	3.8

Table 4. Prediction and identification

VARIABLES	conflict	protest	strike	antijapan
	Panel A			
Regression coefficient				
# Weibo posts	0.670***	0.981***	1.775***	1.054***
	(0.202)	(0.162)	(0.305))	(0.205)
# newspaper articles	0.002*	0.002	0.001	-0.000
	(0.001)	(0.001)	(0.002)	(0.001)
Observations	389,080	389,080	389,080	389,080
R-squared	0.002	0.006	0.007	0.005
	Panel B			
Regression coefficient				
# Weibo posts day before event	0.394***	0.617***	0.792***	0.591***
	(0.136)	(0.141)	(0.197)	(0.131)
# newspaper articles articles	-0.000	0.001	0.000	0.000
day before event	(0.001)	(0.001)	(0.002)	(0.000)
Observations	389,080	389,080	389,080	389,080
R-squared	0.001	0.006	0.005	0.003

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

position	I # posts	II # posts per position	III Nontalk	IV Pct corruption	V Sentiment
Xi Jinping	1,374,780	1,374,780	0.12	0.23	0.72
Wen Jiabao	1,338,882	1,338,882	0.09	0.15	0.77
Li Keqiang	401,451	401,451	0.08	0.14	0.61
Hu Jintao	347,158	347,158	0.13	0.11	0.81
Provincial Governor	728,386	22,072	0.61	1.88	0.65
Provincial Party Secretary	403,074	12,214	0.51	1.91	0.66
City Mayor	3,541,029	10,634	0.52	1.36	0.60
City Party Secretary	718,856	2,159	0.60	2.81	0.56
County Governor	719,634	251	0.49	1.21	0.56
County Party Secretary	324,522	113	0.65	4.40	0.50
Village Chief	1,053,346	25	0.40	0.65	0.66
Village Party Secretary	144,742	3	0.63	4.26	0.39

Table 8. Dependent variable corruption case dummy

VARIABLES	I	II
log(# posts about corruption +1)	3.958***	2.546**
	(0.691)	(1.102)
Observations	21,557	21,557
R-squared	0.010	0.040
Fixed Effects	Year	Year, Prefecture

The regression controls for the log total number of posts +1.
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 10b. Dependent variable is corruption case dummy

VARIABLES	(1)	(2)	(3)	(4)	(5)
# posts mentioning name	-0.0001	-0.0002*			
(2-7 months before first action)	(0.0001)	(0.0001)			
# posts mentioning name and corruption	0.0041***	0.0064***			0.0040***
(2-7 months before first action)	(0.0013)	(0.0016)			(0.0010)
# articles mentioning name	0.0006***	0.0008**			
(2-7 months before first action)	(0.0002)	(0.0003)			
# articles mentioning name and corruption	-0.0008**	-0.0011**			
(2-7 months before first action)	(0.0004)	(0.0005)			
# posts mentioning name			-0.0000	-0.0001	
(12-23 months before first action)			(0.0000)	(0.0001)	
# posts mentioning name and corruption			0.0035**	0.0051**	0.0030
(12-23 months before first action)			(0.0014)	(0.0024)	(0.0019)
# articles mentioning name			0.0003***	0.0005***	0.0005***
(12-23 months before first action)			(0.0001)	(0.0001)	(0.0001)
# articles mentioning name and corruption			-0.0015***	-0.0023***	-0.0023***
(12-23 months before first action)			(0.0005)	(0.0008)	(0.0008)
Observations	680	680	680	680	680
R-squared	0.0275	0.0682	0.0303	0.0740	0.0822
Fixed Effects	No	Case Id	No	Case Id	Case Id

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. Dependent variable: share government users

	I
CPC stronghold	0.282**
	(0.122)
Treaty port	-0.031
	(0.082)
Distance to Beijing	-0.212***
	(0.079)
Population	0.171***
	(0.065)
GDP	-0.400***
	(0.052)
Latitude	0.027***
	(0.008)
Longitude	-0.019**
	(0.007)
Observations	259
R-squared	0.347

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