Geographical Diffusion of Protests: Evidence from China

Ye Wang (yw1576@nyu.edu)

Han Zhang (hz2@princeton.edu)

Mengyang Zhao (zoezhao@sas.upenn.edu)

3/27/2019

Abstract

Protest events often cluster in time and space, reinforcing each other to form cycles of contention. Although scholars have analyzed protest diffusion and protest cycles for decades, there is no systematic evidence on the range and speed of diffusion. Less is known about the process of protest diffusion in an authoritarian setting, which is complicated by the use of repression and cooptation. Exploiting a unique dataset of collective actions occurred in China between 2011 and 2017 and a novel estimator, we find evidence on the diffusion of violent, conventional, and disruptive protests, as well as protests organized by farmers and homeowners, even though the effect's magnitude is at most moderate. In addition, the probabilities for protests to occur concurrently in nearby cities are negatively correlated with each other. It suggests that the diffusion of protests under non-democratic regimes may be largely affected by the preventive actions taken by the government.

Introduction

The occurrence of protest movements are often highly clustered in time and space (Tarrow, 1993). This clustering effect not only results from the similar socio-political environments different protests share, but are also the product of diffusion, when one protest movement fuels another in the nearby locations.

On the one hand, no matter successful or not, protests in one place may ignite the hope for successful mobilization in the minds of people living elsewhere. In this way, protests diffuse by exporting political efficacy to the neighboring areas (Weyland, 2014). On the other hand, modern protest movements create stable repertoires that could be readily recreated in another setting, facilitating information transmission and cross-region coordination (Tilly, 1995). Overall, the closer two locations are, the less the information loss there is, and the more likely citizens in the other location will be motivated. Furthermore, proximate locations or countries tend to share similar socioeconomic background and history, thus existing repertoires can be easily replicated from one to another.

In the US, the Civil Rights Movement, the 2006 immigration reform protests, and the Occupy Wall Street are typical examples where one protest triggers other neighboring protests like wildfire (Andrews and Biggs, 2006; Vasi and Suh, 2016; Zepeda-Millán, 2017). The OWS even sparks solidarity movements in other continents. In the transnational level, although the Arab Spring rapidly descended into a winter of chaos, it still provides valuable lessons on the contagious effects of revolutions. In liberal democracies, the recent Yellow Vests Movement in France diffused to multiple countries in the Europe, again demonstrating the power of ordinary people in spreading movements. Looking back at history, revolutions of 1848 in Europe, worldwide protests in 1968, multiple waves of democratization throughout the 20th century, anti-globalization movements in 1999, and the color revolutions in post-communist Europe and Eurasia all followed similar patterns (Hardt and Negri, 2005; Weyland, 2014).

Protest diffusion constitutes one of the most crucial aspects in the mobilization, development and decline of collective actions. Many theoretical questions could be derived from the rich history of people's struggles all over the world. How persistent is the momentum created by one protest? How far could a protest transmit in space? Do multiple types of events share similar diffusion patterns? In an authoritarian context, protesters usually meet with coercive government reactions ranging from harassment to arrest, which could change or even reverse the logic of diffusion. Accordingly, does protest diffusion differ under a non-democratic setting? These answers remain unclear to social scientists, although they are of both theoretical and practical importance. In theory, protest is considered as one of the major means to constrain the ruler's power in non-democracies (De Mesquita, 2010; Fearon, 2011; Casper and Tyson, 2014), and a primary driving force of democratization (Acemoglu and Robinson, 2000, 2001; Aidt and Franck, 2015). Analyzing protest diffusion in non-democracies could provide new insights into the repression-dissent nexus (DeMeritt, 2016).

However, scholars face three major challenges in studying protest diffusion in non-democratic regimes. The first big challenge is to distinguish diffusion from common shocks, or "contextual factors" (Fowler et al., 2011), which applies to all protests regardless of contexts. That is to say, the reason for us to observe the correlation between protests in different time or locations may simply be that they are caused by the same unobservable confounders.

Second being the data issue. Authoritarian countries tend to censor protest-related information, so there is no readily available dataset to be used by independent researchers. Till now, most national level diffusion studies are based on a single protest campaign in democracies.

Thirdly, authoritarian regime's relatively unrestrained use of repressive power adds another layer of complexity. Fearing the prospect of cross-region and cross-class coordination after a social protest, the government might use preemptive measures to quell the dissatisfaction in the surrounding areas. At the same time, ordinary citizens in the nearby regions are also likely to foresee the heightened repression, strategically refraining themselves from taking to the street. As a result, we might observe the opposite phenomenon of non-diffusion of social protests, at least for events deemed as "politically sensitive".

In this paper, we try to overcome these obstacles by combining a comprehensive dataset containing daily protest events occurred in China between 2011 and 2017 with a novel method proposed by Egami (2018) to identify the diffusion effect of protests. The dataset, CASM-China, is introduced by Zhang and Pan (2018) and constructed via extracting protest-related posts from Sina Weibo, one of China's largest social media platforms. Given its vast territory and intensifying conflict between the global capital, central/local government and ordinary citizens, China offers an ideal setting for studying the diffusion and non-diffusion of multiple types of protests. The dataset includes 142,427 collective action events in total, ranging from worker's strikes to the violate attack of local governments, and covers most area of China.

The method adopted by this paper is developed by Egami (2018) under the framework of stationary directed acyclic graphs (DAGs). The author investigates conditions under which the diffusion effect of an outcome variable in a social network can be identified and proposes a placebo test to examine the validity of the conditions. He also suggests a debiased estimator to obtain the correct estimate when the test fails. Applying the method to the CASM-China dataset, we find that the direct estimate of the diffusion effect— the impact of protest occurrence in neighboring areas one week ago— is weak and insignificant across various types of protest. Nevertheless, the placebo test implies that the estimate is likely to be biased and driven by "negative common shocks", meaning that the probabilities for protests to appear concurrently in contiguous areas are negatively correlated with each other. After the bias is corrected, our results indicate that protests with a particular form (violent, conventional, or disruptive) or organized by specific groups (farmers and homeowners) do diffuse in the geographical space.

In the next section, we briefly summarize the related theoretical literature. Section three introduces our dataset and section four demonstrates the method. Section five presents preliminary results and section six concludes.

Theory

Over the years, scholars on protest diffusion have discovered a wide array of factors contributing to the domestic and international diffusion of protests.

First, grand historical processes give rise to modern forms of social movements that have the ability to be diffused. In the widely acclaimed book *Popular Contention in Great Britain, 1758-1834*, Tilly (1995) documents the gradual transformation of British mass politics, depicting a vivid picture of the rise of new social movements and the accelerating pace of protest diffusion. Due to urbanization and migration, the concentration of capital and the proletarianization of the workforce, increase in state capacity, and the emergence of private associations, parochial and particular protests were replaced by more cosmopolitan and modular ones. Study on protests in the Mid-Qing China confirms the strong influence of state capacity on protest intensity and repertoires. Chinese protesters were able to combine traditional repertories with more modern schemas from the West, creating a unique hybrid protest culture that are activated in multiple historical periods (Hung, 2013).

Geographical proximity, although not the prerequisite, greatly increases the probability of diffusion. Muiznieks (1995) delineated how popular movements in Baltic states catalyzed unrest in Russia. Beissinger (2007) adopts the framework of modular political phenomena to study democratic revolutions in the post-communist region, showing that previous contention sets a standard for later mobilization in other localities with similar political institutions, histories, and cultural affinities. Analyzing the diffusion of xenophobic violence in German, Braun and Koopmans (2009) further confirms that the effect of geographical proximity on diffusion is mediated by social similarity. Gleditsch and Rivera

(2017) reveals that diffusion of nonviolent campaigns are mostly confined to neighboring countries, and the diffusion effect is larger for nondemocratic regimes.

Meso level factors also help explain the diffusion processes. First being the social networks. Using 1871 Paris Commune as an example, Gould (1991) demonstrates that the cross district uprising benefited from the intersection of informal neighborhood ties and the organizational network of the Paris National Guard. In a study of the Temperance movement during the 1870s, scholars find that railroad and telegraph-mediated information about neighboring protest events were main driving forces of the protest diffusion (**?**). Analyzing public protests in the US from 1965-1995, Wang and Soule (2012) indicate that collaboration between social movement organizations (SMOs) constitutes an important platform for tactical diffusion. Second, the role of news media is well recognized. Study on the sit-ins in the US South indicates that movement activists and news media are both vital channels in orchestrating protest Andrews and Biggs (2006). Biggs (2013) argues that mass media help disseminate the protest tactics to a broader audience, initiating the international diffusion of suicide protest. More recent studies on Occupy Wall Street discover that Facebook and Twitter activities greatly facilitate protest diffusion by mediating the effect of spatially proximate protests (Vasi and Suh, 2016), and also by turning repression into online activism which further instigates the offline occupy (Suh, Vasi and Chang, 2017).

Scholars have also explored the micro foundations of protest diffusion. Kuran (1991) has argued that one protest may result in the formation "revolutionary bandwagon" by reducing "preference falsification" among citizens. Kuran claims that, after witnessing the action of some hard-core dissidents, ordinary citizens may feel less pressure to reveal their real political attitude, and join the dissidents to create a larger protest. Weyland (2014) introduces social psychology factors into the diffusion framework. He points out that the diffusion speed of protests is negatively correlated with their probability of being successful: As a society becomes more complicated, collective actions will be mobilized via large organizations such as parties instead of small clubs. The existence of these organizations will reduce the misjudgment of individuals and force them to think twice before action. That is why the Arab Spring diffused less quickly than the 1848 Revolution, but was far more successful in toppling dictators.

Our paper builds upon the works mentioned above, and try to quantitatively investigate rules that govern the geographical diffusion of different types of protests in China. We pay particular attention to protest events organized by workers, farmers and homeowners, which are the three major types of contentious actions in contemporary China. We also differentiate between conventional, disruptive and violent protests, which could generate disparate diffusion patterns.

Data

Conventional approaches to detect protest events often rely on media reports such as newspapers, which renders them less reliable when the context of interest is a non-democratic regime. Since mass protests consist one of the biggest threats to the incumbent, most related information will be blocked and censored by the government, leading to severe selection bias in the available data. The emergence of social media platforms offers researchers new possibilities.

In Zhang and Pan (2018), the authors introduce a seminal system, Collective Action from Social Media (CASM), to extract protest-related contents from posts on social media platforms. The system exploits deep learning algorithms, including convolutional neural network (CNN) and the combined convolutional and recurrent neural network with long short-term memory (CNN-RNN), to isolate posts about real protests from all the posts with relevant keywords. Applying the system to Sina Weibo, one of China's most popular social media platforms, the authors construct CASM-China, a comprehensive dataset on protest events occurred in China between 2011 and 2017.

The authors use Wickedonna, an existing dataset on protests in China, as their source of keywords and training set. Wickedonna records 67,502 protest events described by 240,521 text-based posts and 233,288 images and videos. They were manually collected on a daily basis from various social media platforms (mainly Weibo) by two activists Yuyu Lu and Tingyu Li between 2013 and 2016. Based on the 50 most frequently occurring words in Wickedonna, the authors scraped a vast corpus with approximately 9.5 million posts from Weibo. Posts about real protests are then identified in two stages.

In stage one, the deep learning algorithms are trained on texts and images in Wickedonna and applied to search for posts of interest in the Weibo corpus. The trained model correctly classify 63% of posts as collective action (precision) and correctly identify 79% of collective action posts (recall) on average in out-of-sample validation. The relatively high false positive rate is driven by the misclassification of many posts expressing grievance into protest events. To further improve the algorithms's performance, the authors hire native Chinese speakers in the second stage to hand code 40,505 posts and retrain the model. The ultimate classifier reaches an accuracy rate of more than 90% in out-of-sample validation. 283,427 posts are finally selected by the model, among which 142,427 unique events with time and location are identified. The dataset covers 96% of all the counties in China. Figure 1 (Page 34, Zhang and Pan, 2019) below shows the distribution of the logged aggregated number of protests across Chinese prefectures. Darker color means more protest events. Prefectures without any protest concentrate in ethnic minority regions such Tibet and Xinjiang.

To further validate the reliability of the method, the authors compare the result with other datasets on collective actions in China, such as the Global Database of Events, Language, and Tone (GDELT), the Integrated Conflict Early Warning System (ICEWS), and WiseNews (protests reported by newspapers). It turns out that CASM-China covers more events than any of them, especially those in the rural area. But ethnic and religious conflicts are underrepresented in CASM-China, since the Internet is more strictly regulated in Tibet and Xinjiang. They also show that only a small faction of posts on protests (5.4%)

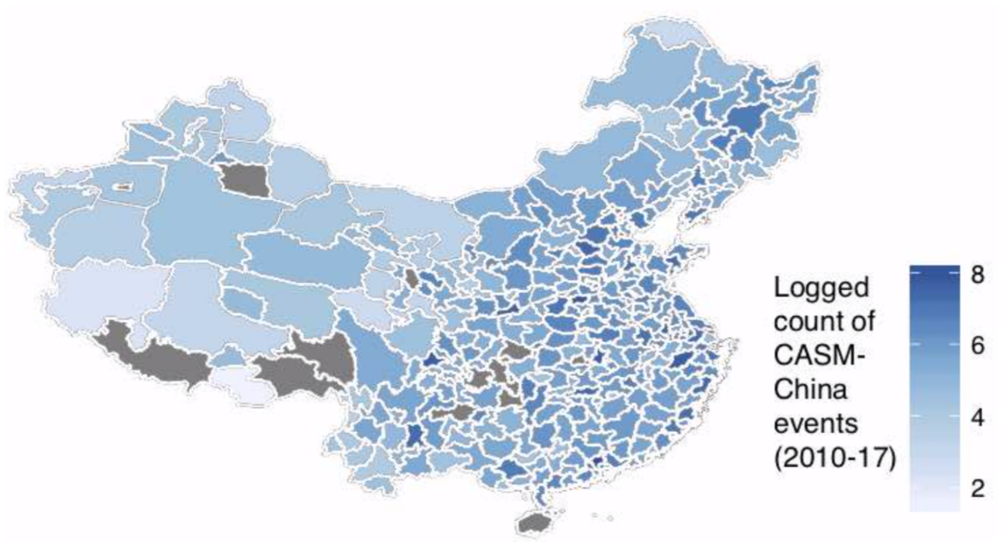

Figure 1: Log count of CASM-China events by prefecture

will be finally censored by the platform, as most posts do not spread virally online. For our purpose, CASM-China has been the best data source available so far.

Events in the dataset can be classified by either their forms or main issues mentioned in the posts. The authors code three mutually exclusive forms of protest: violent, conventional, and disruptive. Violent protests refer to those involving explicit physical conflicts such as attack local officials. Conventional protests consist of more common repertoires, worker's strikes, public demonstrations and petitions included. Disruptive protests are mainly occupy movements and power-cutting actions. The percentages of three forms in the dataset are 24%, 39%, and 37%, respectively. There are also 11 issues coded by the authors, from medical dispute to complaints of veterans. For simplicity, we re-classify all the events based on the main participants and keep only protests organized by workers, farmers, and homeowners. That leaves us with 91,779 events in total, among which 45.6% are organized by workers, 33.4% by farmers, and 37.5 by homeowners.

Method

Identifying the contagion/diffusion effect of specific outcomes (e.g. disease, technology adoption, protest, etc.) in social networks is a well known challenge in statistics and social sciences. Some scholars even believe that it cannot be identified due to the prevalence of homophily and common shocks (Angrist, 2014). The former means that nodes with similar attributes tend to cluster in a network. If there exists a positive relationship between some attribute and the outcome, we may mistake the outcome's clustering for its contagion. The latter indicates that the outcomes of nearby nodes are often affected jointly by the same unobservable factors. Then what we deem as contagion is just bias driven by omitted variables. Yet, to what extent homophily and common shocks confound the contagion effect? Is there any way to calibrate their influences and correct the bias they bring to the estimate? These questions have not been sufficiently addressed in the literature.

The recent paper by Egami (2018) fills up this lacuna under the framework of directed acyclic graphs or DAGs (Pearl, 2009). He first discusses conditions under which the identification of the outcome's contagion effect becomes possible, and proposes a placebo test for the validity of these conditions. A bias-correction estimator is then developed to obtain the correct estimate when the test fails.

Egami assumes that all the relationships between variables of interest could be described by a stationary DAG, in which the subgraph in period t has the same structure as the subgraph in any other period t′. Intuitively, the assumption suggests that the outcome is always affected by the same set of variables— observable or not— even though the value of these variables can change over periods. This is implicitly assumed in most methods for TSCS data estimation, such as the two-way fixed effects models or synthetic control. Figure 2 below displays a stationary DAG with two units and three periods (page 18 of Egami, 2018). Observed characteristics are omitted in the graph. Suppose we are interested in the causal effect of Y_{11} on Y_{22}. The estimation based on selection-on-observables assumption (e.g. IPW) will be unbiased when 1. observed variables are properly controlled, and 2. both arrows emanating from G_t (common shocks) and leading to W (homophily) are absent.

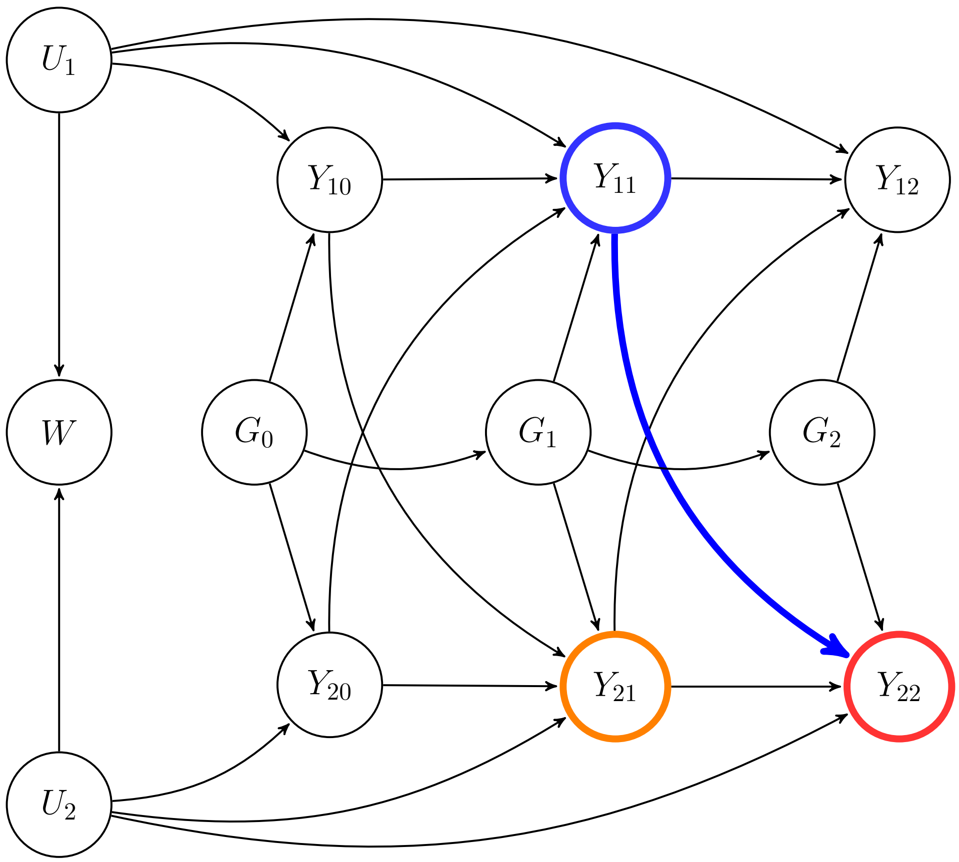

Figure 2: A stationary causal DAG

When either homophily or common shocks exist, there will be a "back-door path" between Y_{11} and Y_{22} (e.g. $Y_{11} \leftarrow G_1 \rightarrow G_2 \rightarrow Y_{22}$) as the source of bias. Meanwhile, Y_{11} and Y_{21} will be connected by the same variables on the path due to the stationary structure of the graph. In other words, if the relationship between Y_{11} and Y_{22} is confounded, so will be the relationship between Y_{11} and Y_{21}. Egami thus suggests that we may test whether the selection-on-observables assumption holds by repeating the analysis with Y_{21} as the dependent variable (the covariates set needs to be changed accordingly). The result from this place test not only indicates the identification assumption's validity, but also provides an estimate of the bias's magnitude. If we further assume that the contagion effect is stable over time, we can obtain a debiased estimate by subtracting Y_{11}'s coefficient in the placebo test from its coefficient in the main analysis.

Following the roadmap proposed by Egami (2018), our analysis proceeds in the next steps. First, we construct a network of Chinese prefectures in which two prefectures are connected if they are contiguous geographically. As we have illustrated, nearby administrative units are similar in culture and socio-economic background, which lowers the cost of protest diffusion. New protesters may easily adopt slogans and strategies from previous protests. Xinjiang and Tibet are dropped from the network as the CASM doesn't include enough ethnic and religious conflicts. We denote the adjacency matrix derived from the network as $\mathbf{W}_{N\times N}$, and the number of protests happened in week t across prefectures as $\mathbf{Y}_t = (Y_{1t}, Y_{2t}, \ldots, Y_{Nt})$. Then the number of protests occurred in neighboring cities in the same week will be $\mathbf{D}_t = \mathbf{W}\mathbf{Y}_{t-1}$. Second, we estimate the following regression model as the main analysis:

$$Y_{it} = \mu + \beta D_{i,t-1} + \gamma X_{it} + \alpha_i + \zeta_t + \varepsilon_{it}$$

where Y_{it} is the outcome variable, the number of protests in prefecture i, week t; $D_{i,t-1}$ is the treatment variable, the total number of protests appeared in neighboring cities one week ago; X_{it} represents time-varying covariates, including lagged dependent variables; α_i and ζ_t are prefecture and week fixed effects, respectively; ε_{it} is the idiosyncratic error term. We assume the linearity of X_{it} and constant treatment effect. The estimated β will capture the diffusion effect when assumptions are valid. For robustness, we also try other ways of coding the variables. For example, both Y_{it} and D_{it} are recoded as 1 when their values are positive and 0 otherwise .

Third, a similar specification will be used for the placebo test:

$$Y_{i,t-1} = \mu + \beta_1 D_{i,t-1} + \beta_2 D_{i,t-2} + \gamma_1 X_{it} + \gamma_2 X_{i,t-1} + \alpha_i + \zeta_t + \varepsilon_{it}$$

Here $Y_{i,t-1}$ becomes the outcome variable and we control for the one-period lag of both D_{it} and X_{it}. If the estimate of β_1 is significant, we may suspect that the identification assumptions have been violated by homophily or common shocks. Then we conduct step four, and obtain the bias-correction estimate $\widehat{\beta^*} = \hat{\beta} - \hat{\beta_1}$. A conservative estimate of $\widehat{\beta^*}$'s standard error is $\sqrt{Var(\hat{\beta}) + Var(\hat{\beta_1})}$. We repeat these steps on different categories of protests in our dataset, including different participants (workers, farmers, homeowners) and different forms (violent, conventional, disruptive). The results are presented in the next section.

Result

Regression estimates

We present results from the main analysis, the placebo test, and the debiased estimator successively in the three graphs of Figure 6. The solid dots represent the point estimates and segments mark 95% confidence intervals. It is clear from the first graph that there exists no significant relationship between the number of protests in prefecture i, week t and the number of protests in its neighboring prefectures in week $t-1$. Moreover, the magnitude of all the estimates is small. For the full sample and those organized by workers, the coefficient is even negative.

Results from the placebo test suggest the existence of bias. Even with the lagged outcome and the lagged treatment being controlled, the number of protests in prefecture i is still negatively correlated with the number of protests occurred over the same period in neighboring prefectures. The estimate is significant across all the categories. As the position of prefectures is pre-fixed in our geographical network, homophily is unlikely the driving force of the bias. Therefore, "negative common shocks" on nearby areas could be a potential explanation. One possibility is that local governments in China respond fast to collective actions happened in neighboring prefectures. Preventive repression or conciliation may be implemented as soon as possible such that local citizens would not copy the action of their neighbors.

One perspective to understand the bias correction is to compare it with the mediation analysis. The total impact of a protest can be decomposed into two parts: the direct one on other protests, which is the effect of interest, and the indirect one that evokes the government's response. Bias correction allows us to subtract the indirect effect from the total effect, such that the diffusion effect can be calculated. After correcting the bias, we obtain significantly positive estimates for the diffusion effect of protests with any particular form (violent, conventional, or disruptive). The same is observed on protests organized by farmers or homeowners. Yet the coefficient for the protest of workers, one of the most severe challenges to the regime, remains insignificant. This finding stands in contrast to arguments in classic theories (e.g. Tilly (1995)) and hints that the dynamics of protest diffusion can have a quite different pattern under non-democratic regimes. The other coefficients are still small in magnitude (see Table 1 below). For example, when the number of occurred protests by homeowners in neighboring prefectures increases by 10, we should expect the number of protests by homeowners to rise by 0.05 in the current prefecture the next week. Suppose there is such an increase for every week in a year, the total number of protests happened in the current location should rise by 52 * 0.05 = 2.6 for the whole year. At most the influence is moderate.

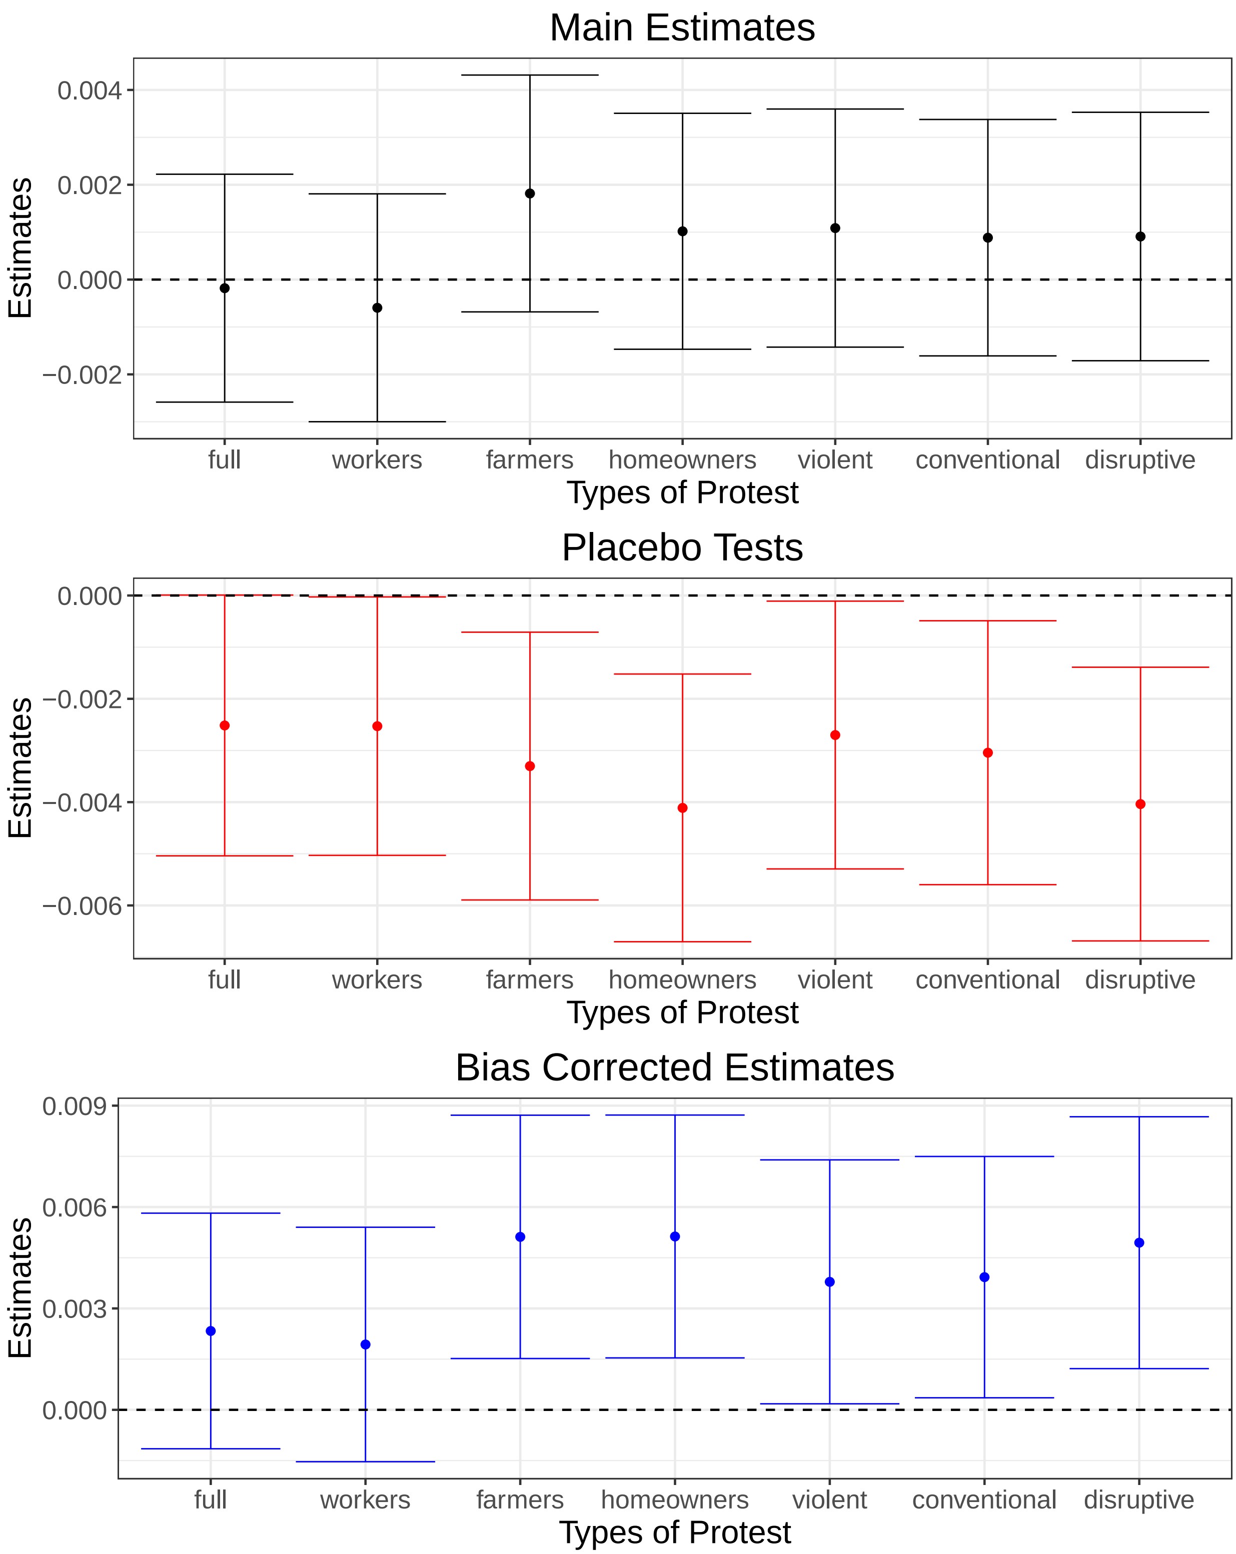

Figure 3: Results of the main analysis, placebo tests, and bias corrected estimates

Table 1: Regression coefficients

	full	*workers*	*farmers*	*homeowners*	*violent*	*conventional*	*disruptive*
Main analysis	-0.0002	-0.0006	0.0018	0.0010	0.0011	0.009	0.009
	(0.0012)	(0.0012)	(0.013)	(0.0013)	(0.013)	(0.013)	(0.013)
Placebo test	-0.0025^{*}	-0.0025^{**}	-0.0033^{***}	-0.0041^{***}	-0.0027^{**}	-0.0030^{***}	-0.0040^{***}
	(0.0013)	(0.0013)	(0.013)	(0.0013)	(0.013)	(0.013)	(0.014)
Bias corrected	0.0023	0.0019	0.0051^{***}	0.0051^{***}	0.0038^{**}	0.0039^{**}	0.0049^{***}
	(0.0018)	(0.0018)	(0.018)	(0.0018)	(0.018)	(0.018)	(0.019)
Prefecture FE	Y	Y	Y	Y	Y	Y	Y
Week FE	Y	Y	Y	Y	Y	Y	Y

Notes: *p<0.1; **p<0.05; ***p<0.01.

Heterogeneity of the estimates

(In progress)

Robustness

(In progress)

Conclusion

Exploring a large dataset containing daily protest events happening in China from 2011 to 2017, we analyze the geographical diffusion of six types of protests using two dimensions: participants and repertoires. Contrary to the common expectation of protests clustering geographically, we discover the opposite pattern that protest intensity in the neighboring cities in the same week are negatively related. This provides implications for the existence of preemptive state repression in China, when the government quickly react to existing protests to prevent the occurrence of similar events.

We use the method suggested by Egami (2018) to control for the mediating effects of the potential state intervention. After the adjustment, we find moderate evidence of protest diffusion for five of the six protest types we examined. The only protest type that lacks a contagion pattern is labor movement. While this seems to contradict previous literature on the strong mobilization tradition of workers in other countries (Nam, 2006), it still confirms the uniqueness of labor unrest in China. It is possible that workers are under a different type of surveillance. As suggested by Fu (2017), labor NGOs and workers may also deliberately repackage collective actions into personal grievances, leading to an underestimation of workers' protests.

In the next steps, we will add covariates to examine the robustness of our findings, and also use another method to estimate the cross province variations in government repression.

References

Acemoglu, Daron and James A Robinson. 2000. "Why did the West extend the franchise? Democracy, inequality, and growth in historical perspective." *The Quarterly Journal of Economics* 115(4):1167–1199.

Acemoglu, Daron and James A Robinson. 2001. "A theory of political transitions." *American Economic Review* pp. 938–963.

Aidt, Toke S and Raphaël Franck. 2015. "Democratization under the threat of revolution: Evidence from the Great Reform Act of 1832." *Econometrica* 83(2):505–547.

Andrews, Kenneth T and Michael Biggs. 2006. "The dynamics of protest diffusion: Movement organizations, social networks, and news media in the 1960 sit-ins." *American sociological review* 71(5):752–777.

Angrist, Joshua D. 2014. "The perils of peer effects." *Labour Economics* 30:98–108.

Beissinger, Mark R. 2007. "Structure and example in modular political phenomena: The diffusion of bulldozer/rose/orange/tulip revolutions." *Perspectives on Politics* 5(2):259–276.

Biggs, Michael. 2013. "How repertoires evolve: the diffusion of suicide protest in the twentieth century." *Mobilization: An International Quarterly* 18(4):407–428.

Braun, Robert and Ruud Koopmans. 2009. "The diffusion of ethnic violence in Germany: The role of social similarity." *European Sociological Review* 26(1):111–123.

Casper, Brett Allen and Scott A Tyson. 2014. "Popular Protest and Elite Coordination in a Coup d'état." *The Journal of Politics* 76(2):548–564.

De Mesquita, Ethan Bueno. 2010. "Regime change and revolutionary entrepreneurs." *American Political Science Review* pp. 446–466.

DeMeritt, Jacqueline HR. 2016. "The strategic use of state repression and political violence.".

Egami, Naoki. 2018. "Identification of Causal Diffusion Effects Using Stationary Causal Directed Acyclic Graphs." *arXiv preprint arXiv:1810.07858* .

Fearon, James D. 2011. "Self-enforcing democracy." *The Quarterly Journal of Economics* 126(4):1661–1708.

Fowler, James H, Michael T Heaney, David W Nickerson, John F Padgett and Betsy Sinclair. 2011. "Causality in political networks." *American Politics Research* 39(2):437–480.

Fu, Diana. 2017. *Mobilizing without the masses: control and contention in China*. Cambridge University Press.

Gleditsch, Kristian S and Mauricio Rivera. 2017. "The diffusion of nonviolent campaigns." *Journal of Conflict Resolution* 61(5):1120–1145.

Gould, Roger V. 1991. "Multiple networks and mobilization in the Paris Commune, 1871." *American Sociological Review* pp. 716–729.

Hardt, Michael and Antonio Negri. 2005. *Multitude: War and democracy in the age of empire*. Penguin.

Hung, Ho-fung. 2013. *Protest with Chinese characteristics: Demonstrations, riots, and petitions in the Mid-Qing dynasty*. Columbia University Press.

Kuran, Timur. 1991. "Now out of never: The element of surprise in the East European revolution of 1989." *World politics* 44(01):7–48.

Muiznieks, Nils R. 1995. "The influence of the Baltic popular movements on the process of Soviet disintegration." *Europe-Asia Studies* 47(1):3–25.

Nam, Taehyun. 2006. "The Broken Promises of Democracy: Protest-Repression Dynamics in Korea 1990-1991." *Mobilization: An International Quarterly* 11(4):427–442.

Pearl, Judea. 2009. *Causality*. Cambridge university press.

Suh, Chan S, Ion Bogdan Vasi and Paul Y Chang. 2017. "How social media matter: Repression and the diffusion of the Occupy Wall Street movement." *Social Science Research* .

Tarrow, Sidney. 1993. "Cycles of collective action: Between moments of madness and the repertoire of contention." *Social Science History* 17(2):281–307.

Tilly, Charles. 1995. *Popular Contention in Great Britain, 1758-1834*. Harvard University Press.

Vasi, Ion Bogdan and Chan S Suh. 2016. "Online Activities, Spatial Proximity, and the Diffusion of the Occupy Wall Street Movement in the United States." *Mobilization: An International Quarterly* 21(2):139–154.

Wang, Dan J and Sarah A Soule. 2012. "Social movement organizational collaboration: Networks of learning and the diffusion of protest tactics, 1960–1995." *American Journal of Sociology* 117(6):1674–1722.

Weyland, Kurt. 2014. *Making Waves: Democratic Contention in europe and Latin america since the revolutions of 1848*. Cambridge University Press.

Zepeda-Millán, Chris. 2017. *Latino mass mobilization: Immigration, racialization, and activism*. Cambridge University Press.

Zhang, Han and Jennifer Pan. 2018. "CASM: A Deep-Learning Approach for Identifying Collective Action Events with Text and Image Data from Social Media." *URL: https://www. princeton. edu/hz2/files/protest method. pdf* .