



Supplementary Materials for

Reverse-engineering censorship in China: Randomized experimentation and participant observation

Gary King,* Jennifer Pan, Margaret E. Roberts

*Corresponding author. E-mail: king@harvard.edu

Published 22 August 2014, *Science* **345**, 1251722 (2014)
DOI: 10.1126/science.1251722

This PDF file includes:

Materials and Methods

Figs. S1 to S7

References and Notes

Contents

1	Introduction	2
2	Topic Details	2
3	Account Blocking	5
4	Examples of the Text of Social Media Posts	6
4.1	Examples of Uncensored critical posts about Xi Jinping	6
4.2	Examples Illustrating Important Topics	7
4.3	Examples Illustrating Discussions of Deep Political Reform	10
5	Mediation Analysis	11
6	More Information on Research Ethics	12
7	Website Screen Shots	13

1 Introduction

This paper includes supplementary materials for this paper:

Gary King; Jennifer Pan; and Margaret E. Roberts. 2014. “Reverse Engineering Chinese Censorship: Randomized Experimentation and Participant Observation.”

Detailed replication information and data for the paper can be found in:

Gary King; Jennifer Pan; Margaret E. Roberts, 2014, ”Replication data for: Reverse Engineering Chinese Censorship: Randomized Experimentation and Participant Observation”, <http://dx.doi.org/10.7910/DVN/26212>, UNF:5:K/LGmB0vjskGYBobxbT+8g== IQSS Dataverse Network [Distributor] V3 [Version]

2 Topic Details

Here we offer details about the collective action and non-collective action events we found and used in the paper. Also included are the two edge case events. We list the events used within each of the three rounds of our experiment, by round.

Round 1

- CA event 1, Tibetan Self-immolation: Pro-government posts attribute the tragedy of her death on the Dalai Lama who is instigating these tragedies, Anti-government posts attribute her death to government policies.
- CA event 2, Protest in Panxu village over illegal land seizure: Pro-government posts say that this sort of protest and violence is wrong and that the villagers are greedy and want money. Anti-government posts say the local officials are unfair to the villagers.
- Non-CA event 1 Corruption Policy: new policy that bribes over 10,000 Chinese yuan will be subject to criminal investigation and penalties. Pro-government posts support this policy because it will reduce corruption. Anti-government posts believe this policy is punishing those who give bribes but the real fault lies with officials

who accept bribes and not those who are forced by the system to give bribes in order to get things done.

- Non-CA event 2, Eliminate Golden Week: people were calling for removal of the 10 day holiday that occurs during China's National Day. Pro-government posts support the 10 day holiday, saying that it stimulates domestic consumption, tourism revenues, stimulates economic development, and allows everyone to relax to promote social harmony. Anti-government posts call for removal of the policy because millions of people traveling at the same time is unsafe and unsanitary and the government should heed the call of the many people who are calling for the government to abolish the Golden Week holiday.
- Non-CA event 3, Rental tax: several cities in China are piloting taxes for renting housing (charging taxes on their rental income), which stimulated a lot of discussion and debate. Pro-government posts support the rental tax because it is income that should be taxed, just as income from salaries and wages are taxed. Anti-government posts criticize the tax saying it will increase already high rental taxes as landlords will push the tax onto renters.
- Non-CA event 4, Yellow Light fines: China promulgated new traffic regulations, which generated debate, especially the part that running yellow lights will incur punishment and fines. This debate prompted the authorities to say that punishment will be in the form of education, not fines or harsher penalties. Pro-government supports the new policy because it will improve transportation safety, and says that education not punishment is what's needed. Anti-government rejects and criticizes the authorities for not upholding the spirit of the law (i.e., education is not punishment).

Round 2

- CA event 1, Dissident Ai Weiwei releases a new album called Divine Comedy: Pro-government criticizes Ai Weiwei for releasing the album. Anti-government supports Ai Weiwei's actions and the album.

- Non-CA event 1, Shanghai Stock Market crash: Steep decline in the Shanghai stock market (the largest single day decline in the past four years). Pro- government says the government has done everything it can to regulate financial markets and this crash is the work of speculators and hackers. Anti-government posts say the stock market crashed and caused hardship to ordinary investors because of bad government interventions, policies, and actions.
- Non-CA event 2 (Corruption), Investigation of Sichuan Vice-Governor Guo Yongxiang: Guo is being investigated for serious breaches of discipline (i.e., corruption). Guo was a member of the Sichuan Province Standing Committee and a Vice Governor. Pro- government says the investigation is good because it will cut down on corruption. Anti-government says all officials are corrupt and Guo is being investigated for other political reasons.
- Edge case 1, Online Protest of Child Abuse. Pro-government posts we wrote criticize Ye Haiyan and this form of protest as unproductive and harmful to social order. Anti-government posts support Ye and criticize a corrupt educational system.

Round 3

- CA event 1, Protests in Xinjing. Pro-government posts calls this an act of terrorism against the Chinese people. Anti-government posts say that this event may be due to forced housing demolition instead of terrorism.
- Non-CA event 1 (Corruption), Li Tianyi Scandal: Li Tianyi is the son of a famous People's Liberation Army performer, Li Shuangjiang. The Beijing police department announced that Li Tianyi and four other young men gang raped a young women on February 17, 2013, and that investigation of Li has been completed. Pro-government posts say the government did a good job arresting Li, even though his father is well connected. Anti-government posts say the government is not doing enough, and asks why the other four participants have not been named.

- Non-CA event 2, Gender Imbalance: new report released by the National Statistics bureau says that by 2020, China will have 30 million “bare branches” (extra men). Pro-government says that is the results of backwardness and preference for boys in rural China. Anti-government says that this is the result of the China’s one-child policy.
- Edge case 1, Hong Kong protest. Pro-government criticizes these protests are trouble-making and disruption to social harmony. Anti-government says the protests are a means of expression for better government and democracy.

3 Account Blocking

In addition to automated review, and content filtering by censorship, some entire accounts are sometimes blocked, which is another form of information control. We did not design our experiment to study blocking, but we are able to glean some important information about it anyway. Under our experimental design, each social media account we set up ultimately had the same number of collective action related posts. However, blocking can occur at any time, and at different times during our experimental protocol, each account had submitted different numbers of collective action related posts. In addition, censorship of collective action posts was not perfect and so we can also leverage these differences as well. Figure 1 gives the basic relationship among sites that use blocking as a tool. It shows that once the percent censored on an account (see the horizontal axis) hits a rate of at least 60-80%, the probability that that account will be blocked (vertical axis) more than doubles.

We also study whether censorship acts as a mediator between collective action posts and blocked accounts. Using the same methods again, we find an average mediation effect of 0.17 with a 95% confidence interval of (0.09,0.25). This means that censorship alone, independent of content and the collective action content of posts, is what alerts the internet service provider to accounts with collective action content, making them more likely to block the offending account from posting further. Blocking thus appears to be a relatively automated process that is calculated from the number of posts that were censored from

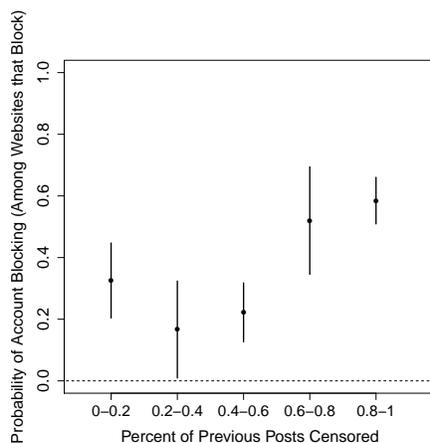


Figure 1: How Blocking is a Function of Prior Censorship

previous attempted posts. It does not seem to be the subject of separate analysis or human judgment in many cases.

4 Examples of the Text of Social Media Posts

We give some posts for the volume burst in two categories: critical posts that are not censored and posts that illustrate that important topics are covered.

Below we include the original chinese text, as well as an English translation. Note, however, that many of these posts are difficult to translate because they use slang, homonyms, and various types of clever wordplays. So we choose to include a literal translation, as well as an explanation for the intended meaning. (A side point is that we translate the food 包子 as bun instead of dumpling, since both are dough products with pork based filling but the bun that Xi ate is steamed, while dumplings are boiled.)

4.1 Examples of Uncensored critical posts about Xi Jinping

三国时郭嘉预测孙策必死于刺客之手，朋友问我旁边跟习大大合影的人如果左手抽刀插入——习大大在福建浙江上海执政的时候从来没有“亲民”，为何到了北京就变得“亲民”了呢？个人认为仅仅是作秀而已

Literal translation: During the Three Kingdoms period, Guo Jia predicted Sun Ce would die at the hands of an assassin, a friend asked me what would have happened if one

of those people taking pictures with Xi at the dumpling shop had pulled out a knife and stabbed him — when Big Xi was governing in Fujian, Zhejiang, and Shanghai he never got close to the people, why is it that now he's in Beijing he's getting close to the people? In my opinion, he's just putting on a show

Explanation: The second part of this passage accuses Xi of faking his care toward the people, specifically “getting close to the people” means getting to know ordinary citizens and showing them you care.

谁要吧把习大大弄来草包一次，肯定比庆丰包子强

Literal translation: who can make big Xi into a grass bun, it'll definitely be better than Feng Qing buns

Explanation: to make into a grass bun means to make a fool of someone. Big Xi is northern slang for Uncle Xi, a very colloquial way of referring to Xi Jinping. The passage means that to make a fool of Xi Jinping would surely be better than Feng Qing buns.

我觉得关注明星的事情没有什么大不了的，总比关心习总吃什么包子穿什么衣服彭妈用什么包然后鼓吹个什么光明盛世国货雄起之类的恶心论调好得多，在那些人眼里他们也不过是消费领导人罢了

Literal translation: I think there's nothing wrong with paying attention to celebrities, and it's infinitely better than paying attention to what kind of bun president Xi is eating, what he is wearing, what handbag Mother Peng (Xi's wife) is using and what bright shiny new product she is promoting, and other sorts of sickening behavior, that shows they (Xi and his wife Peng) are just leaders of consumerism

Explanation: pretty much what the translation says, but the wording, e.g., Mother Peng is so casual that it's clearly making fun of Xi and his wife.

4.2 Examples Illustrating Important Topics

尊敬的习大人，您一顿包子21元，那一天三顿饭光吃包子就算50元好了。一个月30天，单算全吃包子就要1500元。加上老婆孩子全吃包子总共4500元。请问上海一个月5000元的工资能干什么？全家都吃包子都不够！

Literal translation: Respected Xi, you ate a meal of buns for 21 CNY, if you ate buns for all the meals, then that would be 50 CNY a day. There are 30 days in a month, if you only ate buns that would be 1500 CNY. If you add in a wife and a child, that's 4500 CNY a month. My question to you: what can you get with Shanghai's monthly wage of 5000 CNY? It's not even enough for a family just eating buns!

习总书记，能不能给无家可归的人吃个包子？“这无家可归的人能吃到包子吗？？？是什么让她沦落街头？？天朝百姓的包子梦能实现吗？？维稳领导会不会让她吃锤子、刀子？？？”

Literal translation: President Xi, can you give some buns to the homeless? Do the homeless get to eat any buns?? How do these people end up on the streets?? Can the people achieve their bun dreams?? Will leaders who want to maintain social stability force these people (the homeless) to eat hammers and knives???

Explanation: bun dream is a play on Xi's China Dream initiative. Eating hammers and knives refers to the homeless getting beaten up implying China's leaders only care about maintaining stability not improving people's lives.

习近平先生，上回你吃包子，我评论了几句，把我的号给封了。这事，我知道不是你直接干的，但是，我想知道你对此类事件的态度。这算是侵犯公民言论自由权利的行为吗？如果你的回答让我不满意，对不起，我不会投你的票。习近平先生，告诉你一个秘密，如有选票的话，我真投你的票，原因嘛，你毕竟让吃喝玩乐的收敛了一些。但是我选举你的意愿没有得到你的尊重，因为我连选票是什么样子都不晓得，自然我就不高兴了，所以我不想选举你了。看到你讲尊重台湾人民，我又有了选举你的意愿，那前提是你尊重我一回，给一张选票。

Literal translation: Mr. Xi Jinping, last time you ate steamed buns, I wrote a few sentences, and my social media account got shut down. I know you're not directly responsible for that, but I want to know what you think of it (my account getting shut down). Do you regard that as a violation of citizens' rights to freedom of expression? If your answer

doesn't satisfy me, sorry, but I won't vote for you. Mr. Xi Jinping, I want to tell you a secret, if China had elections, I would vote for you, this is because you're kind of reining all that eating, drinking and entertainment. But you don't respect my desire to elect you because I have no idea what a ballot even looks like, so of course I'm not happy so I wouldn't vote for you.

Explanation: Xi has said that Chinese people should have freedom of expression; eating, drinking, and entertainment refers to the corruption of officials.

习总昨天去庆丰包子铺吃包子后，各地一把手会不会掀起去街边小吃店吃饭的高潮呢？天津书记吃起了“狗不理”包子，河北书记吃起了驴肉火烧，陕西书记吃起了肉夹馍，山东书记吃起了煎饼卷大葱，山西书记吃起了刀削面。。。2014年小吃年

Literal translation: After President Xi ate Feng Qing steamed buns yesterday, all of the local top leaders are now eating street foods. Tianjin's party secretary ate "goubuli" buns, Hebei party secretary ate donkey fire (a local dish), Shaanxi party secretary ate Rou jia mou, Shandong party secretary ate onion pancake rolls, Shanxi Party Secretary ate sliced noodles...2014 is the year of local delicacies.

Explanation: China's local leaders all follow the top leader, even when it's something as ridiculous as eating steamed buns and street food.

中国不是奴隶社会，但绝对是奴才社会

Literal translation: China is not a slave society, but it is definitely a society where everyone sucks up to leaders (lackey society).

Explanation: this is a tweet quoting the news that the manager of the Feng Qing shop has preserved the table where Xi sat.

包子就是包子，热度高了还是容易露馅

Literal translation: A bun is a bun, if it's cooked too long, the filling will come out

Explanation: this is in response to a post that says the biggest lie occurs when the liar believes what he is saying, which was in turn a post responding to news of Xi Jinping saying China has tried Constitutionalism, parliamentarism, multiparty presidential system,

and none are feasible for China. Basically someone is calling Xi a liar, and this person uses the bun analogy to make the point that eventually (when the bun is cooked too long) the truth (the filling) will be revealed.

4.3 Examples Illustrating Discussions of Deep Political Reform

想群众成为真正的主人，只有终结集权专政，建立人手一票、多党竞争 的民主制度

To truly give power to the masses means ending the centralized dictatorship, and adopting multi-party democracy where everyone has a vote.

别提反腐了，甚至现在中国的民主、法治、人权都不要再提了，已经在二十多年前彻底死在“履带”之下了！所谓公审，就是演戏，骗老百姓呢罢了。在民主、法治、人权相对健全的国家，治理腐败绝对不是靠国家，有效治理腐 败靠的是多党竞争，如果是姓共的监管姓共的，这不是天大的笑话吗！？醒来 吧国民！

Don't talk any more about anti-corruption, don't even talk about China's current democracy, rule of law, and human rights; those things were completely destroyed over twenty years ago under the treads of tanks! Our so-called court of law is merely for show, to deceive ordinary citizens. In a country with real democracy, rule of law, and human rights, corruption is not fought by the state; effective anti-corruption efforts rely on multi-party competition; the Party (CCP) providing oversight on the Party is an absolute joke. Wake up Chinese citizens!

宪政是多党竞争更是多党监督，在朝可施展抱负在野负监督之责。在 民选国家，公民个人和任何政党，有为才有位有为就有位，没有官位还是公 民，失去执政地位不是失去政治地位和价值。对于一个正常的政党，宪政只会 令其保持生机，绝无令其亡的道理。所谓亡党亡国之忧，罪孽自知罢了。

Constitutionalism is multi-party competition and more importantly multi-party supervision, where parties in office work toward their goals, and parties out of office provide

oversight. In a country with popular elections, citizens and parties with promise gain power, and when out of office, (former) politicians become ordinary citizens, losing political office does not lead one to lose political rights. For any normal political party, losing office doesn't lead to its demise. Those who imply constitutionalism will lead to the demise of the Party and the country are guilty of ignorance.

5 Mediation Analysis

To estimate the influence of automated review in the causal path between collective action and censorship, we follow [Imai et al. \(2011\)](#) in conducting a mediation analysis. We are interested in how the treatment (T), collective action, is mediated through mediator (M), automated review, on an outcome (Y), censorship. Let $Y_i(t)$ denote the potential outcome Y for individual i under treatment T . The causal effect within the potential outcomes framework of the treatment on the outcome is $Y_i(1) - Y_i(0)$, or censorship when the post discusses a collective action event, minus censorship had the post discussed a policy event instead. Let $M_i(t)$ define the potential value of automated review under treatment T for individual i (1 if held for review and 0 if not). The mediation effect of automated review on censorship is: $Y_i(t, M_i(1)) - Y_i(t, M_i(0))$ for treatment value $T = t$. That is, we are interested in the difference between (a) censorship when the post discusses a collective action event and (b) censorship when the post was about a collective action event, but its review status was as if it were not a collective action event. The idea, then, of mediated causal effects is to control the mediator and see what happens to the outcome variable.

To estimate the average mediation effect, we fit two models, an outcome model and a mediator model. The outcome model is $Y_i = \beta_0 + \beta_1 T_i + \beta_2 M_i + \beta_3 T_i M_i$, and the mediator model is $M_i = \theta_0 + \theta_1 T_i$. We allow for interaction between the mediator and the treatment because it is plausible that the mediation effect is stronger for collective action posts. In a simple linear structural equation model, the average causal mediation effect is $\theta_1(\beta_2 + \beta_3)$.¹

In order to accurately estimate the mediation effect, the assumption of sequential ig-

¹We complicate this estimator further, following [Imai et al. \(2011\)](#), to allow for a binary outcome and mediator.

norability must hold. First, the treatment must be ignorable conditional on pretreatment covariates. Because collective action posts are randomized in our experiment, this assumption holds. Second, the mediator must be ignorable conditional on the treatment. We can not think of a covariate that influences both review and censorship, conditional on treatment. Therefore, we believe that this second condition holds.

Following the above procedures, we find and report in the paper that review does not mediate the relationship between collective action posts and censorship. This further consolidates our opinion that review is an ineffective technique in finding what the Chinese government is after — collective action.

6 More Information on Research Ethics

The introduction to Section 2 of the paper discusses IRB approval and the fact that we imposed rules on ourselves that go far beyond those required by the IRB, which only protect research subjects according to the specifics of US law. The main feature of our additional rules was that we attempted to avoid influencing the system we were studying. Adhering to this rule has several purposes. First, it is standard practice in social science to attempt to avoid Hawthorne effects and any other type of investigator-induced change in the system that may confound the intended treatment effect of the experiment. Second, violating this rule would likely increase the probability that the experiment would be interrupted by some of the many social media companies where we submitted posts, or many parts of the Chinese government; this is especially crucial since most large scale social experiments fail for because of “unexpected” political interventions (King et al., 2007). The safety of our large research team was also a continuous concern and the subject of a host of other procedures we followed that we cannot detail here (legitimate researchers are welcome to contact us for more information).

We verified our adherence to the rule, to some degree, by studying whether and how anyone else on Chinese social media responded to our randomly assigned posts. As it turned out, we didn’t find a single case in which one of our posts was commented on by another netizen that indicated that they had a hint that the post was part of an experiment.

We also studied the pattern of posts, censoring, and commenting and found it to be close in a variety of ways to that which we obtained by pure observation.

For our participant observation, we were extremely careful to only ask questions that would normally be asked in China by participants in social media in China. Our questions to customer service professionals included no deception. We actually did set up a social media site, we paid for everything we accessed that had a cost, and we did want to know the answers to the questions we asked. Customer service answered questions for us that were similar or identical to those people in Chinese social media companies ask them regularly. We asked no questions about the people answering our questions or about any other human subject. We followed extensive procedures to avoid standing out as different in any way from the thousands of others making inquiries and doing what we did for other purposes. The underlying motivation of almost anyone (including us) posing questions to customer service was not raised, nor was it relevant to the workers. We kept no records of any name or online ID of anyone who provided us information. We believe it would be exceedingly difficult to find out who we spoke to, but even if someone did, what we asked and what they told should be nearly indistinguishable from numerous other interactions that occur in China all the time.

7 Website Screen Shots

The following are full screen shots of websites that censor, review, and block accounts from posting.

Figure 2: Censorship



Sorry, the host you were looking for does not exist, has been deleted, or is being investigated



Figure 3: Review



Administrator has set up review for new topics, your post will be shown after it has passed review.

Figure 4: Review

论坛 > 生活·消费·服务 > 房产家居 > 我代表广大群众强烈谴责这次新疆鄯善暴力恐怖事件袭击者 ...

Haier 海尔专卖店 贺仪(白电)连续4年全球销量第一 昌吉新邦海尔专卖店“不一样智冷节晚宴”开启 时间: 7月26日-28日

发帖 回复 返回列表

查看: 2 | 回复: 0

[房市楼盘] 我代表广大群众强烈谴责这次新疆鄯善暴力恐怖事件袭击者, 相信政府能全力以赴解决好... (审核中) [复制链接]

发表于 2013-7-4 06:23 | 倒序浏览 | 取消回复通知 1楼 电梯直达

我代表广大群众强烈谴责这次新疆鄯善暴力恐怖事件袭击者, 相信政府能全力以赴解决好此事, 安稳新疆不被那些别有用心
的敌对分子侵犯, 保证祖国疆土完整。作为人民虽然日常生活中还是没有尽如人意, 但我感觉还是幸福的, 至少还能活着,
享受着人生冷暖。

On behalf of ordinary citizens I strongly condemn the perpetrators of the terrorist
attack in Shanshan, and I believe the government will do everything in its power to
adequately resolve this matter...(under review) [copy link]

威望 3
UID 588335
帖子 3
精华 0
银币 40
羊肉串 0
积分 6
注册时间 2013-2-13
最后登录 2013-8-6

论坛游侠
威望 3
UID 588335
帖子 3
精华 0
银币 40
羊肉串 0
积分 6
注册时间 2013-2-13
最后登录 2013-8-6

补充

发帖 回复 返回列表

Figure 5: Review

个人空间
zhangjunlove
7486498 [收藏] [复制]

空间首页 动态 记录 日志 相册 主题 分享 好友 留言板 个人资料

主题	版块/群组	回复/查看	最后发帖
网上这几天有关纪英男的事情好像炒得很凶啊 已忽略	长城杂谈	0 / 2	zhangjunlove 2013-7-12 08:27
在这次建群联系某论坛上发帖 回收站	时事评论	0 / 1	zhangjunlove 2013-7-11 08:40
有一个站被某论坛上建群 回收站	时事评论	1 / 1	zhangjunlove 2013-7-11 08:40
重点少	网友爆料	0 / 450	zhangjunlove 2013-7-4 05:06
新疆维吾尔族这个民族的文化真是和我们汉人不夫一样 回收站	论坛活动	0 / 1	zhangjunlove 2013-7-1 10:30
恢复成以前的原样	行军日记	0 / 0	shankunlove
恢复成以前的原样	行军日记	479 / 0	2013-4-21 00:09
澳门同安西科潘翰村抗说担任土地 回收站	时事评论	80 / 0	zhangjunlove 2013-4-19 09:23
**是相对的, 不是绝对的	媒体焦点	0 / 498	zhangjunlove 2013-1-17 10:50
阅读灯暂不是比较好	热门话题	0 / 578	zhangjunlove 2013-1-17 10:44
乘黄河谷之旅	驴友俱乐部吧	0 / 64	zhangjunlove 2013-1-11 06:19

Post reviewed then published

Post reviewed then deleted

Powered by Discuz! X2
© 2001-2011 Comsenz Inc.

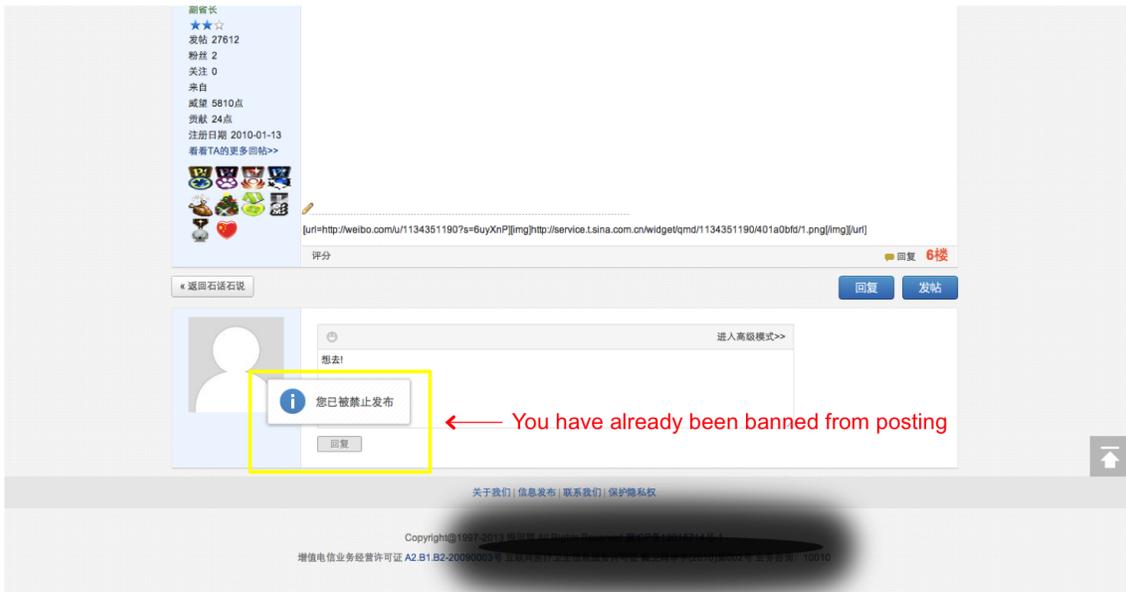
手机版 | 联系我们 | 长城论坛

GMT+8, 2013-8-6 01:48, Processed in 0.029291 seconds(6), 9 queries, Apc On.

Figure 6: Account Blocking



Figure 7: Account Blocking



References and Notes

1. Freedom House, "Freedom of the press, 2012"; www.freedomhouse.org.
2. G. King, J. Pan, M. E. Roberts, How censorship in China allows government criticism but silences collective expression. *Am. Pol. Sci. Rev.* **107**, 326–343 (2013).
[doi:10.1017/S0003055413000014](https://doi.org/10.1017/S0003055413000014)
3. S. L. Shirk, *China: Fragile Superpower: How China's Internal Politics Could Derail Its Peaceful Rise* (Oxford Univ. Press, New York, 2007).
4. S. L. Shirk, *Changing Media, Changing China* (Oxford Univ. Press, New York, 2011).
5. M. K. Whyte, *Myth of the Social Volcano: Perceptions of Inequality and Distributive Injustice in Contemporary China* (Stanford Univ. Press, Stanford, CA, 2010).
6. L. Zhang, A. Nathan, P. Link, O. Schell, *The Tiananmen Papers* (Public Affairs, New York, 2002).
7. A. Esarey, Q. Xiao, Political expression in the Chinese blogosphere: Below the radar. *Asian Surv.* **48**, 752–772 (2008). [doi:10.1525/AS.2008.48.5.752](https://doi.org/10.1525/AS.2008.48.5.752)
8. R. MacKinnon, *Consent of the Networked: The Worldwide Struggle For Internet Freedom* (Basic Books, New York, 2012).
9. P. Marolt, Grassroots agency in a civil sphere? Rethinking internet control in China. In *Online Society in China: Creating, Celebrating, and Instrumentalising the Online Carnival*, D. Herold, P. Marolt, Eds. (Routledge, New York, 2011), pp. 53–68.
10. M. Dimitrov, The resilient authoritarians. *Curr. Hist.* **107**, 24–29 (2008).
11. P. L. Lorentzen, Regularizing rioting: Permitting public protest in an authoritarian regime. *Q. J. Pol. Sci.* **8**, 127–158 (2013). [doi:10.1561/100.00012051](https://doi.org/10.1561/100.00012051)
12. P. L. Lorentzen, China's strategic censorship. *Am. J. Pol. Sci.* **58**, 402–414 (2014).
[doi:10.1111/ajps.12065](https://doi.org/10.1111/ajps.12065)
13. X. Chen, *Social Protest and Contentious Authoritarianism in China* (Cambridge Univ. Press, New York, 2012).

14. E. Malesky, P. Schuler, Nodding or needling: Analyzing delegate responsiveness in an authoritarian parliament. *Am. Pol. Sci. Rev.* **104**, 482–502 (2010).
[doi:10.1017/S0003055410000250](https://doi.org/10.1017/S0003055410000250)
15. G. Distelhorst, Y. Hou, Ingroup bias in official behavior: A national field experiment in China. *Q. J. Pol. Sci.* **9**, 203–230 (2014). [doi:10.1561/100.00013110](https://doi.org/10.1561/100.00013110)
16. D. Bamman, B. O’Connor, N. Smith, Censorship and deletion practices in Chinese social media. *First Monday* **17** (no. 3) (March 2012). [doi:10.5210/fm.v17i3.3943](https://doi.org/10.5210/fm.v17i3.3943)
17. T. Zhu, D. Phipps, A. Pridgen, J. Crandall, D. Wallach, The velocity of censorship: High-fidelity detection of microblog post deletions. In *22nd USENIX Security Symposium* (Washington, DC, 14 to 16 August 2013; www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/zhu).
18. R. MacKinnon, China’s censorship 2.0: How companies censor bloggers. *First Monday* **14** (no. 2) (February 2009). [doi:10.5210/fm.v14i2.2378](https://doi.org/10.5210/fm.v14i2.2378)
19. See supplementary materials on *Science Online*.
20. We also added our own ethics rules, not required by the IRB, which dictate that we avoid, wherever possible, influencing or disturbing the system we are studying (19). The similarity to the Prime Directive in *Star Trek* notwithstanding, this seems like the appropriate stance for scientists attempting to understand the world, as distinct from advocates trying to change it, and in any event is more likely to yield reliable inferences.
21. In the process of setting up the site, they recommended that we hire two or three censors for every 50,000 users. That enables us to back out an estimate of the total number of censors hired within firms at between 50,000 and 75,000, not counting censors within government, 50 Cent Party members, or the Internet police.
22. See (40, 41) for numbers of registered users, which are substantial even if we account for automated sites created by marketing firms (42).
23. K. Imai, G. King, E. Stuart, Misunderstandings between experimentalists and observationalists about causal inference. *J. R. Stat. Soc. Ser. A* **171**, 481–502 (2008).
[doi:10.1111/j.1467-985X.2007.00527.x](https://doi.org/10.1111/j.1467-985X.2007.00527.x)

24. All posts were made to mainland China accounts. Some were submitted from outside China, when feasible, and many from within China. Recent work has noted that overseas accounts are subject to less stringent censorship regulations than mainland accounts (43). This issue does not affect our work because all accounts created and used were mainland China accounts. Users could control account location when creating the account by specifying a location in China, by entering a local mobile number, or by creating the account from a local IP address. We used all of these methods.
25. For each of our three rounds, we wrote 200 posts on non–collective action events (split equally between pro- and anti-government) and 200 posts on collective action events or edge cases (again split equally between pro- and anti-government). Thus, 600 posts submitted relate to non–collective action events, and 600 relate to collective action events or edge cases. We have in total four collective events and two edge cases, and so 400 posts focused on collective action events and 200 on edge cases.
26. J. Crandall, M. Crete-Nishihata, J. Knockel, S. McKune, A. Senft, D. Tseng, G. Wiseman, Chat program censorship and surveillance in China: Tracking TOM-Skype and Sina UC. *First Monday* **18** (no. 7) (July 2013). [doi:10.5210/fm.v18i7.4628](https://doi.org/10.5210/fm.v18i7.4628)
27. J. Fallows, The connection has been reset. *Atlantic* (March 2008); www.theatlantic.com/magazine/archive/2008/03/-the-connection-has-been-reset/306650.
28. K. Imai, L. Keele, D. Tingley, T. Yamamoto, Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *Am. Pol. Sci. Rev.* **105**, 765–789 (2011). [doi:10.1017/S0003055411000414](https://doi.org/10.1017/S0003055411000414)
29. J. Pearl, Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence* (Morgan Kaufmann, San Francisco, 2001), pp. 411–420; <https://dslpitt.org/uai/papers/01/p411-pearl.pdf>.
30. O. Blanchard, A. Shleifer, “Federalism with and without political centralization: China versus Russia” (National Bureau of Economic Research, Cambridge, MA, 2000); www.nber.org/papers/w7616.
31. S. Heilmann, E. Perry, *Mao’s Invisible Hand: The Political Foundations of Adaptive Governance in China* (Harvard University Asia Center, Cambridge, MA, 2011).

32. Y. Qian, G. Roland, Federalism and the soft budget constraint. *Am. Econ. Rev.* **88**, 1143–1162 (1998).
33. Y. Qian, B. R. Weingast, Federalism as a commitment to perserving market incentives. *J. Econ. Perspect.* **11**, 83–92 (1997). [doi:10.1257/jep.11.4.83](https://doi.org/10.1257/jep.11.4.83)
34. G. Egorov, K. Sonin, Dictators and their viziers: Endogenizing the loyalty-competence trade-off. *J. Eur. Econ. Assoc.* **9**, 903–930 (2011). [doi:10.1111/j.1542-4774.2011.01033.x](https://doi.org/10.1111/j.1542-4774.2011.01033.x)
35. G. Salton, *Automatic Text Processing: The Transformation, Analysis, and Retrieval of Information by Computer* (Addison-Wesley, Reading, MA, 1988).
36. D. Kelleher, S. Luz, Automatic hypertext keyphrase detection. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence* (Erlbaum, Hillsdale, NJ, 2005), pp. 1608–1609.
37. For examples see (44).
38. For news coverage of the protests, see (45–48).
39. D. Hopkins, G. King, A method of automated nonparametric content analysis for social science. *Am. J. Pol. Sci.* **54**, 229–247 (2010). [doi:10.1111/j.1540-5907.2009.00428.x](https://doi.org/10.1111/j.1540-5907.2009.00428.x)
40. K. Hong, China’s Twitter-like Sina Weibo service now has over 50 million active users per day. *The Next Web*, 13 August 2013; <http://tnw.co/1fdNFPS>.
41. S. Millward, Tencent Weibo, the ‘other weibo’ that nobody cares about, reaches 540 million users. *Tech in Asia*, 22 January 2013. <http://bit.ly/1byjSNW>.
42. K. W. Fu, M. Chau, Reality check for the Chinese microblog space: A random sampling approach. *PLOS ONE* **8**, e58356 (2013). [Medline doi:10.1371/journal.pone.0058356](https://doi.org/10.1371/journal.pone.0058356)
43. J. Q. Ng, Weibo keyword un-blocking is not a victory against censorship. *Tea Leaf Nation*, 21 June 2013; <http://bit.ly/1kfqNBC>.
44. P. Barefoot, “Principal, get a room with me, spare the schoolchildren!” *China Smack*, 31 May 2013; <http://j.mp/19yuv7E>.
45. Al-Jazeera, Democracy push as Hong Kong marks handover. 1 July 2013; <http://j.mp/145Jvpp>.

46. S. Lee, K. Wong, Hong Kong protests to underscore Leung's record-low appeal. *Bloomberg BusinessWeek*, 28 June 2013; <http://j.mp/13r3v7v>.
47. J. Ngo, July 1 protest is Hong Kong's taste of democracy. *South China Morning Post*, 30 June 2013; <http://j.mp/15PcwBt>.
48. C. Yung, Annual Hong Kong protest to focus ire on leader. *Wall Street Journal*, 28 June 2013; <http://j.mp/13FJB3w>.
49. G. King, J. Pan, M. E. Roberts, Replication Data for: Reverse Engineering Chinese Censorship: Randomized Experimentation and Participant Observation (2014). [doi:10.7910/DVN/26212](https://doi.org/10.7910/DVN/26212)
50. K. Imai, L. Keele, D. Tingley, T. Yamamoto, Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *Am. Pol. Sci. Rev.* **105**, 765–789 (2011). [doi:10.1017/S0003055411000414](https://doi.org/10.1017/S0003055411000414)
51. G. King, E. Gakidou, N. Ravishankar, R. T. Moore, J. Lakin, M. Vargas, M. M. Téllez-Rojo, J. E. Hernandez Ávila, M. Hernández Ávila, H. Hernández Llamas, A “politically robust” experimental design for public policy evaluation, with application to the Mexican universal health insurance program. *J. Policy Anal. Manage.* **26**, 479–506 (2007). [Medline doi:10.1002/pam.20279](https://doi.org/10.1002/pam.20279)