

Testing ‘Control 2.0’: Why Chinese Authorities Allow or Suppress State and Social Media Coverage of Labor Disputes

Christopher Cairns* and Manfred Elfstrom**

Abstract

Rising labor unrest has historically brought political change. How does an authoritarian state manage the impact of worker activism in today's digital age? Using a variety of data sources, this article finds that rather than simply blocking all news of strikes, protests, and riots, China's state media reports those incidents that are impossible to ignore, such as transit strikes, while highlighting “positive” outcomes, e.g., mediation by officials. Chinese social media follows a different path, though, focusing on large-scale conflicts with a nationalist angle, suggesting either limits to authorities' abilities to curate the Internet or a different censorship and propaganda logic.

Introduction

Scholars have long understood working class mobilization to be an important ingredient in political change. Rueschemeyer, Stephens, and Stephens argue that workers—not the middle class—ultimately catalyze democratization.¹ In Huntington's analysis, conversely, labor activism absent mature institutions is a recipe for instability or the emergence of a “praetorian state.”² Historical institutionalists view the specific alliances formed between workers and other classes as crucial in determining their countries' regime choices.³ Outside Western Europe, labor's influence has been documented at “critical junctures” in Latin America;⁴ in Eastern Europe and the former Soviet Union;⁵ and in post-colonial Southeast Asia.⁶

In this article, we examine how the Chinese state is seeking to control news reporting and online discussion surrounding an astonishing wave of worker activism. We use a crowd-sourced and geo-referenced dataset of strikes, protests and riots by Chinese workers; Web searches of

*PhD Candidate, Department of Government, Cornell University. cmc467@cornell.edu

**Postdoctoral Fellow, Ash Center for Democratic Governance and Innovation, Harvard University. manfred_elfstrom@hks.harvard.edu

state media; a collection of postings from the country's popular micro-blogging service, Sina *Weibo*; and one case study to test a theory of Chinese media governance that David Bandurski calls "Control 2.0."⁷ According to this theory, Beijing is increasingly trying to get out ahead of "sensitive" stories, reporting itself on contentious collective actions instead of simply suppressing coverage,⁸ while making sure to portray the state's role in the best possible light. The implications for worker-led change are complex: news of unrest is more readily available than ever before, but underground narratives must contend with a sophisticated official retelling of events.

We hypothesize that three conditions determine whether Chinese authorities will opt for "Control 2.0" or fall back on censorship ("Control 1.0") when confronted with working class contention: 1) the 'publicness' of an incident; 2) whether it touches on cornerstones of the state's legitimacy; and 3) its volatility, i.e. its potential to spur follow-on protests. With regard to newspapers, we find that the bigger an incident is, i.e. the more public it is, the more likely it is to be reported—up to a certain point, when unrest becomes so large-scale that the state's instinct to suppress kicks in again. Given their disruptiveness, transit strikes are similarly more likely to receive state media coverage than other kinds of unrest. Unexpectedly, protests in foreign-invested and state-owned enterprises (SOEs) attract no more or less coverage than conflicts in other firms. But consistent with our theory, newspapers tend to emphasize positive official interventions and avoid volatile tales of police or thug violence.

Our analysis of Weibo commentary reveals that publicness, legitimacy, and volatility function very differently on social media: we observe different effect sizes, significance, and even directions of correlation compared with newspapers. Micro-bloggers are uninterested in government efforts to resolve disputes, fixating instead on incidents at foreign companies.

Moreover, the bigger a protest, the more likely Weibo users are to discuss it. In other words, the government appears unable or unwilling to keep social media conversations within the same parameters as newspapers. However, a negative correlation between police or thug violence and micro-blog coverage suggests that authorities nonetheless exert cruder control by deleting postings and blocking search terms they consider harmful. Whether this tug-of-war between workers and the state and between popular and official narratives will ultimately lead to increased openness or to the effective containment of labor insurgency is unclear.

In the following section, we provide background on China's rising labor protests and on media and censorship in the country. Next, we introduce our main concepts and hypotheses. The following section details our data sources. Then, we estimate two logit models. After discussing our statistical results, we illustrate our argument with a case study of a taxi strike in Lanzhou. Finally, we draw tentative conclusions regarding labor globally in today's media environment.

Rising Labor Unrest

Labor activism has played an important role at key points in modern Chinese history.⁹ Since the mid-1990s, a fresh round of unrest has swept China's shop floors. Between 1996 and 2010, the annual number of mediated, arbitrated, and litigated labor disputes rose from 48,121 to 600,865.¹⁰ Meanwhile, between 1993 and 2005, "mass incidents" jumped ten-fold from 8,700 to 87,000 incidents per year, with approximately a third involving workers.¹¹ Subtle but important changes in the form of labor protest have occurred. In the early 2000s, researchers noted sharp divides between laid off state-owned enterprise (SOE) and migrant workers,¹² between laid off workers in different regions of the country,¹³ between generations of workers, and, within factories, between workers with different social connections.¹⁴ These divisions were said to be driven by unique foreign direct investment policies¹⁵ and by patterns in the allocation of welfare

benefits,¹⁶ among other factors. Workers' demands were also largely regarded as defensive: focused on protecting a tattered Mao-era "socialist social contract" or enforcing minimal legal protections.¹⁷ Now, supported by increasingly sophisticated organizing, workers from different backgrounds seem to have converged around a set of more aggressive demands: for higher wages; attention to the details of workplace life; and simple respect.¹⁸ Scholars at government think tanks have begun warning of serious threats to social order from angry "disadvantaged groups"¹⁹ and the official All China Federation of Trade Unions has scrambled to regain workers' trust.²⁰

Chinese Media Marketization, Censorship, and Public Opinion Guidance

As labor politics in China have evolved, so too have the roles of state newspapers and the Internet. The marketization of state newspapers has dramatically altered the ability of and incentives for journalists to serve the Party-state's interests. As Daniel C. Lynch and Daniela Stockmann argue, the conversion of state-owned media outlets into private companies in the reform period was necessary to ensure their profitability.²¹ In the process, though, incentives were created for publications to deviate from dull coverage of leader meetings and improved harvests and to instead cover topics and use a writing style with greater mass appeal. The result was a product with greater influence over the hearts and minds of Chinese citizens—but one over which propaganda authorities, through supervision and regulation, retained ultimate control. Regarding labor, Daniela Stockmann and Mary Gallagher posit that news reports about citizens using the country's new labor laws "as a weapon" have driven aggrieved workers toward legal remedies, not street protests.²² State publications have simultaneously been pushed to exercise "public opinion supervision" (*yulun jian du*) over corrupt local officials, holding them accountable to superiors.

Since the early 2000s, China's online population has continually increased. In contrast to scholarly discussions about state newspapers' role, debate about the Web's impact in China has centered not on its power as a propaganda venue, but its status as a space to be tamed—and whether the state's legal and technological means of censoring it have outpaced the ability of “netizen” activists to spread counter-hegemonic messages. Cyber optimists see a tug-of-war between approximately equal contestants, with bloggers fighting a “war of position” against official sources²³ in a realm where “authority of all kinds is subject to doubt and ridicule.”²⁴ In this vein, Xiao Qiang describes the viral nature of certain critical stories online as like “water pouring through a hole in a dam” whose speed and force “denies the government agenda-setting power.”²⁵ However, cyber pessimists note that Internet company censors, under close official scrutiny, can remove over 90% of posts on a topic designated for censorship within 24 hours of their posting.²⁶ In an influential article, Gary King, Jennifer Pan and Margaret E. Roberts²⁷ find that the government targets “volume bursts” of posts that “represent, reinforce, or spur social mobilization,” while allowing mere criticisms of authorities. Rebecca MacKinnon contends that online complainants who deploy Albert O. Hirschman's concept of “voice”—remaining loyal to the system while petitioning for changes—are more likely to succeed than those who express a desire for “exit” (challenge the government).²⁸ Is *de facto* freedom of speech slowly being established in China or is discussion unfolding largely on the government's terms?

The Control 2.0 Theory

This paper tests a theory of Chinese media governance that David Bandurski calls “Control 2.0.”²⁹ According to Bandurski, since at least the 2008 Beijing Olympics and that year's uprising in Tibet and Sichuan earthquake, authorities have aimed to provide a greater degree of believability and perceived openness in authorized forums by acknowledging unrest, rather than

suppressing all related discussion. At the same time, they have tried to frame protests in terms favorable to the state. Officials may also, we add, encourage discussions of contention to pressure employers or local officials to act in workers' favor, whether to ensure social stability or to achieve economic goals. Here, we seek to go beyond the existing literature and theorize three conditions under which officials are likely to choose a "Control 2.0" response when dealing with worker actions, rather than mere information suppression ("Control 1.0"). We refer to these conditions as *publicness*, *legitimacy*, and *volatility*.

Publicness

Officials may resort to "Control 2.0" when an incident has become high-profile enough that it can no longer be effectively suppressed from public view. Recent work on censorship suggests that rather than intimidating "netizens" and "chilling" online discussions, *visible* censorship (censorship that netizens observe taking place) actually increases their interest in a topic.³⁰ Additionally, the government has learned, from experiences like its botched handling of the 2003 SARS epidemic, that attempting to suppress big news when the public already has alternative sources of information is futile. Although "mass incidents" involving workers do not have the same public impact as epidemics or other mass disasters, very large or especially disruptive protests may fall into the same category, forcing the government to deal with the problem via skillful messaging.

Legitimacy

Second, the legitimacy of protester grievances matters for what strategy officials choose. Continuing MacKinnon's discussion above, the government is likely to view select groups in Chinese society as interested in exercising "voice" rather than "exit". It may also find certain

groups' claims hard to deny for historical reasons. Particular groups of workers are well-placed to take advantage of the state's own rhetoric.³¹ When dealing with such groups, the state is likely to employ "Control 2.0", as this provides officials a platform for acknowledging what they see as legitimate demands and for portraying themselves as consistent and principled.

Volatility

Finally, protest volatility matters. By this we mean an incident's likelihood of escalating further or spreading across social groups. Certain confrontations feature brutality by elites that, if publically known, could spur further actions by either the original participants, or newcomers.

Hypotheses

Our first hypothesis is that predictions regarding *publicness*, *legitimacy* and *volatility* should apply equally to state newspapers, and social media like Weibo; we have no *a priori* reason to believe that a "Control 2.0" approach better suits one or the other. One empirical implication is that coverage by the two kinds of media should be positively correlated:

H1: Social media discussion of an issue will increase its likelihood of receiving newspaper reporting, and newspaper reporting will increase the amount of social media discussion.

Concerning legitimacy, nationalism in China, by many accounts, is on the rise. Since the Tiananmen Square protests, schools have educated children about China's "century of national humiliation" at the hands of foreigners lasting from the Opium Wars through to "liberation" by the Communist Party in 1949.³² Michael E. Santoro warns of a near-future where "China blames other countries for its economic woes" and "some Chinese begin to characterize Western investment... as being similar to the foreign domination of the nineteenth and twentieth

centuries.”³³ Foreign executives sometimes complain that the state already targets them for criticism:

H2: The government is likely to allow newspaper and social media coverage of disputes in foreign firms.

Also regarding legitimacy, in 1978, state-owned enterprises accounted for 77 percent of industrial output and their profits totaled 14 percent of GDP; by 1996, their share of output had fallen to 33 percent and their profits amounted to less than 1 percent of GDP.³⁴ There are nonetheless reasons to believe that coverage of labor unrest in SOEs is less common than dispute coverage of private enterprises. Ideologically, SOEs still stand for the CCP’s commitment to maintain control of the “commanding heights” of the economy. SOEs are, moreover, often based in company towns, where firm, and local authorities’ interests are virtually indistinguishable:

H3: Incidents involving state-owned enterprises should appear less in newspapers and social media alike.

Concerning the *publicness* of an incident, we believe there are at least two categories of incidents likely to impact the broader public enough that officials are unlikely to censor them:

H4: Incidents involving more workers are more likely to be covered.

H5: Transit strikes—work stoppages involving cabbies, bus drivers, railway workers and truckers—are more likely to be covered.

Finally, regarding *volatility*, we reason that strikes with pro-worker outcomes are less likely to escalate, because they have been resolved and demonstrate to others the state’s impartiality. However, incidents that involve violence against workers are volatile in that they could lead to spiraling violence or cause outrage in the broader population over the lack of protection provided to vulnerable groups:

H6: Incidents with pro-worker outcomes are more likely to be covered.

H7: Incidents where violence occurs are less likely to be covered.

Next, we explain how we proceed to test these hypotheses.

Data Sources & Methodology

Data on Strikes and Protests by Chinese Workers

The starting point for this article is a geo-referenced dataset of 793 strikes, protests, and riots by Chinese workers between 2008 and 2012 called *China Strikes*.³⁵ We focus on the year 2012, as it is the year most thoroughly documented by the dataset (393 incidents). Estimates of the total number of conflicts in this period differ, but although *China Strikes* is the most comprehensive collection available, it still likely only represents a small sample of what is happening.³⁶ To be included in the dataset, an event must be contentious (i.e., go beyond litigation), involve collective action, and relate to an employment dispute (not, for example, ethnic tensions between workers). *China Strikes* draws principally on Chinese and foreign newspapers and blogs. However, the website connected to the dataset also allows visitors to submit their own reports, and many of these tip-offs are included. *China Strikes* has been checked against another mapping project by the Hong Kong-based China Labour Bulletin (CLB),³⁷ and any incidents captured by CLB but initially missed by *China Strikes* have been added to the latter. Although *China Strikes* may be biased toward coastal areas with more Internet users, more foreign journalists, and more commercialized domestic media, the dataset covers a remarkably wide swath of the country, including everything from small cities in Yunnan to Pearl River Delta boomtowns. The occupations represented are also diverse, ranging from car salespeople to shipbuilders. Small gatherings of less than a dozen workers appear alongside mobilizations of thousands. One area of concern may be, though, that more information is known

about some incidents than others (e.g., small town incidents may only be known through brief posts in web forums), resulting in more thorough coding of outcomes in some cases than others. Nonetheless, at present, *China Strikes*, while imperfect, is the most comprehensive record available of the years it documents.³⁸

State Media Coverage

To measure state media coverage of strikes, we searched Baidu.com for each of the 393 incidents from 2012 in the *China Strikes* dataset, using as search terms different combinations of the place and date of the incident, the name of the enterprise or sector in question, the grievances of the workers involved, and the actions reportedly taken by workers, such as strikes, translated as *bagong*, *tinggong*, and *bake* (in the case of teachers' strikes); protests, translated as *kangyi*, *shiwei*, and *tao shuofa*; and riots, translated as *baoluan*. Sometimes searching for grievances without including the (more politically sensitive) actions taken by workers yielded more results. Often, though, searching for grievances without reference to the actions taken only resulted in job forums where workers shared information about company conditions. If the original source for a *China Strikes* incident was already a state media source (or appeared to be a state media report reposted by advocacy groups), then sections of the complete text were used as search terms to speed the process along.

Any state media coverage was coded a 1 and non-coverage, 0. We understand "state media coverage" to be reporting in local and national newspapers or state news agencies (e.g., *Guangzhou Ribao*, *People's Daily*, and *Xinhua*), as well as semi-official news magazines and news websites (e.g., *Caijing* or *Economic Observer Online*). Sources that are not primarily news-oriented, such as industry forums, are not included.³⁹ We found that if state media reported an incident, those reports rose to the top of searches, with web forums, industry websites, and other

informal coverage appearing lower down. Thus, we are unlikely to have missed any official coverage. In all, 42.7% of incidents were captured by state media. Where news we came across deepened our understanding of an incident or corrected flaws in the dataset, we updated *China Strikes* accordingly. The process was therefore somewhat iterative. However, the corrections do not substantially affect our results.

Weibo Data and Censorship

To measure social media content related to disputes, we used a dataset titled *WeiboScope*. The dataset is a sample of the following individuals: a) a group of ten University of Hong Kong researchers who themselves posted frequently on Sina Weibo (often referred to as China’s ‘Twitter’), b) a group of roughly 5,000 “Chinese dissident writers, journalists, and scholars”, and c) about 38,000 “users with an authenticated, or VIP, status and more than 10,000 followers.”⁴⁰ Data are available for the entire year of 2012, and involve daily (or more frequent) crawls of the user timelines of each of the above-mentioned individuals to gather new updates and store these in the dataset. Each row in the dataset consists of one social media post (including reposted content) plus associated meta-data: e.g. a pseudo user ID (the Hong Kong team anonymized user identities). The data files do not include images or videos. For this project, we use only the post text, including embedded retweets but excluding comments. This sample, because it is not random, does not represent all voices on Chinese social media. However, for our purposes its bias toward a more elite segment of society makes it more useful than a random sample, as it speaks to the degree to which news about worker unrest is picked up by individuals not directly involved. In other words, the individuals one would expect to be most likely to post about strikes online, such as labor organizers, journalists and workers themselves, are not (mostly) the same individuals who comprise the above sample. Rather, the sample reflects a wide variety of

intellectuals, scholars, and other people with considerable influence over *Weibo* users at large, who tend to be middle class. It thus documents precisely the potential for cross-group “contagion” that the Party fears.

To prepare the data for analysis, we first ran scans of the entire *WeiboScope* dataset, filtering posts according to a) the names of cities where incidents in the *ChinaStrikes* dataset had occurred that year, b) the week in which each incident occurred, and c) whether the post searched contained one of two keywords: *bagong* (“strike”) or *tinggong* (“stop work”). With a few exceptions, we used the city/municipality name where the strike occurred to locate the incident.⁴¹ This procedure yielded counts of the number of mentions of each keyword, per city, per week. We then matched these counts with known incidents in the *China Strikes* data using each incident’s date and location.

Quantitative Models

This paper measures six different independent variables, each of which tests different aspects of our seven hypotheses, across two models, each with a different outcome variable. The main outcome variable in the first model is an indicator, *newspaper*, for whether a strike was mentioned in one or more newspapers captured by Baidu News search. The second model is a binary outcome, *Weibo mentions*, that captured the probability that Weibo traffic is zero, or nonzero for a given incident. Both models use logistic regression.

Case Study: Lanzhou Taxi Strike

Finally, we conduct a brief case study—a taxi strike in Lanzhou—to illustrate what our results mean practically. The taxi strike is representative of the sorts of incidents we find are subject to a “Control 2.0” approach.

Worker Activism and News Coverage

In both quantitative models, we treat our data as a panel, grouping by geography. Important regional differences exist in China for both the local news media's ability to openly cover contentious issues like labor unrest, and the degree to which Weibo posters are a) active and b) allowed to discuss labor issues uncensored. We rely on Sinologist G. William Skinner's concept of "macroregions,"⁴² which divides China into nine areas based on river basins that cut across provinces. While Skinner's concept is intended to apply mainly to pre-modern China, it remains a good approximation of contemporary differences in economic development (and therefore, labor activity and media coverage of such activity). Grouping the incidents in this manner yielded nine groups, with the smallest containing 12 incidents and the largest containing 149, for an average of 43.6.⁴³ Since the macroregions are all-encompassing within the study's scope (China), rather than drawn from a larger population, we estimate a fixed effect for each region. We then run a fixed-effects logit model that simultaneously tests multiple hypotheses using different explanatory variables. Additionally, we include several measures to control temporally and spatially for the overall level of media activity. First, we include a measure of *strike activity*, which groups the data by weeks and is defined as the count of the number of strikes in the *China Strikes* data occurring in each of the 52 weeks of 2012; we hypothesize that newspaper reporting of strikes might "bandwagon" during periods of high activity. Second, we group observations by province, and assign a group value (common to all observations within a province) to the variable *soc. weight* which measures that province's overall social media traffic as a percentage of all Weibo traffic in the *WeiboScope* data. In essence, we are measuring the level of social media usage and connectivity of a given province, an attribute that proxies for the sophistication level of media consumers in that province.

In this model we also control for the possibility that journalists may be more (less) likely to report incidents with pro-worker outcomes. Three such (non-mutually exclusive) outcomes are 1) if officials intervene on workers' behalf, 2) if workers end up engaging in informal bargaining with management, or 3) if workers receive payments (wages in arrears, unpaid benefits, etc.) they asked for or if the company commits to a desired policy change, such as increased pay for overtime. Model specifications I through III below address these possibilities: Model I controls for a composite measure of all three pro-worker outcomes, Model II considers cases of official intervention, and Model III controls for all three pro-worker outcomes measured individually. Finally, Model IV adds in *violence*. The specification for Model IV, which includes all above-mentioned variables, is below:

$$(1)Y_{itjk} = \beta_o + \beta X + \beta A_t + \beta V_j + U_k + \varepsilon_{itjk}$$

Where Y is the probability of observing newspaper coverage; X is a set of explanatory variables that test our hypotheses, all dummies except for *size* which is categorical; U_k is the macro-regional fixed effect; V_j is the provincial-level measure of media sophistication (*soc. weight*); A_t is the week-level measure of strike activity (*strike activity*); and ε_{itjk} is the individual-level error term. The results, displayed in terms of average marginal effects, are in Table 1:⁴⁴

Table 1: Average Marginal Effects: Newspaper as Outcome Variable

| Explanatory Variable | Model I | Model II | Model III | Model IV |
|------------------------------|---------|----------|-----------|----------|
| <i>Strike activity</i> | -0.009 | -0.015* | -0.009 | -0.020* |
| <i>Soc. Weight</i> | 0.008 | 0.009 | 0.006 | 0.007 |
| <i>Weibo mentions</i> | -0.009 | 0.038 | -0.005 | 0.012 |
| <i>Foreign</i> | -0.097 | -0.054 | -0.108 | -0.021 |
| <i>SOE</i> | -0.113* | -0.087 | -0.105 | -0.079 |
| <i>Size=1</i> | 0.142* | 0.184* | 0.151* | 0.256* |
| <i>Size=2</i> | 0.067 | 0.093 | 0.123* | 0.181* |
| <i>Size=3</i> | -0.106 | -0.083 | -0.077 | -0.033 |
| <i>Transit strike</i> | 0.053 | 0.076 | 0.089 | 0.084 |
| <i>Pro worker</i> | 0.364* | | | |
| <i>Official intervention</i> | | 0.295* | 0.200* | |
| <i>Informal bargaining</i> | | | 0.333* | |
| <i>Payment/policy change</i> | | | 0.238* | |
| <i>Violence</i> | | | | -0.210* |
| Log Likelihood | -128.07 | -141.01 | -122.12 | -151.86 |
| <i>N=281</i> | | | | |

* $p < 0.1$

Notably, *Weibo mentions* is insignificant and near zero in all specifications. This interesting finding immediately challenges our prior assumption (*H1*) that patterns of state media control run in parallel for both newspapers, and social media—if they did, we would expect the presence of *Weibo* traffic to correlate positively with newspaper reporting. Next, *Foreign* and *SOE* are both insignificant except for Model I for *SOE*. Surprisingly, we find no strong evidence to support *H2* or *H3*—at best, the consistent negative signs for *SOE* raise the possibility that SOEs may enjoy limited preferential treatment or at any rate, not be exposed to high levels of strike-related reporting. Nevertheless, our confidence in this possibility is limited—the overall finding is a null one. With regard to *foreign*, journalists seem not to respond (or, be allowed to respond) positively to the nationalist angle of Chinese workers confronting foreign management, despite the prominence of many high-profile disputes with foreign-owned enterprises in our data. Rather than reading into this and the previous *SOE* finding substantively, one conclusion is

simply that great diversity exists across China concerning the types of disputes unfolding in foreign-owned companies, Sino-foreign joint ventures, and SOEs; it is difficult to generalize about the nature or political stakes of these disputes, and this leads to a statistically null result. However, foreign investors' complaint that they are systematically held to a different standard than local firms is disconfirmed.

Strike size, on the other hand, is positive at smaller sizes, significant, and shows a sizeable effect. Increases from small to moderate strike size also have a positive effect on reporting. At yet larger sizes, however, the effect reverses, with coefficients turning negative (although insignificant). It may be that at very large strike sizes, government censors ban reporting of the incident; this would be consistent with media management policies designed to prevent collective action from spreading.⁴⁵ The fact that size positively predicts reporting at smaller sizes, however, suggests a more complex logic consistent with "Control 2.0" (*H4* and *H5*): officials allow some reporting in order to get ahead of an issue and allow the media to pressure management to raise wages and benefits. This finding is further reinforced by the positive and large (.295) marginal effect for *intervention*, which suggests either that officials respond positively to media coverage (or foreknowledge of coming media coverage) by intervening on workers' behalf, or that official intervention, showing as it does the state's commitment to resolving thorny social issues, gets a lot of attention in newspapers (in line with *H6*). Either way, the results do not support the conventional wisdom that strikes are invariably subject to attempts at media suppression. Next, coefficients for *transit strike* are positive.⁴⁶ This, too, is consistent with a "Control 2.0" strategy: transit strikes are hard or impossible to ignore, affecting as they do large portions of urban populations. Although such incidents are often big and the issues raised by cabbies and bus drivers may be "sensitive" because they involve local

government policies—fuel subsidies, transportation monopolies, etc.—authorities likely recognize that suppressing information about transit strikes is impossible. Anyone hailing a cab, for example, can tell if taxis are not running on a particular day. Rather, officials may focus on shaping the content of news reporting, ensuring that they are viewed as problem-solvers. Finally, the coefficient for *violence* is negative and significant—this unambiguously supports *H7* that incidents involving beatings by police and thugs receive less reporting.

Worker Activism and Weibo Discussion

Treating *newspaper* as dependent variable enables us to speak to journalist incentives to report on labor strikes and officials’ incentives to allow such reporting. Taking Weibo posts as the outcome, meanwhile, concerns the degree to which a select, elite group of particularly active and high-profile netizens have incentives to spread such information, and also to the degree to which censors will allow these messages to diffuse. Taking Weibo traffic as binary dependent variable—with a one representing incidents with one or more posts and zero, none—we test the same hypotheses as before. We use the same group-level effects A_t and V_j as before, as well as the macro-region fixed effect U_k .⁴⁷ However, because we now have Weibo as our dependent variable, we must control for weekly fluctuations in Weibo discussions of overall strike activity, including many mentions of the keywords *bagong* and *tinggong* that do not specifically reference an incident in *China Strikes*. We reason that netizens may be paying more attention to labor activity at certain times of the year and that this could affect the post count associated with individual strikes. Thus, we add an additional group-level effect, *Weibo activity* (S_p), to Equation 1 above, where p indexes the total number of strike-related Weibo posts in each week:

$$(2)Y_{itjpk} = \beta_o + \beta X + \beta A_t + \beta V_j + \beta S_p + U_k + \varepsilon_{itjpk}$$

Results are in Table 2 below:

Table 2: Average Marginal Effects: Weibo as Binary Outcome Variable

| Explanatory Variable | Model I | Model II | Model III |
|------------------------------|---------|----------|-----------|
| <i>Weibo Activity</i> | -0.000 | -0.000 | -0.000 |
| <i>Strike Activity</i> | -0.008 | -0.006 | -0.008 |
| <i>Soc. Weight</i> | 0.004 | 0.005 | 0.005 |
| <i>Newspaper</i> | 0.043 | 0.007 | 0.016 |
| <i>SOE</i> | 0.033 | 0.033 | 0.029 |
| <i>Foreign</i> | 0.225* | 0.190* | 0.208* |
| <i>Size=1</i> | 0.027 | 0.026 | 0.016 |
| <i>Size=2</i> | 0.069 | 0.091 | 0.051 |
| <i>Size=3</i> | 0.235* | 0.239* | 0.232* |
| <i>Transit strike</i> | 0.081 | 0.084 | 0.076 |
| <i>Official Intervention</i> | -0.084 | -0.084 | |
| <i>Payment/Policy change</i> | | -0.068 | |
| <i>Informal bargaining</i> | | 0.147* | |
| <i>Violence</i> | | | -0.039 |
| Log Likelihood | -156.58 | -154.50 | -157.28 |
| <i>N=281</i> | | | |

* $p < 0.1$

Several findings are worth noting. The coefficient for *newspaper* is small and insignificant. As with our findings in the previous section, this suggests, at a minimum, that print media and Weibo follow very different patterns regarding coverage of labor disputes: *HI* is disconfirmed, as a strike appearing in Weibo has no predictive value for it appearing in major newspapers, and *vice versa*. *SOE* is also insignificant and close to zero. Similar to results in the previous section, we find no strong evidence that state-owned enterprises receive special attention in Weibo compared with non-SOEs. Again, this null result does not necessarily reflect any particular characteristics of SOEs but rather may speak to the diversity of state-owned companies in China.

With *foreign*, however, results are positive, significant and large across Models I-III. In contrast to newspapers, the Weibo bloggers in our sample seem to latch onto disputes at foreign-

owned or foreign-Sino joint venture companies. The fact that our model controls for dispute size (and therefore, to some extent, company size) and a host of other characteristics increases our confidence that this finding reflects the nationalist sentiment prevalent on Weibo. This sentiment was especially powerful during 2012, as documented elsewhere.⁴⁸ Meanwhile, *size* is insignificant at smaller strike sizes but becomes significant when strikes grow to thousands of workers. This finding likely reflects nothing more than the fact that our sample consists of elite Weibo users, most of whom live in major urban areas and tend to blog about issues of perceived national significance when they write about politics. Yet, it provides an interesting contrast with state newspaper coverage, which tapers off at high levels of protester participation. As with newspapers, *transit strike* has a positive coefficient, though only moderately-sized and insignificant.

Two indicators of pro-worker outcomes—*official intervention* and *payment/policy change*, are insignificant, negative, but not close to zero. Conversely, *informal bargaining* is positive and significant. These results could reflect a disinterest on the part of Weibo users in “positive” stories in which the government neatly resolves problems, coupled with support for worker efforts to “help themselves.” It might also be that unmeasured characteristics lead incidents to both prompt especially detailed reports of the sort likely to capture instances of informal bargaining and to draw micro-bloggers’ attention.⁴⁹ Interestingly, the variable *violence* is negative (although insignificant), suggesting censorship could be playing a role here: talk of fists and bricks simply unnerves authorities. In sum, contrary to our default assumption, Weibo is obviously not subject to the same guidance as official media. Even if violence is still played down (or censored) on social media, the familiar newspaper trope of dutiful officials jumping in to resolve problems does not hold. Read together, the *foreign* result, the one for large strike sizes,

and, to a lesser degree the one for *transit strike*, all suggest that middle- and upper-middle class urban bloggers are sometimes interested in worker struggles—but only large and high-profile incidents, particularly those with a nationalist “hook”, grab their attention. Compared to the picture presented by state coverage, worker activism as refracted through Weibo appears less manageable—more aggrieved and more linked to issues of legitimacy.

Case Study: Lanzhou Bus and Taxi Strikes

On May Day 2012, over 6,000 taxi drivers went on strike for several days in Lanzhou, Gansu province. The drivers' complaints centered on rising compressed natural gas prices but also touched on the high fees they had to pay to their cab companies (approximately 648 USD per month) and competition from illegal “black cabs.” Coming as it did during a national holiday and in a mid-sized city, the strike was bound to be disruptive. Indeed, message boards quickly filled with citizens' questions and complaints. For example, a commentator on the “Gansu Forum” of a car enthusiasts' website posted the question: “I heard that Lanzhou has a taxi strike today??? No wonder it is so hard to hail a cab today...”⁵⁰ Other commentators then chimed in with their own difficulties getting a ride from the train station, etc. A BBS writer at club.autohome.com.cn, meanwhile, expressed support for the cabbies, noting the fact that they were striking on a nominal workers' holiday and proclaiming, “Where there is oppression, there is resistance; where there is exploitation, there is struggle”—and received supportive comments from others.⁵¹ Chat rooms hosted by *Baidu* and other portals also carried discussion. Consistent with our hypotheses and in line with our statistical results, the state media moved to reclaim control of the narrative. *People's Daily* described the strike in detail but dwelt on the local government's efforts to resolve the crisis.⁵² It reported that even before the strike started, a

working group of leaders had been formed by the municipal Communist Party Committee, which in turn gathered together representatives of the Public Security Bureau, Politics and Law Committee, Transportation Bureau, Traffic Police, and others to discuss means of defusing conflict; the result, according to the paper, was a series of proposals aimed at reducing fuel pressures, combating black cabs, and increasing communication with drivers.⁵³ A week after the strike started, a *China Youth Daily* report picked up by *Xinhua* informed readers that the government had promised to open another compressed natural gas station and would establish better channels for communicating with drivers.⁵⁴ The article suggested, though, that thornier issues, especially cab fees and “black cab” competition, remained unresolved.⁵⁵ Other official and semi-official outlets also covered the story, such as *Red.net* and *Henan Shangbao*. Readers were left with the impression that authorities were responsive to the needs of a vulnerable population, even if more work needed to be done. The Lanzhou action received a fair number of mentions (29) in our Weibo dataset. Yet, in a possible indicator of more extensive censorship in the full population of micro-blog posts on the incident, at least two posts were marked as deleted. In sum, rather than trying to erase all discussion of an instance of working class activism that impacted the lives of thousands of ordinary urbanites and therefore could not easily be ignored, the government itself took charge of informing the public about the conflict. However, it made sure to show local officials taking action: holding meetings, putting forward proposals, and enacting reforms, however preliminary. Social media, meanwhile, may have been kept somewhat at bay, keeping the state's narrative the dominant one after an initial flurry of public discussion.

Conclusion

Our findings suggest that the Chinese government has adopted a sophisticated approach to controlling news about and popular discussion of rising labor unrest. Whereas research by

Gary King, Jennifer Pan, and Margaret E. Roberts shows the government's censorship apparatus to be principally directed at any discussion of collective action, we find that with respect to newspapers and within the category of collective action, at least related to worker protests, authorities no longer simply suppress coverage. Instead, they themselves actively report those stories that cannot be ignored—at least up to a certain point: the very largest strikes are hushed, as are violent incidents—and emphasize officials' positive role. Social media is inherently messier to control, though: Weibo users push a more aggrieved and nationalistic narrative. Specifically, micro-bloggers are attracted to incidents that are big, that involve foreign companies, and in which the government plays no helpful role (but in which workers might engage in informal negotiations). But on Sina Weibo, too, in regards to violence against workers, we find hints of what might be official curating of discussion.

This all raises again the perennial question: is the state being pushed to continually cede ground on speech, as cyber optimists claim, or have authorities instead solved the classic “dictator's dilemma” of poor data on discontent⁵⁶ without at the same time allowing incidents to trigger information cascades of citizens revealing their anti-regime preferences?⁵⁷ Our research cannot provide a final answer. However, it is clear that workers in China today, along with potential allies in other classes, face a more complex media environment from that of their forebearers. In getting their story out, they must navigate the contrasting proclivities of state and social media, and although they may reach more people with their tales now than workers could during, say, the Luddite sabotage campaigns of the early English industrial revolution or China's own strike waves during the Mao era, labor's message is inevitably distorted along the way. As China exports its censorship technology to other authoritarian and mixed regimes, such as Iran and Zimbabwe, and as China Central Television is broadcasted globally, offering a propaganda

model to developing world elites, workers elsewhere may face the same challenge of selective media transparency combined with greater control. A careful blend of “Control 2.0” and “Control 1.0” tactics may become the rule not only in China but wherever governments wish to head off a working class challenge—and organizing may consequently have to take new forms.

Appendix A: Measurement in Shenzhen, Guangzhou, Dongguan and Shanghai

Shenzhen, Guangzhou, Dongguan and Shanghai collectively witnessed over a third (111) of all strikes in our 2012 sample. These locations are also more likely to host foreign-owned enterprises and larger companies, and have been the location of a number of high-profile strikes, notably at Taiwanese electronics manufacturer Foxconn’s Shenzhen facility. The cities are therefore both qualitatively and quantitatively important to include. However, due to our method of cross-referencing incidents in the *China Strikes* data and the *WeiboScope* sample—necessary given available resources—these cities also exhibited greater measurement error than other locations, leading us to drop them from the main paper analysis. Here, we report two additional procedures that help confirm that our results hold both with, and without the cities. First, using a random sample of disputes from other locations, we show that minimal measurement error exists outside of these four cities. Second, we attempt to account for the rough magnitude and direction of the measurement error induced by adding these cities in our models.

To confirm that no substantial measurement error affected the results in Tables 1 and 2 of the main paper, we drew a random sample of 16 incidents from among the 393 in the *China Strikes* data, throwing out draws from the above four cities. We then searched the *WeiboScope* dataset, filtering by incident week and city name—this procedure was identical to the one used to code the *Weibo mentions* variable in the main paper. Going beyond the main paper, however, we

then manually read through *all* posts from each of these 16 searches to see if they referenced the relevant strike in *China Strikes*, referred to some other simultaneous strike in the same city, or contained the keyword *bagong* or *tinggong* but did not reference any strike (the latter two cases reflected our concern about measurement error). Our key metric, since *Weibo mentions* was a binary variable coded 1 if an incident contained *any* Weibo posts and 0 otherwise, was whether in our random sample a strike incident for which we got one or more hits for *bagong* or *tinggong* (and thus coded as 1) upon manual reading actually contained at least one relevant post. Using this criterion, we found that 13/16 posts or 81.3% did not contain any measurement error.

Since the above four cities are important sites of labor activity in China, we compare results obtained by including as well as excluding them. If we obtain the same results either way, we can have greater confidence that the results were not due to measurement error (which is less when excluding the cities) or due to lack of variation (since the cities contain a large number and wide variety of strike incidents). Appendix Tables 1 and 2 below again show results with *newspaper*, and *Weibo mentions*:

Appendix Table 1: Average Marginal Effects: *Newspaper* as Outcome Variable

| Explanatory Variable | Model I | Model II | Model III | Model IV |
|------------------------------|---------|----------|-----------|----------|
| <i>Strike activity</i> | -0.006 | -0.012* | -0.008 | -0.015* |
| <i>Soc. weight</i> | 0.007 | 0.007 | 0.006 | 0.004 |
| <i>Weibo mentions</i> | -0.024 | 0.029 | 0.007 | -0.001 |
| <i>Foreign</i> | -0.050 | -0.032 | -0.053 | 0.012 |
| <i>SOE</i> | -0.071 | -0.053 | -0.056 | -0.044 |
| <i>Size=1</i> | 0.131* | 0.153* | 0.132* | 0.272* |
| <i>Size=2</i> | 0.041 | 0.072 | 0.071 | 0.195* |
| <i>Size=3</i> | -0.138* | -0.114 | -0.120 | -0.048 |
| <i>Transit strike</i> | 0.100* | 0.122* | 0.126* | 0.137* |
| <i>Pro worker</i> | 0.374* | | | |
| <i>Official intervention</i> | | 0.352* | 0.280* | |
| <i>Informal bargaining</i> | | | 0.255* | |
| <i>Payment/policy change</i> | | | 0.154* | |
| <i>Violence</i> | | | | -0.165* |
| Log Likelihood | -181.12 | -197.45 | -180.10 | -222.91 |
| <i>N=392</i> | | | | |

* $p < 0.1$

**Appendix Table 2: Average Marginal Effects:
Weibo as Binary Outcome Variable**

| Explanatory Variable | Model I | Model II | Model III |
|------------------------------|---------|----------|-----------|
| <i>Weibo activity</i> | -0.000 | -0.000 | -0.000 |
| <i>Strike activity</i> | -0.004 | -0.003 | -0.005 |
| <i>Soc. weight</i> | 0.014* | 0.014* | 0.015* |
| <i>Newspaper</i> | 0.023 | 0.007 | -0.004 |
| <i>SOE</i> | 0.005 | 0.007 | 0.002 |
| <i>Foreign</i> | 0.126* | 0.103* | 0.116* |
| <i>Size=1</i> | 0.073 | 0.070 | 0.065 |
| <i>Size=2</i> | 0.047 | 0.057 | 0.038 |
| <i>Size=3</i> | 0.145* | 0.146* | 0.152* |
| <i>Transit strike</i> | 0.025 | 0.025 | 0.019 |
| <i>Official intervention</i> | -0.061 | -0.056 | |
| <i>Payment/policy change</i> | | -0.057 | |
| <i>Informal bargaining</i> | | 0.085 | |
| <i>Violence</i> | | | -0.101 |
| Log Likelihood | -209.24 | -207.79 | -209.17 |
| <i>N=392</i> | | | |

* $p < 0.1$

With regard to the first model, our results for strike *size*, various pro-worker outcomes, and *violence* remain essentially unchanged—they remain significant, and while coefficient magnitudes do change, they neither inflate greatly nor decrease to near zero. The one exception is that results for *transit strike* are now significant (though previously they were positive in sign) and do not change drastically in magnitude.

Turning to the second model, the results, again, are not drastically different. *Foreign* is still positive and significant, as is *size* at the highest level. *Informal bargaining* remains positive, but loses significance. Finally, we note that the log likelihoods decrease across both appendix models compared to the main paper—this suggests that adding in the above four cities does increase the noise-to-signal ratio.

¹ Dietrich Rueschemeyer, Evelyne Stephens, and John Stephens, *Capitalist Development and Democracy* (Chicago: University of Chicago Press, 1992).

² Samuel Huntington, *Political Order in Changing Societies* (New Haven: Yale University Press, 1968).

³ Barrington Moore, *Social Origins of Dictatorship and Democracy* (Boston: Beacon Press, 1966); Gregory M Luebbert, “Social Foundations of Political Order in Interwar Europe,” *World Politics* 39, no. 4 (1987), 449–78.

⁴ Ruth Berins Collier and David Collier, *Shaping the Political Arena: Critical Junctures, the Labor Movement, and Regime Dynamics in Latin America* (Princeton: Princeton University Press, 1991).

⁵ David Ost, *Solidarity and the Politics of Anti-Politics* (Philadelphia: Temple University Press, 1990); Stephen Crowley, *Hot Coal, Cold Steel: Russian and Ukrainian Workers from the End of*

the Soviet Union to the Post-Communist Transformations (Ann Arbor, MI: University of Michigan Press, 1997).

⁶ Dan Slater, *Ordering Power: Contentious Politics and Authoritarian Leviathans in Southeast Asia* (New York: Cambridge University Press, 2010).

⁷ E.g., David Bandurski, “Taxi Strikes in China Highlight Changing Press Controls,” China Media Project, 2008, <http://cmp.hku.hk/2008/11/12/1344/>.

⁸ See Gary King, Jennifer Pan, and Margaret E. Roberts, “How Censorship in China Allows Government Criticism but Silences Collective Expression,” *American Political Science Review* 107, no. 02 (May 01, 2013), 326–43.

⁹ Elizabeth J Perry, *Shanghai on Strike: The Politics of Chinese Labor* (Stanford: Stanford University Press, 1993).

¹⁰ P.R.C. Department of Population and Employment Statistics, *China Labour Statistical Yearbook 2011* (Beijing: China Statistics Press, 2011).

¹¹ Andrew Wedeman, “Enemies of the State: Mass Incidents and Subversion in China,” in *APSA Meeting Paper* (Toronto, 2009), <http://ssrn.com/abstract=1451828>; Chih-Jou Jay Chen, “Growing Social Unrest and Emergent Protest Groups in China,” in *Rise of China: Beijing’s Strategies and Implications for the Asia-Pacific*, ed. Hsin Huang Michael Hsiao and Cheng-Yi Lin (New York and London: Routledge, 2009), 87–106; Murray Scott Tanner, “China Rethinks Unrest,” *The Washington Quarterly* 27, no. 3 (2004), 137–56.

¹² Lee, *Against the Law: Labor Protests in China’s Rustbelt and Sunbelt*.

¹³ William Hurst, *The Chinese Worker After Socialism* (New York: Cambridge University Press, 2009).

¹⁴ Ching Kwan Lee, *Gender and the South China Miracle: Two Worlds of Factory Women* (Berkeley: University of California Press, 1998).

¹⁵ Mary E. Gallagher, *Contagious Capitalism: Globalization and the Politics of Labor in China* (Princeton: Princeton University Press, 2005).

¹⁶ Teresa Wright, *Accepting Authoritarianism: State-Society Relations in China's Reform Era* (Stanford: Stanford University Press, 2010).

¹⁷ William Hurst and Kevin O'Brien, "China's Contentious Pensioners," *The China Quarterly* 170 (2002): 345–60; Yongshun Cai, "The Resistance of Chinese Laid-off Workers in the Reform Period," *The China Quarterly* 170 (2002), 327–44; Jianrong Yu, *Kangzhengxing Zhengzhi: Zhongguo Zhengzhi Shehuixue Jiben Wenti* [Contentious Politics: Fundamental Issues in Chinese Political Sociology] (Beijing: People's Publishing House, 2010), 132-154; Feng Chen, "Subsistence Crises, Managerial Corruption and Labour Protests in China," *The China Journal* 44 (2000), 41–63; Lee 2007.

¹⁸ Chris King-chi Chan, *The Challenge of Labour in China: Strikes and the Changing Labour Regime in Global Factories* (Abingdon: Routledge, 2010); Manfred Elfstrom and Sarosh Kuruvilla, "The Changing Nature of Labor Unrest in China," *Industrial and Labor Relations Review* 67, no. 2 (2014), 453–80.

¹⁹ Geremie R. Barmé, "The Five Vermin Threatening China," *The China Story*, 2012, <http://www.thechinastory.org/2012/11/the-five-vermin-五蠹-threatening-china/>.

²⁰ Eli Friedman, *Insurgency Trap: Labor Politics in Postsocialist China* (Ithaca: Cornell University Press, 2014); Chris King-Chi Chan and Elaine Sio-Ieng Hui, "The Development of Collective Bargaining in China: From 'Collective Bargaining by Riot' to 'Party State-Led Wage Bargaining,'" *The China Quarterly* 217 (December 05, 2013), 221–42.

²¹ Daniel C. Lynch, *After the Propaganda State: Media, Politics, and 'Thought Work' in Reformed China* (Stanford: Stanford University Press, 1999); Daniela Stockmann, *Media Commercialization and Authoritarian Rule in China* (New York: Cambridge University Press, 2013).

²² Daniela Stockmann and Mary E Gallagher, “Remote Control: How the Media Sustain Authoritarian Rule in China,” *Comparative Political Studies* 44, no. 4 (2011), 436–67.

²³ Yanqi Tong and Shaohua Lei, “War of Position and Microblogging in China,” *Journal of Contemporary China* 22, no. 80 (2012), 292-311.

²⁴ Guobin Yang, *The Power of the Internet in China: Citizen Activism Online* (New York: Columbia University Press, 2009).

²⁵ Xiao Qiang, “The Rise of Online Public Opinion and Its Political Impact,” in Susan Shirk, ed., *Changing Media, Changing China* (New York: Oxford University Press, 2011), 202-24.

²⁶ Tao Zhu, David Phipps, Adam Pridgen, Jedidiah R. Crandall, and Dan S. Wallach, “The Velocity of Censorship: High-Fidelity Detection of Microblog Post Deletions” (2013), available at <http://arxiv.org/abs/1303.0597>. See also Yuen Yuen Ang, “Authoritarian Restraints on Online Activism Revisited: Why 'I-Paid-A-Bribe' Worked in India but Failed in China,” *Comparative Politics* 47, no. 1 (2013), 21-40.

²⁷ King, Pan, and Roberts, p. 326.

²⁸ Rebecca MacKinnon, *Consent of the Networked: The Worldwide Struggle for Internet Freedom* (New York: Basic Books, 2012); Albert O.Hirschman, *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States* (Cambridge, MA: Harvard University Press, 1970).

²⁹ Bandurski.

³⁰ Margaret Roberts, "Fear or Friction? How Censorship Slows the Spread of Information in the Digital Age," prepared for the Association of Asian Studies Annual Meeting, Philadelphia, PA, March 2014.

³¹ See Kevin J. O'Brien and Lianjiang Li, *Rightful Resistance in Rural China* (New York: Cambridge University Press, 2006).

³² See Peter Hays Gries, *China's New Nationalism: Pride, Politics, and Diplomacy* (Berkeley: University of California Press, 2004); Zheng Wang, *Never Forget National Humiliation: Historical Memory in Chinese Politics and Foreign Relations* (New York: Columbia University Press, 2012); Jessica Chen Weiss, *Powerful Patriots: Nationalist Protest in China's Foreign Relations* (New York: Oxford University Press, 2014).

³³ Michael A. Santoro, *China 2020: How Western Business Can--and Should--Influence Social and Political Change in the Coming Decade* (Ithaca: Cornell University Press, 2009), 134.

³⁴ Barry Naughton, *The Chinese Economy: Transitions and Growth* (Cambridge: MIT Press, 2007), 298-308.

³⁵ Accessible at <https://chinastrikes.crowdmap.com>

³⁶ But see "Shekeyuan: 14 Nianjian Bai Ren Yishang Qunti Shijian 871 Qi Guangdong Jushou [Academy of social sciences: over 871 collective incidents in 14 years with Guangdong the most]," *The Beijing News* via *Xinhua Net* (2014), http://news.xinhuanet.com/youqing/2014-02/25/c_126185554.htm (March 2, 2014).

³⁷ Accessible at <http://www.numble.com/PHP/mysql/clbmap.html>

³⁸ Elfstrom and Kuruvilla.

³⁹ Stockmann (2013) adopts a stricter definition based on the degree to which a publication relies on the market for revenue.

⁴⁰ King Wa Fu, Chung-Hong Chan, and Michael Chau, “Assessing Censorship on Microblogs in China: Discriminatory Keyword Analysis and the Real-Name Registration Policy,” *IEEE Internet Computing*, 17, no. 3 (2013), 42-50.

⁴¹ Where incidents occurred outside any city’s jurisdiction (including incorporated rural lands), we used the county name instead; where both city and county names were unavailable, we used the province name—this occurred in only a couple cases.

⁴² G. William Skinner, ed., *The City in Late Imperial China* (Stanford: Stanford University Press, 1977); G. W. Skinner, Mark Henderson, and Merrick Lex Berman, “Socioeconomic Macroregions of China” (Cambridge, MA: Fairbank Center for Chinese Studies, 2013), <http://hdl.handle.net/1902.1/21766>.

⁴³ One incident, in Tibet, did not fall under any of Skinner’s “macro-regions” and thus dropped out of the regressions.

⁴⁴ The results shown here represent the coefficients obtained after we first dropped all observations from four cities – Shenzhen, Guangzhou, Dongguan and Shanghai – that exhibited high levels of measurement error. The specific issue was that we were only able to link *Weibo* traffic in a given city to strike incidents in the *China Strikes* dataset by checking whether any mentions of “strike” or “work stoppage” (*bagong* and *tinggong*, respectively) surfaced in our *Weibo* dataset during a given city and week, and that therefore corresponded to the date and location of a *China Strikes* record. Without much more extensive keyword searches and manual reading of *Weibo* posts for each of the 393 incidents (a task beyond the scope of this paper), we had no way to confirm absolutely that the incidence of *bagong* and *tinggong* for a given city and week was actually referring to a corresponding *China Strikes* event. However, as we found by drawing a random sample of 16 posts, in most (80%) of cases, our simple coding strategy of

counting keywords by city and week in fact did identify *China Strikes* incidents. The major exception to this was in major urban areas (specifically, the four above) which a) showed high levels of *Weibo* traffic overall, and b) witnessed multiple strikes in the same week, thus making it difficult to isolate what fraction of a given keyword count applied to each particular incident. We are aware that dropping many key observations (a total of 111) from these cities, which are epicenters of labor activity – might negatively affect the magnitude, statistical significance and most importantly, generalizability of our findings. Fortunately, this turns out not to be the case: adding these cities back in slightly alters our results but does not negate them. Thus, while we are unable to measure the four cities precisely, we see no evidence that conditions there differ substantially from the rest of the country. Appendix A details our random sample procedure to verify the lack of noise in observations other than these cities, and reports regressions with the cities added in.

⁴⁵ King, Pan, and Roberts.

⁴⁶ Coefficients reported here are insignificant, but this is likely due to dropping the 111 observations from Shenzhen, Guangzhou, Dongguan and Shanghai. When we add these cities in, coefficients in models I-IV are significant at $\alpha=0.1$ for Model I and $\alpha=0.05$ for Models II-IV.

⁴⁷ As in the previous section, we again drop the 111 observations from Shenzhen, Guangzhou, Dongguan and Shanghai. As a robustness check, in Appendix A we do add observations from these four cities back in and coefficients, except for *Transit strike*, do not change much.

⁴⁸ Christopher Cairns and Allen Carlson. “Real World Islands in a Social Media Sea: Nationalism and Censorship on Weibo during the 2012 Diaoyu/Senkaku Crisis.” Presented at the American Political Science Association 2014 Annual Meeting, Washington, D.C.

⁴⁹ As noted, the *China Strikes* dataset may be especially subject to bias concerning the amount of information available on different protests.

⁵⁰ “*Tingshuo lanzhou jintian you chuzuche bagong??? Nanguai jintian che zheme nanda* [Heard that Lanzhou has a taxi strike today??? No wonder today it is so hard to hail a cab],” *Xcar.com.cn* (2012), <http://www.xcar.com.cn/bbs/viewthread.php?tid=17391066>

⁵¹ “*Lanzhou chuzuche wuyi bagong, zhi wei jiaqi nan?* [Lanzhou's May Day taxi driver strike, only because of fuel difficulties?],” *Club.autohome.com.cn* (2012), <http://club.autohome.com.cn/bbs/thread-a-100005-15270940-1.html>.

⁵² “*Wuyi qijian lanzhou bufen chuzuche yin jiaqinan deng wenti bayun* [During May Day, some Lanzhou taxi drivers strike because of fuel difficulties and other problems],” *People's Daily* accessed via *Jititanpan Wang* (2012), <http://www.jtpt.cn/a/report/news/labor/2012/0502/2881.html>.

⁵³ *Ibid.*

⁵⁴ “*Lanzhou chuzuche tingyun shijian diaocha: fenziqian mei yue jiao 412 yuan* [An investigation of the Lanzhou taxi work stoppage incident: 4125 yuan in cab fees must be paid every month],” *China Youth Daily* accessed via *Xinhua* (2012), http://news.xinhuanet.com/local/2012-05/14/c_111941230_2.htm.

⁵⁵ *Ibid.*

⁵⁶ Ronald Wintrobe, *The Political Economy of Dictatorship* (New York: Cambridge University Press, 1998).

⁵⁷ Timur Kuran, “Now Out of Never: The Element of Surprise in the East European Revolution of 1989,” *World Politics* 44, no. 1 (1991), 7–48.