Appendix

A Ethical considerations

The survey and experiment were approved by the *[Board and Id withheld]*. We followed the "do no harm principle" (Wood) in each step of the survey design and implementation by taking a number of precautions to avoid putting respondents and local staff at risk. By collaborating with Lokniti we ensured that the survey did not violate local norms and customs. Lokniti has vast experience of conducting election surveys, including pre- and post-polls for national, state, and local elections. Our questionnaire was screened by Lokniti to ensure that no questions were asked that could harm respondents or enumerators. Lokniti also trained the enumerators and local fieldwork coordinators, ensuring that the local staff could navigate the environment in which the survey was conducted. Finally, we conducted the interviews based on informed consent. Individuals were informed about the purpose of the study, that they were selected at random, that they could withdraw at any time from the interview, that their identity would not be revealed, and that the survey was conducted by Lokniti on behalf of *[Name withheld]*. We only proceeded if oral consent was given.

B The list experiment

I am going to read you a list of things that people have told us they experienced during the 2019 election campaign. I would like you to tell me how many of these things you have personally experienced. Please, do not tell me which ones, only HOW MANY. If you would like me to repeat the list, I will do so.

- You saw a television add endorsing a particular candidate or party.
- You heard about or attended a campaign rally in your locality.
- **Someone threatened you in any way (or threatened someone in your home) to vote for a particular candidate or party on election day** [Mobilization treatment] OR **Someone threatened you in any way (or threatened someone in your home) to not cast your vote on election day.** [Demobilization treatment]
- A politician personally came to your house to discuss his or her candidacy
- You discussed the elections with friends or family

Box 1: The list experiment

Thinking about the 2019 Lok Sabha elections, we are interested in whether a party, politician or their representatives tried to persuade you to vote in a particular way. Please let us know if a party or politician did any of the following:

- Threatened you (or someone in your home) to vote for a particular candidate or party on election day OR Threatened you (or someone in your home) not to vote on election day.

Box 2: The direct question

C Test of assumptions

The difference between the average number of items chosen in the control list and in the treatment list can only be attributed to the inclusion of the sensitive item (here, either the violent mobilization or demobilization item) if three assumptions hold (Blair and Imai). First, that the assignment to control and treatment lists is random. Second, that the inclusion of the sensitive item does not change responses to any of the control items, the so-called "no design effect" assumption. Third, that individuals answer the sensitive item truthfully. While the latter cannot be tested, but must be assumed, the first two can directly be assessed.

To assess the first assumption, i.e. treatment randomization, we conduct balance tests where we regress the treatment assignment (0=control list, 1=violent mobilization list, 2=violent demobilization list) on a set of key demographic variables (Gerber and Green). These include the respondents age (continous), gender (0=men, 1=women), urban/rural residency (0=rural, 1=urban), socio-economic status (set of four mutually exclusive dummy variables, reference=poor), education level (0=low, 1=medium 2=high), caste (0=other backward classes, 1=scheduled casts, 2=scheduled tribes, 3=other/general castes), and religion (0=hindu, 1=muslim, 2=other).Because we have three experimental conditions (control, violent mobilization, and violent demobilization), we chose multinomial regression models. As Table A1 shows, none of these covariates systematically predicts whether an individual has received the control or any of the treatment lists.

Next, turning to the no design effect assumption, we test whether the inclusion of the sensitive item in the treatment group affects the sum of control items that an individual affirmatively responds to. We follow the approach suggested by Blair and Imai, i.e. testing whether we can reject the null hypothesis of no design effects. As the Bonferroni-adjusted p-values in Table A2 indicate, we cannot reject this null-hypothesis. Hence, our experiment does not seem to violate this assumtpion.[41]

[41]Note that for the mobilization treatment, one of the coefficients is negative. While this is a sign of a potential violation of the no design effect assumption, the large Bonferroni-adjusted p-value indicates that the null hypotheses can, overall, not be rejected.

Table A1: Covariate balance (randomization of treatment)

	Multinomial logit	
Control vs. violent mobilization		
Age	0.000	(0.006)
Women	0.110	(0.157)
Urban	0.048	(0.168)
Scheduled Caste (SC)	-0.029	(0.207)
Scheduled Tribes (ST)	-0.031	(0.265)
Other Backward Classes (OBC)	0.022	(0.213)
Muslim	0.012	(0.217)
Other Religions	0.100	(0.420)
Medium Education	0.347	(0.204)
High Education	-0.099	(0.303)
Upper class	-0.057	(0.433)
Middle class	-0.058	(0.423)
Lower class	-0.000	(0.441)
Constant	-0.240	(0.553)
Control vs. violent demobilization		
Age	0.001	(0.006)
Women	-0.036	(0.157)
Urban	0.079	(0.168)
Scheduled Caste (SC)	0.054	(0.204)
Scheduled Tribes (ST)	-0.139	(0.270)
Other Backward Classes (OBC)	-0.060	(0.215)
Muslim	0.053	(0.218)
Other Religions	0.417	(0.399)
Medium Education	0.101	(0.198)
High Education	-0.385	(0.304)
Upperclass	0.565	(0.504)
Middleclass	0.664	(0.494)
Lowerclass	0.691	(0.508)
Constant	-0.705	(0.608)
Observations	1024	

Standard errors in parentheses.* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Reference categories: gender=men, rural/urban=rural,

cast=other backward classes (obc), religion=hindu, education=low, class=poor.

Table A2: Test for design effects

	Mobilization		**Demobilization**	
	Estimate	SE	Estimate	SE
$pi(Y_i(0) = 0, Z_i = 1)$	0.0112	0.0126	0.0191	0.0113
$pi(Y_i(0) = 1, Z_i = 1)$	0.0000	0.0350	0.0236	0.0340
$pi(Y_i(0) = 2, Z_i = 1)$	-0.0139	0.0438	0.0400	0.0436
$pi(Y_i(0) = 3, Z_i = 1)$	0.0457	0.0331	0.0907	0.0343
$pi(Y_i(0) = 4, Z_i = 1)$	0.0700	0.0159	0.0875	0.0174
$pi(Y_i(0) = 0, Z_i = 0)$	0.0156	0.0077	0.0076	0.0054
$pi(Y_i(0) = 1, Z_i = 0)$	0.1717	0.0268	0.1482	0.0255
$pi(Y_i(0) = 2, Z_i = 0)$	0.3307	0.0397	0.2768	0.0394
$pi(Y_i(0) = 3, Z_i = 0)$	0.2902	0.0395	0.2452	0.0406
$pi(Y_i(0) = 4, Z_i = 0)$	0.0788	0.0271	0.0614	0.0281

Note: Displayed are estimated proportions with standard errors.

Bonferroni corrected p-value for mobilization=1; demobilization=1

D Analysis of item-non responses

The main benefit of list experiments is that they grant anonymity when asking sensitive questions, therefore reducing both self-censorship in the form of item non-response and social desirability bias (Blair and Imai). One observable implication that the experiment reduces self-censorship is that individuals should be more willing to respond to the list experiment than to the direct question. We assess whether this is the case by comparing the number of item non-responses for the sub-sample receiving the mobilization treatment and demobilization treatment separately. In contrast to what we expected, we find that there are more item non-responses for the list experiment than for the direct questions. For those receiving the mobilization treatment (N=360), 103 do not respond to the list experiment ($N_{don'tknow}$=67, $N_{refused}$=36), while 74 do not respond to the direct question ($N_{don'tknow}$=50, $N_{refused}$=24). For those receiving the demobilization treatment (N=360), 97 do not respond to the list experiment ($N_{don'tknow}$=65, $N_{refused}$=32), while 65 do not respond to the direct question ($N_{don'tknow}$=46, $N_{refused}$=21). One potential explanation for this is that the list experiment may be more challenging to answer and requires more concentration, leading to a higher proportion of individuals that do not want to engage with it in comparison to the simple direct question.

While the list experiment did not seem to reduce self-censorship, it did indeed seem to reduce social desirability bias. One way of testing this is to assess whether there are systematic

difference in item non-responses across the experimental conditions, i.e. across individuals that receive the control list vs those that receive any of the lists with the sensitive items. If the experiment did reduce social desirability concerns, we would expect to see no systematic differences in the willingness to respond to the experiment across control and treatment groups. We therefore regress whether or not an individual has responded to the list experiment (1=not responded [do not know or refused], 0=responded) on the individual's treatment status (0=received control list; 1=received mobilization item; 2=received demobilization sensitive item). The coefficient plot in Figure A1 shows the results from three separate logistic regression models using do not know responses, refusals, and a pooled indicator as dependent variables, respectively. Individuals receiving the control list are the reference category for the item non-response variable. The coefficients for treatment assignment do not reach conventional levels of statistical significance, indicating that whether individuals received the control list or any of the two lists with sensitive items does not predict whether it responded or not to the list experiment.

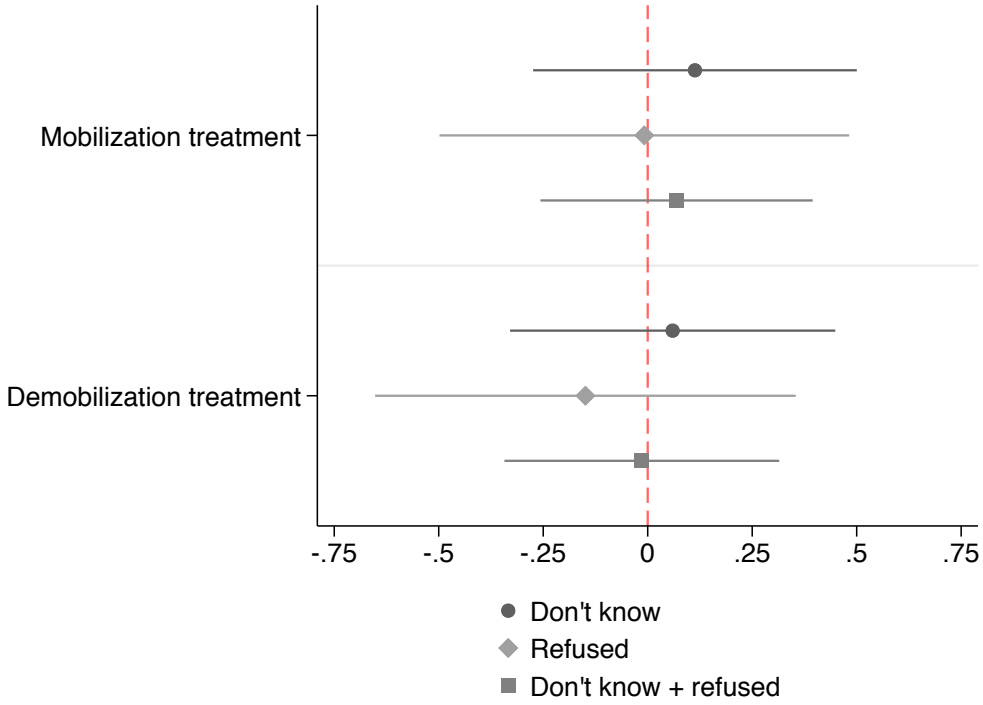

Figure A1: Coefficient plot; treatment assignment (control, mobilization treatment, demobilization treatment) as predictor for item non-responses. Coefficients are reported in log-odds with 95% confidence intervals. $N_{don't know}$=975, $N_{refused}$=887, N_{pooled}=1,080.

E Polling station analysis – data and variable construction

The polling station results for the 2019 elections in India are published online as pdfs, in what is known as 'Form 20'. We downloaded Form 20 for all polling stations, then extracting election results from the pdf documents. Next, we identify geographic coordinates for each polling station and link these to respondent locations in our survey. Recall that we sample 10 districts and nine locations in each district, producing 90 unique respondent locations. To protect the safety of our respondents, we did not record precise coordinates for each respondent, which is why we identify polling stations in the vicinity of these 90 locations. We were unable to identify coordinates for all polling stations in Darjeeling parliamentary constituency, and a few other select polling booths. Despite our best efforts, including contacting the Election Commission of India, we could not determine coordinates for this constituencies. We lose 132 respondents in our survey because of these missing coordinates.

To construct our electoral geography measures for the five closest polling stations, we then first identified a 1.7 kilometer buffer around each location based on the minimum distance between respondent locations, which was 3.4 kilometers. We divided that in half to avoid overlapping buffers. The buffer helps us deal with the fact that the number of polling stations is much higher in urban versus rural locations. By drawing a buffer, we ensure that there will typically be only one or two polling stations in the vicinity of our rural respondent locations, while we can consider a larger number of polling stations for urban locations.

F Probing theory: observable implications

While the findings from both our list experiment and the descriptive case evidence from West Bengal are in line with our theoretical argument that mobilization violence is used for selective in-group policing, while demobilization violence is used against opponents and where mobilization is less feasible due to lack of information and control, they could be produced by different processes that result in the same empirical manifestations. We thus substantiate these findings by probing observable implications that follow from our theoretical argument and that should be supported by empirical evidence if our theoretical reasoning is plausible. In particular, we examine patterns in the sensitivity bias that surround self-reporting of these types of violent intimidation. Discrepancies between what people admit in direct question and what they admit in the list-experiment could provide substantive information about how micro-targeting of voters play out, and thus be used to validate some of our theoretical assumptions.

If our expectation that political actors are using violent mobilization and demobilization for different purposes bears empirical relevance, then we would expect several patterns to emerge in our data. First, we expect that individuals are less likely to underreport mobilization violence compared to demobilization violence. As argued above, mobilization violence will often entail symbolic, yet limited acts of more selective violence, where the purpose is to use fear to prevent defection and consolidate power as a way to exercise control within the perpetrator's own electoral base (e.g. Travaglianti; LeBas). Within this context, denying being subject to mobilization-targeting might itself be seen as an act of defiance. Hence, we believe that respondents targeted with mobilization violence will have incentives to affirm its prevalence to external parties when asked directly, to signal that they are compliant. The context of demobilization violence will often be different: levied less selectively and against the backdrop of more contested political control. This in itself might create incentives to underreport individual exposure when directly asked. In addition, the intended purpose of violence – to induce fear and insecurity for disenfranchisement and withdrawal from the political sphere (e.g. Borzyskowski, Daxecker, and Kuhn) – might predispose individuals to not reveal being subject to this form of violence. This is consistent with what we see in the data. When comparing the difference in reported prevalence between the list experiment and the direct question we see that respondents are much more likely to under-report being subject to demobilization violence in the former, compared to mobilization violence. When asked directly, 7.00% of the respondents report to have been violently demobilized (N=586, Table A3). This is a statistically significant difference of roughly 19 percentage points less than what respondents report based on the list experiment (26.08% - 7.00%). That is, respondents are three times more likely to report demobilization in the list experiment compared to the

direct question. For violent mobilization (N=586), however, the difference between the list experiment (11.30%) and the direct question (8.36%) is only roughly 3%, and statistically indistinguishable. Thus, individuals in West Bengal seem to be much less concerned to report having been intimidated to vote for a particular candidate or party than to have been intimidated to not cast their vote.

Second, we expect that admitting to mobilization violence is least sensitive for those who report that they have voted, because their observable actions are in compliance with the coercive strategies.[42] Those who have been targeted with mobilization violence but reported not to have voted are non-compliant with the coercive threat and thus have larger incentives to under-report also in the direct question. We find evidence in line with this. The under-reporting of being exposed to mobilization violence is, with 0.07% percentage points, lowest for those who also report that they have acted in compliance with the strategy, i.e. voted (Table A3). This difference between the list experiment and direct question is substantially negligibly and statistically not significant. The prevalence of mobilization violence in the list experiment is, with 34.76%, still much higher amongst those that report not having voted, but this is significantly under-reported in the direct question, where only 11.90% report to have been violently prevented from casting their vote (i.e., 3 times less). This difference of 22.85 percentage points is both substantially and statistically significant.

The flip side of this reasoning entails that incentives to under-report among subjects exposed to demobilization violence is particularly high for those who report having voted, as reporting both signals non-compliance. The deviance between the direct and the list experiment questions should, in other words, be lower amongst those targeted with demobilization violence who acted in compliance and abstained from voting. This expectation does only seem to hold partially, as demobilization violence is substantially under-reported both by those who vote and those who did not vote. Nevertheless, the prevalence of demobilization among those who have voted is 4 times less in the direct question (5.57%) compared to the list experiment (22.19%), whereas it is only 2 times less among those who have voted (20% vs. 43.22%). This is in line with what we would expect if respondents have incentives to signal compliance.

[42]Survey responses to voting behavior might themselves be subject to social desirability bias. Yet, independent of whether respondents actually did vote, or just stated that they did, they still have incentives to signal compliance to perpetrator intention.

Table A3: Observable implications: under-reporting across mobilization and demobilization violence

	Mobilization			**Demobilization**		
	% LE	*% Direct*	*Δ LE/Direct*	*% LE*	*% Direct*	*Δ LE/Direct*
Sample average	11.30 (9.68)	8.36 (1.14)	2.93 (1.79)	26.08 (9.72)**	7.00 (1.05)	19.08 (2.18)***
Voted	7.99 (10.34)	8.06 (1.18)	0.07 (1.72)	22.19 (10.41)*	5.57 (0.9)	16.61 (2.14)***
Not voted	34.76 (32.3)	11.90 (5.05)	22.85 (8.96)*	43.22 (28.45)	20.00 (6.4)	23.22 (10.59)*

Note: standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$