

Persistent Political Engagement: Social Interactions and the Dynamics of Protest Movements[†]

By LEONARDO BURSZTYN, DAVIDE CANTONI, DAVID Y. YANG, NOAM YUCHTMAN,
AND Y. JANE ZHANG*

We study the causes of sustained participation in political movements. To identify the persistent effect of protest participation, we randomly indirectly incentivize Hong Kong university students into participation in an antiauthoritarian protest. To identify the role of social networks, we randomize this treatment's intensity across major-cohort cells. We find that incentives to attend one protest within a political movement increase subsequent protest attendance but only when a sufficient fraction of an individual's social network is also incentivized to attend the initial protest. One-time mobilization shocks have dynamic consequences, with mobilization at the social network level important for sustained political engagement. (JEL D72, D74, I23, Z13)

Protests demanding political rights have been a critical driver of economic, social, and political change for centuries (e.g., Acemoglu and Robinson 2012, 2019; Aidt and Franck 2015). While dramatic one-shot events capture public attention (e.g., the Hungarian Revolution of 1956 or Tiananmen Square in 1989), political rights have historically often arisen from successful, long-running *movements*: sequences of events in which sustained political engagement is important. Historically important instances include the women's suffrage movements, the US civil rights movement, or the anti-apartheid movement in South Africa.

*Bursztyn: University of Chicago and NBER (email: bursztyn@uchicago.edu); Cantoni: Ludwig-Maximilians-Universität Munich, CEPR, and CESifo (email: cantoni@lmu.de); Yang: Harvard University and NBER (email: davidyang@fas.harvard.edu); Yuchtman: LSE, CEPR, and CESifo (email: n.yuchtman@lse.ac.uk); Zhang: University of New South Wales (email: jane.zhang@unsw.edu.au). Rohini Pande was coeditor for this article. Helpful and much appreciated suggestions, critiques, and encouragement were provided by Dan Berkowitz, Andrei Shleifer, seminar participants at Amsterdam, Cambridge, Columbia, Harvard, Heidelberg, IIES, Imperial, INSEAD, Mannheim, Northeastern, NUS, Oxford, Pompeu Fabra, Sussex, UC Berkeley, Yale, Zurich, and conference participants (ASSA, NBER SI, SITE, Bruneck, Munich). Raymond Han, Moritz Leitner, Jackson Li, Glen Ng, Aakaash Rao, Vanessa Sticher, and Meggy Wan provided excellent research assistance. Cantoni acknowledges financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 716837). The research described in this article was approved by the University of Munich IRB, Protocol 2017-04; by the Stanford University IRB, Protocol 38481; and by the University of California-Berkeley Committee for Protection of Human Subjects, Protocol ID 2015-05-7571. For further details on IRB approvals, please refer to online Appendix B. The experiment is registered on the AEA RCT registry with ID AEARCTR-0002301.

[†]Go to <https://doi.org/10.1257/aeri.20200261> to visit the article page for additional materials and author disclosure statement(s).

Political movements have received an enormous amount of attention from across the social sciences (e.g., Tilly 1978, Skocpol 1979, McAdam 1982, Goldstone 1991, Kuran 1997). Existing work has argued for the importance of individuals' sustained engagement, working through social structures (Hirschman 1984, McAdam 1986, Tarrow 2005).¹ Up to now, however, there does not exist well-identified—that is, experimental or quasi-experimental—evidence on the causes of individuals' sustained participation in political movements.

Our paper provides the first such evidence. First, we identify the persistent effect of one's own protest participation by randomly indirectly incentivizing Hong Kong university students into participation in an antiauthoritarian protest. We do so by paying subjects for providing us with information about protest crowd size; we thus do not pay for protest turnout per se but behavior conditional on turnout. This allows us to distinguish state dependence—the possibility that participation in one protest causally affects subsequent participation—from serial correlation in preferences.

We next test whether participation by one's *social network* plays a causal role in shaping one's own persistent protest participation by randomizing the intensity of this treatment across major-cohort cells. Finally, we explore mechanisms through which changed social interactions may generate persistent participation. Given that protests are inherently *group* events (see, e.g., Passarelli and Tabellini 2017), we consider changes in friendships, which among other things can affect the social utility arising from protest participation or reduce coordination costs. We also examine changes in preferences and beliefs (beliefs about the world or beliefs about others), as these are central in many models of protest participation (e.g., Tullock 1971, Bueno de Mesquita 2010, Shadmehr and Bernhardt 2011, Edmond 2013, and Barberà and Jackson 2020).

Our context is Hong Kong's ongoing antiauthoritarian movement, demanding political rights from the Chinese Communist Party (CCP).² We study participation in the July 1 marches, yearly protests that represent an important component of Hong Kong's ongoing antiauthoritarian movement. We study the 2017 and 2018 marches: peaceful, modestly sized protests of around 50,000 citizens, aiming to both achieve policy concessions and signal the strength of the movement.

Our study faces a crucial identification challenge: we need to observe *both* exogenous protest participation at the individual level and *independent* exogenous variation in the protest participation of an individual's social network. We design a field experiment to overcome this challenge, leveraging our online surveys with students at Hong Kong University of Science and Technology (HKUST; see Cantoni et al. 2016, 2019). The experiment involves two dimensions of randomization. First, at the individual level, we randomly assign subjects to a condition in which they are indirectly incentivized to attend the 2017 march. Second, to generate exogenous variation in protest participation at the social network level, our design also randomly varies the proportion of treated individuals across major \times cohort cells at

¹ Studying participants in the 1964 Mississippi Freedom Summer Project, McAdam (1986, p. 88) writes that “a prior history of activism and integration into supportive networks acts as the structural ‘pull’ that encourages the individual to make good on his strongly held beliefs.”

² We thus contribute to a growing empirical literature on the political economy of popular dissent in the Greater China region, e.g., Lorentzen (2013); Qin, Strömberg, and Wu (2017); King, Pan, and Roberts (2013); and Zhang (2016).

0 percent, 1 percent, 50 percent, or 75 percent treated. Importantly, these are the only two dimensions of randomization implemented, and both are pre-registered.

We find two main results. First, individual incentives lead to an immediate (2017) increase in protest turnout, and this effect does not vary with how many others in an individual's social network receive incentives. Second, protest participation remains persistently (in 2018) higher but only among treated individuals who are initially treated along with at least 50 percent of their major \times cohort cell. Thus, sustained participation in a political movement is *not* a result of self-selection and serially correlated preferences alone but is to some extent state dependent. In addition, social networks play a crucial role in this state dependence. These results have important implications for the evolution of political movements: a one-time mobilization shock will have *dynamic* consequences, with mobilization of social networks playing an important role in producing sustained political engagement.

We consider several mechanisms through which changed social interactions may produce the persistent protest participation we find among treated subjects in high treatment intensity cells. We begin by presenting evidence that treated subjects in high treatment intensity cells form significantly more new and stronger friendships with people who are politically active—this could directly increase the social utility from protest participation and also increase turnout through other channels. Consistent with changed social interactions reducing coordination costs, we find that treated subjects in high treatment intensity cells are by a large margin the most likely to convert their protest plans into action. We next examine changes in individuals' political preferences and beliefs; while noisy estimates mean we cannot rule out some role for these channels, we do not find compelling evidence that they drive the persistent protest participation we observe among treated subjects in high treatment intensity cells.

Our results contribute to a growing empirical literature on the determinants of protest participation. Much of this work studies individuals' participation in mass movements as a *one-shot* action and thus cannot shed light on the causes of persistent political engagement by individuals (e.g., Enikolopov, Makarin, and Petrova 2020; Manacorda and Tesei 2020; González 2020; Cantoni et al. 2019; Hager et al. 2019a, b). Other work (in particular, Madestam et al. 2013, on the Tea Party protests) identifies the spatial persistence of protests but cannot isolate individual-level persistent behavior or identify its causes. We are able to unpack persistence that has been observed in the aggregate, identifying individual-level persistent behavior that depends also on the behavior of others in one's social network.³

Our own previous work (Cantoni et al. 2019) finds that protest participation in the same Hong Kong setting (although a previous protest) is a game of strategic substitutes. This finding occurs within a single protest when beliefs about the turnout of the broader HKUST student body and the entire Hong Kong population are updated. In contrast to that work, we now study the influence of peers with whom one has relatively strong ties in a dynamic setting. Our work suggests that strong and

³Our work is conceptually related to studies of persistence and social influence in voting behavior (among others, Gerber, Green, and Larimer 2008; Fujiwara, Meng, and Vogl 2016; and DellaVigna et al. 2017), though the dynamics of repeated protest participation may be very different from repeated voting, and the public and social nature of protests may make the role of social interactions distinct.

weak ties may function differently (Granovetter 1973): changes in the participation of the population at large will affect a subject's beliefs about the likelihood that a discrete public good is produced or a government crackdown may occur, potentially generating strategic substitutability. In contrast, friends' participation will have a large effect on the social utility derived from protest participation, the coordination costs of attending, and social image considerations, potentially generating strategic complementarity.

I. Experimental Setting and Design

A. Context: Hong Kong's Antiauthoritarian Movement and the July 1 Marches

In the July 1, 1997, "handover," Hong Kong was transferred from its status as a British colony, with limited democratic political rights but strong protections of civil liberties and respect for the rule of law, to being a special administrative region within the People's Republic of China.⁴ The political institutions of Hong Kong are defined by its quasi constitution—the "Basic Law"—and follow a policy known as "one country, two systems."

The Basic Law left several important dimensions ambiguous, which have been bargained over between the so-called "pan-democracy" and "pro-Beijing" camps since the handover. Every year, the confrontation between Hong Kong citizens and the Chinese government culminates in a protest march held on the anniversary of the "handover" on July 1. Those marches have achieved major policy changes; turnout has varied significantly across years, from less than 20,000 to over 500,000. The repeated nature of the July 1 marches—and their organizers' interest in keeping up high rates of repeated participation—is a feature that the Hong Kong antiauthoritarian protests share with many other political movements.

Our experiment is embedded in the July 1 marches of 2017 and 2018. In both years, protest participation (around 50,000) was modest by historical standards.

B. Overview

Our experimental sample is drawn from the undergraduate student body at HKUST. We recruit subjects through an email sent to the entire HKUST undergraduate student body to participate in a yearly survey on students' preferences (see Cantoni et al. 2016, 2019 for more details); the response rates have ranged between 10 percent and 20 percent. The survey wave in June 2017 includes around 1,100 subjects. Follow-up emails were subsequently sent to experimental subjects between July 2017 and July 2018.

A basic concern regarding self-reported political preferences and behavior is that subjects may not report their participation truthfully. We do not believe that self-censorship is likely in the context of our study. In prior research, we conducted list experiments (also known as the "item count technique") suggesting that subjects

⁴In online Appendix A, we provide a richer description of the political background at the time of our experiment. Note that the implementation on July 1, 2020, of a national security law passed in Beijing has the potential to alter Hong Kong's political landscape, though (as of August 2020) it is still too early for us to know exactly how.

respond honestly to direct questions about sensitive political topics (see Cantoni et al. 2019 for a discussion). More generally, we believe that subjects would have reported their protest participation honestly given the fully legal, peaceful nature of the 2017 and 2018 protests.

It is important to discuss the ethical considerations in conducting our study.⁵ Our research design is based on a careful assessment of ethics. Here we briefly outline salient aspects: (i) IRB approval was received for the study; (ii) no minors are able to participate in the study; (iii) *ex ante*, we assessed a risk level that was minimal, that is, not larger “than those ordinarily encountered in daily life of the general population”: participation in the July 1 marches is unambiguously legal and was peaceful in all years prior to the study; (iv) *ex post*, the assessed risk was minimal, as the marches we studied remained peaceful, with *zero* protesters charged for any offenses across the two years studied; (v) our experiment is tiny relative to the size of the July 1 marches that we study, with treatment affecting total turnout by roughly 0.1 percent.

The timeline of the experiment is as follows (see also online Appendix Figure D.1).

- **June 2017—Baseline Survey and Assignment of Treatment:** We elicit subjects’ own political preferences and beliefs, beliefs about the political preferences and beliefs of others, and planned and past political behavior. We assign and implement the experimental treatment.
- **July 2017—Effect on Protest Participation and Short-Run Impacts on Beliefs and Preferences:** We elicit participation in the 2017 march as well as political preferences and beliefs (short-run treatment effects). Measured beliefs and preferences capture potential mechanisms through which the individual-level or social network-level treatment can shape protest turnout in 2018.
- **June 2018—Long-Run Impacts on Beliefs and Preferences:** We elicit political preferences and beliefs (long-run treatment effects) immediately before the 2018 march. These outcomes again capture potential mechanisms shaping 2018 protest turnout.
- **July 2018—Persistent Effects on Protest Participation and Friendship Formation:** We elicit participation in the 2018 march (our outcome of interest) as well as information on new or stronger friendships formed with politically active individuals. This represents another potential mechanism generating persistent political engagement.

In our study we focus on the 849 subjects for whom we have complete data. The attrition rate is quite low, with over 90 percent retention rates across the multiple waves of the study. In online Appendix Table D.1, we present evidence that the sample that completed all of the study waves looks very similar to the sample of individuals who selected out of the study. We also present all of our analyses re-weighting our experimental sample to match the full sample before attrition, and this has virtually no effect on our findings.

⁵We provide a detailed discussion of ethics and our risk assessment in online Appendix B. All experimental materials (recruitment email, treatment prompts, full survey questions) are provided in online Appendix C.

C. Treatment Design Details

We aim to encourage protest participation without explicitly paying for turnout—directly paying for turnout could potentially generate a set of compliers very different from the typical protest participants we hope to study.⁶ To generate a strong first stage without paying directly for turnout, we pay for behavior *conditional* on turnout: providing information that would help estimate crowd sizes at the protest.⁷

Specifically, within the online survey, individuals randomly selected to be in the treatment group are presented with the following prompt:

Because many students attend the events of July 1, we are asking a subset of survey participants to help us get a better estimate of the July 1 March attendance. ... We would like to ask you to participate in this scientific endeavor. This should take only 5 minutes of your time while you are at the March. ... Once you have uploaded all the information, we will pay you additional HK\$350 for your time and effort.

Subjects in the treatment group received an email the day before the July 1, 2017, march with detailed instructions on how to complete the task. Treated subjects would be able to use a secure link to upload the information we requested. Subjects who upload all requested information and complete the protest participation reporting module would be eligible to receive the bonus payment.

We also want to control for income effects that might arise from our payment in the treatment condition, perhaps generating feelings of reciprocity or otherwise distorting subsequent survey responses in the treatment group. To do so in a politically neutral way, we design a “placebo treatment” that indirectly incentivizes subjects to engage in a very similar activity—traveling to central Hong Kong—for a similar amount of money but *unrelated* to politics (the weekend after the July 1 march). Rather than paying subjects for helping us estimate crowd size at the protest, we pay subjects for helping us estimate metro station crowding. We thereby aim to create a comparison group with identical income effects but no exposure to a political treatment.

Income effects will be comparable between the indirect protest incentive treatment and placebo treatment groups only if take-up rates are similar. As intended, take-up rates in our treatment and placebo treatments are very similar, differing by only around 2 percentage points (online Appendix Figure D.2).

In addition to the random assignment of the treatment (and placebo treatment) at the individual level, we also randomize treatment intensity across relevant social networks. We randomly vary the proportion of study participants receiving the treatment (and placebo treatment) across major \times cohort cells—a relevant social network for university students given the shared course work.⁸ At the cell level, the

⁶“Compliers” in our experiment do not appear to differ significantly from individuals in our sample who had participated in previous protests (online Appendix Table D.2).

⁷Estimating crowd sizes has been conducted by the research team, contributing evidence to a highly contentious debate in Hong Kong (Lin 2018). Using data from our experimental subjects, we estimate that the 2017 march was attended by 26,000–37,000 people—quite similar to the Hong Kong University Public Opinion Programme’s estimates. Refer to online Appendix E for details.

⁸We aim for around 100 cells with 10–20 subjects per cell; when major \times cohort cells are much bigger or smaller, we adjust by merging cells (across majors within cohort) or splitting cells (by gender or residential

treatment intensity is experimentally assigned at a level of 75 percent of subjects in around 35 percent of cells, 50 percent of subjects in 30 percent of cells, 1 percent of subjects in 20 percent of cells, and 0 subjects treated in 15 percent of cells.⁹ The placebo treatment is assigned at the cell level as follows: 0 percent of subjects in approximately 40 percent of cells, 1 percent of subjects in 30 percent of cells, 50 percent of subjects in 25 percent of cells, and 75 percent of subjects in 5 percent of cells. The cell-level intensity of the placebo treatment is cross randomized with the cell-level intensity of the indirect protest incentive treatment, subject to satisfying the adding-up constraint (for example, we could not have a cell with both 75 percent treatment and 75 percent placebo treatment). The result of our cross randomization is that around 45 percent of subjects receive the indirect protest incentive treatment, 20 percent receive the placebo treatment, and 35 percent are pure controls.

In the online Appendix (Tables D.4 and D.5), we present summary statistics and tests of balance at the individual and cell levels. We compare subject characteristics across treatment, placebo treatment, and pure-control subjects as well as between the treatment group and a broader “control group” that pools placebo and pure-control subjects (this is consistent with our pre-analysis plan and supported by our finding that outcomes are nearly identical for placebo treatment and pure-control individuals).

At the individual level, we generally find balance on observables across treatment and control groups, with the exception of gender.¹⁰ At the cell level, we see some systematic differences, with imbalance arising due to our construction of social network cells, which were sometimes defined at the major \times cohort \times gender level. Random assignment generates several high treatment intensity, all-female cells. To address concerns that imbalance affects our estimated treatment effects, we control for cell fixed effects throughout. In addition, we control for gender *interacted* with treatment.¹¹ These analyses suggest that imbalance on observables does not meaningfully affect our results.

II. Main Results: Treatment Effects on Protest Turnout

A. Average Treatment Effects

In Figure 1, panel A, we begin by presenting the short-run (2017) effects of the indirect incentive for protest attendance.¹² In the left-hand graph, one can see that turnout rates in the treatment group are substantially (about 10 percentage points) and statistically significantly higher than in both the pure-control and placebo

address). Online Appendix Table D.3 lists the 98 social network cells that we form.

⁹Due to the small cell sizes, the 1 percent treatment intensity results in cells that have either nobody treated (0 percent) or 1 individual treated (producing a treatment intensity of approximately 10 percent). We present target and actual treatment intensity for each cell in online Appendix Table D.4.

¹⁰This is an important dimension of imbalance to account for, though we do not find evidence that gender is associated with 2017 protest turnout among control subjects (p -value = 0.675).

¹¹Importantly, all of our results that rely on variation across cells (i.e., heterogeneous treatment effects associated with cell treatment intensity) are robust to the inclusion of an interaction between the individual treatment indicator and *any* of the unbalanced cell characteristics identified in Table D.6. See online Appendix Tables D.7 and D.8.

¹²Throughout the analyses presented, we conduct two-sided tests for statistical inference. While deviating from the one-sided tests that we pre-registered, this approach is more conservative.

treatment groups. One can also see that protest attendance rates are very similar (and statistically indistinguishable) in the placebo treatment and pure-control groups. Any income effects contributing to changed protest participation in 2017 are thus unlikely to be large.¹³ To gain power, we pool the pure-control and placebo treatment groups into a larger comparison group that, for concision, we refer to as the “control” group (right-hand graph). Table 1, panel A, column 1, displays the analogous results in regression format, controlling for cell fixed effects. Column 2 adds the interaction of subject gender and the treatment dummy. Regression results suggest around a 10 percentage point increase in 2017 turnout, on average, among treated individuals.

We next examine whether the indirect incentive for protest attendance in 2017 generates long-run (i.e., 2018) average treatment effects on protest participation. Figure 1, panel B, presents the results. In the left-hand graph, we display raw attendance rates across treatment arms. Turnout rates remain substantially—around 5 percentage points—and statistically significantly higher in the treatment group compared to either the placebo or pure-control group. Results are analogous when considering the pooled control group (right-hand graph). Table 1, panel A, column 3, presents regression estimates of the treatment effect in 2018, including cell fixed effects. Column 4 adds the interaction of subject gender and the treatment dummy. We find an approximately 5 percentage point average effect of the incentive treatment on 2018 turnout. We can estimate the average causal effect of 2017 protest attendance on 2018 attendance at the individual level, exploiting variation in 2017 attendance arising from our experimental treatment. Two-stage estimates—from a regression of 2018 turnout on 2017 turnout, instrumented by treatment—show a coefficient of 0.47 (p -value < 0.01); that is, subjects who are randomly indirectly incentivized into protest participation in one year are nearly 50 percent more likely to turnout to protest a full year later when the incentives are no longer in place.¹⁴

B. Heterogeneous Treatment Effects

We next examine the extent to which protest attendance varied in response to *both* individual-level treatment *and* treatment intensity at the social network (major \times cohort) level. Importantly, this is the only dimension of heterogeneity we examine, it is the only dimension of heterogeneity that we included in our pre-analysis plan, and the variation exploited is experimental.

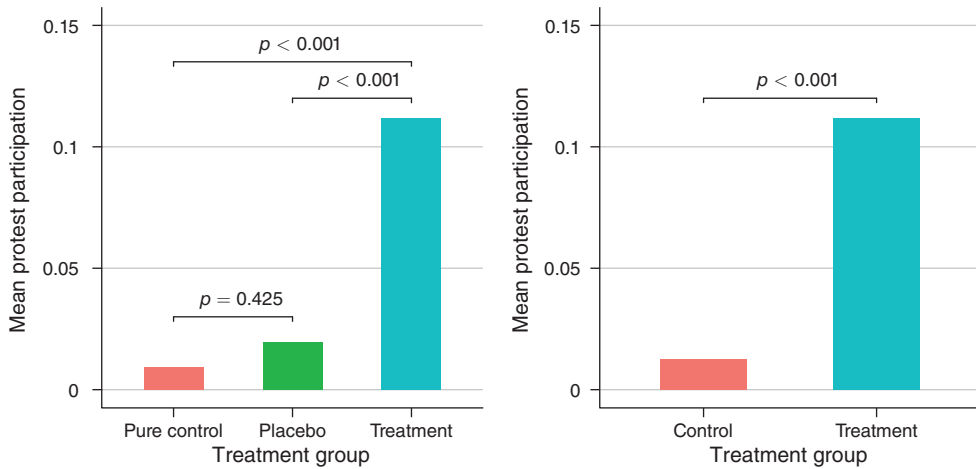
In Figure 2, we plot turnout rates by individual treatment status (treatment versus pooled control) and cell treatment intensity (1 percent treated, 50 percent treated, or 75 percent treated) for 2017 (panel A) and 2018 (panel B).¹⁵ One can see in panel A

¹³The lack of differences between the placebo and pure-control group is also evident in 2018 turnout (see Figure 1, panel B) and across the entire range of survey questions asked in 2017 and 2018 (see online Appendix Table D.8).

¹⁴We benchmark this experimentally induced persistence rate against the naturally occurring one using data we have collected from the HKUST student panel surveys since 2014. The likelihood that a student participates in a July first march in year t , conditional on having participated in year $t - 1$, ranges between 24 percent and 43 percent, slightly lower but not far from the experimental persistence rate (online Appendix Table D.9).

¹⁵In online Appendix Figure D.3 we alternatively plot turnout rates at the *cell* level by individual treatment status and by targeted cell treatment intensity. We also plot the linearly estimated turnout rates as a function of individual treatment status, cell treatment intensity, and their interaction for 2017 and 2018.

Panel A. 2017 participation



Panel B. 2018 participation

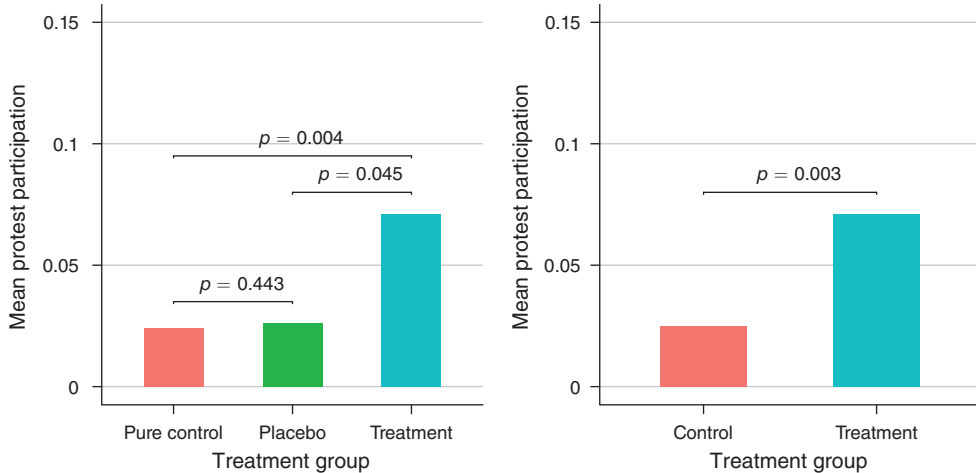


FIGURE 1

Notes: Panel A: participation in July 1, 2017, protest by treatment group. Panel B: participation in July 1, 2018, protest by treatment group. p -values calculated from regressions of protest turnout on treatment group indicators with standard errors clustered at the major \times cohort cell level.

that in 2017 turnout rates are significantly higher among treatment group individuals than control and that the gap in turnout rates between treatment and control subjects is of approximately the same magnitude *regardless* of treatment cell intensity. These results are robust to controlling for cell fixed effects and the interaction of gender with treatment (see Table 1, panel B, columns 1–2). Any complementarities across treated peers within a social network were not very strong in 2017, nor do there seem to have been large spillovers to untreated subjects. It seems that the treatment affected turnout in 2017 very much at an individual level.

In contrast, one can see in Figure 2, panel B, that in 2018 turnout rates are *differentially* higher among treatment group individuals in treatment cells with higher treatment intensity. We find a marginally significant negative treatment effect in

TABLE 1—TREATMENT EFFECTS: PROTEST PARTICIPATION AND PLANS (REWEIGHTED SAMPLE)

	Participation				Plans to participate	
	2017 (1)	2017 (2)	2018 (3)	2018 (4)	2018 (5)	2018 (6)
<i>Panel A. Average treatment effect</i>						
Treatment	0.106 (0.018)	0.094 (0.024)	0.050 (0.016)	0.043 (0.020)	-0.021 (0.023)	-0.028 (0.027)
<i>Panel B. Heterogeneity by cell treatment intensity</i>						
Treatment	0.133 (0.124)	0.114 (0.122)	-0.033 (0.018)	-0.047 (0.030)	-0.104 (0.051)	-0.117 (0.062)
Treatment × 50 percent intensity	-0.028 (0.126)	-0.020 (0.124)	0.062 (0.022)	0.068 (0.025)	0.067 (0.061)	0.073 (0.064)
Treatment × 75 percent intensity	-0.028 (0.127)	-0.021 (0.125)	0.117 (0.036)	0.122 (0.038)	0.110 (0.062)	0.112 (0.065)
DV mean (control group)	0.012	0.012	0.025	0.025	0.100	0.100
DV standard deviation (control group)	0.111	0.111	0.156	0.156	0.299	0.299
DV mean (all)	0.055	0.055	0.045	0.045	0.091	0.091
DV standard deviation (all)	0.229	0.229	0.207	0.207	0.287	0.287
Gender controls	No	Yes	No	Yes	No	Yes
Observations	849	849	849	849	849	849

Notes: Panel A presents estimated coefficients from regressions of protest turnout (or planned turnout) on the individual treatment indicator. Panel B presents estimated coefficients from regressions of protest turnout (or planned turnout) on the individual treatment indicator interacted with major × cohort cell treatment intensity bin indicators (and lower-order terms). Results are shown for 2017 protest turnout (columns 1–2), 2018 protest turnout (columns 3–4), and 2018 planned protest turnout (columns 5–6). Columns 1, 3, and 5 include major × cohort cell fixed effects; in addition, columns 2, 4, and 6 include the interaction between individual treatment status and a gender indicator. Standard errors (reported in parentheses) are clustered at the major × cohort cell level.

the 1 percent treatment intensity cells, modestly greater 2018 protest participation among treated subjects in cells that are 50 percent treated (relative to controls in the same cells), and economically and statistically significantly greater 2018 protest participation among treated subjects in cells that are 75 percent treated (relative to controls in the same cells).¹⁶ One can see in the table of *p*-values reported in Figure 2 that the difference in treatment effects between the 75 percent treated and 1 percent treated cells is highly statistically significant, and the difference between the 75 percent treated and 50 percent treated cells is marginally statistically significant. The difference in treatment effects between the 50 percent treated cells and the 1 percent treated cells is significant as well. These results, too, are all robust to controlling for cell fixed effects and the interaction of gender with treatment (see Table 1, panel B, columns 3–4).¹⁷

¹⁶The negative treatment effect in the 1 percent treated cells may result from sampling variation—estimates become insignificant when we control for the interaction of gender and treatment—or may reflect a particular (negative) experience of 2017 protest participation among treated subjects in low treatment intensity cells that reduces 2018 turnout.

¹⁷In online Appendix Table D.7, we present all of the results in Table 1, panel B, under various alternative specifications. First, we control for the interaction between the treatment dummy and each unbalanced characteristic observed in Table D.6. We also control for the interaction of the treatment dummy and predicted protest attendance. We first predict control group individuals' protest turnout in 2017 using a full set of demographics. Then, using the estimated coefficients from this regression, we predict *all* subjects' turnout based on their demographics. This is a parsimonious way of controlling for relevant subject characteristics without losing too many degrees of freedom. Online

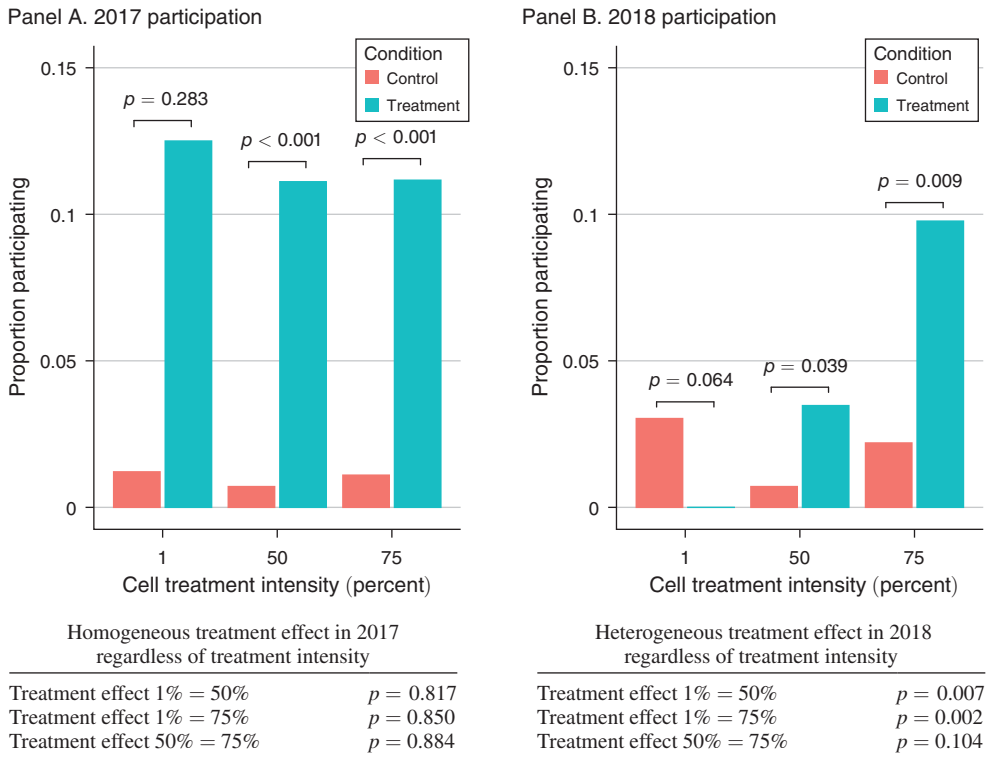


FIGURE 2

Notes: Panel A: participation in July 1, 2017, protest by treatment group and major \times cohort cell treatment intensity. Panel B: participation in July 1, 2018, protest by treatment group and major \times cohort cell treatment intensity. p -values calculated from regressions of protest turnout on interactions of the individual treatment indicator with major \times cohort cell treatment intensity bin indicators as well as lower-order terms. Standard errors clustered at the major \times cohort cell level.

As an additional exercise, we examine treatment effects on *planned* protest participation in 2018 (elicited the week before the July 1, 2018, march) as an auxiliary outcome. While we find no significant average treatment effect on planned participation (Table 1, panel A, columns 5–6), we do find that planned protest participation among treated subjects is greater in major \times cohort cells with higher treatment intensity, matching the pattern observed for actual protest participation (Table 1, panel B, columns 5–6).

The absence of heterogeneous treatment effects by cell treatment intensity in 2017 and their presence for both planned and actual turnout in 2018 suggests that a crucial change took place between the 2017 and 2018 marches. This was specifically among treated individuals within major \times cohort social networks that were more intensely treated (and thus exhibited greater turnout at the 2017 march). We next explore mechanisms related to changed social interactions that might generate sustained political engagement.

Appendix Table D.11 also presents p -values calculated using permutation tests as well as results from a re-weighted sample to account for attrition. Results across these specifications are very similar to those in Table 1.

III. Mechanisms

What explains the persistent engagement of individuals who turn out to protest due to our experimental intervention? Here we consider the possibilities that changed social interactions among treated subjects in major \times cohort cells with high treatment intensity might have shaped subjects' friendship networks, lowered their coordination costs, shaped their political beliefs and preferences, and changed their beliefs about others.

A. *The Formation of New or Stronger Friendships*

How might the variation in treatment intensity at the cell level have generated significant interactions with individual treatment status? Several pieces of evidence are suggestive of the importance of new or stronger friendships formed as a result of march attendance—either at the march itself or thereafter. First, heterogeneity driven by preexisting friendships among treated subjects (prior to 2017) would have made heterogeneous treatment effects in 2017 more likely. We do not find evidence of these. Second, preexisting friendships would have been as common between a treated and control subject as between treated subjects. If attendance in the 2017 march by a treated subject shaped 2018 turnout among her preexisting friends (i.e., those from before the 2017 march), one should see heterogeneity in turnout rates associated with cell treatment intensity in 2018 *among the control group* as well as the treatment group. The fact that we only see differentially large turnout rates in high treatment intensity cells among *treated* subjects suggests that joint attendance at the 2017 march was crucial in shaping turnout in 2018.

We directly elicit changes in subjects' friendships since the 2017 protest in the July 2018 survey.¹⁸ We estimate a regression model analogous to the baseline model estimated in Table 1 but consider as the outcome subjects' reported new or stronger friendships (Table 2, column 1). Indeed, we find patterns of new friendship formation that correspond quite closely to the patterns of 2018 protest attendance: new political friendships are reported significantly more often by treated individuals in the cells with the highest treatment intensity (and new political friendships are actually less common among treated subjects in cells with 1 percent treatment intensity). These new friendships could have directly affected 2018 protest turnout through increased social utility from protest participation or could have stimulated turnout by reducing coordination costs or affecting beliefs or preferences.

B. *The Reduction of Coordination Costs*

One natural role that friends play in shaping protest turnout is in reducing coordination costs. Among subjects who planned to turn out, reduced coordination costs would induce a higher rate of converting their planned protest participation into actual participation. To examine this possibility, in Figure 3, we split our sample of subjects depending on their *planned* 2018 protest participation.

¹⁸While we specifically ask about friendships since the 2017 march, it is possible that some of these friendships were formed after the 2018 march.

TABLE 2—MECHANISMS: POLITICAL BELIEFS, PREFERENCES, BELIEFS ABOUT OTHERS, AND NEW FRIENDSHIPS (REWEIGHTED SAMPLE)

	New friendships	Political preferences		Political beliefs		Beliefs about others	
	2018	2017	2018	2017	2018	2017	2018
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A. Average treatment effect</i>							
Treatment	0.027 (0.020)	0.134 (0.069)	0.093 (0.089)	-0.054 (0.082)	-0.027 (0.089)	0.043 (0.073)	0.015 (0.072)
<i>Panel B. Heterogeneity by cell treatment intensity</i>							
Treatment	-0.036 (0.019)	-0.316 (0.545)	0.155 (0.440)	-0.455 (0.472)	-0.148 (0.177)	-0.424 (0.394)	-0.382 (0.106)
Treatment × 50 percent intensity	0.073 (0.031)	0.414 (0.551)	-0.062 (0.460)	0.362 (0.489)	0.115 (0.215)	0.497 (0.406)	0.521 (0.144)
Treatment × 75 percent intensity	0.058 (0.038)	0.544 (0.556)	-0.069 (0.456)	0.491 (0.483)	0.141 (0.229)	0.489 (0.411)	0.305 (0.149)
DV mean (control group)	0.064	-0.062	-0.052	-0.012	0.005	-0.045	0.005
DV standard deviation (control group)	0.245	0.994	1.039	1.024	1.033	1.016	1.050
DV mean (all)	0.078	-0.011	-0.015	0.002	0.001	-0.015	0.005
DV standard deviation (all)	0.268	0.993	1.007	1.000	1.018	0.998	1.008
Observations	849	849	849	849	849	849	849

Notes: Panel A presents estimated coefficients from regressions of new friendships, indices of preferences, political beliefs, and beliefs about others on the individual treatment indicator. Panel B presents estimated coefficients from regressions of new friendships, indices of preferences, political beliefs, and beliefs about others on the individual treatment indicator interacted with major × cohort cell treatment intensity bin indicators (and lower-order terms). Results are shown for new friendships reported in July 2018 (column 1); for July 2017 preferences, beliefs, and beliefs about others (columns 2, 4, and 6); and for June 2018 preferences, beliefs, and beliefs about others (columns 3, 5, and 7). All regressions include major × cohort cell fixed effects. Standard errors (reported in parentheses) are clustered at the major × cohort cell level. The individual survey questions combined to construct the indices are provided in online Appendix C.1.

We then plot the *actual* participation in 2018 depending on own treatment status and major × cohort cell treatment intensity for subjects who planned to turn out (panel A). We find by far the highest conversion rate of protest plans into action—at over 40 percent—among treated individuals in the highest treatment intensity cells. This may reflect differential information about transportation, meeting times and locations, and differential social pressure as well. Reduced coordination costs might also induce turnout among individuals who did not plan to attend a protest. Indeed, we find that among those subjects who did *not* plan to participate (panel B), there is significantly higher protest turnout among treated subjects in high treatment intensity cells.

C. Changes in Subjects’ Political Preferences and Beliefs

Standard models of protest participation would suggest the importance of changes in expected payoffs from participation arising from changed political beliefs (e.g., about the political climate or incumbent regime) or changed preferences. We consider subjects’ political preferences (e.g., regarding democracy) and beliefs about future political outcomes. We summarize outcomes in each category (preferences and beliefs) by constructing z-score index variables with larger positive values indicating more antiauthoritarian responses, weighting by the inverse covariance

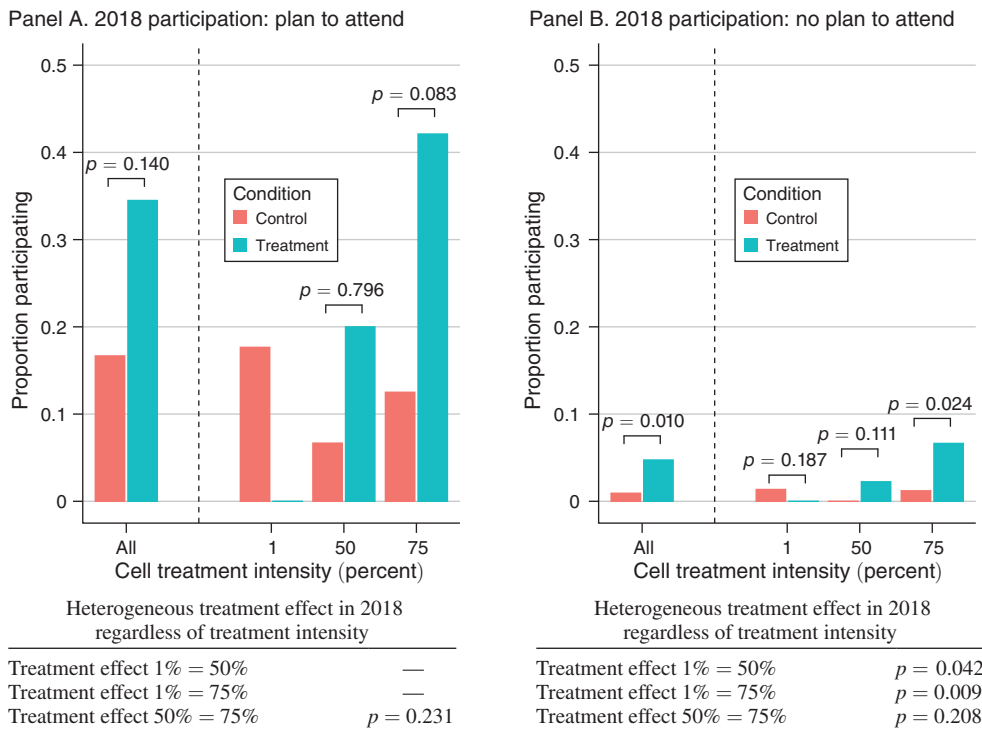


FIGURE 3

Notes: Panel A: participation in July 1, 2018, protest by treatment group and major × cohort cell treatment intensity among subjects who planned to participate in the 2018 protest. Panel B: participation in July 1, 2018, protest by treatment group and major × cohort cell treatment intensity among subjects who did not plan to participate in the 2018 protest. *p*-values calculated from regressions of protest turnout on interactions of the individual treatment indicator with major × cohort cell treatment intensity bin indicators as well as lower-order terms. Statistical tests cannot be conducted among subjects planning to attend the 2018 protest in the 1 percent treatment intensity cells (panel A) because no treated subjects in the 1 percent treatment intensity cells report a plan to attend the 2018 protest. Standard errors clustered at the major × cohort cell level.

of standardized variables, following Anderson (2008).¹⁹ We do so separately for outcomes elicited just after the 2017 protest and just before the 2018 protest, as we pre-register. For completeness, we present the treatment effects on all individual outcome variables in online Appendix Table D.8, adjusting *p*-values for multiple hypothesis testing, following List, Shaikh, and Xu (2019).

In Table 2, columns 2–3, we consider as outcomes subjects’ political preferences in 2017 and 2018 using the baseline specification estimated in Table 1.²⁰ We find a marginally significant shift toward more antiauthoritarian political preferences among treated subjects in 2017, on average. However, we do not find evidence of heterogeneous effects associated with cell treatment intensity (though estimates are noisy). In Table 2, columns 4–5, we examine subjects’ political beliefs in 2017 and

¹⁹The full text of the survey questions entering the indices is provided in online Appendix Section C.1.

²⁰In online Appendix Table D.8, we present all of the results in Table 2 but include a full range of controls. Online Appendix Table D.12 also presents *p*-values calculated using permutation tests as well as results from a re-weighted sample to account for attrition. Results across these specifications are very similar to those in Table 2.

2018 as outcomes. We find very small average treatment effects on political beliefs in both 2017 and 2018. We see some suggestive (albeit noisy) evidence of beliefs moving in an antiauthoritarian direction among treated subjects in high treatment intensity cells in 2017 but not 2018.

Overall, while our estimated treatment effects on political preferences and beliefs are noisy, we do not find compelling evidence matching the heterogeneous treatment effects we observe on protest participation, particularly just prior to the 2018 march.

D. Changes in Subjects' Beliefs about Others

We next examine subjects' beliefs about the political preferences of others. Such beliefs about others may affect strategic considerations in deciding whether to protest (to the extent that they shape subjects' beliefs about other people's protest participation) and could plausibly be affected by the political engagement of subjects' social networks. In Table 2, columns 6–7, we consider as outcomes subjects' beliefs about others in 2017 and 2018. As in the previous section, we construct a z-score index variable with larger, positive values indicating more optimistic (antiauthoritarian) beliefs about others.

We find that in 2018, treated subjects in the high treatment intensity cells are significantly more optimistic about the support of others for Hong Kong's antiauthoritarian movement, while treated subjects in the 1 percent treatment intensity cells are significantly more pessimistic. If such optimism translates into subjects' optimistic beliefs about others' protest participation, then the changed beliefs about others could actually *decrease* the tendency to protest in 2018, as we have previously found that protest participation is a game of strategic substitutes in this context (Cantoni et al. 2019). This points toward previously discussed mechanisms (e.g., social utility or coordination costs) as more plausible explanations for persistent protest turnout among treated subjects in high treatment intensity cells.

IV. Conclusion

Our work provides evidence that social networks play a crucial role in shaping individuals' persistent participation in political movements. The next step is to better understand *how* social interactions affect political engagement. We provide suggestive evidence of the importance of friendship formation and strengthening. Looking ahead, one naturally wonders how important increased joint consumption value from protest participation, changed social image considerations, reduced costs of coordination, and improved information transmission are. We cannot confidently rule out a role for changed political beliefs and preferences. A better understanding of the mechanisms through which social interactions sustain political engagement will not only help interpret patterns of political mobilization but can also inform dynamic models of political movements.

REFERENCES

Acemoglu, Daron, and James A. Robinson. 2012. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. New York: Crown Publishers.

- Acemoglu, Daron, and James A. Robinson.** 2019. *The Narrow Corridor: States, Societies, and the Fate of Liberty*. New York: Penguin Press.
- Aidt, Toke S., and Raphaël Franck.** 2015. "Democratization under the Threat of Revolution: Evidence from the Great Reform Act of 1832." *Econometrica* 83 (2): 505–47.
- Anderson, Michael L.** 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103 (484): 1481–95.
- Barberà, Salvador, and Matthew O. Jackson.** 2020. "A Model of Protests, Revolution, and Information." *Quarterly Journal of Political Science* 15 (3): 297–335.
- Bueno de Mesquita, Ethan.** 2010. "Regime Change and Revolutionary Entrepreneurs." *American Political Science Review* 104 (3): 446–66.
- Bursztyjn, Leonardo, Davide Cantoni, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang.** 2021. "Replication Data for: Persistent Political Engagement: Social Interactions and the Dynamics of Protest Movements." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E122421V1>.
- Cantoni, Davide, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang.** 2016. "The Fundamental Determinants of Anti-Authoritarianism." https://faculty.haas.berkeley.edu/yuchtman/Noam_Yuchtman_files/antiauthoritarian_draft.pdf.
- Cantoni, Davide, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang.** 2019. "Protests as Strategic Games: Experimental Evidence from Hong Kong's Democracy Movement." *Quarterly Journal of Economics* 134 (2): 1021–77.
- Della Vigna, Stefano, John A. List, Ulrike Malmendier, and Gautam Rao.** 2017. "Voting to Tell Others." *Review of Economic Studies* 84 (1): 143–81.
- Edmond, Chris.** 2013. "Information Manipulation, Coordination, and Regime Change." *Review of Economic Studies* 80 (4): 1422–58.
- Enikolopov, Ruben, Alexey Makarin, and Maria Petrova.** 2020. "Social Media and Protest Participation: Evidence from Russia." *Econometrica* 88 (4): 1479–1514.
- Fujiwara, Thomas, Kyle Meng, and Tom Vogl.** 2016. "Habit Formation in Voting: Evidence from Rainy Elections." *American Economic Journal: Applied Economics* 8 (4): 160–88.
- Gerber, Alan S., Donald P. Green, and Christopher W. Larimer.** 2008. "Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment." *American Political Science Review* 102 (1): 33–48.
- Goldstone, Jack A.** 1991. *Revolution and Rebellion in the Early Modern World: Population Change and State Breakdown in England, France, Turkey, and China*. Berkeley: University of California Press.
- González, Felipe.** 2020. "Collective Action in Networks: Evidence from the Chilean Student Movement." *Journal of Public Economics* 188: 104220.
- Granovetter, Mark S.** 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78 (6): 1360–80.
- Hager, Anselm, Lukas Hensel, Johannes Hermle, and Christopher Roth.** 2019a. "Political Activists as Free-Riders: Evidence from a Natural Field Experiment." IZA Discussion Paper 12759.
- Hager, Anselm, Lukas Hensel, Johannes Hermle, and Christopher Roth.** 2019b. "Strategic Interdependence in Political Movements and Countermovements." CESifo Working Paper 7790.
- Hirschman, Albert O.** 1984. *Getting Ahead Collectively: Grassroots Experiences in Latin America*. New York: Pergamon Press.
- Hobbs, William R., and Margaret E. Roberts.** 2018. "How Sudden Censorship Can Increase Access to Information." *American Political Science Review* 112 (3): 621–36.
- King, Gary, Jennifer Pan, and Margaret E. Roberts.** 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107 (2): 326–43.
- Kuran, Timur.** 1997. *Private Truths, Public Lies: The Social Consequences of Preference Falsification*. Cambridge: Harvard University Press.
- Lin, Kurt.** 2018. "Stand and Be Counted." *Monocle* 117: 38.
- List, John, Azeem M. Shaikh, and Yang Xu.** 2019. "Multiple Hypothesis Testing in Experimental Economics." *Experimental Economics* 22: 773–93.
- Lorentzen, Peter L.** 2013. "Regularizing Rioting: Permitting Public Protest in an Authoritarian Regime." *Quarterly Journal of Political Science* 8 (2): 127–58.
- Madestam, Andreas, Daniel Shoag, Stan Veuger, and David Yanagizawa-Drott.** 2013. "Do Political Protests Matter? Evidence from the Tea Party Movement." *Quarterly Journal of Economics* 128 (4): 1633–85.

- Manacorda, Marco, and Andrea Tesei.** 2020. "Liberation Technology: Mobile Phones and Political Mobilization in Africa." *Econometrica* 88 (2): 533–67.
- McAdam, Doug.** 1982. *Political Process and the Development of Black Insurgency, 1930–1970*. Chicago: University of Chicago Press.
- McAdam, Doug.** 1986. "Recruitment to High-Risk Activism: The Case of Freedom Summer." *American Journal of Sociology* 92 (1): 64–90.
- Passarelli, Francesco, and Guido Tabellini.** 2017. "Emotions and Political Unrest." *Journal of Political Economy* 125 (3): 903–46.
- Qin, Bei, David Strömberg, and Yanhui Wu.** 2017. "Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda." *Journal of Economic Perspectives* 31 (1): 117–40.
- Shadmehr, Mehdi, and Dan Bernhardt.** 2011. "Collective Action with Uncertain Payoffs: Coordination, Public Signals, and Punishment Dilemmas." *American Political Science Review* 105 (4): 829–51.
- Skocpol, Theda.** 1979. *States and Social Revolutions: A Comparative Analysis of France, Russia, and China*. New York: Cambridge University Press.
- Tarrow, Sidney.** 2005. *The New Transnational Activism*. Cambridge, UK: Cambridge University Press.
- Tilly, Charles.** 1978. *From Mobilization to Revolution*. Reading, MA: Addison-Wesley.
- Tullock, Gordon.** 1971. "The Paradox of Revolution." *Public Choice* 11: 89–99.
- Zhang, Han.** 2016. "Physical Exposures to Political Protests Impact Civic Engagement: Evidence from 13 Quasi-Experiments with Chinese Social Media." <http://dx.doi.org/10.2139/ssrn.2647222>.

