Appendices

A Variable constructions

- Geolocation: We used the locations identified in CASM. CASM provides locations at the county level if available, and otherwise determines them at the prefecture-city level.
- Time: We used the timestamps of the Weibo posts as the measure of time (i.e., the date). We note that the posts might occasionally describe events that occurred prior to the creation of the post. In the matching procedures discussed below, we conducted sensitivity analyses by extending the time frame, which led to similar results.
- Followers, followees, posts: We included the average number of followers of an account, the number of other users this account is following, and the posts (from the time of account registration until June 2020). The data were obtained by querying the official Application Programming Interface (API) of Weibo.
- Issues: Issues were classified into the following categories: land/rural protests; unpaid wages; homeowner property; fraud/scams; environmental; pension/welfare; taxi drivers; medical; education; veterans. If a protest contained at least one key word describing these issues, it was classified as belonging to the category corresponding to that word. An event could in theory cut across more than one issue area. This procedure was similar to that used in the CASM project, except that slightly different dictionary was used to represent each category. The revised dictionary has reduced the proportion of protests with difficult-to-identify issues from 25 per cent in CASM to 13 per cent in this research. Details are provided in Appendix C.
- Size: The size variable refers to the number of participants in each protest. We used the state-of-the-art crowd counting algorithm, CSRnet, from the computer vision literature to estimate the size of crowds from images.[1] There are two advantages of using images to estimate protest sizes. In CASM-China, about 83 per cent of events have attached images, and most images were from pictures taken at the scene, either by protesters or nearby bystanders who witnessed the protest. In contrast, less than 10 per cent of posts contain a description of the size of the protest.[2] Moreover, descriptions of protest size from social media texts are often inaccurate, while the estimates from images have proven to be as accurate as journalists' reports.[3]

1. Li, Zhang and Chen 2018.

2. The calculation was performed based on a random sample of 3,461 events. Human coders were then instructed to judge whether they could draw a size estimate from the text or images of each event. Based on searching keywords related to quantifiers (e.g., "several hundred"), Goebel collected a list of offline protest events discussed on Weibo and found that they could infer the size of only 10.93% of the protests. See Goebel and Steinhardt 2019. See Goebel and Steinhardt 2019.

3. The reason text-based estimates of size are less accurate is that the majority of texts only provide a rough account of the size of the protest. If a post mentions that several hundred people participated in a protest, only a range estimate of the size (e.g., 100 to 1000 people) can be made, which is not as accurate as image-based estimates which count the exact number of individuals in the picture. See Sobolev et al. 2020.

- Mention of state: This variable was constructed based on mentions of government-related words. Specifically, we collected a list of proper nouns indicating government agencies and officials and checked whether the posts by individuals, media, or government mentioned at least one of these words. More mentions of state-related words indicated willingness to discuss the government. Note that this variable differs from the target variable above. For instance, a government post about protests may mention local government but at the same time may exhibit a more positive tone.

- Police presence: Machine predictions were used to quantify police presence at protests. See Appendix D for details.

- Action forms: Three categories of forms of action were established: peaceful, disruptive, and violent. Dictionaries of violent and disruptive events were constructed using methods similar to those used by the CASM project. A protest was placed in the "violent" category if it contained any of the keywords in that category; a protest was placed in the "disruptive" category if it contained any keywords in that category while containing no keywords in the "violent" category; the remaining events were placed in the "peaceful" category. Details are provided in Appendix C.

- Sentiment: Three commonly used Chinese sentiment dictionaries were concatenated: the National Taiwan University Semantic Dictionary,[4] the Tsinghua-Sentiment-Dictionary,[5] and HowNet Dictionary.[6] This combined dictionary was used to generate sentiment scores using dictionary methods.

B Recognizing government and media accounts

We used both official verification status and usernames to distinguish between the five types of actors on social media as listed in Figure 1. First, for each verified account, Weibo further provide a category label, such as government, media, influencers (famous individuals including self-media), and other organizations (like schools or NGOs).[7] The official verification status was accurate (i.e. lower Type I error), but it suffers from high Type II error. In other word, when Weibo says that an official account is verified and the verification category is a government account or a media account, this information is precise. However, we noted there were not a smaller number of apparent government or news media accounts that had not been verified. We suspect that it is because the owners of these accounts did not want to spend extra time to provide credentials with the Weibo company.

Therefore, we further generated a list of proper nouns used frequently in government, government news media, and other commercial media accounts' usernames. Then we used the following rules to assign Weibo accounts into each of the category below:

- government account: if the verification category indicated it belongs to a government agency, or its username contained at least one phrase in the following list of government

4. http://nlg.csie.ntu.edu.tw/achievements.php#tools
5. http://nlp.csai.tsinghua.edu.cn/
6. Dong and Dong 2003.
7. https://open.weibo.com/wiki/2/users/show. See the "verified_type" field.

names: 司法, 公安, 法院, 交警, 政府, 工会, 政协, 城管, 政务, 高院, 司法, 派出所, 公安, 法院, 交警, 政府, 工会, 派出所, 政务, 高院, 管委会, 中院, 检察, 党支部, 团支部, 政协, 党委, 便民, 警方, 法官, 司法, 公安, 法院, 交警, 政府, 工会, 政协, 城管, 政务, 高院, 司法, 派出所, 公安, 法院, 交警, 政府, 工会, 派出所, 政务, 高院, 管委会, 中院, 检察, 党支部, 团支部, 政协, 党委, 便民, 警方, 执法, 税务, 环保局, 环保部, 监察, 消防, 社保.

- government (news) media: if the verification category indicated it belongs to media, and its username contained at least one phrase in the following list of government news media names: 新华社, 人民日报，环球时报，日报，早报, 晚报，央广，央视，法制报，法院报，司法报, 公安报, 广播, 电视.
- self-media: if the verification category indicated it belongs to *influencers*
- other commercial media: if the verification category indicated it belongs to *media*, or its username contained at least in the list of media names but it was not classified as belonging to either government news media or self-media category defined above. The list of media names we used is: 播报, 宣传, 热线, 交通, 在线, 新闻, 快讯, 直播, 关注, 百态, 掌中, 广视, 那点事, 官微, 交通台, 同城, 交通, 电台, 新鲜事儿, 咨询, 服务, 中心, 乐居, 聚焦, 你好, 发布, 天平, 晚报, 微博, 播报, 宣传, 热线, 日报, 直播, 电视, 广视, 交通台, 交通, 咨询, 服务, 中心, 聚焦, 发布, 人民网, 教育局, 电台, 官方微博, 法律, 城管, 新浪, 新京报, 市场, 民生, 在线, 杂志, 乐居, , 环球市场播报, 播报, 文匯報, 新消息报, 新浪视频, 消防, 荆楚网, 搜房网, 车事, 观察, 輿情, 事事通, 今天, 都市, 楼市, 身边, 草根, 大城小事, 号外, 微生活,V 生活, 纵横, 资讯, 传媒, 事儿, 那些事, 门户.
- individuals: any other accounts.

This approach should distinguish government, self-media, and other commercial media with reasonably high accuracy. The limitations are still Type II error: some self-media are not verified by Weibo as influencers but are still popular Weibo accounts followed by many and posted local news contents regularly. Nevertheless, we find that the above approach is the best solution we could find now.

C Measure issues and action forms

We slightly modified the dictionaries used in CASM-China to classify the issues of protests. This has allowed us to reduce the proportion of posts for which we cannot identify issues or forms. High-quality dictionaries, especially ones that are tailored for a specific corpus such as ours, are not easy to create.

We used state-of-the-art algorithms to generate corpus-specific dictionaries that describe issues and forms of protests.[8] It starts with a set of seed words, and then finds the most similar words (e.g., top 50) for each one. The similarity is then calculated based on the famous word2vec model. Human experts are involved in each step. Applying the algorithm to our setting, we used the original words for each category in CASM as the seed words. We then found the most similar 50 words for each seed word. We used the word2vec model, specifically trained in CASM-China, for 10 million Weibo posts containing protest-related words and another randomly sampled 10

8. Hamilton et al. 2016.

million Weibo posts. Last, we discussed which words to keep and discard. 3897 words were selected for the issue dictionary, and 1269 words were used for the action form dictionary. Given space limitations, we have put the exact word lists on the authors's website.

After we obtained dictionaries for issues and action forms, we used the following dictionary count methods to produce measures for the issue and action form of each protest, and their corresponding reports by the media.

- Issue types: an event was assigned to an issue type if the text of the post contained keywords for that issue type. To reduce false matching, events were only categorized into an issue type if two or more keywords were mentioned.
- Action forms: an event was categorized as one of three action forms (peaceful, disruptive, or violent), if it mentioned three or more keywords for that action form. If an event mentioned keywords from multiple action forms, we took the most violent action form for further analysis.

For issues and action forms, we hired a research assistant to code a random sample of 500 tweets. For issues, the overall false positive rate is 0.027, and the overall false negative rate is 0.099, suggesting a very good performance. Broken down by issues, the false positive rate and false negative rate are respectively:

- Land grabs: 0.036 and 0.070
- Unpaid wages: 0.034 and 0.09
- Homeowners: 0.053 and 0.057
- Fraud/scams: 0.004 and 0.347
- Environmental: 0.012 and 0.095
- Pension/welfare: 0.008 and 0.222
- Taxi drivers: 0.008 and 0
- Medical: 0.075 and 0.091
- Education: 0.024 and 0.026

It can be seen that false positive rates of all issues are below 0.075, suggesting that the Type I error is pretty small. False negative rates of almost all issues were also small, with the exception of protests on commercial frauds and on pension. It suggests that we may have missed some protests on that two particular issues.

For the three action forms, the false positive and false negative rates are respectively:

- peaceful: 0.101 and 0.154
- disruptive: 0.026 and 0.164
- violent: 0.245 and 0.015

It can be seen that the error rates are also in an acceptable range.

D Performances of the Machine Learning Algorithms

A sample of 10,000 events were randomly selected from the entire CASM-China dataset, and research assistants were instructed to code the police presence and targets of each event. Then a Support Vector Machine was used to train machine learning classifiers for these 10,000 events and the trained model was applied to estimate variables for the rest of the events. The parameters were selected using 5-fold cross-validation.

The machine predictions were highly accurate. The average AUC of ROC of the algorithm predicting targets of protests is 0.91.[9] The average AUC of ROC of the algorithm predicting police presence is 0.97.

9. AUC of the ROC curve is a widely used metric to evaluate machine learning's performance. A value of 0.8 is usually considered acceptable and 0.9 is considerable highly accurate.

E Additional results

Figure E.1 shows the same content of Table 6 through a subway graph. The y-axis indicates the ranking, with each word connected by a line. If there was high similarity between the rankings, the line would be flat. Conversely, a line with a big tilting angle would indicate different rankings, which represent distinct reporting styles. For instance, the word police (警察) was the most frequent word used by individuals, and was the third most frequent word used by the news media. However, it was the 21st most frequent word used by government Weibo accounts.

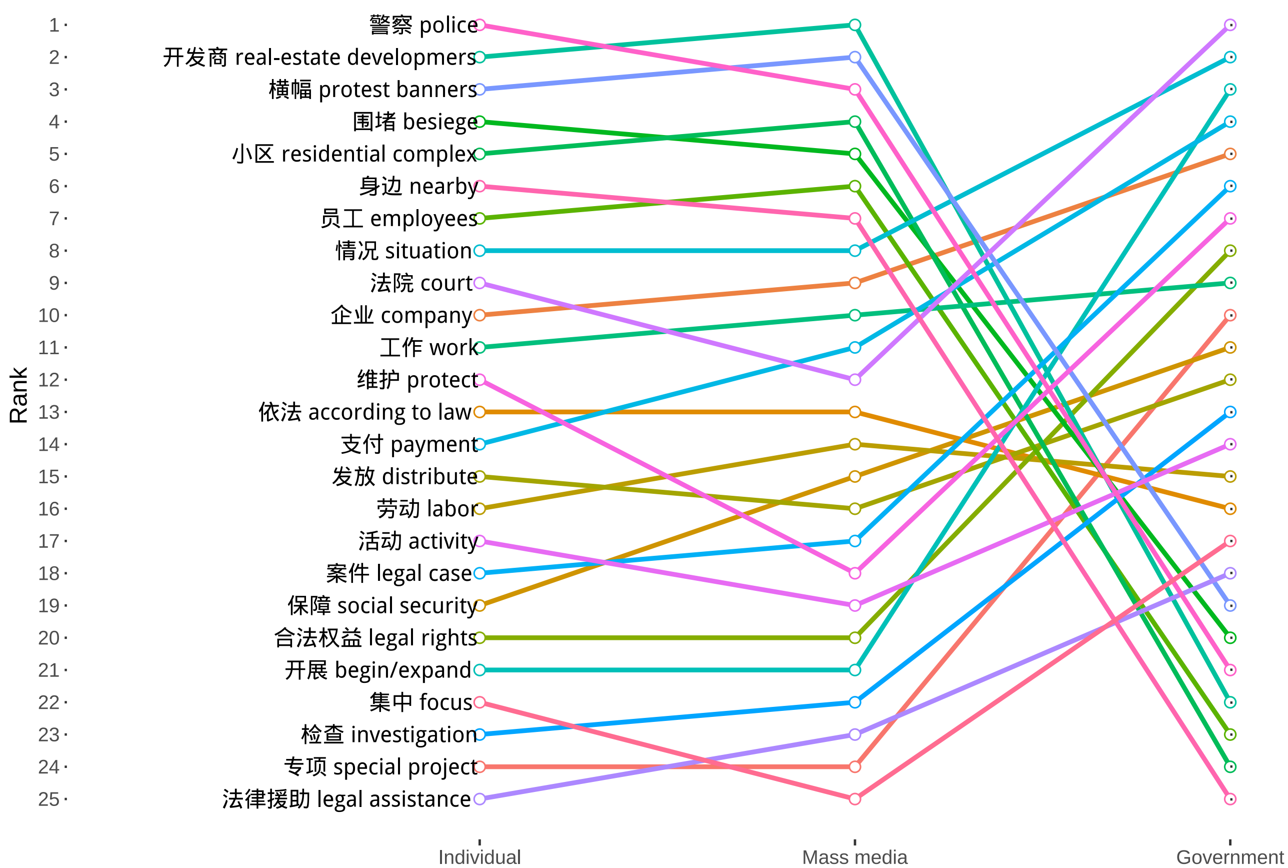

Figure E.1: Subway plot showing frequent words used by individuals, the news media, and the government.

Table E.1: Probability of reporting a protest event by different types of media accounts. Two-way fixed effect regression at province and year level with clustered standard errors at province level.

Model	News Media Logit	News media Quasi-Poisson	Government Logit	Government Quasi-Poisson
Number of posts (log)	0.0525**	0.0585**	0.0028	0.0028
	(0.0187)	(0.0162)	(0.0402)	(0.0402)
Number of followers (log)	0.2538***	0.1995***	0.1892***	0.1892***
	(0.0130)	(0.0114)	(0.0283)	(0.0283)
Number of followees (log)	-0.0823**	-0.0572*	-0.0466	-0.0466
	(0.0285)	(0.0235)	(0.0504)	(0.0504)
Verified	-0.0600	-0.0051	-0.4592**	-0.4592**
	(0.0624)	(0.0505)	(0.1409)	(0.1409)
Size (log)	0.1050***	0.0876***	0.1480	0.1480
	(0.0160)	(0.0132)	(0.0816)	(0.0816)
Issue:				
Land grabs	0.1333*	0.1079	-0.1904	-0.1904
	(0.0677)	(0.0555)	(0.3173)	(0.3173)
Unpaid wages	0.3509***	0.2847***	0.8723***	0.8723***
	(0.0571)	(0.0476)	(0.1808)	(0.1808)
Homeowners	-0.0299	-0.0183	-0.7943**	-0.7943**
	(0.0599)	(0.0503)	(0.2476)	(0.2476)
Frauds	-0.1172*	-0.0931	-0.3373	-0.3373
	(0.0569)	(0.0480)	(0.2496)	(0.2496)
Environment	0.0426	0.0380	0.8310**	0.8310**
	(0.0904)	(0.0768)	(0.2642)	(0.2642)
Pension	0.1487	0.1124	-0.1160	-0.1160
	(0.0795)	(0.0633)	(0.2966)	(0.2966)
Taxi	0.1091	0.0780	-0.1684	-0.1684
	(0.0614)	(0.0483)	(0.2930)	(0.2930)
Medical	0.0462	0.0408	0.1895	0.1895
	(0.0565)	(0.0453)	(0.1767)	(0.1767)
Education	0.0512	0.0356	0.2233	0.2233
	(0.0744)	(0.0603)	(0.2095)	(0.2095)
Veteran	-0.0203	-0.0109	-13.09***	-13.09***
	(0.2179)	(0.1796)	(0.1303)	(0.1303)
Police presence	0.1636**	0.1358*	0.3179	0.3179
	(0.0589)	(0.0494)	(0.1759)	(0.1759)
Action Form:				
Disruptive	0.0361	0.0384	-0.5850*	-0.5850*
	(0.0746)	(0.0607)	(0.2162)	(0.2162)
Violent	-0.0799	-0.0641	-0.2341	-0.2341
	(0.0616)	(0.0515)	(0.1236)	(0.1236)
Target:				
Company; Gov as Mediator	-0.0369	-0.0404	0.0402	0.0402
	(0.0538)	(0.0443)	(0.2152)	(0.2152)
Government	-0.2099**	-0.1868**	-0.1257	-0.1257
	(0.0646)	(0.0549)	(0.1636)	(0.1636)
Sentiment	0.1164	0.1071	0.6728	0.6728
	(0.1498)	(0.1269)	(0.5482)	(0.5482)
Fixed-effects				
province	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes
Observations	54,720	54,720	54,010	54,010

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Table E.2: Probability of reporting a protest event. A dummy variable indicating whether a protest has multiple issues is added as an additional regressor.

Model	Logit	Quasi-Poisson	Logit	Quasi-Poisson
Number of posts (log)	0.0497**	0.0562***	-0.0030	-0.0012
	(0.0178)	(0.0153)	(0.0387)	(0.0378)
Number of followers (log)	0.2557***	0.2007***	0.2048***	0.1971***
	(0.0125)	(0.0111)	(0.0293)	(0.0279)
Number of fans (log)	-0.0812**	-0.0561*	-0.0530	-0.0504
	(0.0287)	(0.0238)	(0.0543)	(0.0532)
Verified	-0.0730	-0.0146	-0.5413***	-0.5176**
	(0.0649)	(0.0529)	(0.1460)	(0.1419)
Size (log)	0.0987***	0.0820***	0.1557	0.1514
	(0.0157)	(0.0129)	(0.0829)	(0.0799)
Multiple issues	-0.0546	-0.0388	-0.1291	-0.1251
	(0.0637)	(0.0511)	(0.2241)	(0.2179)
Issue:				
Land grabs	0.1606*	0.1265	-0.0888	-0.0906
	(0.0818)	(0.0659)	(0.3274)	(0.3239)
Unpaid wages	0.3852***	0.3081***	0.9688***	0.9446***
	(0.0607)	(0.0499)	(0.1860)	(0.1802)
Homeowners	0.0041	0.0059	-0.7530**	-0.7442**
	(0.0588)	(0.0488)	(0.2708)	(0.2673)
Frauds	-0.1003	-0.0825	-0.2756	-0.2681
	(0.0634)	(0.0531)	(0.2777)	(0.2706)
Environment	0.0872	0.0712	0.9879**	0.9664**
	(0.0951)	(0.0794)	(0.3055)	(0.2918)
Pension	0.1928*	0.1429*	-0.0886	-0.0869
	(0.0830)	(0.0657)	(0.3802)	(0.3701)
Taxi	0.1429	0.0990	0.0083	0.0056
	(0.0746)	(0.0584)	(0.2599)	(0.2501)
Medical	0.0508	0.0407	0.2249	0.2204
	(0.0727)	(0.0581)	(0.2122)	(0.2049)
Education	0.1152	0.0824	0.3675	0.3571
	(0.0921)	(0.0733)	(0.2559)	(0.2490)
Veteran	0.1326	0.1106	-12.31***	-13.03***
	(0.1841)	(0.1478)	(0.1398)	(0.1339)
Police presence	0.1368*	0.1140*	0.2999	0.2925
	(0.0576)	(0.0481)	(0.1977)	(0.1931)
Action Form:				
Disruptive	0.0167	0.0240	-0.6674**	-0.6516**
	(0.0748)	(0.0611)	(0.2148)	(0.2109)
Violent	-0.0664	-0.0534	-0.2212	-0.2160
	(0.0582)	(0.0483)	(0.1211)	(0.1172)
Target:				
Company; Gov as Mediator	-0.0169	-0.0241	0.0693	0.0630
	(0.0521)	(0.0430)	(0.2341)	(0.2267)
Government	-0.1982***	-0.1762**	-0.1684	-0.1673
	(0.0599)	(0.0506)	(0.1709)	(0.1665)
Sentiment	0.1623	0.1473	0.6766	0.6649
	(0.1354)	(0.1148)	(0.5699)	(0.5585)
Fixed-effects				
province	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes
Observations	56,376	56,376	55,992	55,992

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

We also add some other examples. We list five exemplary posts by news media accounts:

1. 杭州 16 时将就垃圾焚烧项目相关情况进行说明昨天有部分反对九峰垃圾焚烧项目的群众在杭州余杭中泰乡及附近地区聚集其中有小部分聚集人员封堵杭徽高速公路及 02 省道并出现打砸车辆今天下午 4 点杭州市将在市政府召开新闻发布会就中泰事件相关情况进行说明详细网页链接
2. 媒体称杭州 5 千人聚集抗议垃圾焚烧厂项目杭州余杭区拟建垃圾焚烧厂引发了周围居民的担忧至昨日上午 9 时许在焚烧厂建造地附近已聚集 5 千多人目击者称聚集市民与警方发生冲突一些人受伤余杭官方表态在项目未履行完法定程序和征得大家支持的情况下一定不开工详细网页链接
3. 杭州市政府召开发布会就九峰垃圾焚烧厂项目及 5 月 10 日余杭中泰及附近地区出现聚集堵路打砸事件的情况进行通报
4. 杭州 5 千人聚集抗议垃圾焚烧厂项目官方未征得大家支持不开工至昨日上午 9 时许在焚烧厂建造地附近已聚集 5 千多人目击者称聚集市民与警方发生冲突有人受伤余杭官方表态在项目未履行完法定程序和征得大家支持的情况下一定不开工
5. 9 张图告诉你杭州为何要建垃圾焚烧发电厂目前杭州市公示了九峰垃圾焚烧厂项目规划在余杭区中泰乡新建一个垃圾焚烧发电厂此举遭到附近居民聚集抗议那么为什么杭州要新建垃圾焚烧厂为什么选择了中泰面对垃圾围城的困局普通群众又该怎么办一张图告诉你

The below are two posts by government accounts:

微警讯杭州公布聚集事件违法犯罪嫌疑人照片第二批 5 月 10 日浙江余杭中泰及附近地区出现聚集堵路打砸事件据余杭公安先后公布两批有关该聚集堵路打砸事件的犯罪嫌疑人名单现敦促当事人及时向警方投案自首参看下图群众如有发现请确保自身安全后向公安机关举报电话 110

余杭中泰垃圾焚烧项目杭州市政府召开新闻发布会没有死亡鼓励揭发 1 没有网上传播的死亡事件 2 现场出现打砸围殴等违法犯罪行为 3 现场群众服从民警劝告后自行离去 4 对投案自首的违法犯罪人员依法予以从轻减轻或免除处罚 5 鼓励知情群众揭发他人犯罪行为

References

Dong, Zhendong and Qiang Dong. 2003. HowNet - a Hybrid Language and Knowledge Resource. In International Conference on Natural Language Processing and Knowledge Engineering, 2003. Proceedings. 2003, 820–824.

Goebel, Christian and H. Steinhardt. 2019. “Better Coverage, Less Bias: Using Social Media to Measure Protest in Authoritarian Regimes”.

Hamilton, William L., Kevin Clark, Jure Leskovec and Dan Jurafsky. 2016. Inducing Domain-Specific Sentiment Lexicons from Unlabeled Corpora. In Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, 595–605. Association for Computational Linguistics.

Li, Yuhong, Xiaofan Zhang and Deming Chen. 2018. CSRNet: Dilated Convolutional Neural Networks for Understanding the Highly Congested Scenes. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1091–1100.

Sobolev, Anton, M. Keith Chen, Jungseock Joo and Zachary C. Steinert-Threlkeld. 2020. "News and Geolocated Social Media Accurately Measure Protest Size Variation". *American Political Science Review,* 1–9.