

Predicting Authoritarian Crackdowns

A Machine Learning Approach

Julian TszKin Chan and Weifeng Zhong

MERCATUS WORKING PAPER

All studies in the Mercatus Working Paper series have followed a rigorous process of academic evaluation, including (except where otherwise noted) at least one double-blind peer review. Working Papers present an author's provisional findings, which, upon further consideration and revision, are likely to be republished in an academic journal. The opinions expressed in Mercatus Working Papers are the authors' and do not represent official positions of the Mercatus Center or George Mason University.



MERCATUS CENTER

George Mason University

3434 Washington Blvd., 4th Floor, Arlington, Virginia 22201

www.mercatus.org

Julian TszKin Chan and Weifeng Zhong. “Predicting Authoritarian Crackdowns: A Machine Learning Approach.” Mercatus Working Paper, Mercatus Center at George Mason University, Arlington, VA, January 2020.

Abstract

We have developed a quantitative indicator to predict if and when a series of protests in China, such as the one that began in Hong Kong in 2019, will be met with a Tiananmen-like crackdown. The indicator takes as input protest-related articles published in the *People’s Daily*—the official newspaper of the Communist Party of China. We use a set of machine learning techniques to detect the buildup in the articles of negative propaganda against the protesters, and the method generates a daily mapping between the current date in the Hong Kong protest timeline and the “as if” date in the Tiananmen protest timeline. We call this counterfactual date the Policy Change Index for Crackdown (PCI-Crackdown) for the 2019 Hong Kong protests, showing how close in time it is to the June 4, 1989, crackdown in Tiananmen Square.

JEL codes: C53, C63, D74, D83, K42, N45, P49

Keywords: policy change, machine learning, protest, crackdown, propaganda

Author Affiliation and Contact Information

Julian TszKin Chan
Managing Economist and Lead Data Scientist
Bates White Economic Consulting
julian.chan@policychangeindex.org

Weifeng Zhong
Senior Research Fellow
Mercatus Center at George Mason University
wzhong@mercatus.gmu.edu

Acknowledgments

We appreciate the helpful feedback from Tyler Cowen, Jack Goldstone, Kevin Hassett, Walter Valdivia, Stan Veuger, three anonymous peer reviewers, and seminar participants at the Federal Reserve Board, the National Bureau of Economic Research “Economics of AI” conference, and O’Reilly’s AI conference. We thank NVIDIA Co. for the generous donation of a Titan Xp, which was used for this research. The views expressed here are solely the authors’ and do not represent the views of Bates White, the Mercatus Center, or their other employees. All errors are our own.

© 2020 by Julian TszKin Chan, Weifeng Zhong, and the Mercatus Center at George Mason University

This paper can be accessed at <https://www.mercatus.org/publications/financial-markets/predicting-authoritarian-crackdowns-machine-learning-approach>

Predicting Authoritarian Crackdowns: A Machine Learning Approach

Julian TszKin Chan and Weifeng Zhong

1. Introduction

The year 2019 saw a series of protests in Hong Kong that persist to the time of this writing. The movement began as demonstrations opposing a proposed extradition bill, which, if enacted, would allow the extradition of alleged fugitive offenders in Hong Kong to mainland China. It soon turned into violent confrontations between protesters and the police. As the crisis escalated, the Chinese government ratcheted up its pressure campaign on Hong Kong, spurring speculation of a Tiananmen-like military crackdown on the semiautonomous city.¹

It is challenging to predict the occurrence of such a crackdown because, as the scholarship of political movement has long established, protest activities and the regime's response are highly interactive. To begin with, protesters need to overcome the collective action problem (as discussed in Olson 1965). As Goldstone and Tilly (2001, 187) point out, a regime also faces “considerable hazards” in dealing with protests because it is difficult to pick “the right level of concessions and repression.” A recent experimental study of a 2016 Hong Kong protest (Cantoni et al. 2019) further demonstrates the complexity of the matter. Even if researchers understood the dynamics, the limited number of Tiananmen-like crackdowns in the past—it only happened once—would restrict their ability to quantify such events.

Although fundamental socioeconomic factors that drive the protester–regime interactions may be too complex to sort out, this paper attempts to predict authoritarian crackdowns by

¹ See BBC (2019) for a summary of the 2019 Hong Kong protests.

circumventing this problem through an alternative approach: picking up early warning signs that typically occur before a crackdown takes place.

We start with the observation that propaganda is effective in moving public opinion and mobilizing societal resources.² For that reason, authoritarian regimes heavily use propaganda to prepare the public for their policies. In the days of the 1989 Tiananmen protests leading up to the June 4 crackdown, for example, China’s state-run media consistently escalated their rhetoric about the protests. “Demonstrators” and “protesters” became “rioters,” and the students who used to “have a good heart” now only wanted to “destroy the country’s future.”³ Therefore, although we may not know the political and economic conditions under which the government decides to crack down on protesters, we may still be able to foresee an imminent suppression if we can detect a significant propaganda buildup.

To detect negative propaganda buildups, we have built a machine learning algorithm that takes as input articles published in the *People’s Daily*—the official newspaper of the Communist Party of China. The algorithm first “reads” *People’s Daily* articles in 1989 that are relevant to the Tiananmen protests and learns to predict the date of publication for each sentence of each article—a metadata field already available in the dataset. Because the Tiananmen timeline leading up to the day of the crackdown is largely monotonic in tension, the date variable can be seen as a proxy for the likelihood of crackdown.⁴ Once the algorithm is trained, we deploy it to *People’s Daily* articles in 2019 that are relevant to the Hong Kong protests. Because the algorithm learns

² Communication thinkers have long argued that propaganda “works” (e.g., Lasswell 1927; Lippmann 1922). A review on more recent studies can be found in section 1.1.

³ The newspaper referred to protesting students as “demonstrators” on April 28, 1989 (*People’s Daily* 1989a), “protesters” on May 3, 1989 (*People’s Daily* 1989c), and “rioters” on May 29, 1989 (*People’s Daily* 1989f). It stated on April 29, 1989 (*People’s Daily* 1989b), that the students “have a good heart” but argued on May 17, 1989 (*People’s Daily* 1989e), that their protests would “destroy the country’s future.”

⁴ This likelihood measure, however, is not a statement of probability, which is not feasible because of the nature of rare events.

from the past, the input from 2019 is likely to be mistaken as coming from 1989. This error is exactly what we aim for: each (factual) date in the 2019 Hong Kong timeline is cast back to a counterfactual date in the Tiananmen timeline, giving a daily estimate of how close in time the 2019 Hong Kong protests are to the June 4 crackdown on Tiananmen Square.

This algorithm is similar to our previous work (Chan and Zhong 2019), in which we introduce the Policy Change Index for China (PCI-China), which attempts to predict changes in China's (national) policy priorities by detecting changes in *People's Daily's* editorial emphasis.⁵ In a similar fashion, we call the algorithm developed in this paper the Policy Change Index for Crackdown (PCI-Crackdown).

A natural question about our methodology is whether the two protest events are really comparable, because much in the (factual) world has transpired between 1989 and 2019. We argue, on the basis of our previous work on the PCI-China, that the difference in context does not prevent us from detecting China's propaganda signals effectively. In Chan and Zhong (2019), we make quarterly predictions of changes in China's policy priorities by detecting whether the *People's Daily* has changed the way front-page and non-front-page content is arranged. The PCI-China covers the period from the first quarter of 1951 to the third quarter of 2019. Although the context has changed substantially over the course of nearly seven decades, more so than from the Tiananmen protests to the Hong Kong protests, the PCI-China turns out to be consistently predicting the actual policy changes that took place in China.⁶ Because the *People's Daily* has proven to contain consistent signals on broad policy domains throughout a long horizon, we believe the same will hold true on the issue of crackdowns.

⁵ The index is based on how content is prioritized—that is, whether it is on or off the newspaper's front page.

⁶ The interested reader is referred to the appendix for a more detailed summary on the PCI-China.

Sections 2 and 3 of this study describe the data and methodology, respectively. We have also released the source code for the algorithm and will continue to provide daily updates of the indicator for as long as the 2019 Hong Kong protests last.⁷

Section 4 shows the main results. Using the algorithm, we are able to generate a daily PCI-Crackdown for the 2019 Hong Kong protests. Since the beginning of the movement, the indicator has seen significant fluctuations and has reached as late as May 26, 1989—fewer than 10 counterfactual days from a crackdown. Furthermore, despite the fluctuations, the indicator has stayed well within three weeks from the crackdown date, suggesting that, at least at the time of writing, the political crisis in Hong Kong is far from over.

Section 4 also attempts to provide a partial validation of our method, which is challenging because of the nature of rare events.⁸ First, a Tiananmen-like crackdown on Hong Kong protests by the Chinese government has never happened before, so it is infeasible to validate true positives. However, we are able to show that, when the PCI-Crackdown is close to the June 4 crackdown line, the timing coincides with circumstantial evidence showing that the Chinese government seems ready to move forward with such a crackdown.

Second, although there has not been any true positive, we are able to validate true negatives using the 2014 Hong Kong protests, also known as the Umbrella Movement. The 2014 event was a series of protests that were smaller in scale than the 2019 event, and it did not suffer a Tiananmen-like crackdown by the Chinese government. We show that, consistently, the PCI-Crackdown for the Umbrella Movement is relatively lower, and it declines over time as the protests waned and ended.

⁷ The source code of the project can be found at <https://github.com/PSLmodels/PCI-Crackdown>, and the daily update of the indicator is available at <https://policychangeindex.org>.

⁸ Had there been numerous protest episodes and had some met with military crackdowns, validation would have been easier. Unfortunately—and fortunately—that is not the case.

Although the PCI-Crackdown is specific to the context of China and may not apply to protests and crackdowns elsewhere, the methodology used to construct the indicator can be applied to some other contexts. We discuss that usage in the concluding remarks in section 5.

This paper is related to a body of literature in media economics on the effectiveness of propaganda. The scholarship of political communication has established that propaganda works, especially in countries with weak democratic institutions (e.g., Enikolopov, Petrova, and Zhuravskaya 2011; Lawson and McCann 2005). Gentzkow and Shapiro (2004) show that a broader range of information sources reduces hostility to America in Muslim countries. Yanagizawa-Drott (2014) shows that radio propaganda fueled participation in killings in the 1994 Rwandan genocide. Adena et al. (2015) find that radio propaganda helped the Nazis enroll new members and incited anti-Semitic acts. These findings lend support to our analysis: because propaganda is effective and authoritarian regimes use it heavily, we have opportunities to make inferences on regimes' actions from what they say in the propaganda.

Our study also joins a booming wave of applications that use text analysis and machine learning to examine policy problems in economics and political science.⁹ Some studies gauge US monetary policy by examining the deliberation of policymakers on the Federal Open Market Committee (FOMC), such as the formation of opinion groups in FOMC discussions (Zirn, Meusel, and Stuckenschmidt 2015), the influence of FOMC members on one another (Guo et al. 2015; Schonhardt-Bailey 2013), and the ways external communication affects internal deliberation (Hansen, McMahon, and Prat 2018). Other scholars apply similar techniques to the prediction of lawmaking. Yano, Smith, and Wilkerson (2012) developed a model to predict whether a US congressional bill will survive the committee process, and other algorithms have

⁹ Athey (2017); Gentzkow, Kelly, and Taddy (2019); Mullainathan and Spiess (2017); and Varian (2014) provide excellent surveys on this subject.

been built to predict whether a bill will be voted and enacted into law (Gerrish and Blei 2011; Kraft, Jain, and Rush 2016; Nay 2017). Text analysis and machine learning are also applied to predict court rulings by the US Supreme Court (Agrawal et al. 2017; Katz, Bommarito, and Blackman 2017; Sim, Routledge, and Smith 2016), the German Fiscal Courts (Waltl et al. 2017), and the European Court of Human Rights (Aletas et al. 2016). Finally, we developed the PCI-China, which tries to predict changes in major policy moves in China and detect changes in how the official newspaper prioritizes its content on and off the front page (Chan and Zhong 2019).

2. Data

Our dataset consists of 546 articles published in the *People's Daily* that are relevant to the following three protest episodes:


- *The 1989 Tiananmen protests*: a collection of 155 articles published after April 26, 1989, which marked the beginning of the Chinese government's hard-line stance against student protesters, up to June 4, 1989, the day of the military crackdown.¹⁰
- *The 2014 Hong Kong protests*: a collection of 77 articles published between September 26, 2014, when the first protest against the Beijing-backed electoral reform proposal occurred, and December 15, 2014, the last day of such protests.
- *The 2019 Hong Kong protests*: a collection of 314 articles published between June 9, 2019, the day of the first large protest against the Hong Kong government's extradition bill, and December 2, 2019, the day of this writing, when the protests are still ongoing.

¹⁰ As established in the literature (e.g., Sarotte 2012; Zhang, Nathan, and Link 2001), April 26 marks Deng Xiaoping's decision to respond strongly to student protesters, and the development of the tension later on was largely monotonic. This finding has important implications for the modeling of the PCI-Crackdown.

An article is considered relevant if it directly addresses the respective protests—for example, condemning protesters, praising police forces, or warning foreign governments against interference. For each article, the raw data contain the date of publication, the page the article appeared on, and the full text of its title and body.


Within each episode, the number of relevant articles and their page placement vary over time as the intensity of the movement changes. Figure 1 plots, for all three episodes, the number of relevant articles each day since the beginning of the respective movement. Like the Tiananmen protests, the 2019 Hong Kong protests receive more coverage as the movement intensifies, at times surpassing the Tiananmen level. In contrast, the coverage on the 2014 Hong Kong protests trends down as the movement loses momentum.

Figure 1. Coverage of Protests over Time



Similarly, figure 2 shows the number of those articles that appeared on the front page—the most salient space. Although it remains true that front-page coverage is more common when tension is higher, the two Hong Kong episodes have proved, unsurprisingly, much less prominent issues than were the Tiananmen protests, which took place in the most symbolic location in China’s capital.

Figure 2. Front-Page Coverage of Protests over Time



The contrast between the two Hong Kong samples and the Tiananmen sample in the two figures also illustrates a challenge in the crackdown analysis. That is, just looking at the quantity of coverage—on the front page or on any page—is not sufficient to predict crackdown because

crackdown can occur even if the issue is proportionally less prominent. In other words, one has to go beyond the quantity and examine the textual content of the coverage, to which we now turn.

3. Methodology

This section outlines the PCI-Crackdown’s methodology. Because of the small sample size, we will structure the text data at a sentence level. After training neural network models, we will aggregate the sentence-level predictions to a date level before mapping any other timeline to the Tiananmen timeline.

3.1. Data Structure

Because of the small size of the Tiananmen sample, it is infeasible to train neural network algorithms at the article level. Therefore, we segment all articles into sentences and treat each sentence as a unit of observation in training the algorithm.

At any time t , let A_t denote the set of articles published on that day. Let

$$a_{i_t,t} \in A_t \tag{1}$$

denote the i_t -th article published at time t . Let

$$s_{j_{i_t},i_t,t} \in a_{i_t,t} \tag{2}$$

denote the j_{i_t} -th sentence in the i_t -th article published at time t . To reduce notations, when there is no ambiguity, we simply write a_t as a generic time- t article and s_t as a generic time- t sentence. The raw data of each article contain its title and body. However, for simplicity, we do not treat title sentences and body sentences differently. We also do not use the page number in the analysis because of the different salience levels between Tiananmen and Hong Kong.

We use T^0 to denote the Tiananmen timeline and T' to denote any other generic timelines to map to T^0 , such as the timelines of the two Hong Kong protests. We normalize all timelines such that $t = 0$ represents the beginning of the respective event horizon.

We use stratified sampling (by article) to split the Tiananmen sample into training data and validation data, such that, for each article, 80 percent of the sentences are randomly assigned to the training set and the other 20 percent to the validation set. The training data are used to compute the optimal algorithm under a set of specifications, also known as *hyperparameters*. The validation data are then used to search for the optimal hyperparameters. The best algorithm, therefore, is the one that most closely fits the validation data, not the training data.

3.2. Model

We train a neural network algorithm \hat{f} such that, for each sentence s_t (published at time t) in the training data, the algorithm gives a prediction $\hat{f}(s_t)$ for the time of publication by minimizing the mean absolute error (MAE). That is,

$$\hat{f} \equiv \operatorname{argmin}_f \operatorname{MAE}(t, f(s_t)). \quad (3)$$

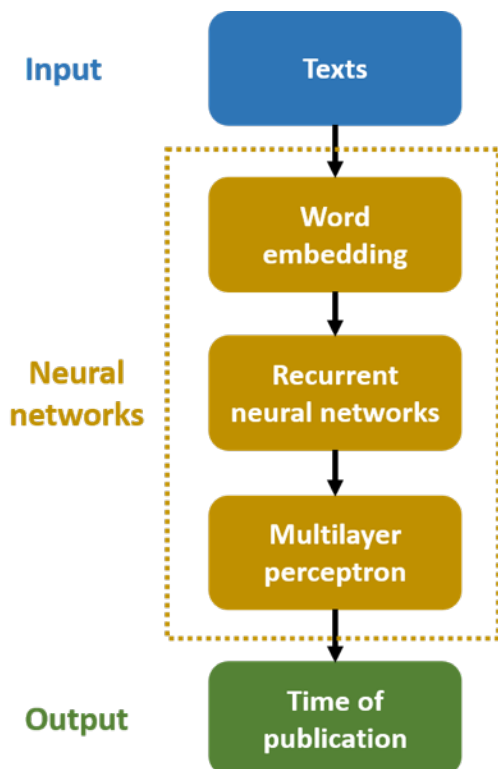
Therefore, the prediction is correct if $\hat{f}(s_t) = t$. In modeling, we treat the time of publication as a continuous variable that happens to have integer realizations in the training sample.

In this step, it is important that a monotonic relationship exist in the training data between the time of publication and the severity of the tension in the movement. As we mentioned before, this assumption is largely satisfied because the Chinese government's attitude toward student protesters was indeed increasingly more negative after April 25, 1989.

We build the structure of the algorithm \hat{f} following Chan and Zhong (2019). It consists of a sequence of neural network models, including a word-embedding layer, a recurrent neural

networks layer, and a multilayer perceptron layer. Figure 3 outlines the structure of the algorithm. From the top to the bottom of the figure, the blue, yellow, and green nodes represent the inputs, neural network models, and output, respectively. An arrow means the (interim) output of the origin node is used as the (interim) input of the destination node. The algorithm takes texts as its input and produces a date predictor as its output. Interested readers are referred to Chan and Zhong (2019) for details about the algorithm's components, which are omitted here.

Figure 3. Model Overview



The text, as a sequence of words, is first fed to a word-embedding layer—a widely used technique in natural language processing pioneered by Mikolov et al. (2013)—that maps each

word or phrase to a numeric vector. In doing so, it reduces the dimensionality of texts while preserving the semantic relationship between words.¹¹

We then feed the outputs of the word-embedding layer into a recurrent neural network, which specializes in processing sequential data (such as sentences and articles) into a vector of hidden variables.¹² In this paper, we implement a type of recurrent neural network called *gated recurrent units*, developed by Cho et al. (2014).

Finally, the hidden variables of the text are passed on to one last layer of multilayer perceptrons, a basic form of neural network model, before generating a date predictor.

Modeling the above layers requires the researcher to optimize the set of hyperparameters that control the complexity of each layer. To find the appropriate hyperparameters, we implement a simulated annealing algorithm (see Kirkpatrick, Gelatt, and Vecchi 1983) to search for the hyperparameters that optimize the classification performance of the validation data.

3.3. Constructing the PCI-Crackdown

The trained algorithm is not sufficient for mapping timelines, because \hat{f} is a sentence-level algorithm, but our goal is to create a date-to-date mapping between timelines. Therefore, we need to aggregate the algorithm's predictions from the sentence level to the date level before creating the mapping.


3.3.1. Aggregation. To put the issue in context, figure 4 plots the fitted time against the actual time for the Tiananmen sample, where each data point represents a sentence. Although the

¹¹ In this paper, we apply the Chinese-language word-embedding developed by Li et al. (2018), which is a neural network algorithm trained on all Chinese words and phrases that have appeared in the *People's Daily*—the same data source as ours—between 1946 and 2017.

¹² See Salehinejad et al. (2017) for a survey on recent advancements in this literature.

sentences generally fall along the 45-degree line, which is what the algorithm was trained to do, substantial noise occurs because not all sentences are central to an article’s meaning.

Figure 4. Fitted Tiananmen Sample



We use a two-step process to aggregate the predictions. First, we summarize each article using its highest-scoring sentences. That is, for a generic article a , we calculate the mean of the $k_1 \in \mathbb{N}$ highest scores among all its sentences, as

$$S_1(a; k_1) \equiv \max_{a' \subseteq a, |a'|=k_1} \left\{ \frac{1}{|a'|} \sum_{s \in a'} \hat{f}(s) \right\}. \quad (4)$$

The rationale for not including all sentences in an article is that, in any natural language document, not all sentences carry the same weight in delivering the meaning of the text. In the

current context, the sentences that are the harshest on the protests can best represent how strong the article’s position is about the protests.

In the second step, we summarize each date using its highest-scoring articles. Similar to the previous calculation, for the generic date t , we calculate the average of the $k_2 \in \mathbb{N}$ highest scores among all its articles, A_t , as

$$S_2(t; k_1, k_2) \equiv \max_{A' \subseteq A_t, |A'|=k_2} \left\{ \frac{1}{|A'|} \sum_{a \in A'} S_1(a; k_1) \right\}. \quad (5)$$

The rationale for not including all articles in a day is similar; the harshest articles are the most representative of the government’s attitude held on that day.

Finally, we have left unspecified the values of parameters k_1 and k_2 . They will be determined in the next subsection, which defines the PCI-Crackdown.

3.3.2. The PCI-Crackdown. The S_2 score defined in equation (5) allows us to estimate the relationship between the score $S_2(t; k_1, k_2)$ at time t and the actual time t in the Tiananmen timeline T^0 . Because the function $S_2(\cdot; k_1, k_2)$ can then be applied to any other timeline T' , any T' can be mapped back to T^0 .

To do that mapping, we first use the aggregated data $\{S_2(t; k_1, k_2), t\}_{t \in T^0}$ to estimate a locally estimated scatterplot smoothing regression model \hat{g} by minimizing the weighted mean squared error (WMSE),¹³ such that $\hat{g}(S_2)$ gives the estimated time of publication that is associated with score S_2 . That is,

$$\hat{g} \equiv \operatorname{argmin}_g \operatorname{WMSE}(t, g(S_2)). \quad (6)$$

¹³ The weight function is a tricube function.

Furthermore, we choose the values for k_1 and k_2 such that they minimize the mean standard error of g —that is, they are chosen to make g most closely fit the aggregated data. We use g^* to denote the function g that is associated with the optimized parameters k_1 and k_2 . Similarly, we use S_1^* and S_2^* to denote the optimized S_1 and S_2 scores.

Once the model g^* is estimated, we are ready to define the PCI-Crackdown, a method to map any other timeline to the Tiananmen timeline.

For any timeline T' , the PCI-Crackdown for time $t \in T'$ is given by

$$\mathcal{C}(t) \equiv g^*(S_2^*(t)). \quad (7)$$


Alternatively, the PCI-Crackdown can be written as $\mathcal{C}(t) = (g^* \circ S_2^*)(t)$, which stresses the fact that the indicator has two components: (1) aggregating the sentence predictions to a date level using S_2^* and (2) mapping the date-level score to the Tiananmen timeline using g^* .

4. Results

Before showing the main results, we visualize the S_2^* score and the estimated g^* curve for the Tiananmen sample in figure 5. Each point in the figure corresponds to one day in the Tiananmen timeline, with the vertical axis representing the actual time t and the horizontal axis representing the aggregate fitted time $S_2^*(t)$. The curve is the estimated function g^* establishing the relationship between the actual time and its associated aggregate score.

For any future protest timeline of interest, the PCI-Crackdown is obtained by first calculating the aggregate score and then finding the counterfactual date using the g^* curve in figure 5.

Figure 5. Fitted Tiananmen Sample (Aggregated)




4.1. The PCI-Crackdown for the 2019 Hong Kong Protests

Figure 6 shows the daily PCI-Crackdown for the 2019 Hong Kong protests from June 15, 2019 (the day the *People's Daily* started to talk about the movement) to December 2, 2019 (the time of this writing). Also plotted in the figure is a set of events that are relevant to the Hong Kong protests, which we will discuss in subsection 4.2.2.

Although the level of tension in the month or so leading up to the Tiananmen crackdown was largely increasing, the tension in the 2019 Hong Kong protests can and did have its ups and downs, as the figure shows. On August 5, 2019, the indicator reaches as high as May 26, 1989, which is fewer than 10 counterfactual days from a crackdown. It spikes to that level again on November 13 and 14, 2019, as violent confrontations between protesters and local police

escalated after the relatively calmer October. Moreover, while the value of the indicator fluctuates over time, it has stayed well within three weeks of the June 4 crackdown line, suggesting that the situation has been rather tense throughout the entire movement.

Figure 6. PCI-Crackdown for the 2019 Hong Kong Protests



We use an example to show that when the indicator jumps up, it indeed corresponds to an escalation of rhetoric in the articles. From August 3 to 5, 2019, the PCI-Crackdown drastically increases from May 10 to 26, 1989—a jump by more than two (counterfactual) weeks in just two days. On August 3, 2019,¹⁴ a newspaper article praising local police said that law enforcement

¹⁴ See *People's Daily* (2019).

was “professional and restrained.” Despite protesters’ violent and illegal behavior, police officers deployed only minimum force and “worked very hard” to maintain “societal order.” Similarly, on May 11, 1989,¹⁵ a newspaper article discussed how to achieve stability as student protests dragged on. The article called for “calm, rationality, restraint, and order,” saying that, despite student strikes and protests, the authorities handled the situation with restraint and “had done a lot to try to restore order.” These two articles, despite coming from two different contexts that are decades apart, share a similar, relatively measured tone.

However, on August 5, 2019, only two days later, the newspaper suddenly soured on Hong Kong. Articles published that day signaled a strong sense of urgency. They argued that the months-long protests had severely affected the entire Hong Kong economy. The retail industry was taking the hardest hit in revenues, while tourism and its related sectors were slammed by the violence as well. The top priority for Hong Kong, the newspaper said, was to immediately restore societal order. Similarly, articles published on May 26, 1989, also sounded urgent. They tied whether the student protests could be swiftly stopped to the fate of the Communist Party, the country, and the well-being of the Chinese people. Again, despite the difference in context, the algorithm links these two dates together by the similar urgency in the articles’ rhetoric.

A similar comparison between articles published on May 26, 1989, and November 13 and 14, 2019, shows the same rhetorical connection, which is omitted here to avoid repetition.

¹⁵ The *People’s Daily* did not publish any article on May 10, 1989, that was relevant to the Tiananmen protests. The May 11 article (*People’s Daily* 1989d) was the closest one.

4.2. Validation

Formally validating the predictive power of the PCI-Crackdown is challenging because there has not been any military crackdown on Hong Kong protests by the Chinese government. In this section, we attempt to tackle this issue with a partial validation.

4.2.1. Other signs of a possible crackdown. Although no data are available that would allow us to validate true positives, we establish circumstantial evidence showing that, when the PCI-Crackdown is high, it tends to coincide with signs that the Chinese military might be ready to move forward with such a crackdown, even if it has not done so yet.


Figure 6 shows a set of events related to the Hong Kong protests. As the timing demonstrates, the dates when the PCI-Crackdown reaches the peak value of May 26, 1989, coincide with the time when unusual moves by the Chinese military were observed, such as the first anti-riot drill near the mainland–Hong Kong border, an explicit warning from the Chinese military that it could reach Hong Kong within 10 minutes, the announcement of a Beijing-backed anti-mask law by the Hong Kong government, and, most recently, the sighting of Chinese soldiers on Hong Kong streets for the first time. In other words, the higher value of an indicator does seem to correspond to higher tension.

Another way to obtain partial validation is to look at how responsive the indicator is to the ups and downs of tension. For example, on October 12, 2019, the United States and China announced that they would soon sign a phase 1 trade deal, an issue that could be seen as giving the United States leverage over China to peacefully resolve the Hong Kong crisis. Consistently, the PCI-Crackdown dropped from May 19, 1989, the counterfactual date before the announcement, to May 8, 1989, the counterfactual date after the announcement.

4.2.2. *The PCI-Crackdown for the 2014 Hong Kong protests.* We are also able to validate true negatives—episodes for which the PCI-Crackdown predicts no crackdown and, consistently, no crackdown took place. The 2014 Umbrella Movement, an earlier series of protests in Hong Kong, is such an example.

Figure 7 shows the daily PCI-Crackdown for the Hong Kong protests from October 1 to December 14, 2014—the first and last days when the *People’s Daily* talked about the Umbrella Movement.

Figure 7. PCI-Crackdown for the 2014 Hong Kong Protests



In comparison to the 2019 episode, the 2014 Hong Kong protests have a generally lower PCI-Crackdown score. For the 2019 episode, the indicator is higher than May 17, 1989, about a

quarter of the time, whereas for the 2014 episode, the indicator is never higher than May 17, 1989. Moreover, the indicator's value trended downward over time, reaching May 5, 1989, in the end. Therefore, the PCI-Crackdown would have predicted a negative, which is consistent with what actually happened in 2014: the protests heated up at the beginning of the movement but soon lost momentum. The protests took an on-and-off pattern several times before ending near the year's end. Moreover, there was never any speculation of a crackdown, and none happened.

5. Conclusion

We have developed the PCI-Crackdown for the ongoing 2019 Hong Kong protests, which attempts to predict on a daily basis how close in time the movement is to a Tiananmen-like crackdown by the Chinese government. We have built this tool relying solely on the text of the *People's Daily* and a set of neural network models. The project is relevant in a practical sense because the official newspaper is still in print, and the political crisis in Hong Kong is far from over at the time of this writing. The indicator, therefore, provides a timely and effective way to monitor the evolution of the protests.

We note that the PCI-Crackdown is specific to China, because the algorithm was trained using Tiananmen protests data. Therefore, the exact model would not apply to protests and crackdowns elsewhere. However, the method of constructing the PCI-Crackdown may have applicability in some other settings. The method leverages the fact that there is a well-defined event with a monotonically increasing severity (i.e., the month leading up to the Tiananmen crackdown) and that the data label to be predicted—the time of publication in the *People's Daily*—is available in the dataset. These two characteristics have allowed us to turn the date variable into a proxy for the severity of the tension, which is abstract and otherwise difficult to

measure. The same method, therefore, may be applicable to other settings in which these two characteristics are present. We leave this interesting possibility to future research.

References

- Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa, and Ekaterina Zhuravskaya. 2015. "Radio and the Rise of the Nazis in Prewar Germany." *Quarterly Journal of Economics* 130 (4): 1885–939.
- Agrawal, Sharan, Elliott Ash, Daniel Chen, Simranjyot S. Gill, Amanpreet Singh, and Karthik Venkatesan. 2017. "Affirm or Reverse? Using Machine Learning to Help Judges Write Opinions." NBER Working Paper, National Bureau of Economic Research, Cambridge, MA, June 29.
- Aletras, Nikolaos, Dimitrios Tsarapatsanis, Daniel Preotiuc-Pietro, and Vasileios Lampos. 2016. "Predicting Judicial Decisions of the European Court of Human Rights: A Natural Language Processing Perspective." *PeerJ Computer Science* 2: e93.
- Athey, Susan. 2017. "Beyond Prediction: Using Big Data for Policy Problems." *Science* 355 (6324): 483–85.
- BBC. 2019. "The Hong Kong Protests Explained in 100 and 500 Words." November 28.
- Cantoni, Davide, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang. 2019. "Protests as Strategic Games: Experimental Evidence from Hong Kong's Antiauthoritarian Movement." *Quarterly Journal of Economics* 134 (2): 1021–77.
- Chan, Julian TszKin, and Weifeng Zhong. 2019. "Reading China: Predicting Policy Change with Machine Learning." AEI Economics Working Paper 2018-11, American Enterprise Institute, Washington, DC.
- Cho, Kyunghyun, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. "Learning Phrase Representations Using RNN Encoder–Decoder for Statistical Machine Translation." In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 1724–34.
- Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya. 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review* 101 (7): 3253–85.
- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. "Text as Data." *Journal of Economic Literature* 57 (3): 535–74.
- Gentzkow, Matthew A., and Jesse M. Shapiro. 2004. "Media, Education and Anti-Americanism in the Muslim World." *Journal of Economic Perspectives* 18 (3): 117–33.
- Gerrish, Sean, and David M. Blei. 2011. "Predicting Legislative Roll Calls from Text." In *Proceedings of the 28th International Conference on Machine Learning*, 489–96.

- Goldstone, Jack A., and Charles Tilly. 2001. "Threat (and Opportunity): Popular Action and State Response in the Dynamics of Contentious Action." In *Silence and Voice in the Study of Contentious Politics*, edited by Ronald R. Aminzade, Jack A. Goldstone, Doug McAdam, Elizabeth J. Perry, William H. Sewell, Sidney Tarrow, and Charles Tilly, 179–94. Cambridge, MA: Cambridge University Press.
- Guo, Fangjian, Charles Blundell, Hanna Wallach, and Katherine Heller. 2015. "The Bayesian Echo Chamber: Modeling Social Influence via Linguistic Accommodation." In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 315–23.
- Hansen, Stephen, Michael McMahon, and Andrea Prat. 2018. "Transparency and Deliberation within the FOMC: A Computational Linguistics Approach." *Quarterly Journal of Economics* 133 (2): 801–70.
- Katz, Daniel M., Michael J. Bommarito II, and Josh Blackman. 2017. "A General Approach for Predicting the Behavior of the Supreme Court of the United States." *PLoS One* 12 (4): 1–18.
- Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi. 1983. "Optimization by Simulated Annealing." *Science* 220 (4598): 671–80.
- Kraft, Peter, Hirsh Jain, and Alexander M. Rush. 2016. "An Embedding Model for Predicting Roll-Call Votes." In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2066–70.
- Lasswell, Harold D. 1927. *Propaganda Technique in the World War*. Oxford: Knopf.
- Lawson, Chappell, and James A. McCann. 2005. "Television News, Mexico's 2000 Elections and Media Effects in Emerging Democracies." *British Journal of Political Science* 35 (1): 1–30.
- Li, Shen, Zhe Zhao, Renfen Hu, Wensi Li, Tao Liu, and Xiaoyong Du. 2018. "Analogical Reasoning on Chinese Morphological and Semantic Relations." In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 138–43.
- Lippmann, Walter. 1922. *Public Opinion*. New York: Harcourt, Brace.
- Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. "Efficient Estimation of Word Representations in Vector Space." In *Proceedings of the International Conference on Learning Representations*, 1–12.
- Mullainathan, Sendhil, and Jann Spiess. 2017. "Machine Learning: An Applied Econometric Approach." *Journal of Economic Perspectives* 31 (2): 87–106.

- Nay, John J. 2017. "Predicting and Understanding Law-Making with Machine Learning." *PLoS One* 12 (5): 1–14.
- Olson, Mancur. 1965. *The Logic of Collective Action*. Cambridge, MA: Harvard University Press.
- People's Daily*. 1989a. "Beijing University Students Take to the Street." April 28.
- . 1989b. "Authorities Hold Talks with Capital University Students." April 29.
- . 1989c. "Government Understands Students' Patriotism, but Protests Will Not Solve Problems." May 3.
- . 1989d. "Maintaining Stability Together." May 11.
- . 1989e. "Ten University Presidents' Open Letter to Students on Hunger Strike." May 17.
- . 1989f. "Beijing Residents and Students Visit Injured Soldiers in Hospital." May 29.
- . 2019. "Professional, Restrained Police Force Confident about Defending Rule of Law and Stability." August 3.
- Salehinejad, Hojjat, Sharan Sankar, Joseph Barfett, Errol Colak, and Shahrokh Valaee. 2017. "Recent Advances in Recurrent Neural Networks." arXiv:1801.01078.
- Sarotte, M. E. 2012. "China's Fear of Contagion: Tiananmen Square and the Power of the European Example." *International Security* 37 (2): 156–82.
- Schonhardt-Bailey, Cheryl. 2013. *Deliberating American Monetary Policy: A Textual Analysis*. Cambridge, MA: MIT Press.
- Sim, Yanchuan, Bryan Routledge, and Noah A. Smith. 2016. "Friends with Motives: Using Text to Infer Influence on SCOTUS." In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 1724–33.
- Varian, Hal R. 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives* 28 (2): 3–28.
- Waltl, Bernhard, Georg Bonczek, Elena Scepankova, Jörg Landthaler, and Florian Matthes. 2017. "Predicting the Outcome of Appeal Decisions in Germany's Tax Law." In *Electronic Participation*, edited by Peter Parycek, Yannis Charalabidis, Andrei V. Chugunov, Panos Panagiotopoulos, Theresa A. Pardo, Øystein Sæbø, and Efthimios Tambouris, 89–99. Cham, Switzerland: Springer.
- Yanagizawa-Drott, David. 2014. "Propaganda and Conflict: Evidence from the Rwandan Genocide." *Quarterly Journal of Economics* 129 (4): 1947–94.

- Yano, Tae, Noah A. Smith, and John D. Wilkerson. 2012. "Textual Predictors of Bill Survival in Congressional Committees." In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 793–802.
- Zhang, Liang, Andrew J. Nathan, and E. Perry Link. 2001. *The Tiananmen Papers*. New York: Public Affairs.
- Zirn, Cäcilia, Robert Meusel, and Heiner Stuckenschmidt. 2015. "Lost in Discussion? Tracking Opinion Groups in Complex Political Discussions by the Example of the FOMC Meeting Transcriptions." In *Proceedings of Recent Advances in Natural Language Processing*, 747–53.

Appendix: Summary of the PCI-China

In our previous work (Chan and Zhong 2019), we construct the PCI-China to predict China’s major policy moves by detecting changes in the *People’s Daily* editorial emphasis. The idea is that if the Chinese government uses propaganda to prepare the public for its policies, an algorithm that can detect changes in propaganda would, effectively, predict future changes in policy.

Because editorial emphasis is an abstract concept, we proxy for it using a front-page classifier that can tell whether an article appears on the front page—the most prominent space in a newspaper—depending on the text of articles in a certain period.


We start with a quarterly rolling window, with the length of five years. For each five-year window, we build a front-page classifier to “read” the text of the articles published within that time and learn to identify front-page articles. Whatever patterns the algorithm learns in this step would constitute a fairly good understanding of the editorial emphasis during the five years in question. We then deploy the same front-page classifier to the quarter following the five-year window. If the editorial emphasis is more or less the same, the algorithm should perform just as well. But if the performance is very different, the editorial emphasis will have changed.

We define the difference in performance between the two classification tasks—the one for the five-year period and the one for the following quarter—as the PCI-China (for that particular quarter). If the indicator’s value is close to zero, the two classification tasks are performing similarly, indicating a stable editorial emphasis. In contrast, if the indicator’s value is high, that result suggests a shift in editorial emphasis.

Figure A.1 shows the PCI-China from the first quarter of 1951 to the third quarter of 2019, together with the set of events that the literature considers important to the history of the

Chinese economy.¹⁶ For an indicator to be predictive in this context, it would have to have two properties: (1) in terms of timing, it should spike before an actual policy change occurs, and (2) in terms of substance, that spike represents different classification performance between the quarter in question and the previous five years, so the content of the articles misclassified by the algorithm should be consistent with the nature of the actual policy change preceded by the spike.

Figure A.1. PCI-China and Major Events in China



Note: The PCI-China series is a predictor of policy changes. A spike in the PCI-China signals an upcoming policy change, whereas a vertical bar marks the occurrence of an actual policy change labeled by the respective event.

¹⁶ For details of the labeled events, see appendix A of Chan and Zhong (2019).

As the figure shows, spikes in the PCI-China typically precede the labeled major events. Moreover, as we have demonstrated in section 5.2 of Chan and Zhong (2019), the substance of those spikes is consistent with the policy changes they precede.