

How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument

GARY KING *Harvard University*

JENNIFER PAN *Stanford University*

MARGARET E. ROBERTS *University of California, San Diego*

The Chinese government has long been suspected of hiring as many as 2 million people to surreptitiously insert huge numbers of pseudonymous and other deceptive writings into the stream of real social media posts, as if they were the genuine opinions of ordinary people. Many academics, and most journalists and activists, claim that these so-called 50c party posts vociferously argue for the government's side in political and policy debates. As we show, this is also true of most posts openly accused on social media of being 50c. Yet almost no systematic empirical evidence exists for this claim or, more importantly, for the Chinese regime's strategic objective in pursuing this activity. In the first large-scale empirical analysis of this operation, we show how to identify the secretive authors of these posts, the posts written by them, and their content. We estimate that the government fabricates and posts about 448 million social media comments a year. In contrast to prior claims, we show that the Chinese regime's strategy is to avoid arguing with skeptics of the party and the government, and to not even discuss controversial issues. We show that the goal of this massive secretive operation is instead to distract the public and change the subject, as most of these posts involve cheerleading for China, the revolutionary history of the Communist Party, or other symbols of the regime. We discuss how these results fit with what is known about the Chinese censorship program and suggest how they may change our broader theoretical understanding of "common knowledge" and information control in authoritarian regimes.

INTRODUCTION

Social media in China appears as vibrant and extensive as in any Western country, with more than 1,300 social media companies and websites, and millions of posts authored every day by people all over the country. At the same time, the Chinese regime imposes extensive and varied controls over of the entire system (Brady 2009; Cairns and Carlson 2016; Knockel et al. 2015; MacKinnon 2012; Ng 2015; Shirk 2011; Stockmann 2013; Stockmann and Gallagher 2011; Yang 2009). Which social media companies are prevented from operating in China is easy to see (the so-called Great Firewall of China), and the scholarly literature now offers considerable evidence on how and why

they censor certain individual social media posts that have appeared on the web or filter them out before appearing. In both cases, the censorship apparatus allows a great deal of criticism of the regime, its officials, and their policies (which can be useful information for the central government in managing local leaders) but stops discussions that can generate collective action on the ground (King et al. 2013, 2014).¹

According to numerous speculations by scholars, activists, journalists, officials in other governments, and participants in social media, the Chinese regime also conducts "astroturfing," or what we might call *reverse censorship*, surreptitiously posting large numbers of fabricated social media comments as if they were the genuine opinions of ordinary Chinese people. The people hired for this purpose are known formally as Internet commentators (网络评论员), although more widely as 50c party members (五毛党), so called because they are rumored to be paid 50 cents (5 jiao, 角, or about \$0.08) to write and post each comment (Tong and Lei 2013). We show that this rumor turns out to be incorrect; however, we adopt this widely used term to denote social media comments posted at the direction or behest of the regime, as if they were the opinions of ordinary people.²

Gary King is Albert J. Weatherhead III University Professor, Institute for Quantitative Social Science, 1737 Cambridge St., Harvard University, Cambridge MA 02138; GaryKing.org (King@Harvard.edu).

Jennifer Pan is Assistant Professor, Department of Communication, 450 Serra Mall, Building 120, Stanford University, Stanford CA 94304; jenpan.com (jpl1@stanford.edu).

Margaret E. Roberts is Assistant Professor, Department of Political Science, University of California, San Diego, Social Sciences Building 301, 9500 Gilman Dr., #0521, La Jolla, CA 92093-0521; MargaretRoberts.net (meroberts@ucsd.edu).

Our thanks to Neel Guha, Peter Dyrud, Yingjie Fan, and many others for superb research assistance; Danielle Allen, Peter Bol, Becky Fair, Chase Harrison, Iain Johnston, Franziska Keller, Blake Miller, Jean Oi, Samantha Ravich, Brandon Stewart, Daniela Stockmann, Andy Walder, Yuhua Wang, Chaodan Zheng, and Yun Zhu for helpful comments; and DARPA (contract W31P4Q-13-C-0055/983-3) and the National Science Foundation (grant 1500086) for research support. Authors are listed alphabetically. Data and information necessary to replicate the results in this article appear in King et al. (2017).

Received: August 26, 2016; revised: November 6, 2016; accepted: April 11, 2017.

¹ Although we make general statements about Chinese censorship, it is important to note that censorship in China is by no means monolithic in its operations or outcomes. Censorship is largely carried out by Internet content providers, and regulations over the flow of information continues to evolve. As a result, there is variation in the precise details of censorship in China by platform, geography, and over time.

² Thus, 50c party members are distinct from "volunteer 50c members" (自干五; known as "bring your own grainers"), who express proregime or anti-Western sentiment online without being paid by

The nearly unanimous view of journalists (and human rights activists) is that 50c party posts strongly argue with and debate against those who criticize the government, its leaders, and their policies. This is also the view of many scholarly publications discussing this activity (and was our view as well prior to the research reported here). We systematically summarize these views, including a quantitative analysis of social media posts openly accused of being written by 50c party members. Unfortunately, until now, no method has existed for detecting 50c party members, which posts they write, their content, the size of the operation, or why they write them, and so the best anyone could do was base claims on intuition, logic, occasional anecdotes, rumors, and leaked government directives.

In this article, we offer the first systematic empirical evidence for the content of 50c party posts, their authors, and the government's strategic objectives.³ We begin by analyzing an archive of emails leaked from the Internet Propaganda Office of Zhanggong, a district of Ganzhou City in Jiangxi Province. These emails give explicit details of the work of numerous 50c accounts in this district. Although in the public domain and reported in the press (e.g., Henochowicz 2014; Sonnad 2014), the structure of the archive is complicated, too large to understand by traditional qualitative methods, and in formats (and attachments) far too diverse to make standard methods of automation feasible. As such, it has never before been systematically analyzed, and little of it has been explored. We have developed an approach to analyze this dataset and have extracted more than 43,000 known 50c party posts and their authors from it.

We first characterize the patterns in these data via their network and time series structures. Then we systematically analyze the *content* of the 50c posts in our leaked archive and extrapolate to the rest of China in stages. We then use this methodology to study the content of the posts and finally infer the goals behind this massive government program and how it may reveal broader government strategies. We validate our 50c party member predictions with a novel sample survey of predicted 50c party members, as well as several unusual gold standard evaluations that we develop to validate our validation. We estimate and reveal the size of what turns out to be a massive government operation that writes approximately 448 million 50c posts a year. We also discuss our assumptions, interpretations, and what might go wrong with our evidence and inferences.

At every stage, our results indicate that prevailing views about the 50c party are largely incorrect. We

the government; the “little red flowers” (小粉红), unpaid red guards who also attack opponents of the regime online; the “American Cent Party” (美分党), who express Western democratic values and criticize the Chinese communist regime online; and the “Internet water army” (网络水军), which refers to for-hire astroturfers working for and advancing the interests of companies and other actors willing to pay their fees. None are known to be organized as an entire collective. Of course, political parties do not exist in China, and so, despite the name, the 50c party is not a political party. For an excellent overview of the purpose and tactics of the volunteer 50c party, see Han (2015a).

³ Because China is a single-party regime, we use the terms *government* and *regime* interchangeably to refer to those in power.

show that almost none of the Chinese government's 50c party posts engage in debate or argument of any kind. They do not step up to defend the government, its leaders, and their policies from criticism, no matter how vitriolic; indeed, they seem to avoid controversial issues entirely. Instead, most 50c posts are about cheer-leading and positive discussions of valence issues. We also detect a high level of coordination in the timing and content in these posts. A theory consistent with these patterns is that the strategic objective of the regime is to distract and redirect public attention from discussions or events with collective action potential.

The theoretical implications of our findings are presented later in the article. We give a unified parsimonious summary of Chinese government internal information control efforts and show how these findings may cause scholars to rethink the notion of “common knowledge” in theories of authoritarian politics more generally. Finally, we conclude and then give a summary of what we might have missed and how scholars can follow up on this work.⁴

WHAT WE THINK WE KNOW

We summarize here views about the 50c party of (1) journalists, (2) academics, and (3) social media participants accusing others of being 50c party members. The dominant view of most is that 50c party members engage in “hand-to-hand” verbal combat, making specific, directed arguments that support the government, its leaders, and their policies, and opposing arguments to the contrary; they do this by engaging in debate with and criticism of China's enemies, including those who oppose it inside the country and from abroad. For (1) and (2) we offer brief literature reviews; for (3) we find and analyze posts accused by others of being 50c. In that section, we also introduce and validate a scheme to categorize 50c posts on the web; we use it in this section to understand the posts accused of being written by 50c members and then for many other purposes throughout this article.

Although the difficulties of collecting data on an inherently secret operation means that most prior literature includes “no successful attempts to quantify regime-sponsored commentary in China” (Miller 2016), the work cited in this section involves considerable effort and creativity, and even a few clever efforts to guess what might be 50c posts before turning to try to explain or predict the guesses (Miller 2016; Han 2015b). For example, Han (2015b) uses information from leaked censorship directives and local media reports of the training of online commentators in an online “guerrilla ethnography.” Still, the lack of ground

⁴ Appendix B briefly describes the (unintended) events and implications of the Chinese government's unexpected decision to respond to an early version of this article through an editorial in a state-owned newspaper. In so doing, they appear to give a rare admission of the existence of the 50c party and confirmation of our empirical results. In doing this, they also provide important information about government strategy that we discuss. (Our Online Supplementary Appendix provides additional details and a translation of the original editorial.)

truth means that for identification of 50c posts, Han has to rely on anecdotal evidence and intuition (e.g., whether posts “smell strongly of official propaganda”). In other works, sophisticated unsupervised statistical techniques have been used but still generated “no evidence of large-scale Wumao [50c] activity on Weibo” (Yang et al. 2015). As these authors make clear, little solid empirical evidence exists about the content and extent of 50c party posts.

Journalists. The popular press describes 50c members as “undercover progovernment Internet commenters” (Keating 2011) who “set out to neutralize undesirable public opinion by pushing pro-Party views through chat rooms and web forums” (Bandurski 2008, 41). They “shape online public opinion” by labeling “critical opinion leaders as traitors of the country” (Lam 2012). Prominent dissident Ai Weiwei said, “If you oppose the US and Japan [online], you are a member of the 50 cents army” (Strafella and Berg 2015, 154). The 50c party members “combat hostile energy,” defined as posts that “go against socialist core values” or “are not amenable to the unity of the people.” Such information should be “resolutely resisted, proactively refuted, and eagerly reported to Internet authorities” (Haley 2010). Through active engagement of opposition views, they try to “sway public opinion” (China Digital Space 2016; Ng 2011), “influence public opinion . . . pretending to be ordinary citizens and defending or promoting the government’s point of view” (Lam 2013), or “steer conversations in the right direction” (Economist 2013). Estimates by journalists of the size of the 50c party is between 500,000 and 2 million (Philipp 2015).

Academics. Academics have indicated that between 250,000 and 300,000 paid 50c party members write pseudonymous posts directed by the Chinese government (Freedom House 2009; Barr 2012; Greitens 2013). Because of the absence of systematic scholarly research on the subject, academics express a wider range of possibilities (and uncertainties) for what 50c party members write about. However, in most cases their conclusions mirror those of the journalists, that 50c party members generate proregime commentary and argue with its critics. Deibert and Rohozinski (2010, 54) describe 50c party members who “patrol chatrooms and online forums, posting information favorable to the regime and chastising its critics.” They “mix control and activism on line . . . making favorable comments, and generally pushing discussion toward pro-Party lines” (Greitens 2013, 265). They are an “army of online commentators . . . promoting the Chinese Communist Party’s line on sensitive subjects” (Bremmer 2010; see also Hassid 2012). They “facilitate state propaganda and defuse crises” (Han 2015b), “post comments favorable towards the government policies” (Tang et al. 2012, 299), “defending the government” and “fighting” those who “criticize the government” (Zhang et al. 2014, 1889), and, for example, “attack calls for the country to launch a ‘jasmine revolution’” (Bambauer 2013, 29).

Social Media Participants. Participants in social media regularly characterize 50c party members by openly accusing others of being members themselves. To systematically characterize their views, we obtained a random sample of 9,911 social media posts from 2010 to 2015 that contain the word “五毛党” (“50c party”). From these data, we drew a sample of 128 posts written by people accused in other posts of being 50c party members.

We then sorted these “accused 50c posts” into one of six categories, using a categorization scheme we will use throughout this article. With two independent Chinese language coders and 200 randomly selected posts from the 9,911 posts, we measured the intercoder reliability of the categorization scheme at 93% agreement (see Appendix A for details). Two of the categories, comprising 65% of the accused 50c posts, represent the views of academics and journalists, and include (1) taunting of foreign countries (which is 29% of this sample) and (2) argumentative praise or criticism (36% of the sample). Taunting includes denigrating favorable comparisons of China compared to other, usually Western, countries and taunting of prodemocracy or pro-West values or opinions. Argumentative praise or criticism involves engaged argument and debate about controversial (nonvalence) issues, criticism of opponents of the government, or praise of the leaders.

The categorization scheme also includes (3) nonargumentative praise or suggestions (22% of the sample) and two categories that everyone agrees are not what 50c party members are writing about, (4) factual reporting (8%), and (5) cheerleading (at 5%). Nonargumentative praise or suggestions includes discussion of noncontroversial valence issues, such as improving housing or public welfare, or praise of government officials, but does not debate or take opposing viewpoints. Category (3) does not threaten the regime in any way, and indeed Chen et al. (2016) show that local governments openly discuss nonargumentative valence issues with others on government websites.)

Factual reporting involves descriptions of government programs, events, initiatives, or plans. Cheerleading includes expressions of patriotism, encouragement and motivation, inspirational slogans or quotes, gratefulness, discussions of aspirational figures, cultural references, or celebrations. (Appendix A also includes a sixth “other” or “irrelevant” category, but we remove this so that the percentages from the first five categories add to 100%.)

Thus social media participants accusing others of being 50c party members agree with journalists and most scholars that the content of 50c posts is basically antidisestablishmentarianism—arguing with those who oppose with the regime, its leaders, or their policies.

We now go a final step and study the identities of accused (as distinct from actual) 50c party members, which can be difficult because such accusations occur on comment or discussion threads where participants are anonymous. However, by careful and extensive cross referencing of profile information across multiple platforms, we were able to unearth personal details for a handful of these individuals. Their backgrounds

vary greatly, but in each case it seems obvious that they are highly unlikely to be real 50c party members. For example, among those accused of being 50c party members include Zhou Xiaoping (周小平), a blogger well known for his anti-West and nationalist sentiment, and He Jiawei (何家), a blogger known for critiques of the Chinese government who posts on Boxun, a site hosted outside of China devoted to covering topics such as Chinese government human rights abuses. Other well-known figures accused of being 50c party members include Lin Yifu (林毅夫), a Peking University professor who was chief economist and senior vice president of the World Bank from 2008 to 2012. In none of these cases are these people likely to be 50c party members. However, those accused of being 50c party members also include figures not connected to politics, such as (in our data) a comedian, a lawyer, and a marketing executive.

It appears that the evidence base of those accusing others of being 50c party members is no better than that of academics or journalists. Although the prior beliefs of all three groups about the content of 50c party posts are almost the same, little evidence supports their claims.

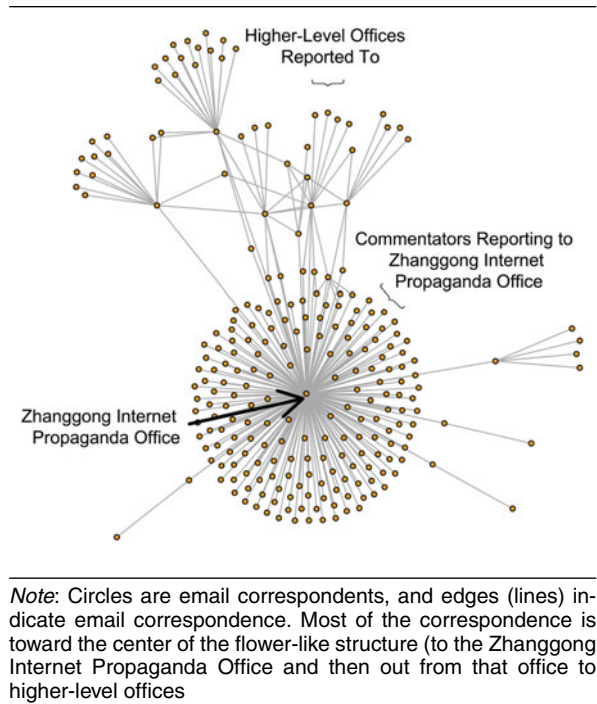
LEAKED INTERNET PROPAGANDA OFFICE COMMUNICATIONS

1. Data and Methods. The problem in the literature has been that “detecting the Wumao [50c party] is difficult because there is no ground truth information about them” (Yang et al. 2015). We are fortunate to be able to change this situation. In December 2014, anonymous blogger “Xiaolan” (<https://xiaolan.me/>) released an archive of all 2013 and 2014 emails to, and some from, the account of the Internet Propaganda Office (网宣办), a branch of the propaganda department, of Zhanggong District. Zhanggong District is a country-level administrative unit (with a population in 2013 of 468,461) that is part of the moderate-sized Ganzhou City, located in Jiangxi Province. The emails reported activities of Internet commentators, including numerous 50c posts from workers claiming credit for completing their assignments, and many other communications. The hack was widely reported, and the archive of emails has been publicly available since (Henochoicz 2014).

The archive’s large size, complicated structure, numerous attachments, diverse document formats (screen shots, Word, Excel, PowerPoint, raw text, text as part of other emails, etc.), multiple email storage formats, and many links to outside information has made digesting much of it impossible either for individuals reading and coding by hand or for existing methods of automated text analysis. Journalists managed to pull out a few examples to write newspaper articles, but no systematic analysis has been conducted of these data.

To systematize this richly informative (and essentially qualitative) data source, we developed and applied a variety of methods and procedures, from large-scale hand coding, to specially tuned and adapted methods of named entity recognition to methods of automated text analysis and extraction. Because of the considerable effort and resources necessary, we


FIGURE 1. Network Structure of Leaked Email Correspondents



have made structured and easy-to-access forms of these data, along with other replication information, publicly available in Dataverse so that others may follow up (see King et al. 2017).

From this work, we identified 2,341 emails sent from February 11, 2013 to November 28, 2014. Of these, 1,208 contained the text of one or (usually many) more 50c posts. In all, from these emails and their attachments we harvested 43,757 known 50c posts that form a basis for our subsequent analyses and, as a training set, help identify other 50c posts. (Although we have the name, direct contact information, and often photographs of many of the people discussed in this article, we have no academic reason to make this information more public than it already is and therefore do not do so. Other data and replication information is available in our Dataverse archive; see King et al. 2017). We conduct rigorous evaluations of our claims in subsequent sections. For now, we characterize the content with several separate descriptive analyses.

2. Structure. We portray the overall structure of communications in these emails with the network diagram in Figure 1. Each circle is a specific email account, and each line denotes where one or more emails was sent from and to. The large flower-like shape at the bottom represents 50c party members sending in copies of their posts to the Zhanggong District Internet Propaganda Office (章贡区网宣办), claiming credit for completing their assignments. This office then reports up to other offices (see the lines out from the center of the flower shape), including the speaker of Zhanggong People’s Court News office (江西省赣州市章贡区人民法院新

FIGURE 2. Time Series of 43,757 Known 50c Social Media Posts with Qualitative Summaries of the Content of Volume Bursts

闻发言人) and the District Party Office Information Department (区委办信息科).

3. Identifying 50c Party Members. Next, most of the scholarly literature describes 50c party members as ordinary citizens hired for very low piecemeal wages. We found instead that almost all 50c workers in our sample are government employees (consistent with some arguments by Han [2015b]). Of the 43,757 posts, only 281 were made by individuals or groups that we could not identify (the content of these posts were very similar to those we could identify). The remaining 99.3% were contributed by one of more than 200 government agencies throughout the Chinese regime's matrix organizational structure (of geographic representation by functional area) in Zhanggong District, including 9,159 posts (20.9% of the 43,757 total) made directly by the Zhanggong Internet Propaganda Office, 2,343 (5.4%) by the Zhanggong District Bureau of Commerce (区商务局), 1,672 by Shuixi Township (水西镇, one of several townships in Zhanggong), and 1,620 by Nanwai Subdistrict (南外街道, one of several subdistricts in Zhanggong). Others come from functional bureaus in Zhanggong District (e.g., 体育局 Sports Bureau, 人保局 Bureau of Human Resources and Social Security, 地税局 Bureau of Taxation, 法院 Zhanggong District court), the government offices of Zhanggong's subdistricts and townships (e.g., 沙河镇 Shahe Township, 赣江街道 Ganjiang Subdistrict), functional departments in each subdistrict or township (水西镇党政办 Shuixi Township Party Office), and administrative offices of neighborhoods and villages in Zhanggong's

townships and subdistricts (e.g., 南外街道东阳山社区 Dongyang Shan neighborhood of the Nanwai Subdistrict, 水西镇何乐村 Hele village of the Shuixi Subdistrict).

Of the 50c posts in this archive, 29.98% did not contain a URL or a description of the site where the content was posted. Of the remainder, 53.38% of the 50c posts were comments on government sites (GanzhouWeb, Newskj, DajiangWeb, JidanWeb, JiangxiWeb, CCTVWeb, RenminWeb, JiujiangWeb, QiangGouWeb), and 46.62% were on commercial sites. Of the 50c posts on commercial sites, 53.98% were on Sina Weibo, 32.10% on Tencent Weibo, 10.75% on Baidu Tieba, and 2.69% on Tencent QZone, with the rest in the long tail receiving less than 1% each.

We also found no evidence that 50c party members were actually paid 50 cents or any other piecemeal amount. Indeed, no evidence exists that the authors of 50c posts are even paid extra for this work. We cannot be sure of current practices in the absence of evidence, but given that they already hold government and Chinese Communist Party jobs, we would guess that this activity is a requirement of their existing job or at least rewarded in performance reviews.

4. Coordination and Content. We now offer a first look at the 43,757 posts from the 50c party we unearthed. We do this by plotting a daily time series of counts of these posts in Figure 2. The most important finding in this graph is that the posts are far from randomly or uniformly distributed, instead being highly focused into distinct volume bursts. This suggests a high level

of coordination on the part of the government. Indeed, often the most influential patterns in most social media are the bursts that occur naturally when discussions go viral. The government's manufactured bursts mirror these naturally occurring influential patterns, but at times of the government's choosing. Bursts are also much more likely to be effective at accomplishing specific goals than a strategy of randomly scattering government posts in the ocean of real social media. (We also looked extensively for evidence that 50c posts were created by automated means such as bots, but the evidence strongly indicates to the contrary that each was written by a specific, often identifiable, human being under direction from the government.)

Although we conduct rigorous, quantitative analyses of the content of 50c posts in the sections to follow, here we provide a feel for the content of the posts by labeling the largest volume bursts in this set (with numbers corresponding to those in the figure). The labels are brief summaries we chose from reading numerous posts, a process we found easy and unambiguous. The following list gives the first indication that the focus of these posts is on cheerleading, possibly for purposes of distraction, rather than engaged argumentation and debate:

1. *Qingming (Tomb Sweeping Day)*: More than 18,000 posts about veterans, martyrs, how glorious or heroic they are, and how they sacrificed for China.
2. *China Dream*: More than 1,800 posts about President Xi Jinping's "China Dream." Potentially a reaction to the April 2013 *People's Daily* piece instructing municipal governments to carry out China Dream propaganda campaigns (see <http://j.mp/chinadream>).
3. *Shanshan Riots*: 1,100 posts, immediately following Shanshan riots in Xinjiang. At 5:30 p.m., Zhanggong County sent an email to itself (probably BCC-ing, many others), highlighting three popular posts about Xinjiang and identifying this as a terrorist incident. At 8:00 p.m. on the same day, Zhanggong County sent an email to Ganzhou City, to which it reports having created hundreds of 50c posts, seemingly to distract from the riots, about China Dream, local economic development, and so forth.
4. *18th Party Congress, 3rd Plenum*: More than 3,400 posts related to the 3rd plenary session of the Chinese Communist Party's 18th Congress, which discussed plans for deepening structural reform.
5. *"Two Meetings"*: More than 1,200 posts about Ganzhou's People's Congress and Political Consultative Committee meetings, and policies to be discussed at the two meetings, including factual reporting of environmental issues, one child policy, rural issues, as well as growth and development.
6. *Early May Burst*: 3,500 posts about a variety of topics, such as mass line, two meetings, people's livelihood, and good governance. Immediately followed the Urumqi railway explosion.
7. *Praise for Central Subsidy*: More than 2,600 posts celebrating the second anniversary of "Central Soviet Areas Development policy" (若干意见), subsidies from the central government to promote the

development of region where the original Chinese Communist Party bases were located (including the region where Zhanggong is located); at the same time, the local government held an online Q&A session for citizens.

8. *Martyr's Day*: 3,500 posts about martyrs and the new Martyr's Day holiday, celebrating heroes of the state.

Although we cannot know for certain the exact cause or intended purpose of each burst of 50c party posts, Figure 2 is consistent with a strategy of distraction. For example, several bursts follow events with "collective action potential" (i.e., actual or potential real-world crowd formation and related activities; see p. 6 of King et al. [2013] for a precise definition). These events include the Shanshan riots and the early May burst following the Urumqi railway explosion. Other bursts occur during national holidays when people are not working, which tend to be prime time periods of political unrest. Indeed, the Qingming festival, or Tomb Sweeping Day, has historically been a focal point of protests in China and, for this reason, was largely banned during the Maoist era. In recent years, Qingming, a day on which people pay respects to the dead, has drawn attention to sensitive events, such as the deaths of those in the 1989 Tiananmen crisis (Johnson 2016). The central regime and Jiangxi provinces have both issued notices about the Qingming festival as a period when local governments need to increase their vigilance to prevent protest (see <http://j.mp/jiangxi> and <http://j.mp/MinistryCivil>). Similarly, political meetings are periods when government and party officials believe that protests are more likely to take place. During these periods, officials gather and attention is focused on the activities of the regime; as such, successful protests can garner greater attention. Prior to these meetings, measures such as a preemptive redistribution and preemptive repression are put into place to decrease the likelihood of social mobilization (Pan 2015; Truex 2016).

5. *The Purpose of 50c Posts*. Although our leaked archive includes specific directions to 50c workers, it does not reveal whether these directions originate from Zhanggong or from higher levels of the government or party. This, and the nearly infinite phenomena that we might identify as potential precipitating events, prevents us from determining the immediate cause of every burst of 50c activity. However, our inference about distraction being the goal of the regime is consistent with directions to 50c party members in emails from the Zhanggong propaganda department. They ask 50c members to "promote unity and stability through positive publicity" (坚持 团结稳定鼓劲、正面宣传为主) and "actively guide public opinion during emergency events" (积极稳妥做好突发事件舆论引导). In this context, "emergency events" are events with collective action potential.⁵

⁵ For example, a website developed by the Ministry of Public Security and Ministry of Education to help young people better understand

We now turn to a more systematic analysis of these posts, their accounts, and others like them beyond Zhanggong.

CONTENT OF 50c POSTS

We now reveal the content of 50c party posts across China by estimating the distribution of these posts over the five main content categories introduced previously (with details in Appendix A). We do this in five separate analyses and datasets that successively expand the initial set of posts from Zhanggong to larger and larger areas across the country.

Ex ante, we do not know how 50c party activity in Zhanggong might differ from that in other counties. Originally part of the Jiangxi Soviet, established in 1931 by Mao Zedong, Zhu De, and other leaders, Zhanggong has a rich revolutionary history. These and other factors may make it unusual. However, directives from the central government, or common interests of different counties in keeping their populations in check, may keep the purpose and content 50c party activity in different counties aligned. As it turns out, for each of the five separate analyses, and in the survey validation in the next section, we find very similar results, with 50c party posts largely comprised of cheerleading and distraction rather than engaged argument. In other words, the patterns found in the leaked data from Zhanggong District do extrapolate.

We conclude this section with a sixth part, reporting on an event that occurred during our observation period that provides strong evidence of coordination across counties and very clear top down control.

1. Leaked 50c Posts. We first analyze the 43,757 50c social media posts that we harvested from the leaked archive from Zhanggong. These posts were made by numerous authors on many different social media sites, including national-level platforms run by private sector firms such as Sina Weibo and Baidu Tieba, as well as government forums at the national, provincial, prefectural, and county levels. To study these data, we began by hand coding a random sample of 200 posts into our categories (again with high intercoder reliability).

One result is immediately apparent: the number of posts from this sample that fall in the categories “taunting of foreign countries” or “argumentative praise or criticism” is exactly *zero*. This is an important surprise,

safety issues (<http://j.mp/EmergEvents>) explains: “Every emergency event involves the self-interest of a particular group of people, leading to psychological pressure and change among this group, and understandably leading to concern and worry. Especially for emergency events of a societal nature [as distinct from natural disasters], most are organized by a small group of people, who through their publicity seeking and encouragement get others involved. Recently, emergency events due to issues like territorial disputes, land requisitions, and housing demolition in certain areas are often organized by one person and involve many, making collective events.” (任何一类突发事件, 都必然要涉及一部分人的切身利益, 使其产生心理压力和变化, 引起人们的关注和不安也属正常。尤其是社会性的突发事件, 多是由少数人操纵, 通过宣传鼓动把一些群众卷到事件中来。近期, 在一些地方因地界、征地、拆迁安置而发生的突发性事件, 往往是一人纠合, 数人响应, 使其具有聚众性。)


as it is essentially the opposite of the nearly unanimous views espoused by scholars, journalists, activists, and social media participants. This result would be highly unlikely to have resulted from (binomial) sampling error if the true share of the full set were even as large as a few percentage points (at 5%, which would still be a major surprise, the probability of seeing the sample that we obtained is essentially zero). To push even further, we did extensive searches and reading among the remaining posts and finally found a few that fit this category (see the examples in Categories 2 and 3 in Appendix A), but the overall result is that 50c party posts are extremely rare in these categories.

We thus infer that the leaked posts contain very little taunting of foreign countries or argumentative praise or criticism; we verify this by formally estimating all category proportions in the entire set of posts. Using a text-analytic method known colloquially as ReadMe (named for the open source software that implements it), we estimate the category proportions directly, without having to classify each post into a category (Hopkins and King 2010). This is fortunate, as individual classifiers that manage to achieve high (but imperfect) levels of the percentage correctly classified may still generate biased estimates of the category proportions. For example, an estimate indicating that zero country dyad-years, since WWII, were at war achieves a predictive accuracy of about 99.9%, but aggregating these classifications yields an obviously biased (and useless) estimate of the prevalence of war. In contrast, ReadMe does not give individual classifications, but it has been proven to give approximately unbiased and consistent estimates of the category proportions, which here is the relevant quantity of interest. The other advantage of ReadMe in this context is that its statistical assumptions are met by our sampling procedures.

The estimated proportions of 50c posts by category for all datasets appear in Figure 3; the results for our first dataset (of all posts found in the leaked emails in Zhanggong) are represented by a histogram, formed by the set of solid disks (●) for the point estimate, and solid line for the confidence interval, for each of the categories. Other results, to be described in the following in order from left to right within each category, also appear in the same graph.

The categories in Figure 3 are arranged so that the two emphasized in the literature appear on the left and our main empirical results on the right. For this analysis, the results indicate that approximately 80% fall within the cheerleading category, 14% in nonargumentative praise or suggestions, and only tiny amounts in the other categories, including nearly zero in argumentative praise or criticism and taunting of foreign countries. Clearly, these results clearly indicate that 50c posts are about cheerleading, not argumentation.

2. Posts from Leaked 50c Weibo Accounts. One possibility that we now consider is whether 50c party members differentially reported cheerleading posts back to the propaganda department, even though they posted about topics at the behest of the regime from other categories as well. To study this question, we constructed

FIGURE 3. Content of Leaked and Inferred 50c Posts, by substantive category (with details in Appendix A) and analysis (given in the legend)

a second dataset by first identifying all Weibo social media accounts revealed in the leaked email archive. We chose Weibo because it is the most widely used social media site that enables mass distribution, and we were able to obtain access in the manner we needed it. We then found these accounts on the web and kept all 498 Weibo accounts that made at least one post. Finally, we downloaded all social media posts from these accounts, yielding a set of 167,977 known—but not previously leaked—posts from 50c accounts.

We drew a random sample (stratified by account) of 500 of these 167,977 social media posts and coded them into our categories as a training set. In this randomly selected training set, like the last, we find no evidence of taunting of foreign countries, although we did find a handful of posts in the category of argumentative praise or criticism, constituting only 3% of the posts. As earlier, we then used (a stratified sample and) ReadMe to estimate the five category proportions for the set of all posts. The results, reported in the second bar of the histogram in Figure 3, are very similar to that from the first dataset. The point estimates (portrayed as solid triangles, ▲, with confidence intervals as dashed lines) indicate that again the bulk of 50c posts from leaked accounts are cheerleading (51%), 20% in factual reporting, 23% in nonargumentative praise or suggestions, and only 6% in argumentative praise or criticism.

3. Partitioning Leaked Accounts for Extrapolation. We designed our third analysis to further explore the leaked data and to prepare the ground for extrapolation. The key idea here was to partition the Sina Weibo accounts (from Analysis 2) into those easy to identify

outside the leaked archive (which we do for Zhanggong in Analysis 4 and in other counties in Analysis 5) and those more difficult to identify. We developed an algorithm to distinguish these two account types and then showed that we only need to extrapolate the first type because they post the same types of content.

To find a useful partition, we began by studying the structure of the 498 known 50c Weibo accounts and their 167,977 social media posts. In each type, we often found many commercial posts, which fall in our “other” category (see the Appendix); since we remove and condition on this category for all analyses, we do not define account types on this basis either. The first type of account, which we call *ordinary*, is used by apparently ordinary people in China to post about their children, funny videos, commercial advertisements, sports teams, pop stars, personal opinions, and many other subjects. Embedded within the stream of these posts are those which these authors indicate in their communication with the propaganda department to be 50c party posts. The second type, which we call *exclusive* accounts, is (aside from commercial posts) almost exclusively devoted to 50c posts. Near as we can tell, via extensive cross checking with external data sources, ordinary accounts are genuine, registered in the name of a person (usually a government employee) posting on it, whereas exclusive accounts are pseudonymous, designed solely to fool those who see it. In both cases, the 50c posts on these accounts are those directed by the government rather than necessarily reflecting the opinions of ordinary people.

Distinguishing between ordinary and exclusive accounts in our leaked archive is easy (the number of real 50c posts reported to the propaganda department,

as a proportion of all posts on the account, is a direct measure), but our goal is to extrapolate to other counties where we have no known 50c posts. Thus we need a formal partitioning algorithm to sort accounts into these two categories without needing the inside information that we have from our extraordinary leaked data. Moreover, since our goal is to determine the content of 50c posts, we must be able to discern whether an account was written by a 50c party member without using the text of the posts.

To develop this partitioning algorithm, we followed the logic of “Bayesian falling rule list” methodology, which is accurate and also highly interpretable (Letham et al. 2015). The interpretability also enabled us to combine qualitative knowledge with modern machine learning, as well as to make choices that were much easier to apply outside of Zhanggong. With this approach as a guide, we found that two simple rules are sufficient to partition our 498 50c accounts into exclusive and ordinary. First, we obtained candidate 50c accounts by collecting all accounts that comment on or forward any post on the Zhanggong government’s Weibo account (<http://weibo.com/u/3880516376>). Second, we narrowed this to accounts with 10 or fewer followers. The result is our definition of exclusive accounts. These two simple, interpretable rules are highly plausible and consistent with what is known about social media. After all, accounts that engage with government websites and have no more than a handful of followers are likely used for a very specific purpose. (Because of how Weibo differs from platforms like Twitter, users of Weibo accounts with few followers can still be highly influential by commenting on other more popular accounts.)

We now show that the 50c posts appearing on exclusive and ordinary accounts have essentially the same types of content, where we can verify both. To do this, we applied our partitioning algorithm to the set of 498 known 50c accounts from our archive and then compared the content of ordinary and exclusive accounts. We found that 202 (41%) are exclusive accounts and the remaining 296 (59%) are ordinary accounts. This partition of the data is neither right nor wrong (and thus statistics like “percent correctly classified” do not apply), but it is useful only to the extent that using only the exclusive posts causes no bias. Thus, we estimate and compare the distribution of posts within the ordinary and exclusive account types across our five content categories. To do this, we applied ReadMe within each partition and compared the results.

Fortunately, the results are very close to each other and (as a result) to the overall results we presented previously. This implies that bias is unlikely to be induced by narrowing our search outside our leaked archive to exclusive accounts. Point estimates for the category proportions appear in Figure 3 (marked as \boxtimes and \times , in red). For both, the bulk of 50c posts appear in the cheerleading category (46% for exclusive accounts and 58% for ordinary accounts). In contrast, the sum of taunting of foreign countries and of argumentative praise or criticism is very small (5% for exclusive and 11% for ordinary).

4. Unleaked 50c Posts in Zhanggong. We now use the results about ordinary and exclusive accounts (from Analysis 3) and expand our extrapolation beyond the 50c posts in the leaked archive (from Analysis 1) and new leaked Sina Weibo posts that we found from the accounts identified in the leaked archive (from Analysis 2). The key for this extrapolation is that all three of these analyses yielded very similar estimates of the distribution of 50c posts across our five categories of interest. We thus now narrow our extrapolation to Weibo posts from exclusive accounts, which are easier to find, even though we strongly expect 50c posts to be made in many different platforms, including those run by private firms, and different levels of government.

In this section, we focus on previously unidentified 50c posts in Zhanggong. To do this, we chose exclusive accounts (by applying the two rules from the previous section). With this procedure, we found 1,031 accounts, of which 829 accounts are not mentioned in our leaked archive. We then found and scraped all 22,702 social media posts available from the front page of each of these accounts. Each front page has up to 45 separate posts. We then analyzed these posts with ReadMe, as earlier.

Results from this analysis appear in Figure 3 (with point estimates represented by \times). The result again is very similar to previous analyses: 57% of the posts made on these accounts engaged in cheerleading, 16% engaged in factual reporting, 22% engaged in nonargumentative praise and suggestions, about 4% in taunting of foreign countries, and essentially zero in argumentative praise or criticism.

5. Unleaked 50c Posts in Counties with County Government Weibo Accounts. We now extrapolate to counties across China. To do this we started with all 2,862 counties (and county-level divisions). We then took as our target of inference 50c behavior in 1,338 of these counties that were structured same way as Zhanggong, with a propaganda department that has a public website. We then drew a simple random sample of 100 of these counties and identified all exclusive accounts and a sample of their social media posts.⁶

To be more specific, for each county government Weibo account, we collected all 151,110 posts, randomly sampled up to 200 posts of these, identified all outside Weibo accounts that commented on or forwarded any one, downloaded all metadata from those accounts, and subsetted to those with 10 or fewer followers. We then downloaded the first page, comprising up to 45 social media posts from each account, as our candidate 50c posts.

Figure 3 provides our results (with point estimates represented as a diamond, \diamond). Again, we find very similar results, highly focused on cheerleading and distraction rather than argumentation and criticism: 64%

⁶ Many of the remaining 1,524 counties have Weibo accounts run by government bureaus and agencies (e.g., the public security department and civil affairs department) but not by the county government. Our informal study of these counties revealed no systematic differences from those we studied, but following up with systematic study in these counties would be a good topic for future research.

of the posts made on these accounts are categorized as cheerleading, 18% in factual reporting, 9% nonargumentative praise and suggestions, 4% in taunting of foreign countries, and only 4% in argumentative praise or criticism.

6. Coordination and Top Down Direction. The analyses thus far suggest a high level of coordination in the timing (see Figure 2) and content (see Figure 3) of 50c party activity. Here we offer evidence that these efforts may be directed from the highest levels of the regime.

In late February 2014, Chinese president Xi Jinping led the first meeting of the Central Leading Group for Internet Security and Informatization. The meeting was also attended by two other top leaders, Li Keqiang, China's premier, and Liu Yunshan, head of the Chinese Communist Party propaganda department. During this meeting, President Xi stressed the need for government officials to "have a good grasp of the timing, degree, and efficacy of online public opinion guidance so that online spaces are clear and unclouded" (把握好网上舆论引导的时、度、效,使网络空间晴朗起来) (Xi 2014). Xi's phrase *public opinion guidance* is the official term for Chinese Communist Party policies and practices designed to control or influence public opinion, which includes "traditional" guidance such as Chinese Communist Party control of the press, as well as newer types of opinion guidance for social media such as 50c party activity, censorship, and the Great Firewall. President Xi repeatedly stressed in the meeting the need for the regime to build infrastructure and a solid foundation for ensuring "Internet security" (which refers to cybersecurity more broadly in addition to public opinion guidance).

As this event occurred near the middle of our data, we can look for evidence that it had an effect. Thus we calculate that over the 2 years we observed in Zhanggong, 50c party members created an average of 7.7 social media accounts per month. Yet 156 accounts were created the month of the meeting and 39 the month after. Similarly, in our predicted data, an average of 19 accounts were created per month. Yet they created 41 accounts in the month following and 174 in the month after. We interpret these strong patterns as evidence that governments all across China responded directly to Xi's call.

VERIFICATION BY DIRECT SURVEY

We now attempt to go an extra step to verify the accuracy of our extrapolation presented earlier to predicted 50c party members across China. To do this we take the unusual step, in this context, of conducting a sample survey of predicted 50c party members, along with gold standard elements designed to validate this method of validation.⁷

⁷ We had full IRB approval for our study (although we cannot make public our exact question wording in Chinese, sample size, certain statistics, or the original data). We also added our own additional ethics rules not required by the IRB designed to further protect the identities of our respondents and to keep our large research team

1. Design. We began by creating a large number of pseudonymous social media accounts. This required many research assistants and volunteers, having a presence on the ground in China at many locations across the country, among many other logistically challenging complications. We conducted the survey via "direct messaging" on Sina Weibo, which enables private communication from one account to another. With IRB permission, we did not identify ourselves as researchers and instead posed, like our respondents, as ordinary citizens. Since information in our archive appears to indicate that government monitoring of 50c party member activities occur only through voluntary self-reporting up the chain of command, our survey questions and the responses are effectively anonymous, which are conditions that have been shown to make respondents more sincere in responding to sensitive questions (Tourangeau et al. 2013).

We drew a random sample of social media accounts that we predicted earlier to be 50c and asked each whether the owner of that account was indeed a 50c party member (in a special manner described in the following). Of course, the difficulties of interpreting these answers is complicated by the fact that our survey respondents are conducting surreptitious operations on behalf of the Chinese government designed to fool participants in social media into thinking that they are ordinary citizens, and we are asking them about this very activity. In most cases, the government is also their employer, and so they have ample incentives to not comply with our requests or to not comply sincerely.

We addressed these uncertainties with two entire additional surveys designed to provide internal checks on our results, as well as a carefully worded survey question in our anonymous survey context. In most surveys, researchers are left trusting the answer, perhaps after a stage of pretesting or cognitive debriefing. In our survey, we are in the unusual position of being able to go further by offering a gold standard validation, where for some respondents we know the outcome to the question that we are posing. In other words, we ask the same question of a random sample of *known* 50c party members from our Zhanggong leaked archive. If the results of our survey of predicted 50c party members give similar results as this survey, then we should have more confidence in the results.

We also fielded a third entire survey that approximates the opposite gold standard by asking those known *not* to be 50c party members. To do this, we drew a random sample from Weibo accounts across China among those who do not engage with government Weibo accounts and have more than 10 followers. Our results would be confirmed if the percentage who say they are 50c in this sample are significantly lower than those who acknowledge being 50c in our predicted 50c sample. A tiny fraction of these accounts may actually be 50c, but that would merely bias the results

safe. Our rules followed the principle articulated in footnote 20 of King et al. (2014) of trying to avoid influencing the system that we were studying, which has the added advantage of reducing the chance for bias.

against the test of our hypothesis of the difference in means from our set of predicted 50c members.

The final way we reduce uncertainty is in the design of our survey question. We followed best practices in designing survey questions about sensitive topics, including adjusting the perceived social environment (Näher and Krumpal 2012) and using familiar language and positive “loading” of sensitive questions (Groves et al. 2011). We also studied a large volume of social media interactions, both via automated means (King et al. 2016) and by direct reading, and found a way within the cultural context to ask the question so that it would be more likely to elicit a sincere answer. We also pretested our survey on an independent sample. Although preserving the confidentiality of our respondents and research team makes us unable to share the exact text of our question, here we report a similar version in English, which will also enable us to explain its features:

I saw your comment, it's really inspiring, I want to ask, do you have any public opinion guidance management, or online commenting experience?

To avoid interfering or influencing the system that we are studying, and to avoid putting our respondents in an uncomfortable position, the question discusses online propaganda in positive terms. We used the terms “opinion guidance management” and “online commenting,” which is the terminology the government uses to discuss these tasks. We avoided terms like “50c,” which has negative connotations for some. Instead of asking someone to “out” themselves as a 50c party member, we asked for advice on where the person learned to write in such a motivating, inspiring manner, thus avoiding generating defensiveness on part of the respondent.

2. Results. High-quality web surveys have response rates of about 3.5% (Pew Research Center 2014). The response rate for our survey was almost twice that at 6.5%, which although small on an absolute level is encouraging given our more challenging environment. In addition, unlike most web surveys, we were able to perform some checks for selection bias because we collected available information on our entire target sample before administering our survey question. This information to test for selection bias included variables such as the number of followers, gender, year of creation, average number of posts for each month, and enabling geolocation; we also observed each of these variables within the five separate data sources mentioned previously. Most tests that we conducted indicated statistically insignificant differences between respondents and nonrespondents. The few differences that appeared were negligible compared to the large effect sizes that we present in the following. As might be expected, the data contained some evidence that 50c party members are less likely to respond to our question than non-50c party members, which has the effect of making it more difficult to confirm our hypothesis. Although we could weight the following results by the

TABLE 1. Survey About 50c Status

50c Status	Origin	Yes (%)
Predicted 50c	Across China	59
Known 50c	Leaked Zhanggong archive	57
Known “Not” 50c	Random sample	19

Note: The first line is from our survey; the second two are gold standard evaluation surveys. The difference between the first and second lines is not statistically significant; the difference between the first and the third is statistically significant (both at $\alpha = 0.05$).

differences we found, they are small enough that we chose to present the raw, unprocessed data instead.

The results for our three surveys appear in Table 1. Overall, we found that 59% of our predicted 50c party members admitted to being 50c party members. If we are correct that they are all 50c party members, then the remaining 41% gave an insincere answer, which would not be surprising given that doing so is essentially their job. To test this, we used our gold standard sample of known 50c party members revealed in our leaked Zhanggong email archive. In this sample, 57% admitted to their 50c party status. The two percentage point difference between these two figures is not statistically significant (at $\alpha = 0.05$), suggesting that indeed all respondents in our predicted sample are 50c.

Also as a test, we use our gold standard sample that approximates those known to not be 50c party members. In this sample, only 19% said that they were 50c; the substantial 40 percentage point difference between this figure and that from our predicted 50c party member sample (59%) is very large and statistically significant, revealing a strong signal of actual 50c party membership among our predicted 50c sample. (Near as we can tell, if we had asked much more directly whether our respondents were 50c party members, those who were not would have responded with angry denials. This would have had the advantage of dropping the 19% figure nearer to 0%, but it would likely also have threatened our entire project. The survey would also have failed, because then few or no actual 50c party members would have answered our survey question.) Overall, the results from this survey strongly support the validity of the predictions of 50c party membership conducted previously.

SIZE OF THE 50c PARTY

In this section, we study how widespread 50c activity is across the country. Overall we find a massive government effort, where every year the 50c party writes approximately 448 million social media posts nationwide. About 52.7% of these posts appear on government sites. The remaining 212 million posts are inserted into the stream of approximately 80 billion total posts on commercial social media sites, all in real time. If these estimates are correct, a large proportion of government website comments, and about 1 of every 178 social media posts on commercial sites, are fabricated by the

government. The posts are not randomly distributed but, as we show in Figure 2, are highly focused and directed, all with specific intent and content. The rest of this section explains how we estimate these numbers. Throughout, in lieu of the possibility of formal standard error calculations, we offer transparent assumptions that others can easily adjust to check sensitivity or improve as more information is unearthed.

1. Number of Social Media Posts. To understand the context into which 50c posts are inserted, we began by estimating the total number of Chinese social media posts nationwide. As of December 2012, netizens were posting approximately 100 million messages a day, or 36.5 billion a year, on Sina Weibo alone (Zhao et al. 2014), which is one of at least 1,382 known social media sites (King et al. 2013). In our data, the ratio of Sina Weibo posts to all posts is 1.85, meaning that an estimate of the total number of posts on all platforms is $(1.85 \times 36.5 \text{ billion}) = 67.5 \text{ billion}$. However, this requires the strong assumption that 50c party members use specific commercial social media platforms in the same proportions as the entire user population. We therefore used the detailed survey from iiMedia Research Group (2014) and calculated the ratio of total posts to Sina Weibo posts to be 2.10 and the total number of posts per year to be about 80.4 billion. This is an underestimate because it is based on microblogs and ignores blogs, but blogs probably number in the millions, which is rounding error on this scale.

2. Number of 50c Posts in Zhanggong. Among the 43,757 confirmed 50c posts, 30,215 were made during a 365-day period between February 11, 2013 (the first day on which we observed a 50c post) and February 10, 2014. We have evidence of at least 1,031 exclusive (Sina Weibo) accounts in Zhanggong, including 202 accounts in the leaked archive and 829 that we identified outside the archive (by following the rules presented previously).

In our archive, a 50c party member needing to make a post chooses an exclusive account on Weibo $(689/43,757 =) 1.57\%$ of the time compared to all other choices (an ordinary account on Weibo or another social media site). We assume that this ratio is approximately the same for nonleaked 50c posts in Zhanggong, which in turn implies that the ratio of total 50c posts to 50c posts in the archive is the same as the ratio of total exclusive accounts to exclusive accounts in the archive. As such, an estimate of the total number of posts in Zhanggong in 2013 is $(30,215 \times 1,031/202 =) 154,216$.

3. Number of 50c Posts in Jiangxi Province. Zhanggong is an urban district of Ganzhou City within Jiangxi Province. According to the 2014 China Internet Network Information Center's *Statistical Report on Internet Development in China*, the 2013 Internet penetration of urban residents was 62.0% and of rural residents was 27.5% (CNNIC 2014). According to the National Bureau of Statistics of China, 48.87% of the 45.22 million people in Jiangxi Province lived in urban areas, or 22.10 million, with 23.12 million living in rural areas (National Bureau of Statistics of China 2014).

We first compute the number of 50c posts per Internet user in Zhanggong, which is $(154,216/468,461 \times 0.62 =) 0.531$. We then assume that this rate is roughly the same in Jiangxi and then scale up. Thus, we estimate the total number of 50c posts in Jiangxi during 2013 as $(0.531 \times [0.62 \times 22.1M + 0.275 \times 23.1M]) = 10.65 \text{ million}$.

4. Number of 50c Posts in China. Finally, to scale this result to all of China, we assume that the ratio of 50c posts to Internet users in other parts of China is roughly the same as in Jiangxi. This ratio of posts per Internet user is $(10.65M/14.68M =) 0.7255$. Applying this assumption to the country as a whole reveals the presence of $(0.7255 \times 617.58M =) 448.0 \text{ million}$ 50c posts in China during 2013 (see CNNIC 2014).

WHAT MIGHT BE WRONG?

Inferences in this article depend on the veracity of the leaked archive that we analyze. The size and extraordinary complexity of this archive makes it highly likely to be real. There are no signs of it having been generated by automated means, and fabricating it by hand to mislead would have been a monumental task. We also verified numerous external references from the data—to specific individuals, email addresses, phone numbers, government departments, programs, websites, social media accounts, specific posts, etc.—and every one checks out. Nevertheless, we have no information about how the leak actually occurred.

Chinese government astroturfing efforts may exist that do not follow the model that we unearthed in Zhanggong. For example, based on anecdotal evidence that we came across, it is possible that the public security bureaucracy and Communist Youth League may also be involved in fabricating social media content. It is possible that other organizations may hypothetically follow different rules and practices, perhaps varying in different places, and may generate 50c posts with different types of content. Determining whether it is must wait for new evidence to be unearthed. Perhaps the window that this article opens on this large and previously opaque government program may help others discover different aspects of it in China, and eventually in other related authoritarian regimes.

We have observed that the content of 50c party posts across China is largely about cheerleading, and to a lesser extent nonargumentative praise or suggestions and factual reporting. Since humans have highly limited attention spans, and the volume of information competing for their attention is growing quickly in the digital age, huge bursts of irrelevant posts about cheerleading will certainly be distracting to at least some degree. We are not able to quantify how distracting these posts are in practice or, as a result, the overall effectiveness of 50c strategy. Our results do suggest some interesting experiments that could be run by future researchers.

We have also gone another step and inferred that the purpose of 50c activity is to (1) to stop arguments (for which distraction is a more effective than

counterarguments) and (2) to divert public attention from actual or potential collective action on the ground. As inferences, these are, by definition, more uncertain than observations, and so we now briefly consider five alternative possible interpretations of our evidence.

First, perhaps 50c activity is a simple extension of the traditional functions of the propaganda system and not always focused on collective action. This point is definitely possible, that propaganda workers engage in cheerleading because they are not motivated to excel and because they are guided by what Han (2015b) describes as a “persistent state propaganda logic” that contravene covert activity. However, the cheerleading that we identify departs from the traditional focus of the Chinese Communist Party propaganda department on guiding the content of media and shaping public opinion (Brady 2009; Lynch 1999). In addition, we have offered clear evidence that most 50c posts from our data appear in highly coordinated bursts around events with collective action potential—either after unexpected events or before periods of time such as the Qingming festival and political meetings when collective action is perceived by the regime to be more likely. Of course it may also be that these bursts of 50c posts have different purposes depending on the need as perceived by the regime.

Second, it may be that cheerleading about (essentially) irrelevant topics merely creates a general sense of positiveness that transfers over to positiveness about other things including the regime. This may well be true, but such an effect is not likely to be large. This hypothesis would, however, be testable by experiment, perhaps even in a lab setting.

Third, might the purpose of 50c posts be to dilute negative opinion through generally positive cheerleading? In fact, this is unlikely, as 50c posts are about irrelevant issues and thus do not change the balance of positive versus negative comments. It is true that 50c posts do change the percentage of negative comments as a proportion of all posts, but more research is needed to determine how 50c posts interact with characteristically bursty and highly variable social media posts about every possible issue unrelated to politics and whether the influx of 50c comments to change the percentage of negative comments as a proportion of all posts has any tangible effect on public beliefs and perceptions.

Fourth, perhaps the point of 50c activity is to signal to the people that they are under surveillance. Although when sent through censorship a signal like this may be effective in getting people to self-censor their posts and other activities, which posts are 50c is not known to the Chinese people and so this strategy, if it exists, is unlikely to be successful.

Finally, we might ask whether some of the few posts appearing in the empirically small categories of nonargumentative praise or factual reporting might actually be sarcastic, backhand ways of making arguments. This is possible, but our methods are human led and computer assisted, and thus such sophisticated and subtle arguments would have to confuse our human coders and yet still not mislead Chinese social media partic-

ipants. In fact, even in the unlikely situation where 100% of these posts were misclassified from argumentative praise or suggestions, most would still be cheerleading and our conclusions would remain largely unchanged.

THEORETICAL IMPLICATIONS

The empirical results offered earlier seem clear, but what do they suggest about the overall strategy of the Chinese government or for authoritarian regimes in general? We first explain these results by generalizing prior findings on (human) censorship and (automated) filtering, all led by the same propaganda department in the same government as the 50c party (King et al. 2013, 2014). We then extend these ideas to the authoritarian literature in general.

1. China. One way to parsimoniously summarize existing empirical results about information control in China is with a theory of the strategy of the regime. This theory, which as with all theories is a simplification of the complex realities on the ground, involves two complementary principles that the Chinese regime appears to follow, one passive and one active. The passive principle is *do not engage on controversial issues*: do not insert 50c posts supporting, and do not censor posts criticizing, the regime, its leaders, or their policies. The second, active, principle is *stop discussions with collective action potential* by active distraction and active censorship. Cheerleading in directed 50c bursts is one way the government distracts the public, although this activity can be also be used to distract from general negativity, government-related meetings and events with protest potential, and so forth. (Citizens criticize the regime without collective action on the ground in many ways, including even via unsubstantiated threats of protest and viral bursts of online-only activity—which, by this definition, do not have collective action potential and thus are ignored by the government.)

These twin strategies appear to derive from the fact that the main threat perceived by the Chinese regime in the modern era is not military attacks from foreign enemies but rather uprisings from their own people. Staying in power involves managing their government and party agents in China’s 32 provincial-level regions, 334 prefecture-level divisions, 2,862 county-level divisions, 41,034 township-level administrations, and 704,382 village-level subdivisions, and somehow keeping in check collective action organized by those outside of government. The balance of supportive and critical commentary on social media about specific issues, in specific jurisdictions, is useful to the government in judging the performance of (as well as keeping or replacing) local leaders and ameliorating other information problems faced by central authorities (Dimitrov 2014a–c; Wintrobe 1998). As such, avoiding any artificial change in that balance—such as from 50c posts or censorship—can be valuable.

Distraction is a clever and useful strategy in information control in that an argument in almost any human discussion is rarely an effective way to put an end to

an opposing argument. Letting an argument die, or changing the subject, usually works much better than picking an argument and getting someone's back up (as new parents recognize fast). It may even be the case that the function of reasoning in human beings is fundamentally about winning arguments rather than resolving them by seeking truth (Mercier and Sperber 2011). Distraction even has the advantage of reducing anger compared to ruminating on the same issue (Denson et al. 2012). Finally, since censorship alone seems to anger people (Roberts 2014), the 50c astroturfing program has the additional advantage of enabling the government to actively control opinion without having to censor as much as they might otherwise.

2. Authoritarian Politics. For the literature on authoritarian politics in general, our results may help refine current theories of the role of information, and particularly what is known as common knowledge, in theories of revolutionary mobilization. Many theories in comparative politics assume that autocrats slow the spread of information critical of the regime to minimize the development of common knowledge of grievances which in turn may reduce the probability of mobilization against the regime. The idea is that coordination is essential to revolution, and coordination requires some common knowledge of shared grievances (Chwe 2013; Egorov et al. 2009; Hollyer et al. 2014; Persson and Tabellini 2006; Tilly 1978).

In contrast, our results suggest that the Chinese regime differentiates between two types of common knowledge—about specific grievances, which they allow, and about collective action potential, which they do a great deal to avoid. Avoiding the spread of common knowledge about collective action events (and not grievances) is consistent with research by Kuran (1989, 1991), Lohmann (1994), and Lorentzen (2013), who focus specifically on the spread of information about real-world protest and ongoing collective action rather than the generic spread of common knowledge more broadly.

The idea is that numerous grievances of a population ruled autocratically by nonelected leaders are obvious and omnipresent. Learning of one more grievance, in and of itself, should have little impact on the power of a potential revolutionary to ignite protest. The issue then appears not to be whether such grievances are learned by large enough numbers to foment a revolution. Instead, we can think of creative political actors, including those aspiring to lead a revolution or coup, as treating issues, ideologies, events, arguments, ideas, and grievances as “hooks on which politicians hang their objectives and by which they further their interests,” including interests that entail initiating or fostering a political uprising (Shepsle 1985). If one hook is not available, they can use another.

By this logic, then, common knowledge of grievances is already commonplace, and thus allowing more information about them to become public is of little risk to the regime or value to its opponents. Since disrupting discussion of grievances only limits information that is otherwise useful to the regime, the leaders have little

reason to censor it, argue with it, or flood the net with opposing viewpoints. What is risky for the regime, and therefore vigorously opposed through large-scale censorship and huge numbers of fabricated social media posts, is posts with collective action potential.

CONCLUDING REMARKS

Academics and policymakers have long been focused on contested physical spaces, over which military wars have been or might be fought. For example, in the South China Sea, the Chinese regime is presently building artificial islands and the United States is conducting military exercises, both highly expensive shows of power. As important as this focus may be, we believe that scholars and policymakers should focus considerably more effort on the Chinese Internet and its information environment, which is a contested virtual space, one that may well be more important than many contested physical spaces. The relationship between the government and the people is defined in this space, and thus the world has a great interest in what goes on there. We believe that considerably more resources and research should be devoted to this area. Whatever the appropriate relationship between governments and their people, a reasonable position is that it be open and known. This is an area where academic researchers can help. By devoting great effort, they can open up this knowledge to the world. It is our hope that others follow up on the research reported here.

More specifically, most journalists, activists, participants in social media, and some scholars have, until now, argued that the massive 50c party is devoted to engaging in argument that defends the regime, its leaders, and their policies. Our evidence indicates the opposite—that the 50c party engages in almost no argument of any kind and is instead devoted primarily to cheerleading for the state, symbols of the regime, or the revolutionary history of the Communist Party. We interpret these activities as the regime's effort at strategic distraction from collective action, grievances, or general negativity, and so forth.

It also appears that the 50c party is mostly composed of government employees contributing part time outside their regular jobs, not, as has been claimed, ordinary citizens paid piecemeal for their work. This, nevertheless, is still an enormous workforce that, we estimate, produces 448 million 50c posts per year. Their effectiveness appears maximized by the effort we found of them concentrating the posts into spikes at appropriate times and by directing about half of the posts to comments on government websites.

Appendix A. CATEGORIZATION SCHEME

Our categorization scheme for social media posts includes the six categories below, along with examples of each. Non-Chinese speakers should be aware when reading these examples that the Chinese language, even on social media, tends to be quite flowery and formal, with frequent creative, and often (to English speakers) stacy-sounding, wordings.

(1) Taunting of Foreign Countries Favorable comparisons of China to other countries; insults to other countries; taunting of pro-democracy, pro-West, pro-individual liberties, or pro-capitalist opinions within China. Examples from leaked Zhanggong 50c posts:

- 去年,奥巴马在香格里拉会议上力邀23国参与围堵中国时这样说道:“中国有13亿人,他们越崛起,我们就会越没饭吃,因为地球资源供给是有上限的。所以为了我们能继续过现在的生活,就必须遏制中国的发展。” [Last year, at the Shangri-la Dialogue where Obama invited 23 countries to participate in the containment of China, he said: “China has 1.3 billion people, the faster China rises, the more difficult it will be for us to live, because the earth’s resources are limited. For us to remain at our current living standard, we must contain China’s development.”]
- 中国的崛起大势已经不可阻挡。美国一边公开宣称不是中国死就是西方亡,一边又拼命告诉中国民众:你们的政府有问题啊,必须推翻它,然后你们就能过上比现在更好的日子。——请问,还有比这更可笑和自相矛盾的逻辑吗? [China’s rise is now inevitable. On one hand, the US publicly asserts that if China does not perish the West will wither; on the other hand it tells the Chinese people that your government is problematic: you have to overthrow it so you can live a better life. Is there a more ridiculous and contradictory logic than this?]

(2) Argumentative praise or criticism Comments on controversial, Pro/Con (non-valience) issues, as well as claims of wrongdoing or unfairness; praise (usually of the government) or criticism (usually of opponents of the government); taking a position or explaining why a particular viewpoint is correct or (more often) wrong. These posts are often part of a debate, in opposition to a previous post. Examples from leaked Zhanggong 50c posts:

- 我亲爱的朋友们,翻一下你的微薄,你会发现系统已经自动帮你添加了诸如薛蛮子、李开复、作业本、韩寒、李承鹏等各种民粹微薄,这是标准的强制灌输和洗脑手段,建议你取消关注 [My dear friends, you if you go through your Weibo, you’ll discover that the system automatically had you follow Xue Manzi, Li Kaifu, Zuo Yeben, Han Han, Li Chengpeng and other populist Weibo users. This is a typical tactic of indoctrination and brainwashing, I suggest you unfollow them.]
- 李开复说纽约60万美金一套别墅,比北京便宜多了,但他不会告诉你那套所谓的别墅其实是个仓库,而且离纽约市区开车需要四个多小时 [Li Kaifu says that you can buy a villa for \$600,000 USD in New York, much cheaper than in Beijing. But what he doesn’t tell you is that this so-called villa is actually a warehouse, which is more than a four hour drive from New York City.]

(3) Non-argumentative Praise or Suggestions Noncontroversial valience issues which are hard to argue against, such as improving housing and public welfare; praise of current government officials, programs, or policies. It doesn’t respond to alternative, opposing viewpoints, and it includes positive sentiment. It is distinguished from cat-

egory (2) in that it praises something specific such as the government, its officials, government programs, or initiatives, but does not take issue with another post. Includes a small number of constructive suggestions for what government policies might include (i.e., added benefits rather than critical complaints). It does not argue against a specific viewpoint, but just says “it would be nice if the government did X,” which usually the government is already in the process of implementing. Some examples of known Zhanggong 50c posts:

- 政府...做了好多实事,其中解决了好大一部分的人住房。 [The government has done a lot of practical things, among which is solving a significant part of the housing problem]
- 土坯房改造政策,让我们村的人搬出了坯房都住上了小洋楼,村里发生了翻天覆地的变化,真是太感谢了。 [The policy of renovating mud-brick houses has allowed villagers to move out of mud-brick dwellings into small, Western-style buildings. The village has been transformed, we are so grateful]
- 希望中央支持力度更大 [We hope the central government provides us with even more support]
- 希望能有更多类似《若干意见》的好政策! [We hope there will be more good policies like “Various Opinions” (the abbreviated name of an economic development policy)]
- 期待书记带领我们...特别是在教育、医疗卫生方面,争取更多政策为百姓谋取更大福祉。 [We look forward to the leadership of our party secretary... We hope that he can carry out more policies that will benefit the people in different aspects, especially in education and health care.]

(4) Factual Reporting Descriptions of current government programs, projects, events, or initiatives, or planned or in progress initiatives. Does not include any praise of these programs or events (which would be category (3)), just that they are occurring. Reporting on what government, government officials are doing. Some examples of known Zhanggong 50c posts:

- 清明三天假期7座小客车继续免费 [During the Qingming festival three-day holiday, [the freeway] will remain free to 7-seater buses]
- 6月27日,江西省委作出向龚全珍同志学习的决定,号召全省党员干部深入学习龚全珍坚定信念、永葆本色的坚强党性,服务群众、一心为民的质朴情怀,忠于职守、矢志不渝的执着追求,淡泊名利、无私奉献的高尚情操,勤俭节约、艰苦朴素的生活作风。 [On June 27, the Jiangxi provincial committee promulgated an opinion to learn from comrade Zhen Gongquan, calling on all provincial party members and cadres to study Zhen Gongquan’s firm conviction, staunch support of the Party’s spirit, service to the masses, straightforward dedication to the people, devotion to duty, abiding dedication, indifference to fame and fortune, selfless dedication to moral character and hardwork.]
- 1月16日,江西省委常委、赣州市委书记史文清将通过中国赣州网与网民在线交流,倾听网民意见、建议和诉求。 [On January 16, Jiangxi Party Committee

Member and Ganzhou City Party Secretary Shi Wenqing will communicate with netizens on the China Ganzhou Web, to hear comments, suggestions, and demands from netizens.]

(5) Cheerleading for China Patriotism, encouragement and motivation, inspirational quotes and slogans, inspirational quotes from government officials, thankfulness, gratefulness, inspiration or thankfulness for historical and aspirational figures or events, and cultural references and celebrations (e.g., describes traditions, actions, suggestions for the community). Excludes positive sentiment toward particular government leaders or specific policies (which would be category (3)), but includes positive sentiment or general praise toward life, historical figures, model citizens (e.g., Lei Feng; Gong Quanzhen, a model teacher; Guo Chuhui, a patriotic villager), or China in general. Some examples of known Zhanggong 50c posts:

- 众多革命先烈们的英勇奋斗，缔造了我们今天的幸福生活！向英雄致敬。 [Many revolutionary martyrs fought bravely to create the blessed life we have today! Respect these heroes.]
- 向所有为中华民族繁荣富强做出伟大贡献的先人们致敬！人民英雄永垂不朽 [Respect to all the people who have greatly contributed to the prosperity and success of the Chinese civilization! The heroes of the people are immortal]
- 接过父辈、祖辈血染的红旗，坚定不移地跟着党走！ [[I will] carry the red flag stained with the blood of our forefathers, and unswervingly follow the path of the CCP!]
- 我们自己要更加努力，不等不靠，主动上前。 [We all have to work harder, to rely on ourselves, and to take the initiative to move forward.]
- 爱我中华 [I love China]
- 大家的日子都过好了，中国梦就实现了！ [[If] everyone can live good lives, then the China Dream will be realized!]
- 赣州加油哦 [Way to go Ganzhou]

(6) Other Irrelevant posts that are entirely personal, commercial (such as ads), jokes, or empty posts that forward information not included. This category is removed and conditioned on in all analyses in this article.

Appendix B. AN UNINTENDED “SURVEY” OF THE CHINESE GOVERNMENT

We describe here a rare tacit confirmation of the existence of the 50c party, as well as an apparent admission to the accuracy of our leaked archive and the veracity of our empirical results, all unexpectedly offered by the Chinese government in response to our work.

Due to a set of unusual and unintended circumstances, an early draft of this article received considerable international attention, so much so that the *Global Times* wrote an editorial about it (this is a newspaper published by the *People's Daily*, the CCP's primary mouthpiece; see Wade 2016).⁸ Although

this editorial is not an official statement of the Chinese government, it is reasonable to interpret it as a close approximation, or at worst of a faction, of the government. (We offer a translation of the editorial, along with a contextual explanation of its content, in the Supplementary Appendix to this paper.)

The main purpose of the editorial is to strongly defend the government's unique system of public opinion guidance (see item 6 in Section 4). The editorial claims that “Chinese society is generally in agreement regarding the necessity of ‘public opinion guidance’” (其实对于“引导舆论”的必要性, 中国社会总体是认同的). To understand the government's position and perspective, it is helpful to use the viral discussion of our paper in social media, following its unexpected news coverage, to test the editorial's claim. To do this, we downloaded posts from two sources, comments on the *Global Times* site and a broader sample from Weibo responding to the editorial. We used ReadMe, as above, to analyze each corpus separately.

We would expect more support for public opinion guidance from comments on a nationalist newspaper website, and much less support (than the regime acknowledges) from a more general population (consistent with Roberts 2014). Indeed, this is just what we found. Our estimates indicate that 82% of the comments on the newspaper's website which expressed an opinion supported China's system of public opinion guidance (with 15% critical). Yet, among the likely broader audience found on Weibo, only 30% were supportive (with 63% critical), clearly contradicting the editorial's rosy view of the government's popularity.

The fact that the regime's central strategy for controlling the dynamic and highly contested social media space lacks universal support likely made the regime feel it all the more urgent to defend public opinion guidance in this forum. Authoritarian regimes like China, with strong international and military power, are usually focused on threats to their rule from their own people rather than, in this case, the international press (or scientific community). Confirming the following four points central to our article (as opposed to denying their previously surreptitious behavior) was of incidental relevance to government leaders but served the purpose of enabling them to engage the discussion and explicitly defend their information control practices.

First, although the *Global Times* has English and Chinese editions, with many articles published in both languages, the editorial about our paper was published only in Chinese. That is, even though it objected to how the story was covered in the international press, the CCP was primarily addressing its own people. This seems to be a regular strategy of the regime and is consistent with our interpretation of their main perceived threats being their own people rather than Western powers.

when a reporter from a major international news outlet somehow obtained a copy and contacted us about publishing a story about it. We asked him to hold off because we had not finished revising, nor of course was it even under review yet. He refused, explaining that our results were too important and he did not want to be scooped by his journalistic competitors. Freedom of the press obviously gave us no recourse, and so we gave in, answered his questions, and posted the paper on our web sites. Within a few days, over 5,000 news outlets across the globe published stories on our paper (according to Google news). The reporter who broke the story turned out to be right, as his next fastest competitor published only a few hours after he did, followed soon by many others.

⁸ We shared an early draft of this article privately with colleagues and others to solicit comments. We were in the process of revising

Second, the editorial appears to admit to the existence of the 50c party and at least tacitly confirms the veracity of our leaked archive. They made these admissions apparently in order to turn the conversation into an explanation for their people about why public opinion guidance is essential. They also use the editorial to explain that traditional public opinion guidance is no longer sufficient to prevent the increase in viral messaging under control of those outside the government, which can spark or fuel collective action. Due to the rise of social media, the editorial says the government has “no choice” but to implement stronger information control practices designed for this new form of communication, such as 50c party activity. In other words, the 50c party exists but the Chinese people should not be focused on it.

Third, in a forum that regularly expresses opinions, including disapproval and disagreement, the editorial began with a summary of our empirical results, and took no issue with any of our conclusions.⁹ Thus, for all practical purposes, the editorial constitutes the answer to a simple sample survey question: That is, instead of asking 50c party members about their status as we do in Section 5, we (inadvertently) asked the Chinese government whether they agreed with our results, and they effectively concurred. Although social scientists often conduct interviews of individual public officials, we are grateful for the unusual, if not unprecedented, chance to pose questions to an organ of the Chinese government and have it respond, for all practical purposes, as a government, or at least in a way that represents it.

Finally, in the editorial, the government also acknowledges that the purpose of public opinion guidance is to constrain or stop the spread of “hot button issues” that go viral online or “grassroots social issues” that have collective action potential. This also confirms a central point of our work.

SUPPLEMENTARY MATERIALS

To view supplementary material for this article, please visit <https://doi.org/10.1017/S0003055417000144>.

Replication files can be found at <https://doi.org/10.7910/DVN/QSZMPD>.

REFERENCES

- Bambauer, Derek E. 2013. “Censorship v3.1.” *IEEE Internet Computing* 17(3): 26–33.
- Bandurski, David. 2008. “China’s Guerrilla War for the Web.” *Far Eastern Economic Review* 171(6): 41–44.
- Barr, Michael. 2012. “Nation Branding as Nation Building: China’s Image Campaign.” *East Asia* 29(1): 81–94.
- Brady, Anne, -Marie. 2009. *Marketing Dictatorship: Propaganda and Thought Work in Contemporary China*. Lanham, MD: Rowman & Littlefield.
- Bremmer, Ian. 2010. “Democracy in Cyberspace.” *Foreign Affairs* 89(6): 86–92.

⁹ The editorial did take one issue with our work (in addition to disagreeing with how the international press covered it), which was to imply that we had confused 50c party posts with those from other groups unauthorized by the Chinese government, but involved in similar activities. In fact, the only posts we analyzed and called 50c were those officially generated by the government. In case this was a misunderstanding, we clarify this point by adding what is now Footnote 2 to our paper.

- Cairns, Christopher, and Allen Carlson. 2016. “Real-World Islands in a Social Media Sea: Nationalism and Censorship on Weibo During the 2012 Diaoyu/Senkaku Crisis.” *China Quarterly* 225: 23–49.
- Chen, Jidong, Jennifer Pan, and Yiqing Xu. 2016. “Sources of Authoritarian Responsiveness: A Field Experiment in China.” *American Journal of Political Science* 60(2): 383–400.
- China Digital Space. 2016. “Fifty Cents.” China Digital Times. <http://j.mp/cdt50cent>.
- Chwe, Michael Suk-Young. 2013. *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton, NJ: Princeton University Press.
- China Internet Network Information Center (CNNIC). 2014. Statistical Report on Internet Development in China. Beijing, CN: CNNIC. <http://j.mp/StatDevC>.
- Chwe, Michael Suk-Young. 2013. *Rational ritual: Culture, coordination, and common knowledge*. Princeton, NJ: Princeton University Press.
- Deibert, Ronald, and Rafal Rohozinski. 2010. “Liberation vs. Control: The Future of Cyberspace.” *Journal of Democracy* 21(4): 43–57.
- Denson, Thomas F., Michelle L. Moulds, and Jessica R. Grisham. 2012. “The Effects of Analytical Rumination, Reappraisal, and Distraction on Anger Experience.” *Behavior Therapy* 43(2): 355–64.
- Dimitrov, Martin K. 2014a. “Internal Government Assessments of the Quality of Governance in China.” *Studies in Comparative International Development* 50(1): 50–72.
- Dimitrov, Martin K. 2014b. “Tracking Public Opinion Under Authoritarianism.” *Russian History* 41(3): 329–53.
- Dimitrov, Martin K. 2014c. “What the Party Wanted to Know: Citizen Complaints as a ‘Barometer of Public Opinion’ in Communist Bulgaria.” *East European Politics and Societies* 28(2): 271–95.
- Economist. 2013. “Cat and Mouse: How China Makes Sure its Internet Abides by the Rules.” *The Economist*. <http://www.economist.com/news/special-report/21574629-how-china-makes-sure-its-internet-abides-rules-cat-and-mouse>.
- Egorov, Georgy, Sergei Guriev, and Konstantin Sonin. 2009. “Why Resource-Poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data.” *American Political Science Review* 103(4): 645–68.
- Freedom House. 2009. “China and the Internet: An Uphill Fight for Freedom.” *Harvard International Review* 31(2): 68–73.
- Greitens, Sheena Chestnut. 2013. “Authoritarianism Online: What Can We Learn from Internet Data in Nondemocracies?” *PS: Political Science and Politics* 46(02): 262–70.
- Groves, Robert M., Floyd J. Fowler, Jr, Mick P. Couper, James M. Lepkowski, Eleanor Singer, and Roger Tourangeau. 2011. *Survey Methodology*. Vol. 561. John Wiley & Sons: Please provide publisher location for Groves reference..
- Haley, Usha. 2010. “China’s Fifty Cent Party for Internet Propaganda.” *Huffington Post*. http://huffingtonpost.com/usha-haley/chinas-fifty-cent-party-f_1_b_749989.html.
- Han, Rongbin. 2015a. “Defending the Authoritarian Regime Online: China’s ‘Voluntary Fifty-Cent Army.’” *China Quarterly* 224: 1006–25.
- Han, Rongbin. 2015b. “Manufacturing Consent in Cyberspace: China’s ‘Fifty-Cent Army.’” *Journal of Current Chinese Affairs* 44(2): 105–34.
- Hassid, Jonathan. 2012. “Safety Valve or Pressure Cooker? Blogs in Chinese Political Life.” *Journal of Communication* 62(2): 212–30.
- Henochowicz, Anne. 2014. “Thousands of Local Internet Propaganda Emails Leaked.” *China Digital Times*. <http://j.mp/leakedEmails>.
- Hollyer, James R., B. Peter Rosendorff, and James R. Vreeland. 2014. “Measuring Transparency.” *Political Analysis* 22(4): 413–34.
- Hopkins, Daniel, and Gary King. 2010. “A Method of Automated Nonparametric Content Analysis for Social Science.” *American Journal of Political Science* 54 (1): 229–47. <http://j.mp/jNFDgl>.
- iiMedia Research Group. 2014. “2013 年中国移动社会化分享年度报告.” China Mobile Social Share Annual Report (内部版). <http://iimedia.cn/37109.html>.
- Johnson, Ian. 2016. “The Presence of the Past: A Coda.” In *The Oxford Illustrated History of Modern China*, ed. Jeffrey N. Wasserstrom. Oxford, UK: Oxford University Press, 301–324.

- Keating, Joshua. 2011. "China's '50-Cent' Party Takes on the Jasmine Revolutions." *Foreign Policy*. <http://foreignpolicy.com/2011/03/01/chinas-50-cent-party-takes-on-the-jasmine-revolutions/>.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107: 1–18. <http://j.mp/LdVXqN>.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2014. "Reverse-Engineering Censorship in China: Randomized Experimentation and Participant Observation." *Science* 345 (6199): 1–10. <http://j.mp/1KbwkJJ>.
- King, Gary, Patrick Lam, and Margaret Roberts. In Press. "Computer-Assisted Keyword and Document Set Discovery from Unstructured Text." *American Journal of Political Science*. Copy at <http://j.mp/2nxUa8N>.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2017. "Replication Data for: How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument." Harvard Dataverse, V1, [UNF:6:/3HIELdmrcyZm+v5mx0OJg==]. DOI: <https://doi.org/10.7910/DVN/QSZMPD>, URL: <dx.doi.org/10.7910/DVN/QSZMPD>.
- Knockel, Jeffrey, Masashi Crete-Nishihata, Jason Q. Ng, Adam Senft, and Jedidiah R. Crandall. 2015. "Every Rose Has Its Thorn: Censorship and Surveillance on Social Video Platforms in China." In *Proceedings of the 5th USENIX Workshop on Free and Open Communications on the Internet (FOCI '15)*.
- Kuran, Timur. 1989. "Sparks and Prairie Fires: A Theory of Unanticipated Political Revolution." *Public Choice* 61(1): 41–74.
- Kuran, Timur. 1991. "Now Out of Never: The Element of Surprise in the East European Revolution of 1989." *World Politics* 44(1): 7–48.
- Lam, Oiwan. 2012. "China: Challenging the 50 Cent Party." *Global Voices*. <https://globalvoices.org/2012/03/04/china-challenging-the-50-cent-party/>.
- Lam, Oiwan. 2013. "China Beefs Up '50 Cent' Army of Paid Internet Propagandists." *Global Voices*. <https://advoc.globalvoices.org/2013/10/17/china-beefs-up-50-cent-army-of-paid-internet-propagandists/>.
- Latham, Benjamin et al. 2015. "Interpretable Classifiers Using Rules and Bayesian Analysis: Building a Better Stroke Prediction Model." *Annals of Applied Statistics* 9(3): 1350–71.
- Lohmann, Susanne. 1994. "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989–1991." *World Politics* 47(1): 42–101.
- Lorentzen, Peter L. 2013. "Regularizing Rioting: Permitting Public Protest in an Authoritarian Regime." *Quarterly Journal of Political Science* 8(2): 127–58.
- Lynch, Daniel C. 1999. *After the Propaganda State: Media, Politics, and "Thought Work" in Reformed China*. Stanford, CA: Stanford University Press.
- MacKinnon, Rebecca. 2012. *Consent of the Networked: The Worldwide Struggle For Internet Freedom*. New York, NY: Basic Books.
- Mercier, Hugo, and Dan Sperber. 2011. "Why Do Humans Reason? Arguments for an Argumentative Theory." *Behavioral and Brain Sciences* 34(02): 57–74.
- Miller, Blake Andrew Phillip. 2016. "Automatic Detection of Comment Propaganda in Chinese Media." Available at https://papers.ssrn.com/sol3/Papers.cfm?abstract_id=2738325.
- Näher, Anatol-Fiete, and Ivar Krumpal. 2012. "Asking Sensitive Questions: The Impact of Forgiving Wording and Question Context on Social Desirability Bias." *Quality and Quantity* 46(5): 1601–16.
- National Bureau of Statistics of China. 2014. *China Statistical Year Book*. Beijing, CN: China Statistics Press. <http://j.mp/ChinaSY>.
- Ng, Jason. 2011. "Wu Mao." *Blocked on Weibo blog*. <https://blockedonweibo.tumblr.com>.
- Ng, Jason Q. 2015. "Politics, Rumors, and Ambiguity: Tracking Censorship on WeChat's Public Accounts Platform." *Munk School of Global Affairs*. <https://citizenlab.org/2015/07/tracking-censorship-on-wechat-public-accounts-platform/>.
- Pan, Jennifer. 2015. "Buying Inertia: Preempting Social Disorder with Selective Welfare Provision in Urban China." PhD dissertation. Harvard University, Cambridge, MA.
- Persson, Torsten, and Guido Tabellini. 2006. "Democracy and Development: The Devil in the Details." *American Economic Review* 96: 319–24.
- Pew Research Center. 2014. "The Political Typology: Beyond Red vs. Blue." *Pew Research*. <http://www.people-press.org/2014/06/26/the-political-typology-beyond-red-vs-blue/>.
- Philipp, Joshua. 2015. "Leaked Emails Show Chinese Regime Employs 500,000 Internet Trolls." *Epoch Times*. <http://j.mp/twomill>.
- Roberts, Margaret Earling. 2014. *Fear, Friction, and Flooding: Methods of Online Information Control*. Dissertation. Harvard University.
- Roberts, Margaret E. 2015. "Experiencing Censorship Emboldens Internet Users and Decreases Government Support in China." Unpublished Working Paper. <http://margaretroberts.net/wp-content/uploads/2015/07/fear.pdf>.
- Shepsle, Kenneth A. 1985. "Comment of Why the Regulators Chose to Deregulate." In *Regulatory Policy and the Social Sciences*, ed. Roger Noll. Berkeley, CA: University of California Press, 231–39.
- Shirk, Susan L. 2011. *Changing Media, Changing China*. New York, NY: Oxford University Press.
- Sonnad, Nikhil. 2014. "Hacked Emails Reveal China's Elaborate and Absurd Internet Propaganda Machine." *Quartz*. <http://j.mp/Sonnad>.
- Stockmann, Daniela. 2013. *Media Commercialization and Authoritarian Rule in China*. Cambridge, UK: Cambridge University Press.
- Stockmann, Daniela, and Mary E. Gallagher. 2011. "Remote Control: How the Media Sustain Authoritarian Rule in China." *Comparative Political Studies* 44(4): 436–67.
- Strafella, Giorgio, and Daria Berg. 2015. "'Twitter Bodhisattva': Ai Weiwei's Media Politics." *Asian Studies Review* 39(1): 138–57.
- Tang, Min, Laia Jorba, and Michael J. Jensen. 2012. "Digital Media and Political Attitudes in China." In *Digital Media and Political Engagement Worldwide: A Comparative Study*, ed. Eva Anduiza. Cambridge, UK: Cambridge University Press, 221–39.
- Tilly, Charles. 1978. *From Mobilization to Revolution*. New York, NY: McGraw-Hill.
- Tong, Yanqi, and Shaohua Lei. 2013. "War of Position and Microblogging in China." *Journal of Contemporary China* 22(80): 292–311.
- Tourangeau, Roger, Frederick Conrad, and Mick Couper. 2013. *The Science of Web Surveys*. Oxford, England: Oxford University Press.
- Truex, Rory. 2016. "Focal Points, Dissident Calendars, and Preemptive Repression." *SSRN*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2802859.
- Wade, Samuel. 2016. "Who Does Global Times Speak For?" *China Digital Times*. <http://j.mp/GT-CCP>.
- Wintrobe, Ronald. 1998. *The Political Economy of Dictatorship*. New York, NY: Cambridge University Press.
- Xi, Jinping. 2014. "Overall Plan Involving All Parties for Innovation and Development to Strive to Build Our Country into an Internet Power." <http://j.mp/XiJingSph>.
- Yang, Guobin. 2009. *The Power of the Internet in China: Citizen Activism Online*. New York, NY: Columbia University Press.
- Yang, Xiaofeng, Qian Yang, and Christo Wilson. 2015. "Penny for Your Thoughts: Searching for the 50 Cent Party on Sina Weibo." In *Proceedings of the 9th International AAAI Conference on Web and Social Media (AAAI '15)*.
- Zhang, Qingpeng, Dominic DiFranzo, and James A. Hendler. 2014. "Social Networking on the World Wide Web." In *Encyclopedia of Social Network Analysis and Mining*. Berlin, Germany: Springer, 1879–92.
- Zhao, Juanjuan, Weili Wu, Xiaolong Zhang, Yan Qiang, Tao Liu, and Lidong Wu. 2014. "A Short-Term Trend Prediction Model of Topic over Sina Weibo Dataset." *Journal of Combinatorial Optimization* 28(3): 613–25.