

# When Autocracies Threaten Citizens With Violence: Evidence from China<sup>\*</sup>

Brett L. Carter<sup>†</sup>      Erin Baggott Carter<sup>‡</sup>

March 26, 2018

## Abstract

When do autocrats threaten their citizens with violence? Are these threats effective? We propose a theory of propaganda-based threats in autocracies that rests on insights from experimental psychology. Our basic insight is that even credible threats of violence are costly, and so they should be reserved for moments when mass protests are most likely. Since threats of violence are employed sparingly, we also expect them to be effective. We test the theory with data from China, the world's most populous autocracy. We employ a series of computational techniques to collect all 164,707 articles published between 2009 and 2016 in the *Workers' Daily*, China's state-run newspaper which focuses on domestic issues and targets a non-elite audience. We identify every explicit threat of violence it issued. We show that propaganda-based threats of violence are deployed overwhelmingly around the anniversaries of ethnic separatist movements in Tibet and Xinjiang provinces, and are tightly correlated with pro-regime propaganda. We then employ an instrumental variables estimation strategy to show that propaganda-based threats render protests less likely.

Word Count: 10,864

---

<sup>\*</sup>We thank an anonymous research assistant for excellent work and Victor Shih for helpful feedback. Authors appear alphabetically.

<sup>†</sup>Assistant Professor, School of International Relations, University of Southern California. [blcarter@usc.edu](mailto:blcarter@usc.edu).

<sup>‡</sup>Assistant Professor, School of International Relations, University of Southern California. [baggott@usc.edu](mailto:baggott@usc.edu).

# 1 Introduction

The dynamics of autocratic politics have undergone important shifts since the end of the Cold War. As the rate of elite coups has declined (Goemans and Marinov 2014), mass protests have emerged as a chief threat to the world’s autocrats. As they have, scholars have sought to understand their dynamics: who protests (Rosenfeld 2017), when (Carter and Carter 2017a), how they organize (Howard and Hussain 2013), and which tactics are most effective (Chenoweth and Stephan 2011; Beber, Roessler and Scacco 2014; Enos, Kaufman and Sands 2017). Scholars have also sought to understand how the world’s autocrats attempt to inoculate themselves: when autocrats incarcerate dissidents (Truex 2016), under what circumstances repression is circumscribed by external constraints (Aronow, Carnegie and Marinov 2012), and how autocrats employ censorship and propaganda to manipulate their citizens’ informational environments (King, Pan and Roberts 2013, 2017; Carter and Carter 2017b).

In this paper, we ask two questions: When do autocrats use their propaganda apparatuses to threaten their citizens with violence? Are these threats effective? We develop a theory of propaganda-based threats of violence that rests on insights from experimental psychology. We argue that even credible threats of violence are costly. For propaganda apparatuses that aim to persuade citizens of the regime’s merits – and hence invest in acquiring some amount of credibility – threats of violence undermine those efforts. For propaganda apparatuses that aim to signal to citizens the regime’s repressive capacity – for instance, by broadcasting absurdly positive propaganda that everyone knows to be false – the force of a threat may diminish in how often it is issued by the regime. Threats of violence may also give special salience to moments that are already politically sensitive: moments when citizens are inclined to protest a regime’s past crimes. For all these reasons, we expect propaganda-based threats of violence to be used sparingly, at moments when the regime is most concerned about popular protests. Since we expect threats to be deployed to maximize their effects on citizen behavior, we also expect threats to be effective.

To answer these questions, we focus on China, the world’s most populous autocracy. We employ a series of computational techniques to collect all 164,707 articles published in the *Workers’ Daily* (工人日报), China’s state-run newspaper that focuses on domestic issues and targets a non-elite audience, between 2009 and 2016. We then identified every explicit threat of violence against citizens. We find that explicit threats are overwhelmingly associated with the anniversaries of ethnic separatist movements in Tibet and Xinjiang provinces: the Tibetan “Liberation,” Tibetan Rebellion, Serf Emancipation Day, and Xinjiang Uprising. This is consistent with our theoretical expectations, but also adds nuance. Explicit threats of violence appear to be targeted disproportionately at political out-groups: ethnic or regional identity groups that are excluded from a regime’s core constituency. Huang (2015) suggests that the Chinese government’s pro-regime propaganda is itself interpreted as threatening by citizens. Consistent with this, we find that explicit threats of violence and spikes in pro-regime propaganda are strongly correlated.

We employ a novel identification strategy that plausibly identifies the causal effect of propaganda-based threats on the probability of popular protest. Our identification strategy rests on two features of contemporary Chinese politics. First, regime propaganda is set at the national level, but must occasionally respond to local conditions, such as the anniversaries of ethnic separatist movements in Tibet and Xinjiang. As a result, provinces that are geographically and culturally distant from each other are occasionally “treated” with propaganda content that responds to conditions in the other’s locale. Second, China’s ethnic diversity and sprawling size renders these ethnic separatist anniversaries salient only in certain provinces: Tibet, Xinjiang, and a handful of others. We argue that these ethnic separatist anniversaries in Tibet and Xinjiang plausibly condition the rate of protest in geographically and culturally distant provinces through the propaganda-based threats that the regime issues in anticipation of them. We find that propaganda-based threats have a plausibly causal effect on nationwide protest levels. By doubling the number of references to “stability” or “harmony,” China’s propaganda apparatus halves the number of protests over the course of the subsequent week.

This identification strategy rests on an exclusion restriction: the dates of these ethnic separatist anniversaries, salient in some province  $i$ , should condition the rate of protest in geographically and culturally distant province  $j$  only through their effects on propaganda-based threats. We provide a range of evidence that this is indeed the case. First, we drop from our analysis both Tibet and Xinjiang provinces, as well as seven other provinces that are geographically contiguous, contain substantial numbers of co-ethnics, or contain substantial numbers of other ethnic minorities. After dropping these nine provinces, we are left with a sample that nonetheless includes 88.5% of Chinese citizens. Second, we present preliminary results from a pilot survey that shows the dates of these ethnic separatist anniversaries are generally unknown outside Tibet and Xinjiang, which suggests that they lack political salience elsewhere. Third, we show that the government’s repressive apparatus does not repress protests that emerge outside Tibet and Xinjiang during these separatist anniversaries *at higher rates* than other days, which suggests that the security apparatus is not on high alert in geographically and culturally distant provinces on these anniversaries.

This paper proceeds as follows. Section 2 surveys recent scholarship on autocratic propaganda, the psychological effects of threats, and our theoretical framework for understanding propaganda-based threats of violence. Section 3 introduces our corpus of *Workers’ Daily* articles and describes our measure of propaganda. Section 4 shows that propaganda-based threats are driven overwhelmingly by the anniversaries of ethnic separatist movements in Tibet and Xinjiang. Section 5 provides evidence that explicit threats of violence in the *Workers’ Daily* have reduced protest rates in geographically and culturally distant provinces. Section 6 concludes with suggestions for future research. The online appendix provides additional information about our measure of propaganda-based threats and a series of validation and robustness checks.

## 2 Theoretical Framework

### 2.1 Models of Autocratic Propaganda

Formal theorists have provided a range of insights into how the world’s autocrats employ propaganda. The canonical models entail a set of citizens and an autocrat, who can employ some combination of repression, propaganda, and financial transfers to discourage mass protests and ultimately retain power.<sup>1</sup>

Citizens are generally uncertain about something: whether the autocrat implements sound public policy (Guriev and Treisman 2015; Egorov, Guriev and Sonin 2009; Gehlbach and Sonin 2014; Chen and Xu 2015), or whether the autocrat has a substantial repressive apparatus (Edmond 2013; Huang 2015). In some cases, citizens are also uncertain about the beliefs of their neighbors: about whether their neighbors are prepared to revolt (Little 2015). Propaganda is generally modeled as a public signal, which is designed to somehow manipulate these beliefs and, ultimately, discourage citizens from attempting to force the autocrat from power, either at the ballot box or through protests.

Carter and Carter (2017*b*) refer to these competing approaches as “propaganda as persuasion” and “propaganda as domination.” The former aims to persuade citizens of the regime’s merits; the latter to signal its capacity to repress them. These two approaches have long historical antecedents. Joseph Goebbels, architect of Nazi Germany’s propaganda apparatus, wrote that “propaganda becomes ineffective the moment we are aware of it” (Doob 1950). As George Orwell wrote, forcing citizens to consume propaganda that everyone knows to be false serves to dominate them, and to broadcast the regime’s capacity to dominate them to others. In turn, the regime’s capacity for repression becomes common knowledge. This insight underlies Wedeen (1999)’s account of Syria under Hafez al-Assad and Huang (2015)’s work in contemporary China.

Empirically, Carter and Carter (2017*b*) show that the extent to which an autocrat is bound by nominally democratic political institutions determines which propaganda strategy he employs, and that these competing strategies have profound implications for the substance of propaganda itself. When autocrats employ propaganda to signal their “strength” or repressive capacity, as in Huang (2015) and Edmond (2013), its coverage of the regime is absurdly positive. Indeed, its absurdity gives it power. Again, by forcing citizens to consume information that everyone knows to be false – and to be seen doing so by their neighbors – autocrats broadcast their capacity for repression, creating common knowledge among citizens. When autocrats employ propaganda to persuade readers of the regime’s merits, as in Egorov, Guriev and Sonin (2009), Gehlbach and Sonin (2014), Chen and Xu (2015), Guriev and Treisman (2015), and Little (2015), autocrats craft propaganda to be plausible. They do so either by explicitly mixing fact with fiction, or otherwise investing in the credibility of the propaganda apparatus.

---

<sup>1</sup>There is also a relatively large literature on censorship. See, for instance, Egorov, Guriev and Sonin (2009), Lorentzen (2014), and Shadmehr and Bernhardt (2015).

These competing approaches to propaganda have different implications for how we understand propaganda-based threats. When propaganda is designed to dominate citizens – and to signal their domination to other citizens – absurdly positive coverage of the regime itself constitutes the threat. When propaganda is designed to persuade citizens of the regime’s merits, propaganda is not threatening at all. It strives for credibility by cultivating the appearance of genuine news, mixing verifiable facts with useful fictions.

## 2.2 Mechanisms

We define a threat as some explicit claim by the government about the likelihood that anti-regime behavior will be met with violence. In that respect, it is analytically distinct from the sort of absurd pro-regime propaganda that serves to dominate citizens. We propose three mechanisms that should render explicit propaganda-based threats uncommon.

### 2.2.1 Propaganda as Persuasion: Source Derogation

First, insofar as pro-regime propaganda is meant to persuade citizens of the regime’s merits, as in models by Gehlbach and Sonin (2014), Guriev and Treisman (2015), Egorov, Guriev and Sonin (2009), we should expect threats to be rare. The reason is that publishing threats alongside content that is designed to persuade citizens should invalidate the persuasive content. Put simply, a regime that threatens citizens with violence for asserting basic rights does not privilege their interests. This mechanism is fully consistent with Bayesian rationality. If citizens are strictly rational, then, in response to threats, citizens should update their beliefs about the government’s merits.

This mechanism is widely accepted among psychologists, who have long sought to understand the range of factors that compel individuals to resist persuasion.<sup>2</sup> Many of these accounts rest on the theory of psychological reactance, which suggests that humans have an innate desire for autonomy and independence. When that freedom is somehow threatened or eliminated, they are compelled to maintain – or to restore – the opinion or behavior that was threatened (Brehm and Brehm 1981; Burgoon et al. 2002; Rains 2013). Psychologists refer to these as “threats to freedom,” which are defined broadly. These threats may be mild: messages that are clearly designed to limit an individual’s autonomy to craft and maintain their own beliefs. These threats may also be more strident: messages that use the imperative form or otherwise leave no option for refusal (Kronrod, Grinstein and Wathieu 2012).

Messages that constitute threats to freedom routinely compel individuals to contest a message’s *source*. This is commonly referred to as “source derogation.” Smith (1977) documents source derogation along three dimensions: objectiveness, expertness, and trustworthiness. Put simply, when exposed to threatening information, individuals evaluate the source of the message “as someone less expert, as less objective, and as less trustworthy.”

---

<sup>2</sup>For one prominent account about when citizens are open to persuasion, see Cialdini (2006).

The source derogation mechanism has clear implications for propaganda-based threats in autocracies. Gehlbach and Sonin (2014) and Carter and Carter (2017*b*) make clear that building credibility with readers is difficult: It requires conceding regime malfeasance or policy errors. Once acquired, this reputation for credibility can be spent on persuading citizens of “useful fictions.” Put simply, obtaining the capacity to persuade is costly insofar as it requires conceding damaging information. Autocratic propaganda apparatuses that seek to persuade should be reticent to suffer source derogation, which would squander their credibility.

When propaganda aims to persuade, threats should be rare.

### **2.2.2 Propaganda as Persuasion or Domination: The Boomerang Effect**

Messages that constitute threats to an individual’s freedom may elicit a second response: a “boomerang effect,” in which individuals contest the message’s *content*. Individuals do so by engaging in less (more) of the encouraged (discouraged) behavior (Baek, Yoon and Kim 2015). As a result, messages that include threats to freedom may not only prove ineffective, but also lead to the opposite of the desired results (Clee and Wicklund 1980; Ringold 2002).

The boomerang effect also has clear implications for understanding propaganda-based threats in autocracies. If citizens perceive the threat as a warning against mass protests in response to an event or grievance, then the boomerang effect suggests that the threat itself could endow that event or grievance with some additional or special salience. In turn, the explicit threat may help catalyze the episode of collective action that it was meant to discourage. This, indeed, is one explanation for why Chinese censorship emboldens China’s online activists to engage in even more dissent than they would otherwise (Roberts 2018). Likewise, citizens who experienced censorship became more likely to purchase censorship evasion software and to consume anti-regime content online (Hobbs and Roberts 2018).

### **2.2.3 Propaganda as Domination: Desensitization**

Even if absurdly positive propaganda makes the regime’s repressive capacity common knowledge among citizens, there is no *ex ante* reason to believe that propaganda apparatuses employ pro-regime propaganda and explicit threats in the same way, at the same time, or to the same effect. Indeed, when pro-regime propaganda is explicitly designed to threaten citizens, as is the case for Huang (2015), it is possible that regimes may wish to employ explicit threats of violence as well. The reason is that the force of absurdly positive propaganda may diminish over time. In turn, explicit threats of violence may provide a useful and vivid reminder of the regime’s capacity for violence.

Psychologists refer to this mechanism as *desensitization*, defined by Carnagey, Anderson and Bushman (2007) as “a reduction in emotion-related physiological reactivity to real violence.” Desensitization has been widely documented by experimental psychologists. The canonical experiment

entails exposing treatment and control groups to some manner of violence via some media platform, and then measuring the physiological responses of both groups to some subsequent exposure to violent stimulus. To ensure that the desensitization mechanism is robust, psychologists have varied a range of features in the canonical experiment: the medium of exposure to violence (television, movies, and video games), the kind of violence to which subjects are exposed and desensitized (fictitious or real-life), age groups (children or adult), and the measurement of the physiological response (heart rate or galvanic skin response). The results are clear. Exposure to violence in the past renders people less sensitive to violence in the future.<sup>3</sup> The link between repeated exposure to violence and subsequent desensitization is so widely accepted that it now constitutes part of military training (Grossman and Degaetano 1999).

We are unaware of any experimental studies that test whether individuals become desensitized to *threats of violence* in the same way as to *actual violence*. One analog, however, may be how citizens in Western democracies have responded to threats of terrorism. Although this literature is small, scholars have observed important variation. Waxman (2011)'s central argument is that, in response to *sustained* terrorist attacks during the Palestinian Intifada, Israeli citizens

basically became accustomed to terrorism and adapted accordingly. ...The effects of ongoing, chronic terrorism may significantly differ from the effects of a one-off terrorist attack. ...A society can gradually grow accustomed to chronic terrorism, and consequently, its impact declines. In short, societies can effectively become habituated to terrorism and learn to cope with it.

For Waxman (2011), American citizens were so affected by the September 11, 2001, attacks precisely because they had no precedence. Waxman (2011)'s argument was anticipated by Wardlaw (1989), who wrote that repeated terrorist attacks would ultimately result in a society's indifference to them.

This observation is increasingly common in the popular press, motivated by a sense that ISIS terrorist attacks are no longer as unsettling to citizens in Western countries as they were. In *Slate*, two political scientists dubbed it "terrorism fatigue": "ISIS attacks," they wrote, are losing their ability to terrify" (Amarasingam and Clarke 2017). *USA Today* put it succinctly: "What does it mean if terrorism no longer terrorizes us?"<sup>4</sup> *Vox* declared, simply: "We're becoming numb to terrorism. Terrorists are getting worse at the one thing they seek to do: terrify people."<sup>5</sup>

### 2.3 Hypotheses

Each of these mechanisms suggests that propaganda-based threats of violence should be employed sparingly, during moments of genuine crisis, when popular protests are most dangerous for the

---

<sup>3</sup>The literature on desensitization is large. See, among others, Lazarus et al. (1962), Linz, Donnerstein and Adams (1989), Thomas et al. (1977), Thomas (1982), Cline, Croft and Courier (1973), Bartholow, Bushman and Sestir (2006), and Carnagey, Anderson and Bushman (2007).

<sup>4</sup><https://www.usatoday.com/story/news/2017/06/05/london-terror-we-becoming-apatetic-resilient/369387001/>

<sup>5</sup><https://www.vox.com/world/2017/11/2/16591706/terror-attack-new-york-city-terrorism>

regime. The source derogation effect suggests that, insofar as propaganda aims to persuade, threats invalidate its capacity to do so. The desensitization effect suggests that repeated threats over time diminish their force. The boomerang effect suggests that threats may endow politically salient grievances or moments with even more significance to citizens. For all these reasons, threats of violence should be reserved for when autocrats most need them: when popular protests most profoundly threaten the autocrat’s survival. This constitutes our first hypothesis, and it focuses on *temporal* variation within countries.

Our second and third hypotheses focus on cross-country variation in how often threats are employed. The source derogation effect suggests the rate of propaganda-based threats should *decrease* when autocrats employ their propaganda apparatuses primarily to *persuade* their citizens. The reason, from above, is that explicit threats of violence against citizens undermine the regime’s claims of serving their citizens’ interests. Accordingly, in a setting like China, where the purpose of propaganda is to signal the regime’s capacity to dominate its citizens, we expect the rate of threats to be higher than autocracies that use propaganda to persuade. Put differently, the protest likelihood threshold for issuing a threat should be lower in autocracies, like China, that use propaganda primarily to dominate rather than to persuade. Third, and as a corollary, we also expect propaganda-based threats to be more explicit – and less veiled – in autocracies where propaganda serves to dominate rather than persuade. The reason, again, is that as the explicitness of a threat increases, the propaganda apparatus further undermines its ability to persuade citizens of the regime’s merits.

Finally, since propaganda-based threats should be employed sparingly, we expect them to generally be effective. That is, we expect citizens to generally believe that threats of violence are credible. To be sure, whether propaganda-based threats reduce the likelihood of protest depends on a range of factors: how frustrated citizens are, how mobilized they are, how vulnerable they believe the regime is. But if autocrats reserve threats of violence for moments of genuine crisis, then citizens should know this, and it should have some effect on their behavior.

In this paper, for reasons of data availability, we restrict attention to the first and fourth hypotheses: that explicit threats of violence should be driven by a society’s calendar of collective action, and that threats of violence should render protest less likely.

## 3 Data

### 3.1 Measuring Propaganda-Based Threats

To measure propaganda-based threats, we combine our substantive knowledge of Chinese politics with a series of computational techniques. We focus on the government’s *Workers’ Daily* newspaper,



which reports a daily circulation of 960,000.<sup>6</sup> For context, the *New York Times* has a weekday circulation of 551,500 and a Sunday circulation of 1,120,000.

### 3.1.1 The *Workers' Daily*

The *Workers' Daily* was established by Mao Zedong in 1949, just months after the PRC was founded. It targeted Chinese workers, who were then central to the CCP's ideological legitimacy.<sup>7</sup> To do so, the *Workers' Daily* employed colloquial language, and its editors routinely rewrote *People's Daily* and Xinhua content. The *Workers' Daily* also covered issues of more direct relevance to labor, occasionally even editing the more "politicized" content in the *People's Daily* to be of genuine interest to readers.<sup>8</sup> Reflecting this, on the eve of the Cultural Revolution, the *Workers' Daily* had a daily circulation of 400,000.<sup>9</sup> Although the *Workers' Daily* ceased publication during the Cultural Revolution, its circulation figures rebounded during the early reform era (1978-1992). It focused on labor issues and occasionally highlighted the regime's policy shortcomings in areas like manufacturing waste, misuse of bonuses, deprivation of leisure time, and worker safety.<sup>10</sup> Partly because of its reputation for genuine news, during 1980s the *Workers' Daily* reached a daily circulation of 2,440,000.

China's economic take-off since 1992 has transformed the status of workers: from a politically important and socially mainstream class to a relatively weak and marginalized one. Accordingly, the *Workers' Daily* has shifted its target audience. In addition to its original readership of industrial laborers, union leaders, and SOE employees, the *Workers' Daily* now caters to migrant workers, intellectuals, and civil servants who are members of unions.<sup>11</sup> Its content reflects this shift. One study examined a random sample of 868 articles on labor issues published between 1979 and 2008.<sup>12</sup> Prior to 1992, just 15% of articles cited government or party sources; following 1992, some 43% did. Prior to 1992, nearly 8% of articles cited labor sources; since 1992, only 1% have. Prior to 1992, 22% of articles included direct quotes attributed to workers; since 1992, less than 9% have.<sup>13</sup>

---

<sup>6</sup>Its reported daily circulation of 960,000 is exceeded by its number of followers on Weibo, China's social media platform. There, the *Workers' Daily* counts 1,400,000 subscribers.

<sup>7</sup>According to a widely-cited quote, in the decision to start *Workers' Daily* in July 1949, All China Federation of Trade Union (中华全国总工会) stated “(工人日报的诞生) 是中国工人阶级的历史上破天荒的一件大事! 这是中国工人阶级多年来在工人阶级自己的政党——中国共产党、中国工人阶级的英明领袖——毛主席的领导下英勇奋斗, 流血牺牲所换来的重大革命成果之一”。

<sup>8</sup>张刃. 贴近大众是报纸生存与发展的必由之路——我理解的工人日报的特色 [J]. 新闻三昧, 2009, (06):4-7.

<sup>9</sup>Ibid.

<sup>10</sup>At times the *Workers' Daily* even defended workers' interests against government officials. When the “No. 2 Bohai” drilling ship sank in November of 1979 and caused the death of 72 workers, the *Workers' Daily* took the lead in investigating and reporting the episode. The reports on “No. 2 Bohai” led to the sentence of a bureau chief (局长), the firing of a minister (部长), and a recorded mistake (记大过) of a vice prime minister. (张刃. 贴近大众是报纸生存与发展的必由之路).

<sup>11</sup>Ibid.

<sup>12</sup>For the purpose of their study, they narrowly define workers as the ones who have urban hukou, are dependent on wages for subsistence, directly engage in industrial production and are SOE employed laborers.

<sup>13</sup>夏倩芳, 景义新. 社会转型与工人群体的媒介表达——《工人日报》1979-2008年工人议题报道之分析 [J]. 新闻与传

The *Workers’ Daily* has retained its reputation for relatively legitimate news content, but its target audience is now far broader than at any point in its 70 year history. For this reason, we focus on the *Workers’ Daily* rather than the *People’s Daily*, which targets an elite rather than popular audience.

### 3.1.2 Identifying Propaganda-Based Threats: Supervised Learning

We identify threats of violence in the *Workers’ Daily* with supervised learning methods. After scraping the entire *Workers’ Daily* online archive between 2009 and 2016 – which yielded 164,707 unique articles – we read several thousand articles to create a list of all possible coverage topics. We identified 29 topics, which appear in Table 1. Then, we randomly sampled 500 articles from the *Workers’ Daily*. We refer to these 500 articles as our training set, and the remaining articles as our test set. We labeled each training set article with as many topic labels as appropriate.

Table 1: Topic List

corruption	culture	democracy and rights	economy
education	electoral politics	environment	government action
international cooperation	international news	law enforcement	legal system
media	military	nation	natural disaster
obituary	protest	public health	religion
science	social	sports	stability
terrorism	traffic	war	weather
youth			

Note: The blue label reflects our gold standard set of threats: articles which focus on social stability maintenance. The red label reflects our broader measure of threats: articles that discuss law enforcement capabilities and the internal security apparatus. Our classifier identified 16 and 45 of these published in the *Workers’ Daily* between 2009 and 2016, respectively.

We identified two topics that routinely included threats to citizens. They appear in Table 1 in blue and red, respectively. The first is widely regarded by observers of Chinese politics as profoundly threatening: allusions to “social stability maintenance” (维稳). “The term,” Huang (2015, 426) writes, “is broadly understood in China as a code word for maintaining the stability of the existing regime.”<sup>14</sup> Within China, it is associated with Deng Xiaoping’s response to the Tiananmen Square

播评论,2008,(00):174-182.

<sup>14</sup>See also Chen (2013), Steinhardt (2016), Yue (2012), Benney (2016), Yang (2017), Sandby-Thomas (2011), and Wang and Minzner (2015).

massacre: “stability overrides everything.”<sup>15</sup> We regard articles labeled as “social stability” as constituting our gold standard set of threats. As we discuss below, this set has no false positives – each article labeled social stability is profoundly threatening – but may exclude some articles that are designed to threaten citizens as well. That is, it may produce a sample with false negatives: articles that threaten in some other way.

To accommodate the possibility of false negatives, we identify a second topic, which highlights the capacity of the government’s law enforcement and internal security apparatus. By describing the loyalty of police forces and the government’s efficient response to crime, the propaganda apparatus may communicate to citizens that anti-regime behavior will be detected and quickly repressed. To be clear, however, this may generate some false positives: by labeling as threats some articles that constitute legitimate reporting on crime. We refer to this as our “law enforcement” set of threat articles.

After labeling our training set, we processed our corpus for text analysis and implemented a multi-label topic model.<sup>16</sup> Between 2009 and 2016, our classifier identified 16 social stability articles (a rate of 2.3 per year) and 45 law enforcement articles (a rate of 6.4 per year).

We validate the performance of the classifier in two ways. First, we asked a native speaker to read a random sample of 300 articles, and, to each, apply all appropriate topic labels. We then compared these human classification decisions with the results of the topic model. We did so by employing a binary classification decision *for each topic in each article*: whether, for instance, a social stability tag or a law enforcement tag was appropriately applied. We then computed an overall accuracy rate by topic. Figure 1 reports these accuracy rates by topic. For each topic along the *y*-axis, the cells are shaded according by accuracy rate. The legend on the right visualizes our accuracy gradient, with red representing random accuracy and blue representing perfect accuracy. The accuracy rates appear in each cell, rendered in gray. These accuracy rates are high across topics, and especially so for threat labels. Our classifier identified articles about social stability with a 95% accuracy rate and articles about law enforcement with an 89% accuracy rate.

We also validate the topic model using a series of qualitative techniques: by inspecting most common words by topic label and the text of the articles themselves. We include both in the online appendix. Again, the topic model appears to have substantial validity.

---

<sup>15</sup>“稳定压倒一切,” from Deng Xiaoping, ‘The United States should take the initiative in putting an end to the strains in Sino –American relations’, Selected Works of Deng Xiaoping, vol. 3 (Beijing: Renmin Chubanshe, 1993), p. 331. After Tiananmen, Deng instructed state-run media to emphasize the importance of social stability in their coverage. In 1992, the CCP announced that the media should follow three principles in shaping public opinion: unity (团结), stability (稳定), and morale (鼓精) (Zhonggong zhongyang guanyu jiaqiang he gaijin xuanchuan sixiang gongzuo (3 September 1992), 2180.

<sup>16</sup>Before classifying articles, we employed standard text preprocessing practices. We removed numbers, symbols, and punctuation. We used the `jieba` segmenting algorithm to split chunks of Chinese text into constituent words. We used the `multilabel` classifier in Python’s `sklearn` module. The multi-label classifier employs a linear support vector classifier one-versus-rest model to assign as many topics as appropriate to test set articles.




Figure 1: Validation of our classifier model.

### 3.1.3 Identifying Propaganda-Based Threats: Word Frequency Counts

We identify propaganda-based threats in a second way: by measuring frequency counts for two terms, “stability” (稳定) and “harmony” (和谐), closely associated with repression in China. “Stability” is routinely employed by government officials and the propaganda apparatus to underscore the regime’s commitment to maintaining social stability. “Harmony” is so widely identified as a threat that it has entered colloquial speech. The phrase “being harmonized” (被和谐了) is used as a euphemism for being detained, arrested, or censored. Critically, these words are seldom employed *not* in a threatening context.

Text from articles themselves makes this clear. One typical article, from 2010, summarized a working conference attended by President Hu Jintao and other senior officials. To strengthen the rule of law, the article concluded, it was necessary to “actively participate in the comprehensive management of social order and safeguard national security and social stability ...[and] maintain a harmonious social management order.” In February 2011, another article claimed this:

Strengthening and innovating in social stability management is an inevitable requirement for continuing to build on the important strategic opportunity for our country in promoting the cause of the Party and the state, an inevitable requirement for building a harmonious socialist society, and a necessary requirement for safeguarding the fundamental interests of the overwhelming majority of the people. ...In a country like ours, the task of social stability management is increasingly arduous. The fundamental purpose of our efforts to strengthen and innovate social management is to maintain social order, promote social harmony, ensure that the people live and work in peace and contentment, and create a favorable social environment for the development of our Party and our country.

In February 2013, still another article applauded the government’s efforts to “establish and perfect

a risk assessment mechanism for social stability ...[and] to improve the national security strategy, exercise vigilance, and resolutely prevent separatist movements and infiltration and subversive activities by hostile forces.”

These word frequency counts are particularly attractive for measurement purposes. The topic labels described in Section 3.1.2 yield a dichotomous outcome: whether some article  $j$  focused on either social stability or law enforcement. It includes no information about an article’s length, or what *share* of the article sought to threaten citizens with violence. We show below that the *Workers’ Daily* generally publishes no more than one such article per day. By contrast, from the empirical distributions in Figure 2, these word frequency counts are more continuous: combined, ranging from 0 to 188 per day, with a mean daily count of 17.3. As a result, they permit us to measure just how threatening was *Workers’ Daily* content on day  $t$ .

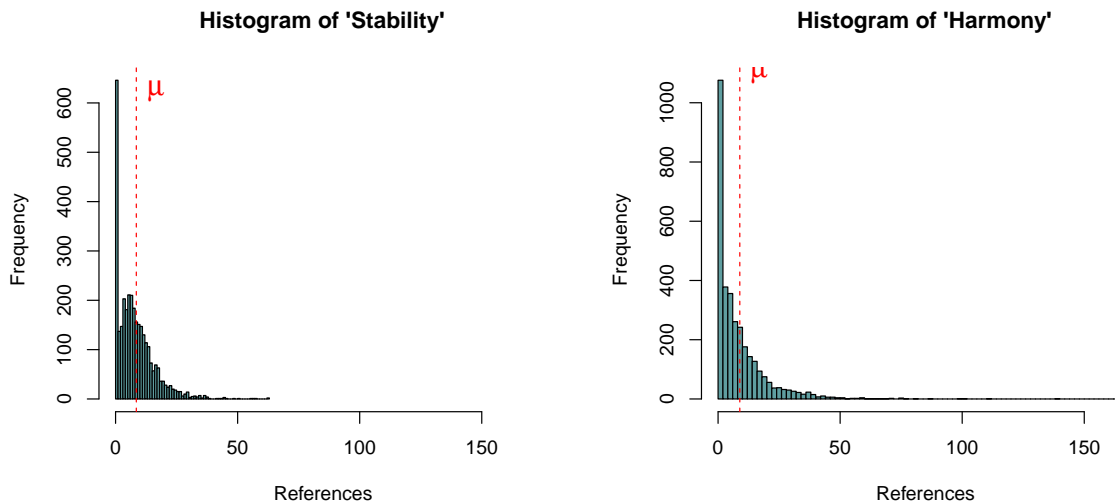


Figure 2: Histogram of word frequency counts. The left panel gives the distribution of “stability” references per day. The right panel gives the distribution of “harmony” references per day. Mean values are shown in red.

### 3.2 Collective Action in China

We use protest data from Manfred Elfstrom and the China Labour Bulletin (CLB), a non-governmental organization in Hong Kong that advocates for labor rights in China. Drawing on international, domestic, and social media, they maintain a geocoded dataset of all known strikes and protests. Elfstrom’s dataset covers the period between 2006 and 2012, while the CLB dataset covers the period between 2011 and 2016.<sup>17</sup>

<sup>17</sup>An earlier version of these data covering 2008-2012 was analyzed by Distelhorst and Hou (2017). Note that social media reports are an accepted way to measure unrest in China, as the government does not release official data on

To maximize coverage, we merged the Elfstrom and CLB datasets. We did so after confirming that their respective data generating processes – their coding rules and their sources – are essentially identical, as all available information suggests they are. We did this by exploiting the fact that the two datasets overlap for 2011 and 2012. Figure 3 reports the number of protests, by day, between 2006 and 2016, with the Elfstrom daily average in blue and the CLB daily average in red. As expected, the two datasets are essentially identical in 2011 and 2012. We also conducted a Granger test to confirm that their protest records are statistically indistinguishable; we report the results in the online appendix. Accordingly, we create the variable  $Protests_{it}$ , which it records the number of protests in province  $i$  on day  $t$ . For the period between 2006 and 2010, we use Elfstrom’s data; for 2011 through 2016, we use CLB data. As an additional precaution, we also employ year-level fixed effects to accommodate unobserved differences in the data generating process.

Figure 3 suggests that the rate of popular protest across China has risen exponentially since 2006. In part, this may reflect the data collection process. By relying on social media reports, the Elfstrom and CLB measures may under report protest events prior to the recent boom in internet penetration and social media use. There are good reasons to believe, however, that this increase does not only reflect rising social media rates. The Chinese government reported roughly 10,000 protests in 1994 and 80,000 protests in 2008. Then, since its own data indicated rising popular discontent, the government stopped releasing it. One Chinese sociologist estimated that, in 2010, there were 180,000 protests across the country. The time trend in Figure 3 reflects this rise in popular discontent.

Figure 4 presents two empirical distributions: of the number of protests each day, and the number of protests over the course of the coming week. The mean values for the two distributions – 2.9 and 20.2, respectively – are shown in red. We use both protest measures as outcome variables in the statistical models below. The online appendix includes a range of descriptive statistics about protests by province. In short, we find some evidence that, at the province level, protests are correlated with economic output, which may reflect higher levels of social media use or urbanization rate. Likewise, protests are also recorded disproportionately in Guangdong, close to CLB headquarters.

## 4 Calendars of Threats and Collective Action

### 4.1 Descriptive Statistics

We visualize the life cycle of regime threats and collective action in Figure 5. The top panel focuses on threats. We index each calendar day as  $d \in \{1, 365\}$ , and then, for each day  $d$  along the  $x$ -axis, the left and right  $y$ -axes report our measures of threats. The left  $y$ -axis reports the total number of times, per day, between 2006 and 2016 that the *Workers’ Daily* mentioned “stability” (red) and “harmony” (orange). We also report the number of times the *Workers’ Daily* mentioned some other protests (Cai 2010; Wallace and Weiss 2015).

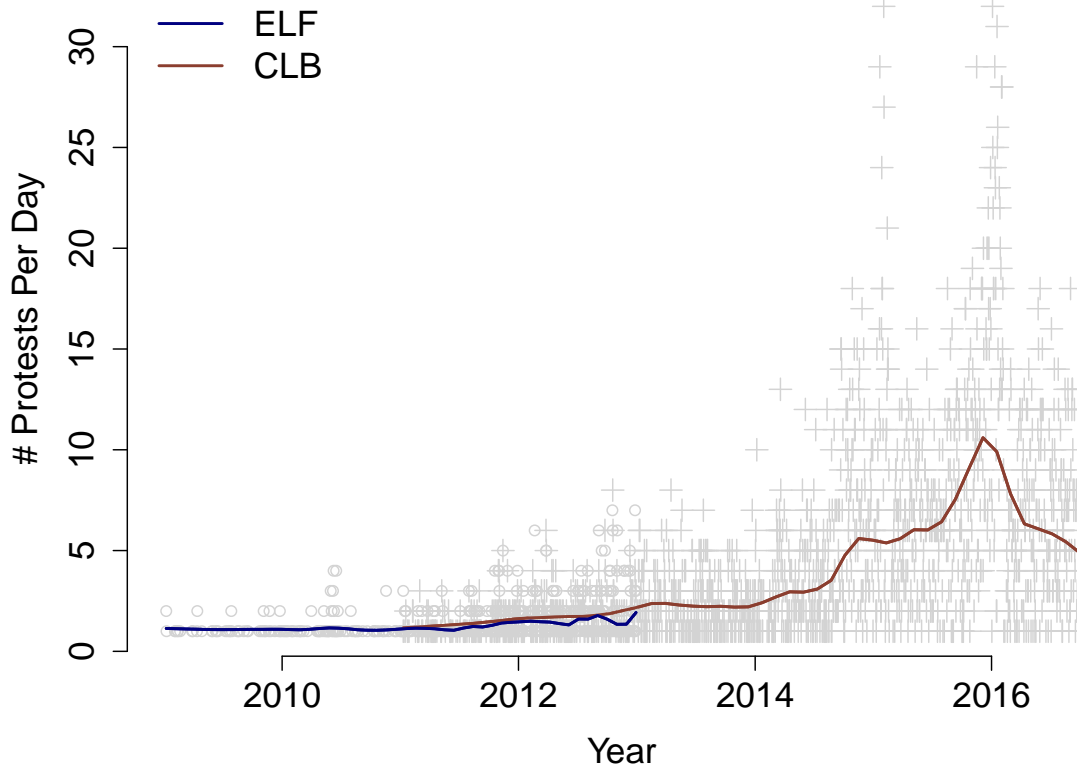


Figure 3: Protests over time. Data from Manfred Elfstrom are shown in blue and cover the 2006-2012 period. Data from the China Labour Bulletin are shown in red and cover January 2011 through September 2016. We aggregate these two measures into a single variable,  $Protests_{it}$ , which uses Elfstrom’s data for the 2006-2010 period and the CLB’s data thereafter.

terms associated with social stability maintenance: “social stability” (green), and “police” (blue). The right  $y$ -axis reports the number of articles, again by day, between 2006 and 2016 that are classified as social stability (plus sign) and law enforcement (times sign).

The bottom panel visualizes the life cycle of collective action in China. Again, we index each calendar day as  $d \in \{1, 365\}$ , and then, for each day  $d$  along the  $x$ -axis, we compute the total number of protests across the country between 2006 and 2016. The two dashed horizontal lines indicate daily protest levels equal to the mean plus one or two standard deviations, respectively.

These descriptive statistics make clear that China’s propaganda apparatus does not issue threats randomly, or frequently. Rather, as the theory in Section 2 suggests, propaganda-based threats appear to cluster around politically sensitive moments, when citizens are engaged in politics and




Figure 4: Histogram of protests. The left panel gives the distribution of protests per day. The right panel gives the distribution of protests per week. Mean values are shown in red.

aware of their shared frustrations with the regime. We label two sets of politically sensitive dates above both panels: the anniversaries of ethnic separatist movements (orange) and scheduled political meetings and the anniversaries of failed pro-democracy movements (both in gray).

As expected, our continuous and dichotomous measures of threats are correlated. We estimate a series of bivariate regressions of the form

$$\text{Continuous Threat Measures}_t = \alpha + \beta + (\text{Dichotomous Threat Measures}_t) + \epsilon \quad (1)$$

which give the additional number of references to “stability” and “harmony,” by day, when the *Workers’ Daily* publishes social stability and law enforcement articles. The results appear in Table 2. The baseline daily number of references to “stability” and “harmony” are between 8.5 and 9. When the *Workers’ Daily* publishes a social stability article, the mean number of references to “stability” rises by nearly 7 words; the number of references to “harmony” rises by 4 words. When the *Workers’ Daily* publishes a law enforcement article, the mean number of references to “stability” and “harmony” rises by between 0.4 and 2 words.

### Anniversaries of Ethnic Separatist Movements

Modern China has experienced two major separatist movements – one in Tibet and another in Xinjiang – which have simmered for decades. These separatist movements have accumulated four anniversaries that are now focal moments for popular protest *in those provinces* and, as a result,






Figure 5: The life cycles of threats and protest in China. In the top panel, the left  $y$ -axis gives the total number of times that keywords – social stability, police, stability, and harmony – were published on day  $t$ , and right  $y$ -axis gives the total number of articles labeled as “social stability” or “law enforcement” published on day  $t$ . In the bottom panel, the  $y$ -axis records the total number of protests on day  $t$ . The light blue lines gives the mean protest level plus one standard deviation. The dashed navy line indicates the two standard deviation threshold. Ethnic separatist anniversaries appear in orange; pro-democracy anniversaries and major political moments appear in gray.

government repression. That is, these anniversaries are deeply salient – for citizens and for the regime – within Tibet and Xinjiang, but not elsewhere. Table 3 describes these anniversaries, and

Table 2: Correlation Between Term Counts and Topic Models

	<i>Dependent variable:</i>			
	“Stability”	“Harmony”	“Stability”	“Harmony”
	(1)	(2)	(3)	(4)
Social Stability	6.667*** (1.937)	4.050 (2.871)		
Law Enforcement			1.846 (1.162)	0.387 (1.720)
Constant	8.458*** (0.135)	8.887*** (0.200)	8.465*** (0.136)	8.902*** (0.201)
Observations	3,287	3,287	3,287	3,287
R <sup>2</sup>	0.004	0.001	0.001	0.00002

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01

we provide more detail in the online appendix.

The descriptive statistics in Figure 5 suggest that these separatist anniversaries compel the government to issue propaganda-based threats. Between 2009 and 2016, 4 of 16 social stability articles and 9 of 45 law enforcement articles were published in the week centered upon each of our four separatist anniversaries. Because separatist anniversaries are so rare, this implies a much higher rate of threats during these periods. Figure 6 shows these relative rates. The left panel reports topic model data. Social stability articles are four times as likely during separatist windows, which we define as the seven day period centered on a separatist anniversary. Law enforcement articles are three times as likely during separatist windows. The right panel reports term frequency counts. Between 2009 and 2016, references to “stability” and “harmony” were some 20% more common during separatist windows. These threats may be directed principally at the ethnic minority populations in Tibet and Xinjiang, but citizens across China are nonetheless exposed to them.

### **Anniversaries of Failed Pro-Democracy Movements and Political Meetings**

Carter and Carter (2017*b*) find that the anniversaries of China’s failed pro-democracy movements constitute focal moments for protest. They identify five, in particular: Tiananmen, Democracy Wall, Constitution Day, Charter 08, and NPC Direct Elections. Carter and Carter (2017*b*) find more limited evidence that major political meetings like the quinquennial Party Congress and the annual leadership retreat to a beach resort outside Beijing also occasion protest. Figure 5 reinforces this. There is strong evidence that the government prepares for these focal moments in advance: by arresting dissidents, increasing censorship, and flooding social media with pro-regime content (Truex 2016; King, Pan and Roberts 2013, 2017). We find some evidence that the government may issue propaganda-based threats during these moments as well.

Table 3: Ethnic Separatist Anniversaries

Date	Anniversary	Description
May 23	Tibetan “Liberation”	After founding the PRC in 1949, Mao was keen to reassert state authority in Tibet, which had enjoyed de facto autonomy in the nineteenth and early twentieth centuries due to the collapse of the Qing Dynasty. Mao launched a military intervention and on May 23, 1951, forced Tibetan leaders to sign the Seventeen Point Agreement, which stated that Tibet was part of China. <sup>18</sup>
March 10	Tibetan Rebellion	On March 10, 1959, the Tibetan Rebellion began in Lhasa, the capital. The PLA put down the rebellion on March 20, forcing the Dalai Lama to flee to India. 7,000 Tibetan refugees followed him that spring. The exile community celebrates the day as Tibetan Uprising Day. The anniversary has occasioned Tibetan protests since, as well as increased levels of repression in Tibet. <sup>19</sup>
March 28	Serf Emancipation Day	CCP authorities commemorate the crushing of the Tibetan Rebellion as Serf Emancipation Day. On March 28, 1959, Premier Zhou Enlai dissolved the Tibetan government and assigned policymaking authority to the Preparatory Committee of the Tibet Autonomous Region. The decision was publicized nationally by Xinhua. <sup>20</sup>
July 5	Xinjiang Uprising	Following news of a factory altercation between Uygher and Han workers in Guangdong, on July 5, 2009, protests broke out in Urumqi, the capital of Xinjiang. The protest drew over 1,000 and became violent, targeting Han people and businesses. Officially, 197 were killed and 1,721 injured. Unofficial tolls are much higher. Over 1,000 Uyghers were arrested, 30 of whom received death sentences. Mosques were closed and communications networks cut. Local police surveillance, armed patrols, and citizen brigades have increased during subsequent anniversaries (Freedom House N.d.).




Figure 6: Rates of propaganda-based threats. The left panel reports topic model data. The top two bars restrict attention to our law enforcement threat definition. The bottom two bars restrict attention to our social stability threat definition. The  $x$ -axis gives the percent of days on which a threat occurred according to the windows defined along the  $y$ -axis. Anniversary windows reflect a seven day period centered on the anniversary itself. The right panel reports term frequency data. The  $x$ -axis gives the number of times a given term was referenced per day, during ethnic separatist windows and otherwise.

## 4.2 Specification

To probe the determinants of threats more systematically, we estimate a baseline model of the form

$$y_t = \alpha + \beta_1 (\text{Ethnic Separatist Anniversary Window}_t) + \beta_2 (\text{Pro-Democracy Anniversary Window}_t) + \phi X_t + \gamma_s + \epsilon \quad (2)$$

where  $t$  indexes day,  $s$  indexes year, and vector  $X_t$  gives day-level covariates. To accommodate unobserved characteristics by year, we include a full set of year-level fixed effects, given by  $\gamma_s$ . Our dichotomous outcome variable,  $Threat_t$ , assumes value 1 if the *Workers' Daily* published an explicitly threatening article on day  $t$  and value 0 otherwise. Our more continuous outcome variable,  $Count_t$ , counts the number of times that the *Workers' Daily* published either “stability” or “harmony” on day  $t$ . We use logit models for dichotomous outcomes and negative binomials for count outcomes.

Our explanatory variables of interest are *Ethnic Separatist Anniversary Window<sub>t</sub>* and *Pro-Democracy Anniversary Window<sub>t</sub>*. Although we expect threats to cluster around these anniversaries, we are agnostic about *how close* to the anniversaries themselves the propaganda apparatus will issue threats. To accommodate the possibility that threats may be issued several days before or after the anniversary, we also construct a series of anniversary *windows*, ranging from the anniversary plus/minus one day to the anniversary plus/minus five days.

We are unaware of any features that may be correlated with these anniversary windows and propaganda-based threats, save one. We include in  $X_t$  a measure of underlying political instability, which accommodates the possibility that past protests *themselves* generate propaganda-based

threats, rather than the calendar. In particular, we include the variable *Recent Protests<sub>t</sub>*, which counts the number of protests that occurred across the country on day  $t - 1$ .<sup>21</sup> The year-level fixed effects, given by  $\gamma_s$  accommodate any unobserved annual features, such as changes in underlying economic conditions, changes in the data generating process, and changes in the government’s political strategy. In the online appendix we estimate a series of models that include a range of year-level covariates: GDP, population, rural population share, sex ratio, and urban unemployment rates.

### 4.3 Results

The results appear in Tables 4 through 7. Tables 4 and 5 use social stability and law enforcement topic labels, respectively, as the outcomes of interest; Tables 6 and 7 use “stability” and “harmony” word counts as the outcomes of interest. The results are consistent with our theoretical expectations, similar across models, and virtually identical across outcome variables.

The Chinese government is far more likely to issue propaganda-based threats around the anniversaries of ethnic separatist movements. The effect is large and substantively meaningful. From Tables 4 and 5, the daily odds that the government publishes a social stability or law enforcement article are between 330% and 530% greater around ethnic separatist anniversaries than on any other day of the calendar year. From Tables 6 and 7, each day within an ethnic separatist anniversary window witnesses between 14% and 26% more mentions of “stability” and harmony.” By contrast, we find no evidence that the anniversaries of failed pro-democracy movements are associated with propaganda-based threats. These results are substantively unchanged when we include year-level covariates as well as year-fixed effects in the online appendix, though the sample size is somewhat smaller. Tables 6 and 7 also suggest that protests on day  $t - 1$  compel the government to issue propaganda-based threats.

In short, even for a government as repressive as the Chinese Communist Party, explicit threats of violence are rare. These results are consistent with our theoretical expectations, but they add important nuance. Carter and Carter (2017b) find that the anniversaries of failed pro-democracy movements are associated with higher levels of protests across China. In a sense, then, a key question is why the anniversaries of ethnic separatist movements elicit propaganda-based threats, but the anniversaries of failed pro-democracy movements do not. Our data do not permit us to answer this question, though we regard it as an important direction for future research.

Our theory suggests two possible answers. First, the regime may simply view ethnic separatist anniversaries as more threatening than pro-democracy anniversaries. While pro-democracy anniversaries sometimes occasion protest spikes, ethnic separatist anniversaries almost always do *in the province in which they are salient*. Since killing citizens is costly, the government would prefer

---

<sup>21</sup>Note that the results are substantively unchanged if we control for the number of protests that occurred across the country over the preceding week.

to avoid doing so, and hence reserve threats for these moments of particular political instability. This may be why the government was so slow to actually threaten protesters in the run-up to the Tiananmen Square massacre, even as thousands of citizens joined in the weeks prior. In this reading, the propaganda apparatus avoided issuing threats lest they render citizens even more angry. In turn, the infamous *People's Daily* editorial of April 26, 1989, represented a final effort to discourage Tiananmen Square protesters before engaging in mass violence. And, indeed, the *People's Daily* editorial elicited a boomerang effect, as tens of thousands of citizens joined students in the Square to protest the regime's willingness to threaten them.

Second, the government – and perhaps the world's other autocrats – reserve explicit threats of violence for political out-groups. Intuitively, insofar as explicit threats of violence may alienate their recipients, repressive governments may reserve them for groups of citizens on whom their survival does not rest. In China, for instance, Tibetans and Uyghers constituted just 1.4% of the population in 2016. While the government may prove unable to suppress a nationwide pro-democracy movement, it likely can rule the small Tibetan and Uygher communities with force. Accordingly, in Tibet and Xinjiang, repression is widespread. The government routinely conducts military patrols; imposes curfews, internet, and cellular shutdowns; and incarcerates Tibetans and Uyghers in mass. In Xinjiang, the government has even banned the Quran, the name “Mohammed,” and “unreasonably long” beards. At least 120,000 Uyghers are being held in re-education camps (Phillips 2018). Given the depth of hostility among Tibetans and Uyghers for the Beijing government, threats of violence are also less costly than they otherwise might be. While threats directed at ethnic Han citizens may elicit source derogation and the boomerang effect, the government is already maximally unpopular in Tibet and Xinjiang.

#### 4.4 Threats and Pro-Regime Propaganda

As we discussed in Section 2.1, scholars increasingly understand the Chinese government's propaganda strategy as a signal to citizens that the government's capacity for repression is so substantial that it can compel them to consume propaganda content that everyone knows to be false. Huang (2015) refers to this as “hard propaganda,” and Carter and Carter (2017b) refer to it “domination propaganda.” Carter and Carter (2017b) find that the Chinese government's propaganda apparatus is among the most absurdly pro-regime in the world.

One implication of our theory in Section 2 and the results in Section 4.3, then, is that propaganda spikes should be correlated with propaganda-based threats. To probe this, we use the measure of pro-regime propaganda developed by Carter and Carter (2017b). For every reference to the Chinese Communist Party on day  $t$ , we extracted the 10 words before and after the reference, a string known as a “concordance segment.” Drawing on standard semantic dictionaries, we measured how fulsome or critical were these 20 words.<sup>22</sup> The variable *Positive Coverage<sub>t</sub>* constitutes our measure of pro-

---

<sup>22</sup>We segmented each word in our corpus, dictionaries, and list of identifiers. We removed numbers, symbols, and

Table 4: When Autocrats Issue Propaganda-Based Threats: Social Stability Topic

	<i>Dependent variable:</i>					
	Social Stability Topic					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Ethnic Separatist Anniversaries</i>						
Anniversary	-15.275 (2,841.667)					
1 Day Window		0.591 (1.040)				
2 Day Window			1.663*** (0.586)			
3 Day Window				1.265** (0.584)		
4 Day Window					1.363** (0.554)	
5 Day Window						1.209** (0.564)
<i>Pro-Democracy Anniversaries</i>						
Anniversary	-15.065 (2,774.881)					
1 Day Window		-16.038 (2,641.730)				
2 Day Window			-15.872 (2,047.488)			
3 Day Window				-15.934 (1,740.651)		
4 Day Window					-0.116 (1.054)	
5 Day Window						0.482 (0.788)
<i>Other Covariates</i>						
Recent Protests	-0.596 (0.610)	-0.590 (0.611)	-0.615 (0.618)	-0.577 (0.617)	-0.589 (0.619)	-0.602 (0.619)
Constant	-4.754*** (0.580)	-4.769*** (0.583)	-4.938*** (0.598)	-4.886*** (0.597)	-5.020*** (0.612)	-5.081*** (0.621)
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,281	3,281	3,281	3,281	3,281	3,281

Note:

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01

Table 5: When Autocrats Issue Propaganda-Based Threats: Law Enforcement Topic

	<i>Dependent variable:</i>					
	Law Enforcement Topic					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Ethnic Separatist Anniversaries</i>						
Anniversary	-16.265 (2,892.280)					
1 Day Window		0.291 (0.732)				
2 Day Window			1.336*** (0.402)			
3 Day Window				1.174*** (0.386)		
4 Day Window					1.301*** (0.353)	
5 Day Window						1.118*** (0.356)
<i>Pro-Democracy Anniversaries</i>						
Anniversary	-16.238 (2,804.682)					
1 Day Window		-16.242 (1,619.225)				
3 Day Window			-0.839 (1.018)			
3 Day Window				0.301 (0.536)		
4 Day Window					0.382 (0.490)	
5 Day Window						0.607 (0.431)
<i>Other Covariates</i>						
Recent Protests	-0.267 (0.390)	-0.261 (0.390)	-0.269 (0.393)	-0.267 (0.394)	-0.278 (0.396)	-0.290 (0.396)
Constant	-3.895*** (0.382)	-3.895*** (0.383)	-4.028*** (0.390)	-4.097*** (0.394)	-4.184*** (0.401)	-4.213*** (0.404)
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,281	3,281	3,281	3,281	3,281	3,281

*Note:*

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01



Table 6: When Autocrats Issue Propaganda-Based Threats: “Stability” Count

	<i>Dependent variable:</i>					
	Stability					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Ethnic Separatist Anniversaries</i>						
Anniversary	0.235*					
	(0.139)					
1 Day Window		0.128				
		(0.082)				
2 Day Window			0.162**			
			(0.064)			
3 Day Window				0.172***		
				(0.055)		
4 Day Window					0.230***	
					(0.049)	
5 Day Window						0.243***
						(0.045)
<i>Pro-Democracy Anniversaries</i>						
Anniversary	-0.111					
	(0.142)					
1 Day Window		-0.022				
		(0.082)				
2 Day Window			-0.013			
			(0.064)			
3 Day Window				0.013		
				(0.055)		
4 Day Window					0.040	
					(0.050)	
5 Day Window						0.053
						(0.047)
<i>Other Covariates</i>						
Recent Protests	0.184***	0.182***	0.181***	0.182***	0.183***	0.185***
	(0.039)	(0.039)	(0.039)	(0.039)	(0.039)	(0.039)
Constant	2.666***	2.665***	2.659***	2.651***	2.635***	2.621***
	(0.042)	(0.042)	(0.042)	(0.042)	(0.042)	(0.042)
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	
Observations	3,272	3,272	3,272	3,272	3,272	3,272

Note:

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01

Table 7: When Autocrats Issue Propaganda-Based Threats: “Harmony” Count

	<i>Dependent variable:</i>					
	Harmony					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Ethnic Separatist Anniversaries</i>						
Anniversary	0.261 (0.169)					
1 Day Window		0.057 (0.100)				
2 Day Window			0.140* (0.078)			
3 Day Window				0.159** (0.067)		
4 Day Window					0.161*** (0.060)	
5 Day Window						0.147*** (0.055)
<i>Pro-Democracy Anniversaries</i>						
Anniversary	-0.113 (0.172)					
1 Day Window		0.013 (0.100)				
2 Day Window			0.052 (0.078)			
3 Day Window				0.090 (0.067)		
4 Day Window					0.106* (0.061)	
5 Day Window						0.093
<i>Other Covariates</i>						
Recent Protests	0.167*** (0.047)	0.165*** (0.047)	0.165*** (0.047)	0.165*** (0.047)	0.165*** (0.047)	0.165*** (0.047)
Constant	2.539*** (0.050)	2.539*** (0.051)	2.531*** (0.051)	2.521*** (0.051)	2.513*** (0.051)	2.509*** (0.051)
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	
Observations	3,272	3,272	3,272	3,272	3,272	3,272

*Note:*

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01

regime propaganda, and it gives the number of fulsome words, less critical words, among the 20, summed for day  $t$ . Since we are interested in *spikes* in pro-regime propaganda, we define a “high propaganda” day as occurring when the observed level of *Positive Coverage<sub>t</sub>* is greater than or equal to the sample mean plus  $\kappa$  standard deviations. To ensure robustness, we let  $\kappa$  range from [0.25, 2.0]. We then estimate the model

$$y_t = \alpha + \beta (\text{High Propaganda Day}_t) + \epsilon \quad (3)$$

where  $t$  indexes day.

Table 8 reports results for social stability and law enforcement topic threats; Table 9 reports results for “stability” and “harmony” counts. Across outcome variables and  $\kappa$  threshold values, the results suggest that the Chinese government’s propaganda apparatus employs absurdly positive pro-regime propaganda and threats of violence against citizens at roughly the same moments. Most strikingly, from Table 9, high propaganda days are associated with roughly four additional references to “stability,” and roughly six additional references to “harmony.”

This is consistent with our theory of threats and existing research that understands absurdly positive pro-regime propaganda as itself implicitly threatening. They are complements, with explicit threats of violence reserved for moments of most profound political tension.

## 5 Do Threats of Violence Work?

### 5.1 Selection Bias

Determining whether threats reduce popular protests is complicated by the fact that governments issue threats strategically. As we show above, the government is far more likely to threaten violence against citizens during politically sensitive moments, when citizen frustrations are most likely to coalesce into mass protests, and in response to protests on day  $t - 1$ . Empirically, then, there are likely to be two competing effects: a negative effect on protest due to the force of the threat, and a positive effect on protest due to the underlying political tensions that compelled the government to issue the threat. As a baseline, we estimate a “naive” model of the form

$$\text{Protest}_{t+1:t+7} = \alpha + \beta (\text{Threat}_t) + \phi X_t + \psi W_s + \gamma_s + \epsilon \quad (4)$$

where  $t$  indexes day and  $s$  indexes year. The vectors  $X_t$  and  $W_s$  include a range of day- and year-level covariates, respectively. To accommodate unobserved characteristics by year, we include a full set of year-level fixed effects, given by  $\gamma_s$ . Since our outcome variable is effectively continuous – the number of protests across the country over the next seven days – we employ OLS. We use our

---

punctuation from the corpus before generating the concordance segments from which we extracted our measure of tone. See Grimmer and Stewart (2013) and Lowe et al. (2010).

Table 8: Threats and Propaganda: Topic Models

		<i>Dependent variable:</i>							
		Social Stability Topic							
		$\kappa = 0.25$	$\kappa = 0.5$	$\kappa = 0.75$	$\kappa = 1.0$	$\kappa = 1.25$	$\kappa = 1.5$	$\kappa = 1.75$	$\kappa = 2.0$
High Propaganda		0.006* (0.003)	0.011*** (0.004)	0.013*** (0.004)	0.016*** (0.005)	0.018*** (0.006)	0.015** (0.007)	0.006 (0.007)	0.008 (0.008)
Constant		0.004*** (0.001)	0.003*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.005*** (0.001)	0.005*** (0.001)
Observations		3,287	3,287	3,287	3,287	3,287	3,287	3,287	3,287
R <sup>2</sup>		0.001	0.003	0.003	0.003	0.003	0.001	0.0002	0.0003

		<i>Dependent variable:</i>							
		Law Enforcement Topic							
		$\kappa = 0.25$	$\kappa = 0.5$	$\kappa = 0.75$	$\kappa = 1.0$	$\kappa = 1.25$	$\kappa = 1.5$	$\kappa = 1.75$	$\kappa = 2.0$
High Propaganda		0.003 (0.005)	0.012* (0.006)	0.015** (0.007)	0.016** (0.008)	0.024** (0.010)	0.025** (0.012)	0.019 (0.012)	0.024* (0.013)
Constant		0.013*** (0.002)	0.012*** (0.002)	0.012*** (0.002)	0.013*** (0.002)	0.013*** (0.002)	0.013*** (0.002)	0.013*** (0.002)	0.013*** (0.002)
Observations		3,287	3,287	3,287	3,287	3,287	3,287	3,287	3,287
R <sup>2</sup>		0.0001	0.001	0.001	0.001	0.002	0.001	0.001	0.001

*Note:* \*p<0.1; \*\*p<0.05; \*\*\* p<0.01

Table 9: sord Frequency Counts

<i>Dependent variable:</i>										
"Stability"										
	$\kappa = 0.25$	$\kappa = 0.5$	$\kappa = 0.75$	$\kappa = 1.0$	$\kappa = 1.25$	$\kappa = 1.5$	$\kappa = 1.75$	$\kappa = 2.0$		
High Propaganda	4.585*** (0.338)	4.114*** (0.399)	3.689*** (0.470)	3.522*** (0.554)	3.671*** (0.670)	3.035*** (0.766)	3.119*** (0.817)	2.978*** (0.870)		
Constant	7.639*** (0.146)	7.967*** (0.142)	8.161*** (0.140)	8.269*** (0.139)	8.336*** (0.137)	8.393*** (0.137)	8.403*** (0.137)	8.417*** (0.137)		
Observations	3,287	3,287	3,287	3,287	3,287	3,287	3,287	3,287		
R <sup>2</sup>	0.053	0.031	0.018	0.012	0.009	0.005	0.004	0.004		
<i>Dependent variable:</i>										
"Harmony"										
	$\kappa = 0.25$	$\kappa = 0.5$	$\kappa = 0.75$	$\kappa = 1.0$	$\kappa = 1.25$	$\kappa = 1.5$	$\kappa = 1.75$	$\kappa = 2.0$		
High Propaganda	6.113*** (0.503)	6.041*** (0.591)	6.461*** (0.692)	6.138*** (0.818)	5.551*** (0.992)	5.576*** (1.133)	6.090*** (1.207)	6.234*** (1.285)		
Constant	7.773*** (0.217)	8.139*** (0.211)	8.331*** (0.207)	8.522*** (0.205)	8.674*** (0.203)	8.729*** (0.202)	8.736*** (0.202)	8.753*** (0.202)		
Observations	3,287	3,287	3,287	3,287	3,287	3,287	3,287	3,287		
R <sup>2</sup>	0.043	0.031	0.026	0.017	0.009	0.007	0.008	0.007		

Note:

\*p<0.1; \*\*p<0.05; \*\*\* p<0.01

standard set of threat measures.

The results appear in Table 10, and they provide strong evidence of competing effects. The estimated coefficients for social stability and law enforcement topics hover around 0. For “stability” and “harmony” word counts, the estimated effect on protest levels is negative and statistically significant, but substantively almost meaningless. For the number of protests over the next week to fall by just one, these results suggest, the Chinese propaganda apparatus would have to publish 15 additional references to “stability” and 20 additional references to “harmony.” The *Workers’ Daily* makes just eight references to “stability” and “harmony” alike per day, with standard deviations of eight and 11, respectively.

## 5.2 Identification Strategy

The results in Section 4 offer an opportunity to address the possibility of selection bias, and hence obtain plausibly causal estimates of the effect of propaganda-based threats on popular protest. We treat the theory in Section 2 and the empirical results in Section 4 as a model of the treatment assignment mechanism: of how the Chinese government chooses when to threaten citizens. We then use this treatment assignment mechanism as the foundation for an instrumental variables (IV) estimation strategy.

In particular, we exploit China’s ethnic diversity and its sprawling size to identify a third variable, which predicts propaganda-based threats *and* conditions the probability of popular protest only through its effect on threats. Our key idea is that the content of the *Workers’ Daily* is set at the national level, but must occasionally respond to local conditions. As a result, provinces that are geographically and socially distant from each other are occasionally “treated” with propaganda content that is directed primarily at the other province’s readership. In practice, this identification strategy entails identifying events that generate mass protests in province  $i$ , induce the government to issue propaganda-based threats to citizens in province  $i$ , and are otherwise unrelated to mass protests in some geographically or culturally distant province  $j$ . We argue that the ethnic separatist anniversaries identified in Section 4 provide a suitable instrument.

## 5.3 The Exclusion Restriction

This identification strategy rests on an exclusion restriction. The anniversaries of ethnic separatist movements in province  $i$  must only reduce the probability of protests in province  $j$  through their effects on propaganda-based threats. We believe this is the case, though we consider two potential violations.

Table 10: Naive Estimates of Threats and Protest

	<i>Dependent variable:</i>			
	Protests <sub>t+1:t+7</sub>			
	(1)	(2)	(3)	(4)
Social Stability	-0.374 (2.769)			
Law Enforcement		0.557 (1.547)		
“Stability”			-0.070*** (0.026)	
“Harmony”				-0.049*** (0.017)
Recent Protests	2.663*** (0.529)	1.667*** (0.074)	1.688*** (0.075)	1.686*** (0.075)
Ethnic Separatist Anniversary Window	-2.605*** (0.795)	-2.043*** (0.731)	-1.788** (0.734)	-1.886*** (0.730)
Pro-Democracy Anniversary Window	6.625*** (1.186)	5.368*** (1.090)	5.315*** (1.089)	5.270*** (1.089)
log GRP	-691.054*** (24.250)	-527.759*** (23.377)	-527.381*** (23.344)	-532.450*** (23.390)
log Pop	7,879.511*** (805.076)	5,948.561*** (744.139)	5,940.678*** (743.129)	5,977.090*** (742.986)
Rural Population Share	-3,123.816*** (425.469)	-2,489.962*** (391.731)	-2,454.449*** (391.434)	-2,510.460*** (391.149)
Sex Ratio	23,873.710*** (8,053.456)	19,833.380*** (7,395.003)	18,922.770** (7,392.446)	20,308.630*** (7,384.682)
Urban Unemployment Rate	-1,210.566*** (266.884)	-976.255*** (245.258)	-945.443*** (245.186)	-998.679*** (244.980)
Constant	-104,429.000*** (12,005.540)	-80,357.270*** (11,079.350)	-79,453.290*** (11,068.900)	-81,042.020*** (11,063.640)
Year Fixed Effects	Yes	Yes	Yes	Yes
Observations	2,549	2,549	2,549	2,549
R <sup>2</sup>	0.736	0.777	0.778	0.778

*Note:*

\* p&lt;0.1; \*\* p&lt;0.05; \*\*\* p&lt;0.01

### 5.3.1 Potential Violation 1: Repression During separatist Anniversaries

First, it is possible that the Chinese government chooses to deploy additional security forces in province  $j$  during ethnic separatist anniversaries in province  $i$ , or place existing security forces in province  $j$  on high alert. This may be especially likely if the relevant ethnic minority in province  $i$  is represented non-trivially in province  $j$ . In turn, the repressive apparatus – not propaganda-based threats – may be responsible for any observed decline in popular protests.<sup>23</sup>

In practice, this is a significant concern in three sets of provinces: Xinjiang and Tibet, their neighboring provinces (Qinghai, Gansu, Sichuan, Yunnan, Ningxia, and Inner Mongolia), and Beijing. Xinjiang and Tibet are home to these ethnic separatist movements, and a range of observational evidence suggests that the government dramatically increases its security presence during these ethnic anniversaries. Accordingly, we drop them from all models.

The neighboring provinces of Qinghai, Gansu, Sichuan, and Yunnan have several Tibetan and Uygher communities along the border, and on several occasions have witnessed episodes of political instability that emanated from these ethnic separatist movements. In 2014, Xinhua News blamed Uygher separatists for a knife attack at Kunming Rail Station in Yunnan province that killed 31 and injured over 140; no group claimed responsibility. In 2008, Tibetan protests spilled into Qinghai, Sichuan, and Yunnan (NYT 2008); in 2012 Tibetans protested in Qinghai and Gansu (Greene 2012); and in 2012 and 2014, respectively, Tibetan monks self-immolated in Sichuan and Qinghai (Greene 2012; Rudolph 2014). In these provinces, the security apparatus treats these ethnic separatist anniversaries as politically sensitive. In 2017, for instance, the CCP blocked internet in the 10 western Sichuan counties that neighbor Tibet in the weeks surrounding the Tibetan Rebellion anniversary (Finney 2017). Accordingly, we also drop Qinghai, Gansu, Sichuan, and Yunnan from the sample.

We drop three other provinces from the sample as well. Ningxia contains the Hui ethnic group, which is China’s second largest Muslim population. While the Hui are ethnically, linguistically, and culturally distinct from Uyghers and do not have separatist aims, they may be more sympathetic to Uygher concerns than most Han citizens. Inner Mongolia has few Muslims but a relatively large ethnic minority community, which may be more sympathetic to Uygher concerns than most ethnic Hans. Separatist anniversaries may also have special significance in Beijing, but for different reasons. Uygher terrorists conducted bombings in Beijing in 2013, and so Beijing residents may bear some lingering frustration towards the government for permitting a security lapse or for failing to

---

<sup>23</sup>It is well known that separatist anniversaries prompt state repression *inside* Tibet and Xinjiang. For instance, shortly before the 50th anniversary of the Tibetan Rebellion, Zhang Qingli, the CCP chief in Tibet, told local riot police that “We must keep a watchful eye, and with clenched fists, constantly be on the alert. ...We must resolutely and directly strike at criminal elements who dare to stir up incidents. We must foil the separatist schemes of the Dalai clique” (FlorCruz 2009). The CCP regularly imposes a province-wide security alert, holds military parades, and bans foreign travel in Tibet in March, for the anniversary. Likewise, in 2013, one Xinjiang CCP member told reporters that “during each anniversary, he had to patrol the neighborhood or visit residents for a whole month to ensure stability” (Choi and Zuo 2013).



indemnify affected residents. Anecdotal reports suggest that anti-terrorism precautions in Beijing were heightened after the 2009 Uygher bombing in Urumqi, the capital of Xinjiang (Yang Fan 2014).



Figure 7: Map of China. Provinces with separatist movements appear in red. Our analysis excludes all provinces in blue. These include the neighboring provinces of Gansu, Qinghai, Sichuan, and Yunnan, where there are Uygher and Tibetan communities. This set also includes Ningxia and Inner Mongolia, where there are relatively large Muslim and minority communities that might be more sympathetic with separatist causes. Finally, this set includes Beijing, which might draw separatist displays because of its symbolic importance as the capital. Our analysis includes the remaining provinces, shaded in white.

The map in Figure 7 visualizes our final province sample. Tibet and Xinjiang appear in red; the seven other provinces we exclude appear in green. In 2016, these nine provinces accounted for just 11.5% of Chinese citizens. Our sample, in white, accounted for 88.5% of Chinese citizens. In these provinces, there is no evidence that the security apparatus increases its presence during our four ethnic separatist anniversaries. To test this, we ask whether the Chinese government is more likely to repress protests that emerge during ethnic separatist anniversaries than on other days of the calendar year. The basic idea is this: If the security apparatus is on high alert during these ethnic separatist anniversaries, then any protests that emerge should be repressed at a higher rate. The model is

$$\Pr(\text{Repression}_j = 1 | \text{Protest}_j = 1) = \alpha + \beta(\text{Ethnic Separatist Anniversary Window}_j) + \phi X_j + \gamma_s + \epsilon \tag{5}$$

where  $j$  indexes protest,  $\gamma_s$  gives year-level fixed effects, and the vector  $X_j$  gives a range of protest-level controls. The results appear in Table 11, and they suggest that protests that emerge during ethnic separatist anniversaries are indeed no more likely to be repressed than otherwise. Rather, the best predictor of whether a protest will be repressed is the number of protests that have emerged in the preceding week.

In short, there is no evidence that the Chinese government's repressive apparatus is on high alert *outside* the nine provinces that we exclude from analysis. As a result, we view this violation of the exclusion restriction as unlikely. Since these nine provinces account for just 11.5% of the population, this geographic restriction comes at relatively little cost in forgone population coverage.

Table 11: State Repression of Popular Protests During Ethnic Separatist Anniversaries

	<i>Dependent variable:</i>			
	Repression			
	(1)	(2)	(3)	(4)
Recent Protests	0.056** (0.022)	0.054** (0.022)	0.028 (0.029)	0.028 (0.029)
Ethnic Separatist Anniversary Window	-0.100 (0.103)	-0.091 (0.107)	-0.238* (0.140)	-0.237 (0.147)
Pro-Democracy Anniversary Window	-0.057 (0.077)	-0.057 (0.077)	0.024 (0.097)	0.022 (0.097)
Minority Province		0.006 (0.068)		0.065 (0.098)
Minority Province $\times$ Separatist Window		-0.101 (0.375)		-0.026 (0.479)
Log GRP			-8.334** (3.948)	-8.409** (3.949)
Log Population			155.476*** (57.048)	156.451*** (57.070)
Rural Population Share			-0.934*** (0.273)	-0.972*** (0.279)
Sex Ratio			0.018*** (0.007)	0.018*** (0.007)
Urban Unemployment Rate			0.047 (0.052)	0.042 (0.052)
Constant	-1.489*** (0.197)	-1.488*** (0.197)	-1,752.055*** (632.531)	-1,762.768*** (632.763)
Year Fixed Effects	Yes	Yes	No	No
Observations	9,262	9,262	5,336	5,336

*Note:*

\*  $p < 0.1$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

### 5.3.2 Potential Violation 2: Citizen Awareness of Separatist Anniversaries

We consider a second potential violation of the exclusion restriction. It may be possible that citizens in province  $j$  intentionally choose not to protest during ethnic separatist anniversaries because they do not want to be perceived as sympathetic to the cause of a separatist movement in province  $i$ .

To assess how aware Chinese citizens are of these ethnic separatist anniversaries, we fielded an online pilot survey on **Zhubajie/Witmart**, China’s largest crowdsourcing website. It was founded in 2006 and claims over 16 million users, making it 30 times larger than Amazon’s Mechanical Turk.<sup>24</sup> We recruited a total of 31 respondents living in mainland China between March 18 and 20, 2018.<sup>25</sup> Our survey was modeled after Li, Shi and Zhu (2018) and Carter, Johnston and Quek (2018). It included an introduction to the purpose of the study, participation rules, and a Qualtrics link. We then asked respondents to identify the date of 13 national holidays in a text response box. Among these, four were our separatist anniversaries. Others included typical political and cultural holidays like Labor Day, Youth Day, National Day, and Singles’ Day.<sup>26</sup> We then asked a battery of demographic questions, including sex, ethnicity, birth year, province of residence, educational attainment, occupation, religion, and party membership.

No respondents employed the same IP address, indicating that repetitive survey participation was not a problem. Nor did any respondents fail our attentiveness check embedded two-thirds of the way through the survey, a question that asked them to identify the date of New Year’s Day. We instructed respondents not to look up dates on the internet. This is crucial for the validity of our survey. If respondents performed such searches, then the results would suggest that Chinese citizens have more awareness of separatist anniversaries than they actually do. To reduce the likelihood of this possibility, we dropped 6 respondents (19.4%) who took more than 10 seconds to answer each of our 25 questions.<sup>27</sup> This group was over twice as likely to correctly identify separatist anniversaries.

The results are visualized in Figure 8, and they suggest, however tentatively, that the overwhelming majority of respondents are unaware of these ethnic separatist anniversaries. For each anniversary along the  $y$ -axis, the  $x$ -axis records how accurate was each respondent relative to the correct date; 0 days corresponds to a correct response, 364 days corresponds to a maximally incorrect response, and the right column indicates respondents who answered “Don’t know.” No respondents correctly identified the date of the Tibetan Rebellion. Just one respondent – from Inner Mongolia, among the nine provinces we exclude from our sample – correctly identified the date of the Xinjiang Uprising. One respondent correctly identified Serf Emancipation Day. Two

---

<sup>24</sup>Li, Shi and Zhu (2018) find that **Zhubajie/Witmart** workers are representative of Chinese netizens, who are more likely to be male, young, and highly educated.

<sup>25</sup>Note to readers: Yes, our pilot sample is small. The final survey will obviously be much larger.

<sup>26</sup>We did not employ lunar holidays since these have different Gregorian dates each year, making them harder to identify.

<sup>27</sup>On average, respondents took 3.6 minutes to complete the 25-question survey, a rate of 8.6 seconds per question.

respondents correctly identified Tibetan Liberation Day. The aggregate accuracy rate across anniversaries was just 4%. The aggregate accuracy rate rises to just 5% if we count as accurate errors of a week or less. By contrast, respondents easily identified typical holidays. Every respondent correctly identified Labor Day and National Day, and 90% of respondents correctly identified Youth Day.




Figure 8: Survey Data.

#### 5.4 IV Specification and Results

We regard the variation in propaganda-based threats predicted by ethnic separatist anniversaries as plausibly exogenous *outside* of the nine provinces that we exclude from our analysis. In estimating the effect of propaganda-based threats on citizen behavior, our outcome variable is the total number of protests across the country – save for the excluded provinces identified above – in the week following a threat. This accommodates the possibility that *Workers’ Daily* content may take several days to fully disseminate, or for citizens to update their beliefs following a threat.<sup>28</sup> Accordingly, our second stage model is

$$\text{Protests}_{t+1:t+7} = \alpha_P + \beta_P \left( \widehat{\text{Threat}}_t \right) + \phi_P X_t + \psi_P W_s + \gamma_{P_s} + \epsilon_P \quad (6)$$

<sup>28</sup>In the online appendix, we show that our results are not sensitive to the size of this protest window.

where  $t$  indexes day and  $s$  indexes year. The endogenous regressor  $\widehat{\text{Threat}}_t$  is the propaganda-based threat level predicted by the first stage, given by

$$\begin{aligned} \text{Threat}_t = & \alpha_T + \beta_T (\text{Ethnic Separatist Anniversary Window}_t) \\ & + \phi_T X_t + \psi_T W_s + \gamma_{Ts} + \epsilon_T \end{aligned} \tag{7}$$

In both models, the vector  $X_t$  includes two day-level covariates: whether day  $i$  falls within a pro-democracy anniversary window, and the number of protests that occurred across the country on day  $t - 1$ . The vector  $W_s$  includes year-level covariates observed at the national level: logged GRP, logged population, rural population share, sex ratio, and urban unemployment. Our  $\gamma_s$  parameters give year fixed effects.

We use as the endogenous regressor our word count measures of threats – references to “stability” and “harmony” by day – which are both generally continuous and well above 0. So too is our outcome measure of nationwide protests, from the right panel of Figure 4. Accordingly, we implement the IV estimator using two-stage least squares. In the online appendix, we employ social stability and law enforcement articles as endogenous regressors; the results are substantively unchanged.<sup>29</sup>

The results appear in Table 12. Models 1 through 3 use the daily number of references to “stability” as the endogenous regressor; Models 4 through 6 use the daily number of references to “harmony.” The bottom panel gives the first stage estimates; the top panel gives the second stage estimates. The models pass standard weak instrument tests at the 1% level and Wu-Hausman cotnsistency tests at the 10% level. The first stage results are generally consistent with those in Tables 6 and 7. Ethnic separatist anniversaries are strongly correlated with threats, as are protests on day  $t - 1$ .

In the second stage, we find strong evidence that propaganda-based threats reduce the daily rate of protest in provinces that are geographically and culturally distant from Tibet and Xinjiang. For a unit increase in the predicted number of references to “stability” and “harmony,” the number of protests across the country over the subsequent week falls by between  $-0.65$  and  $-1.6$ . These results are visualized in Figure 9. When the *Workers’ Daily* makes no reference to “stability,” the predicted number of protests nationwide is roughly 48. When it publishes the mean number of references to “stability” – roughly 9 per day – the predicted number of protests falls to roughly 40. At 16 references, or the mean plus 1 standard deviation, the predicted number of protests falls to roughly 30. We estimate similar effects for “harmony.”

These effects are substantively meaningful. To halve the number of protests over the course of the following week, China’s propaganda apparatus has to publish eight additional references to

---

<sup>29</sup>Since social stability and law enforcement articles occur so seldom, the predicted values from the first stage models are generally quite small. As a result, the coefficients on the endogenous regressors in the second stage models are implausibly large.

“stability” or “harmony” on a given day. This represents just a doubling of its baseline daily rate of 8.5 “stability” references and 8.9 “harmony” references. The government’s propaganda apparatus issues threats sparingly. When it does, we find, citizens take them seriously.




Figure 9: Simulation from IV Model.

## 5.5 Robustness Checks

We report a range of robustness checks in the online appendix. We regard two as particularly important. First, the nationalization of Uygher terrorism after 2013 makes a violation of the exclusion restriction through the citizen awareness channel somewhat more likely. That year, Uygher separatists conducted a terrorist attack in Beijing. Other scholars have confirmed that “Uygher terrorism” was virtually never discussed in the national media before 2013, but increasingly so after 2013. In the online appendix, we show that our results are robust to focusing only on the 2008-2012 period.

Second, our results are robust to dropping Guangdong, a populous southern province. A factory brawl between Han and Uygher workers in Guangdong sparked the 2009 riot in Urumqi, Xinjiang that is among our set of ethnic separatist anniversaries. Guangdong also experienced a Uygher terrorist attack in 2015. Though we find no evidence that Guangdong’s provincial security apparatus is more repressive during secessionist windows, citizens in Guangdong may be more aware of Uygher secessionist anniversaries.

## 6 Conclusion

As the rate of elite coups has declined, popular movements have emerged as a chief threat to the world’s autocrats. These popular movements coalesce into mass protests according to a well-defined calendar: on the anniversaries of historically sensitive events. Autocrats know this, and attempt

Table 12: IV Results

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Second Stage</i>						
“Stability”	-1.440** (0.588)	-1.297*** (0.389)	-0.645*** (0.224)			
“Harmony”				-1.569** (0.713)	-1.404*** (0.540)	-0.788** (0.344)
Recent Protests		5.273*** (0.857)	3.421*** (0.604)		5.475*** (1.154)	3.769*** (0.854)
Pro-Democracy Window			5.224*** (1.110)			4.295*** (1.502)
Log GDP			-586.569*** (23.132)			-674.912*** (53.339)
Log Population			6,226.620*** (749.228)			6,888.957*** (1,004.186)
Rural Population Share			-2,545.334*** (406.282)			-3,231.862*** (526.250)
Sex Ratio			12,776.340 (7,988.475)			29,023.670*** (10,015.320)
Urban Unemployment Rate			-769.645*** (263.968)			-1,426.668*** (354.847)
Constant	29.352*** (5.005)	18.899*** (5.589)	-76,415.730*** (11,300.110)	31.106*** (6.368)	17.966*** (6.819)	-97,364.010*** (15,506.310)
<i>First Stage</i>						
Separatist Window	2.241*** (0.506)	2.322*** (0.446)	3.275*** (0.549)	2.061*** (0.750)	2.142*** (0.667)	2.674*** (0.850)
Recent Protests		1.541*** (0.311)	1.555*** (0.365)		1.567*** (0.466)	1.711*** (0.566)
Pro-Democracy Window			-0.445 (0.819)			-1.548 (1.270)
Log GDP			-23.941 (16.751)			-131.684*** (25.963)
Log Population			263.361 (556.122)			1,055.917 (861.965)
Rural Population Share			428.347 (293.960)			-520.795 (455.626)
Sex Ratio			-12,613.990** (5,564.284)			10,298.850 (8,624.402)
Urban Unemployment Rate			408.652** (184.394)			-499.289* (285.803)
Constant	8.318*** (0.140)	14.117*** (0.359)	8,533.339 (8,293.738)	8.749*** (0.208)	12.410*** (0.537)	-19,602.970 (12,854.940)
Year Fixed Effects	No	Yes	Yes	No	Yes	Yes
Observations	3,287	3,287	2,556	3,287	3,287	2,556
Second Stage R <sup>2</sup>	-0.076	0.493	0.676	-0.334	0.165	0.490
First Stage R <sup>2</sup>	0.006	0.233	0.090	0.002	0.215	0.109
F Statistic	19.610***	99.317***	28.035***	7.539***	89.503***	34.691***

Note:

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01  
Excludes Tibet, Xinjiang, Gansu, Qinghai, Sichuan, Yunnan, Inner Mongolia, Ningxia, and Beijing.

to steel themselves in advance: by incarcerating opposition leaders, amplifying their censorship efforts, and engineering social media campaigns.

In this paper, we show that autocrats also issue propaganda-based threats of violence. These threats are issued sparingly, however, and, in China, timed to coincide with the anniversaries of ethnic separatist movements that occurred in Tibet and Xinjiang. Drawing on a large psychology literature, we advance a theoretical framework to explain variation – both temporal and cross-national – in how propaganda-based threats are employed. We identify three mechanisms that render threats costly: source derogation, which renders individuals less trusting of those who threaten them; the boomerang effect, which may give an event that induced a threat special salience for the threatened; and desensitization, which diminishes the force of a threat with the frequency with which it is issued. Across countries, our theory suggests, autocrats whose authority rests purely on their capacity for repression – and who, therefore, employ absurd pro-regime propaganda to signal their capacity for repression to citizens (Huang 2015; Carter and Carter 2017*b*) – are more likely to issue propaganda-based threats. We regard these cross-country determinants as an important avenue for future research.

We employ a novel identification strategy that plausibly identifies the causal effect of propaganda-based threats on the probability popular protest in Chinese provinces. We exploit two features of contemporary Chinese politics. First, regime propaganda is set at the national level, but must occasionally respond to local conditions, such as the anniversaries of ethnic separatist movements in Tibet and Xinjiang. Second, China’s ethnic diversity and sprawling size renders these ethnic separatist anniversaries salient only in certain provinces: Tibet, Xinjiang, and a handful of others. We argue that these ethnic separatist anniversaries in Tibet and Xinjiang plausibly condition the rate of protest in geographically and culturally distant provinces *only* through the propaganda-based threats that the regime issues in anticipation of them. Indeed, there is no evidence that the regime’s security apparatus *in other provinces* is on high alert during these anniversaries. Likewise, most citizens appear to be unaware of the dates on which these anniversaries occur. We find that propaganda-based threats have a plausibly causal effect on nationwide protest levels. By doubling the number of references to “stability” or “harmony,” China’s propaganda apparatus can halve the number of protests over the course of the subsequent week.

Though propaganda-based threats appear to reduce the rate of popular protest, it is nonetheless possible that threats impose other costs on repressive governments. Roberts (2018), for instance, finds that censorship emboldens China’s online activists to engage in more dissent. Might being threatened with violence compel citizens to engage in more anti-regime activity in the future? Do threats of violence issued by the government damage social capital, perhaps by increasing social distrust? As in Shih and Trinh (2018), do threats affect citizens’ perceptions of political risk, and thus their economic decisions? If explicit threats of violence reduce the probability of popular protests, then surely they have other effects as well, which scholars would do well to understand.



## References

- Amarasingam, Amarnath and Colin P. Clarke. 2017. "Terrorism Fatigue." Washington: Slate Magazine.
- Aronow, Peter M., Allison Carnegie and Nikolay Marinov. 2012. "The Effects of Aid on Rights and Governance: Evidence from a Natural Experiment." Available at SSRN: <http://ssrn.com/abstract=2199131>.
- Baek, T. H., S. Yoon and S. Kim. 2015. "When Environmental Messages Should Be Assertive: Examining the Moderating Role of Effort Investment." *Int. J. Advert* 34:135–157.
- Bartholow, Bruce D., Brad J. Bushman and Marc A. Sestir. 2006. "Chronic Violent Video Game Exposure and Desensitization to Violence: Behavioral and Event-Related Brain Potential Data." *Journal of Experimental Social Psychology* 42(4):532–539.
- Beber, Bernd, Philip Roessler and Alexandra Scacco. 2014. "Intergroup Violence and Political Attitudes: Evidence from a Dividing Sudan." *Journal of Politics* 76(3):649–665.
- Benney, Jonathan. 2016. "Weiwen at the Grassroots: China's Stability Maintenance Apparatus as a Means of Conflict Resolution." *Journal of Contemporary China* 25(99):389–405.
- Brehm, S. S. and J. W. Brehm. 1981. *Psychological Reactance: A Theory of Freedom and Control*. New York: Academic Press.
- Burgoon, M., E. Alvaro, J. Grandpre and M. Voulodakis. 2002. *The Persuasion Handbook: Developments in Theory and Practice*. Thousand Oaks: Sage chapter Revisiting the Theory of Psychological Reactance: Communicating Threats to Attitudinal Freedom.
- Cai, Yongshun. 2010. *Collective Resistance in China: Why Popular Protests Succeed or Fail*. Stanford University.
- Carnagey, Nicholas L., Craig A. Anderson and Brad J. Bushman. 2007. "The Effect of Video Game Violence on Physiological Desensitization to Real-Life Violence." *Journal of Experimental Social Psychology* 43:489–496.
- Carter, Erin Baggott and Brett L. Carter. 2017a. "Anniversary Politics: Focal Moments and Collective Action in China." University of Southern California.
- Carter, Erin Baggott and Brett L. Carter. 2017b. "Fighting for Citizens' Minds: Autocratic Propaganda and Political Institutions." Book Manuscript, University of Southern California, <http://www.erinbcarter.org>.
- Carter, Erin Baggott, Iain Johnston and Kai Quek. 2018. "Heterogenous Chinese Nationalism."

- Chen Jian. 2006. "The Tibetan Rebellion of 1959 and China's Changing Relations with India and the Soviet Union." *Journal of Cold War Studies* 8(3):54–101.
- Chen, Jidong and Yiqing Xu. 2015. "Information Manipulation and Reform in Authoritarian Regimes." *Political Science Research and Methods* 5(1):163–178.
- Chen, Jie. 2013. *A Middle Class without Democracy: Economic Growth and the Prospects for Democratization in China*. Oxford: Oxford University Press.
- Chenoweth, Erica and Maria J. Stephan. 2011. *Why Civil Resistance Works: The Strategic Logic of Nonviolent Conflict*. New York: Columbia University.
- Choi, Chi-Yuk and Mandy Zuo. 2013. "Xinjiang Calm Ahead of Anniversary of Deadly Riots." *South China Morning Post* p. July 5.
- Cialdini, Robert B. 2006. *Influence: The Psychology of Persuasion, Revised Edition*. New York: Harper Business.
- Clee, M. A. and R. A. Wicklund. 1980. "Consumer Behavior and Psychological Reactance." *J. Consum. Res* 6:389–405.
- Cline, Victor B., Roger G. Croft and Steven Courier. 1973. "Desensitization of Children to Television Violence." *Journal of Personality and Social Psychology* 27(3):360–365.
- Distelhorst, Greg and Yue Hou. 2017. "Constituency Service under Nondemocratic Rule: Evidence from China." *Journal of Politics* 79(3):1024–1040.
- Doob, Leonard W. 1950. "Goebbels' Principles of Propaganda." *The Public Opinion Quarterly* 14(3):419–442.
- Edmond, Chris. 2013. "Information Manipulation, Coordination, and Regime Change." *Review of Economic Studies* 80(4):1422–1458.
- 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):645–668.
- Enos, Ryan, Aaron Russell Kaufman and Melissa Sands. 2017. "Can Violent Protest Change Local Policy Support? Evidence from the Aftermath of the 1992 Los Angeles Riot." Harvard University.
- Finney, Richard. 2017. "Internet Blocked in Tibetan Areas of China in Run-up to Sensitive Anniversary." *Radio Free Asia* p. March 10.

- FlorCruz, Jaime. 2009. "Tensions High on Tibet Anniversary." *CNN* p. March 10.
- Freedom House. 2009. "Human Rights Abuses Intensify in Tibet Ahead of Key Anniversary." <https://freedomhouse.org/article/human-rights-abuses-intensify-tibet-ahead-key-anniversary> p. March 6.
- Freedom House. N.d. "Repression and Surveillance of Uighers Still a Concern on Second Anniversary of Urumqi Crackdown." <https://freedomhouse.org/article/repression-and-surveillance-uighurs-still-concern-second-anniversary-urumqi-crackdown>. Forthcoming.
- Gehlbach, Scott and Konstantin Sonin. 2014. "Government Control of the Media." *Journal of Public Economics* 118:163–171.
- Goemans, Hein and Nikolay Marinov. 2014. "Coups and Democracy." *British Journal of Political Science* 44:799–825.
- Greene, Scott. 2012. "Self-Immolations, Protests Continue in Western China." *China Digital Times* p. March 16.
- Grimmer, Justin and Brandon Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21(3):267–297.
- Grossman, Dave and Gloria Degaetano. 1999. *Stop Teaching Our Kids to Kill: A Call to Action Against TV, Movie, and Video Game Violence*. New York: Crown.
- Guriev, Sergei and Daniel Treisman. 2015. "How Modern Dictators Survive: An Informational Theory of the New Authoritarianism." NBER Working Paper.
- Hobbs, William and Margaret E. Roberts. 2018. "How Sudden Censorship Can Increase Access to Information." Working Paper.
- Howard, Philip N. and Muzammil M. Hussain. 2013. *Democracy's Fourth Wave? Digital Media and the Arab Spring*. New York: Oxford University.
- Huang, Haifeng. 2015. "Propaganda as Signaling." *Comparative Politics* 47(4):419–444.
- 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(2):326–343.
- King, Gary, Jennifer Pan and Margaret E. Roberts. 2017. "How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, not Engaged Argument." *American Political Science Review* . Harvard University.

- Kronrod, A., A. Grinstein and L. Wathieu. 2012. “Go green! Should Environmental Messages Be So Assertive?” *J. Mark.* 76:95–102.
- Lazarus, Richard S., Joseph C. Speisman, Arnold M. Mordkoff and Leslie A. Davison. 1962. “A Laboratory Study of Psychological Stress Produced by a Motion Picture Film.” *Psychological Monographs: General and Applied* 76(34):1–35.
- Li, Xiaojun, Weiyi Shi and Boliang Zhu. 2018. “The face of internet recruitment: Evaluating the labor markets of online crowdsourcing platforms in China.” *Research and Politics* pp. 1–8.
- Linz, Daniel, Edward Donnerstein and Steven M. Adams. 1989. “Physiological Desensitization and Judgments about Female Victims of Violence.” *Human Communication Research* 15(4):509–522.
- Little, Andrew T. 2015. “Propaganda and Credulity.” Cornell University.
- Lorentzen, Peter. 2014. “China’s Strategic Censorship.” *American Journal of Political Science* 58(2):402–414.
- Lowe, Will, Kenneth Benoit, Slava Mikhaylov and Michael Laver. 2010. “Scaling Policy Preferences from Coded Political Texts.” *Legislative Studies Quarterly* XXXVI(1):123–155.
- NYT. 2008. “Tibetan Riots Spread Outside Region.” *New York Times* p. March 16.
- Phillips, Tom. 2018. “China ‘Holding at Least 120,000 Uighers in Re-Education Camps.’” *Guardian* p. January 25.
- Rains, S. A. 2013. “The Nature of Psychological Reactance Revisited: A Meta-Analytic Review.” *Hum. Commun. Res* 39:47–73.
- Ringold, D. J. 2002. “Boomerang Effects in Response to Public Health Interventions: Some Unintended Consequences in the Alcoholic Beverage Market.” *J. Consum. Policy* 25:27–63.
- Roberts, Margaret E. 2018. *Censored: Distraction and Diversion Inside China’s Great Firewall*. Princeton: Princeton.
- Rosenfeld, Bryn. 2017. “Reevaluating the Middle-Class Protest Paradigm: A Case-Control Study of Democratic Protest Coalitions in Russia.” *American Political Science Review* 111(4):637–652.
- Rudolph, Josh. 2014. “Nun Becomes 133rd Tibetan to Self-Immolate in Protest.” *China Digital Times* p. April 1.
- Sandby-Thomas, Peter. 2011. *Legitimizing the Chinese Communist Party since Tiananmen: A Critical Analysis of the Stability Discourse*. Routledge.

- Shadmehr, Mehdi and Dan Bernhardt. 2015. "State Censorship." *American Economic Journal: Microeconomics* 7(2):280–308.
- Shih, Victor and Duy Trinh. 2018. "Protective Umbrella: How Factional Ties Attenuated the Impact of an Anti-Corruption Campaign in China." University of California, San Diego.
- Smith, M. J. 1977. "The Effects of Threats to Attitudinal Freedom as a Function of Message Quality and Initial Receiver Attitude." *Commun. Monogr.* 44:196–206.
- Steinhardt, H. Christoph. 2016. "State Behavior and the Intensification of Intellectual Criticism in China: The Social Stability Debate." *Modern China* 42(3):300–336.
- Thomas, Margaret Hanratty. 1982. "Physiological Arousal, Exposure to a Relatively Lengthy Aggressive Film, and Aggressive Behavior." *Journal of Research in Personality* 16(1):72–81.
- Thomas, Margaret Hanratty, Robert W. Horton, Elaine C. Lippincott and Ronald S. Drabman. 1977. "Desensitization to Portrayals of Real Life Aggression as a Function of Television Violence." *Journal of Personality and Social Psychology* 35(6):450–458.
- Truex, Rory. 2016. "Focal Points, Dissident Calendars, and Preemptive Repression." Princeton University, <https://static1.squarespace.com/static/5431e6ebe4b07582c93c48e3/t/56f898d92eeb8183c2c859b8/14591+Temporal+Logic+of+Repression+v+1.15.16.pdf>.
- Wallace, Jeremy L. and Jessica Chen Weiss. 2015. "The Political Geography of Nationalist Protest in China: Cities and the 2012 Anti-Japanese Protests." *China Quarterly* 222:403–429.
- Wang, Yuhua and Carl Minzner. 2015. "The Rise of the Chinese Security State." *China Quarterly* 222:339–359.
- Wardlaw, Grant. 1989. *Political Terrorism: Theory, Tactics and Counter-Measures*. 2nd ed. New York: Cambridge University.
- Waxman, Dov. 2011. "Living with Terror, not Living in Terror: The Impact of Chronic Terrorism on Israeli Society." *Perspectives on Terrorism* 5(5-6):4–26.
- Wedeen, Lisa. 1999. *Ambiguities of Domination: Politics, Rhetoric, and Symbols in Contemporary Syria*. Chicago: The University of Chicago.
- Yang, Dali L. 2017. "China's Troubled Quest for Order: Leadership, Organization and the Contradictions of the Stability Maintenance Regime." *Journal of Contemporary China* 26(103):35–53.
- Yang Fan. 2014. "Beijing Steps Up Security Ahead of Urumqi Riot Anniversary." *Radio Free Asia* p. July 3.
- Yue, Xie. 2012. "The Political Logic of Weiwen in Contemporary China." *Issues & Studies* 48(3):1–41.