The Political Economy of Social Media in China

Bei Qin, David Stromberg, and Yanhui Wu*

Abstract

This paper studies the role of Chinese social media in three areas: collective action, monitoring of politicians, and the gauging of public opinion. Our study is based on a data set of 13.2 billion blog posts from Sina Weibo – the most prominent Chinese microblogging platform – over the period of 2009-2013. In contrast to the previous studies that conclude that collective action events are censored in Chinese social media, we find millions of posts discussing protests, strikes, and demonstration. Moreover, we find that microblog posts 1) are highly informative in predicting collective action events and corruption charges; and 2) have a significant and positive effect on the incidence of strikes and protests, although their effects on the incidence of large-scale massive conflicts and government-sanctioned demonstrations (e.g., anti-Japan) are muted. Finally, we find that the number of government microblog accounts, based on machine learning estimation, is larger in areas with a higher level of internet censoring and where newspapers have a stronger pro-government bias. Overall, our findings suggest that the Chinese government regulate social media to balance threats against regime stability against the benefits of 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

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 the 186th among 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 primarily study outcomes in three areas: collective action, the monitoring of politicians, and the gauging of public sentiment. These areas are central to understanding the pros and cons of social media to an authoritarian regime (Egorov et al. 2009, Lorentzen, 2014). As starkly illustrated by the role of social media in the Arab Spring uprisings against regimes in Tunisia, Egypt, and Libya, social media's role in facilitating collective action poses a potential threat to authoritarian regimes. However, the bottom-up information flows in media allow a regime to identify and solve social problems before they become regime threatening. This role of media has been an important part of the propaganda policy of the Communist Party in China (CPC hereafter), 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 a central regime is to monitor local governments and officials. In a large country like China, many political and economic decisions are decentralized to local governments. The implementation of policies and evaluation of local politicians require substantial efforts to collect information other than internal reports which are likely to be distorted. A prominent example is that in recent years, "online anti-corruption" has become an important complement to the "offline anti-corruption" campaigns launched

[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.

by the Chinese central government.

We examine four questions related to the effects of social media in the above three areas. First, how much do the microblog posts cover collective action events, people's grievances with public policies and explicit criticism of political leaders? 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 it became 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 all types of 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

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

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

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

in addition to commenting and re-posting. Because of its easy access and use, Sina Weibo soon became the most popular microblog in China. By 2010, 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 implemented largely 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. However, 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. King et al. (2013) find that the Chinese government allows criticism of officials and bureaucrats but prohibits information about collective action. Bamman (2012), Fu (2013) and Zhu (2013) find that internet censoring in China focuses on political and minority group issues. Zhu et al. (2013) find that the implementation of censorship is speedy: 30% of deletions happen in 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 taken out. This crucially depends on self-censoring: what people dare blog. The self-censoring in turn depends on the policing of government and the punishments of users.

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 politicians may use their own internet police to pursue their own goals. For example, out of career concern, they may suppress negative information about the regions under their administration, even if this information is tolerable or encouraged

[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.

by the central government. Local governments' control of social media is largely through the punishment of users who 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 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 levels 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 social outcomes that bloggers and the central regime have a common interest. This could include addressing local problems such as corruption, and small-scale strikes and protests. A caveat is that these effects could be muted by the 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 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, a stricter enforcement of regulation may be observed during periods of higher political tensions.

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

3 Data

Our primary data of Sina Weibo posts were collected by Weibook Corp. Since September 2009, this company has executed a massive data collection strategy, downloading the posts from active users. First, they identify users as authentic active persons according to individual information and the interaction with friends. In total, they locate 200-300 active million users. Second, according to the number of followers, they categorize users into 6 tiers. They download the microblogs of the top tier at least daily, and the downloading frequency decreases with tiers with the second to third tiers every 2-3 days and the last tier on a weekly basis. For a user 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 automatic censoring that is immediate, the data contains posts that are later censored. For each post, we obtain the information on the content, the posting time and user information (including self-reported location).

In total, the data set contains 13.2 billion posts from 2009 to 2013. For the sub-period of 2009-2011, we construct an independent measure of the volume of posts on Sina Weibo.[12] According to our estimation, the Weibook data contains approximately 95% of the posts on Sina Weibo. As illustrated in Figure 1, the blue line indicates 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 Weibook data base, we pull out 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 level, 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 with a very high frequency. We only collect a 10 percent random sample of the posts mentioning these words. The second group of key words refer to specific events that we have recorded, including events in censoring directives issued by the Chinese media control authorities and a large number of massive events from 2009 to 2013. In total, our 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 then remove stopwords, punctuations, URLs, usernames and non-Chinese characters other than meaningful English abbreviations such as "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 occur-

[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 the public timeline API was denied access.

rences).

4 Collective Action

Social media can be used to coordinate and organized collective action events. However, because of extensive policing and censoring, it is an open question whether they play such a role in China. To play a role, first, the relevant content must be present in social media. Second, social media coverage must precede, or at least be concurrent with, the events. Finally, social media entry should 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-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 massive conflicts between governments and the public. The second category is protests, 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 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 censoring in China, King et al. (2014) study four collective action events in China and 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.)

However, we find a large number of posts covering even the most sensitive collective action events in our classification. From our data set, we identify 382 thousand posts in the conflict category and over 2.5 million posts in the protest category. To describe the content, we characterize "hot topics" in posts about collective action. These topics are identified by words used more often in collective action posts than in the entire sample.[13]

[13] More precisely, we compare the frequency of each word in a given category with the overall

Table 1 presents the results. 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 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 in further detail, we investigate a random sample of 1,000 posts for each category. We manually code whether and how the posts actually cover a particular type of event; see Table 2. The share of posts that are actually about the events ranges from 50.4 percent for the anti-Japan events to 31.2 percent for strikes. To estimate the total number of posts in a given category, we should multiply the shares with the total number of posts in a 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 listed below:

- "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]

To investigate whether a regular user who posts this type of content is punished, we look at the 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, indicating that 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 Sina Weibo account, whose IP address is hidden and which contains few other posts that

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.

[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).

can be used to identify them. Then, 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 data patterns that indicate censoring or self-censoring. The collective action events in our data are substantial enough to catch media attention, and bloggers would write extensively about these events if they were not worried about being punished. For a significant number of events, we find no blogs posts at all. Table 3 tabulates two binary indicators for the presence of an event and any post about the event, respectively, at the prefecture-day level. For example, all the 46 cases of anti-Japan events are covered, presumably because these events are supported by the government. The number of events uncovered 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 is, the more likely it is censored and self-censored. Because we do not find much evidence that people self-censor when posting about these topics, censoring is a more likely explanation.

4.2 Predicting and Identifying Events

A key question is whether collective action events are caused or amplified by posts in social media. 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 are preceding or at least concurrent with the event.

4.2.1 Individuals

We study whether blog posts cover a forthcoming or ongoing event in a set of random samples of 1,000 posts. Compared to the less sensitive events, the more sensitive events (e.g. conflict and protests) are covered more in the format of general comments and in retrospect; see Table 2. In the conflict category, 1 in 1,000 posts covers a forthcoming event, and 11 covers ongoing events. For protests, the numbers are nearly doubled, and for strikes while anti-Japan demonstrations the numbers are much higher.

We next investigate whether social media activity predicts the collective action events. The last two lines of 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. For example, the large-scale conflicts are only covered in 6 posts on the event days, compared to 2,000 posts per day for the government sanctioned anti-Japan events. The average number of posts is much higher on the day and the day before a collective action events than on other days. This holds for all types of collective action. 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 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[15] and also controls for the total number of Sina Weibo posts in the entire Weibook data at the prefecture-month level. Standard errors are clustered by prefectures.

Table 4 shows that microblogs are highly significant in predicting where and when the collective action events take place. In contrast, newspaper coverage is uninformative. The fact that people start discussing events before they happen indicates that Sina Weibo may be used to organize or at least coordinate on the collective action events. This is the case even for the most sensitive "conflict" events, which pose a potential threat to the regime.

4.2.2 Government

We now turn the tables and ask: how useful are these data 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 Chinese government is risk-averse and has a high cost of missing regime-threatening events. Using terminology from information retrieval, 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 thus wants a classifier with high recall, but also desires a classifier with high precision, given the high cost of investigate a large number of false positives.

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 thus our data represent the same information that the government will use. Table 3 shows that there are anti-Japan posts in all 46 anti_Japan events. No anti-Japan event is under the social media radar, and the simple classifier has a perfect recall. On the other hand, this classifier has terrible precision, 46/(125,359+46). To find the 46 events, the government would have to investigate over 125 thousand prefecture-days. This is very costly.

We explore machine-learning methods to improve the classifier. In particular, we use a method similar to the one that is developed by Sasaki et al. (2010) to detect earthquakes based on twitter flows. This approach combines a Support Vector Machine (SVM) with regression analysis. Such an approach dramatically increases the precision (by 97%), but at a cost of missing 4 anti-Japan events. We are not able to increase precision further without reducing the recall. Hence,

[15] Newspaper coverage is taken from the WiseNews archives. For the 2010-2012 sample, this includes 62 general-interest newspapers.

a risk-averse regime may be forced to use very costly methods. Although it is possible that the Chinese government can develop better classification methods, the improvement is challenging given the fact that machine-learning techniques have not been very successful in identifying other large-scale events, such as flu epidemics (Lazer et al., 2014).

More generally, the above methods can be used for data collection. For example, it is easy to identify collective action events that are not in our original data. We investigate the 100 prefecture-dates where the predicted probability of a strike is 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 are not in our original data (as well as 14 strikes that are in our data).

4.3 Social Media Effects

Much has been written on the role of social media in triggering and coordinating protest activity during, for example, the Arab Spring. Nevertheless, there has been little systematic analysis of the impact of social media on protest activity (that we are aware of).[16] So far, we have provided evidence that Chinese social media contain an extensive and vivid discussion of collective action events, and that blog posts about the events precede and predict these events. While this suggests that posts on social media have helped coordinate the events, the evidence is far from conclusive. Note that collective action events could be also be affected by blog posts entirely different from those discussed above. For example, collective action events could be sparked by information about injustices, corruption, ethnic tensions or collective action events in other areas.

We explore these potential types of effects 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., increases in the use of Sina Weibo.

Table 5 shows the results from regressing a dummy for each type of event on the Weibo penetration (defined as $log\ (\#posts\ per_capita + 1)$). 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 above estimates suggest that the use of Sina Weibo increases strikes and perhaps protests, but not conflicts and anti-Japan events. The results are sensible. People do not need social media to organize anti-Japan events,

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

which are sanctioned by the government and can be organized other information channels, such as newspapers. The strikes and protests are of a small scale, confined within a small region, and do not threaten the central regime. The highly sensitive conflicts pose a potential threat to the regime and hence coverage of these is suppressed.

Our findings suggests that Chinese social media are unlikely to be regime threatening. The regime has contained the effect on Sina Weibo to protests and strike and conflict events are not 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.

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.[17] 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 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.[18] With the 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 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 "pe-

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

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

titioners"), or to different punishments, such as "reeducation through labor", "imprisoning", "police" and "specialized hospital".[19]

Furthermore, 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 forthcoming petitioning cases, 469 about past petitioning cases, and 406 about petitioning in general. There is little evidence that users who post petitioning content are punished. Users who post on this topic continue to be as active as other users, implying that their accounts are not closed. The petition posts are not generated from special accounts with few posts. Rather, they are 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 explains 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

It is unclear how much discussion of corruption among political leaders to expect in Chinese microblogs. On one hand, the central government and the regular social media users have a common interest in cracking down on corruption. On the other hand, corruption of local politicians reflects badly on the CPC and may undermine its legitimacy. For this reason, the central government may want to silence discussion about corruption. In addition, local politicians have strong incentives to suppress corruption coverage in their jurisdiction, because this may jeopardize their careers. Consequently, users may self-censor in order to avoid being identified and punished by the powerful local politicians.

To study the 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 major political positions at the central, provincial, prefectural, county, and village levels. In total, we obtain 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, 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 a higher level are more extensively discussed, and executive positions are more covered than the party secretaries are.

In the second category, we search for words that are widely used to describe corrupt behavior and wrongdoings and punishment of officials (e.g., bribery, cronyism, and the misuse of power). We find over 5.3 million posts in this

[19] 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.

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 them are just general comments on corruption. Of the 419 posts that talk about specific cases of corruption, 293 are written after the government has 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.

- "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!"
- "Billions of money went to 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 around 668,000 posts that discuss specific instances of corruption before government action. This is a wealth of information for upper-level governments that are seeking 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 do not fear 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 some posts that explicitly criticize the top national leaders, although the posts do not contain 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 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 are actually discussing particular instances of corruption, rather than, for example, discussing government anti-corruption policy. Using a widely used text classification algorithm - Support Vector Machines (SVM),[20] we estimate the probability that each post discusses a specific case of corruption. 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.[21] 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). To obtain estimated classification probability, we 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 probability of a post being in our corruption category and also mentioning a political position. 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 just report the government stance on corruption. The share of 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.

To get a broader measure of how people feel about their leaders, we also conduct a sentiment analysis on 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 sentiment words from the number of positive sentiment words. Column V in Table 7 shows the average sentiment in posts mentioning each political position.

Figure 3 plots the percentage of corruption posts in column IV against the average sentiment in column V. Sentiment is most positive for the national

[20] 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).

[21] 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).

politicians, who are also infrequently mentioned in connection with corruption. County and village party secretaries have the most negative sentiment as well as 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 for a few years. These village officers are less likely to be corrupt because they have strong career concerns and may have less effective control of resources and lack social networks.

5.2 Predicting Corruption Charges

We finally relate this coverage to 185 corruption cases involving high-ranking Chinese government or party leaders. Among these cases, 39 are 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 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 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 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 the corruption charge. Table 9 shows that the corrupt officials are mentioned in 49 posts on average within the 2-7 months before the corruption charge, while the non-corrupt officials are mentioned in 44.4 posts. However, the corrupt officials appear much more frequently in posts that mention our corruption words (3.9 compared to .4). A similar pattern is found in the posts 12-23 months before the charge. Given the substantial difference in the number of posts, it is not surprising that the 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 an official's name and corruption.[22] Columns II, IV and V include dummy variables for case indicators, which are

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

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.

A significant number of 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 investigate 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 into features of the microblog posts, outside of our corruption category. Of course, it is possible that there are false negatives in this group: corrupt officials who have not been charged.

To sum up, we find a massive volume of posts discussing corruption in Chinese microblogs. These posts help to identify political positions, regions, times and individuals whose incidence of corruption is likely to be high. This suggests that the improved monitoring of lower-level officials outweighs the bad publicity from corruption coverage, to the regime. It also suggests that local politicians are unable to impose self-censoring on the users.

6 Propaganda

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

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 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. However, these correlations should be absent if government accounts were used for neutral information. We will test these hypotheses using the measure of Bamman et al. (2012) of deletions of social media posts across regions and the measure of bias in Chinese newspapers in Qin et al. (2013). The latter measure is based on nine content categories, including leader mentions, cites of the official CPC news agency, and the coverage of regime critical stories.

Informative propaganda may be more effective on audiences that share the propagandist's view, while the effect of propaganda can be negative when the audience holds opposing views to message sender. 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 the support for Nazi policies in places with historically low anti-Semitism. Similarly, in a laboratory experiment, DellaVigna et al. (2014) find 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 lower consumers' valuation of social media. To the extent that the 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.

6.1 Identifying Official Accounts

We identify official government and newspaper accounts on Sina Weibo. 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. Similarly, we identify the official accounts for the main parts of Chinese legal system: People's Court and Procurate at all administrative levels. In this way, we identify a total of a set of 1,043 government accounts in our data.

For newspapers, we identify the official accounts representing the 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. To obtain the newspaper accounts, we first search the user names for newspaper names from the comprehensive newspaper directory used by Qin et al. (2013).[23] We identify 538 newspaper accounts. In our data, these identified government and newspaper users contribute 72,691 posts.

We investigate how well we can identify government 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.[24] Of the 786 users who were classified by the SVM to be government/newspaper users, 640 (81%) actually were. We compute the probability that each account is a government account using the methodology described in Section 5.

[23] We further compare the newspaper location from the newspaper directory with that from Sina Weibo. If this location information is matched, the corresponding accounts are coded as the official newspaper accounts. If the locations do not match, we then manually check by locating the microblog of the user in Sina Weibo by the user ID in our data.

[24] 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).

We then use the word frequencies of all 30 million users in our data to compute the probability that each user is a government organization or a newspaper. We compute the average such probability by province. This gives us a measure of the share government users on Sina Weibo in each province.

6.2 The Correlates of Government Content

We now investigate the geographical distribution of the share government users on Sina Weibo. Figures 4 and 5 plot the estimated share of government users against the share of deleted posts (from Bamman et al.) and the measure of media bias in Party Daily Newspapers (from Qin et al. 2013). 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 for in both cases). 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 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 a variable called CCPstronghold, which equals the share of counties in the prefecture touched by the Long March 1933-1935 or that were a part of a CCP Soviet before 1949. In the case that the Long March went through or the CCP Soviet located in the metropolitan center of the prefecture, CCPStronghold is set as one. Some 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 lower in areas with a high level of GDP and higher in CPC strongholds. The latter result is consistent with the view that propaganda is more effective in areas where audience shares the same ideology. The estimated share of government users also appears higher in areas closer to Beijing and in areas that are more populous.

7 Conclusions

We analyze a large data set of blog posts from 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 the 200 politicians who were charged with corruption with a sample of similar politicians who 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 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. "The Power of the Street: Evidence from Egypt's Arab Spring". No. w20665. National Bureau of Economic Research, 2014.

[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, 09.

[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. "Why resource-poor dictators allow freer media: A theory and evidence from panel data." American Political Science Review 103.04 (2009): 645-668.

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

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

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

[15] 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.

[16] Human Rights Watch, "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, 2013, 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, 1999.

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

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

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

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

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

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

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

[27] 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

[28] (Yuanxi and Jing, 2009)

[29] 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.

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

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

Figure 1: Estimated number of Sina Weibo post by Weibook and hei-ya data

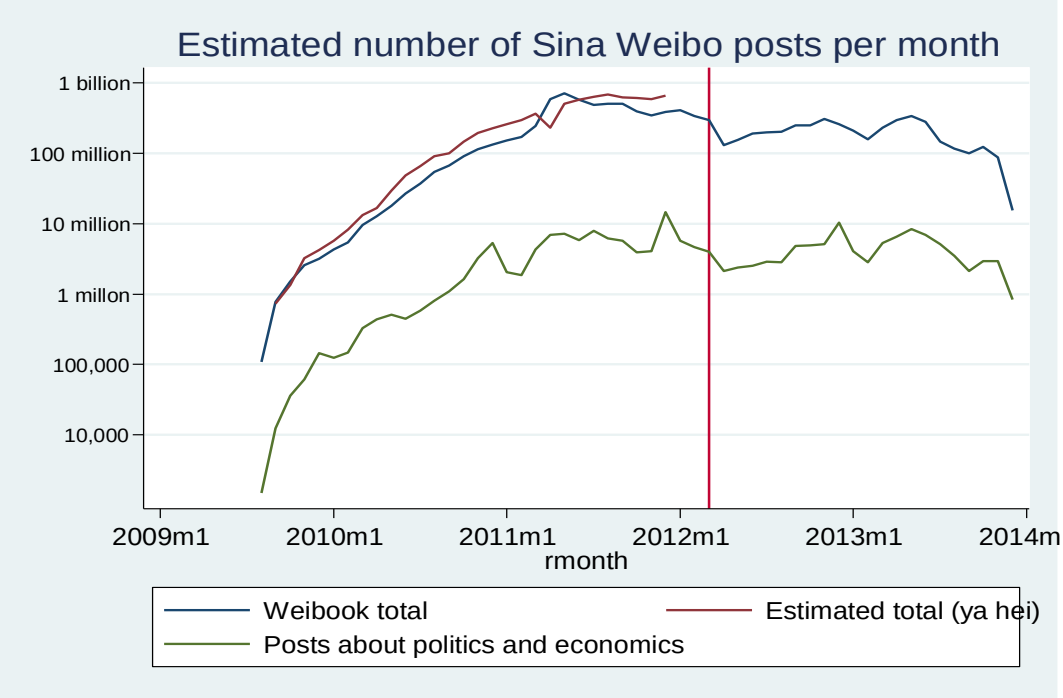
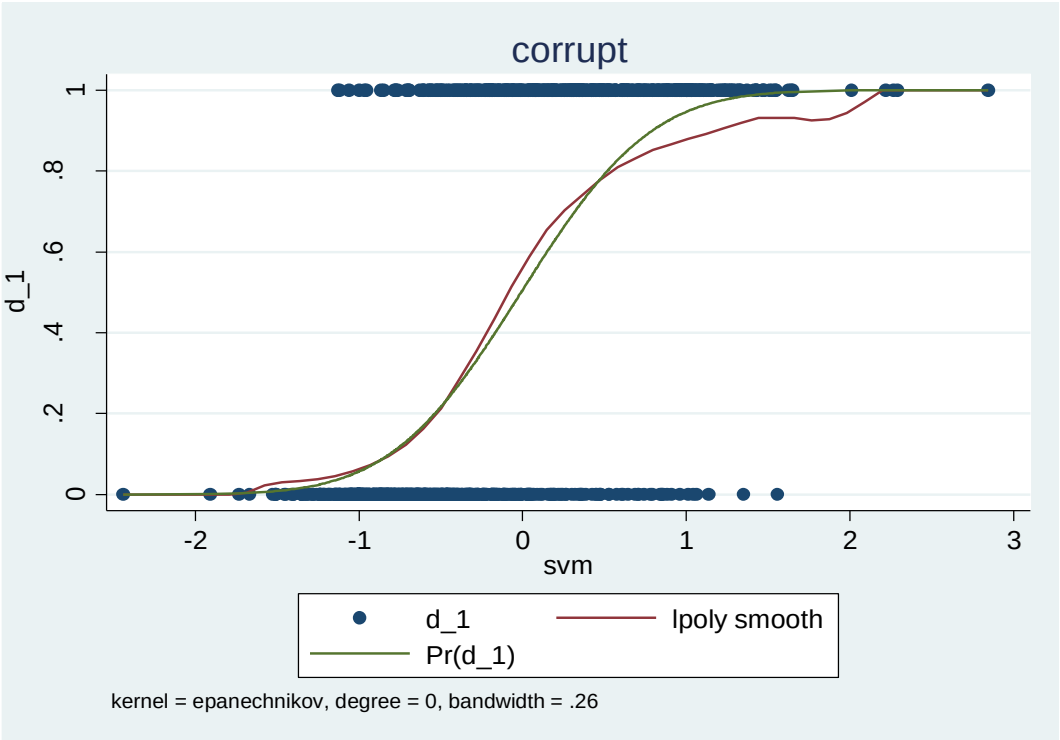


Figure2: The predicted probability and the non-parametric average share corruption posts



Note: The predicted probability is predicte by the probit regression, with the dependent variable as the indicator for the post being about corruption and the independent variables as SVM classification parameter. The red line shows the smoothed value obtained by a kernel-weighted local polynomial regression of the indicator of the post being about corruption on SVM classification parameter.

Figure 3 Sentiment of posts against the percent posts about corruption

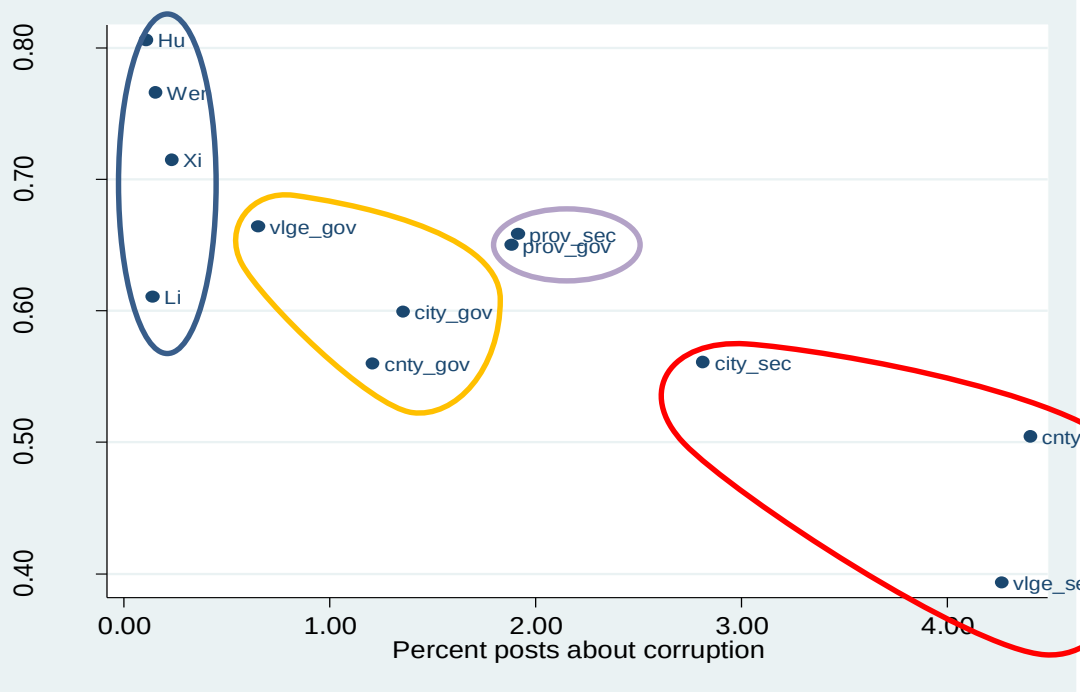
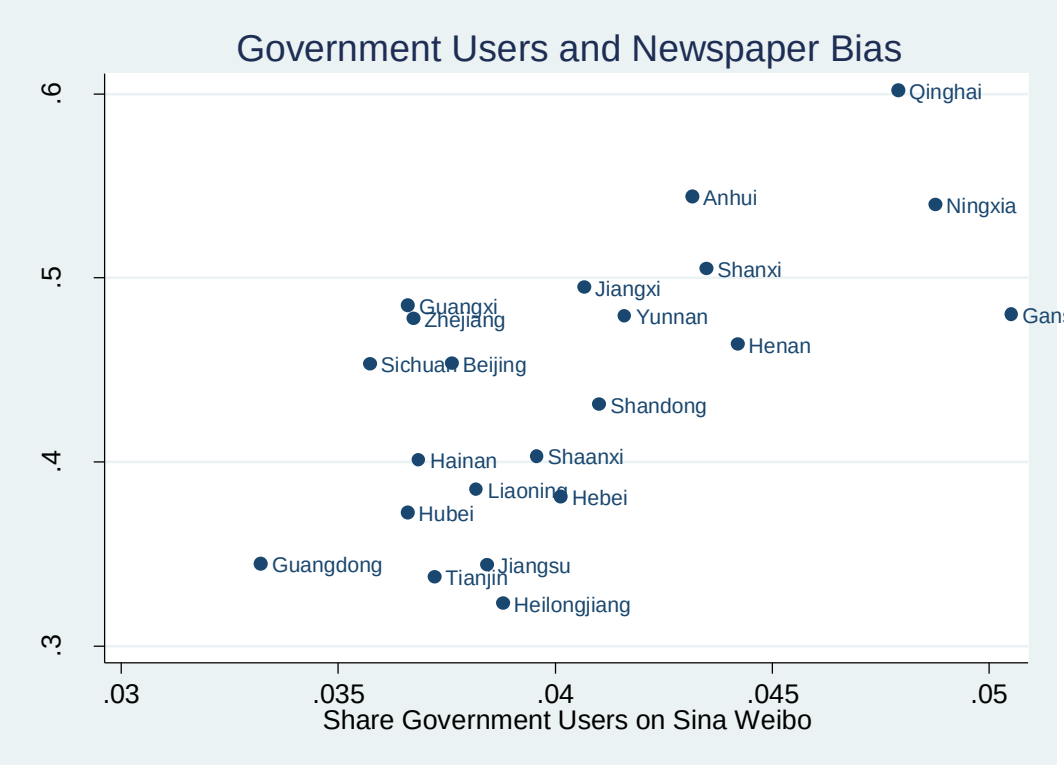
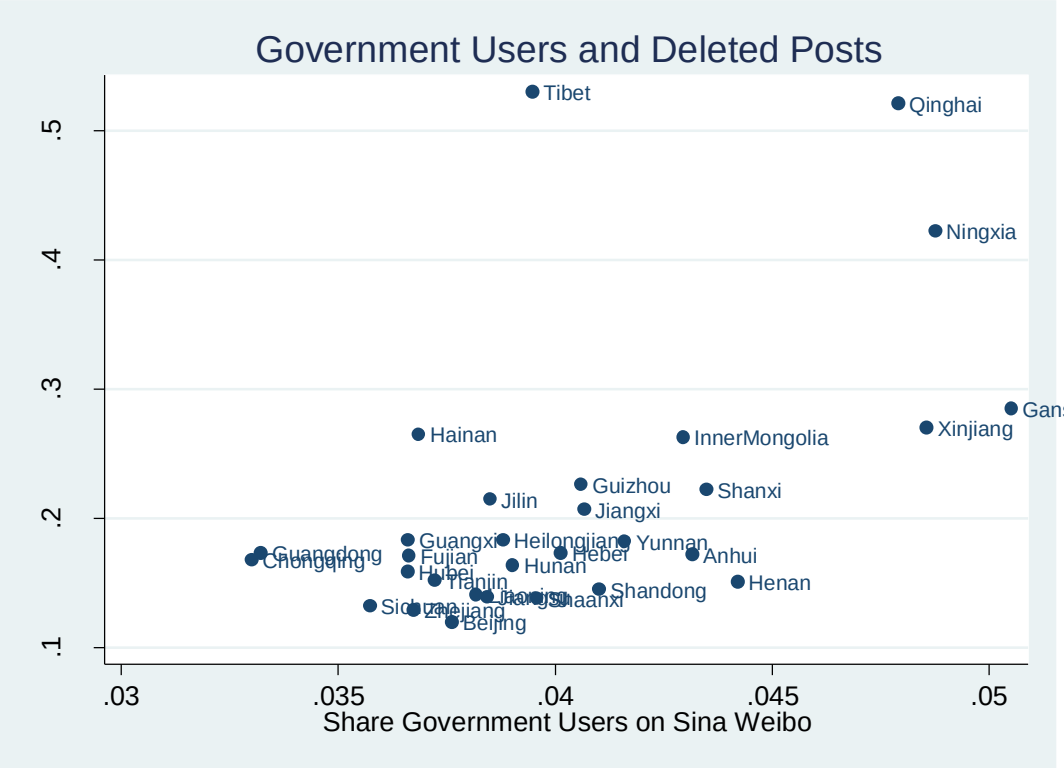


Figure 4 Newspaper bias and the share of government users on Sina Weibo



Note: Each dot represents one province in China.

Figure 5 Share of deleted posts and the share of government users on Sina Weibo



Note: Each dot represents one province in China.

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

	Total	Random 1000 post sample				
		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

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

	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	0.72
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