# Pre-analysis plan
# Attitudes on police reform: Evidence from the Cooperative Election Survey

Mitch Downey, Institute for International Economic Studies, Stockholm University
John Kuk, Department of Political Science, University of Oklahoma
Matthew Nanes, Department of Political Science, Saint Louis University

# Contents


Introduction ..... 4

Research strategy ..... 4

Design and empirical strategy ..... 5

- Making use of baseline information ..... 5
- [CC22_334grid] {grid} Do you support or oppose each of the following proposals? ..... 6
- Analyzing binary outcomes ..... 8
- Experiment 1: Information about policing policies in place in large cities ..... 8
  - Motivation ..... 8
  - Design ..... 8
  - Substantive hypotheses ..... 10
  - Controls ..... 12
  - Empirical implementation ..... 13
  - Additional heterogeneity and tests for the mechanism ..... 16
- Experiment 2: Information about voters' views ..... 17
  - Motivation ..... 17
  - Design ..... 18
  - Substantive hypotheses ..... 20
  - Empirical implementation ..... 21
- Experiment 3: Information about legislative behavior ..... 24
  - Motivation ..... 24
  - Design ..... 24
  - Hypotheses ..... 26
  - Empirical implementation ..... 26
  - Robustness: Excluding Florida ..... 28
  - Additional heterogeneity and tests for the mechanism ..... 28
  - Adjusting for earlier experiments ..... 29
- Experiment 4: Priming the salience of state and local politics ..... 31
  - Motivation ..... 31
  - Design ..... 31
  - Hypotheses ..... 32
  - Empirical implementation ..... 33
- [CC22_363] {single} Do you intend to vote in the 2022 general election on November 8th? ..... 33
  - Additional heterogeneity and tests of the mechanism ..... 36
- Adjusting for inattention ..... 39
- Prioritizing hypotheses ..... 40

References.......................................................................................................................... 41

## Introduction

This project is broadly motivated by the fact that people often take cues from their parties, which can lead them to form policy views they personally have not thought about very much. However, given the importance of partisanship and political views for individuals' identity and self-conception, many people go to great lengths to psychologically defend and justify those policy views. As a result, attempts at overt persuasion by presenting contradictory evidence often backfire as people dig in to their existing beliefs.

Our main question, then, is how might people be convinced to revisit a previously-held belief, particular when it was formed out of partisan identity more than a genuine personal preference? Our overarching idea is that this can be achieved by providing informational corrections which are *tangential* to the targeted belief. In other words, rather than trying to directly correct respondents' views on a policy question, our aim is to provide respondents with new information about questions related to policy, while allowing them ample space to draw their own conclusions and inferences.

In pursuing this aim, we sought to explore one particular issue in depth, rather than seeking to cover a broad range of issues. We sought to target a specific issue where these information corrections are likely to be effective. Specifically, in light of the nationalization of politics, we expect the greatest divergence between individuals' (national) partisan identity and their own personal lived experience to occur in issues where i) the policy in question has a tangible and noticeable impact on most people's lives, ii) state and local governments are especially important, iii) there is wide variation in policies and practices across place, iv) the national salience is high, and v) political elites have connected the issue to broader cultural conflicts. We believe that policing is the most important example of such an issue.

We hypothesize that in this context, modest information interventions can have strong effects when the information relates the policy in question to national-level politics and party identification. We believe this information can affect individuals' policy preferences, partisan loyalties, interest group perceptions, and voting behavior. We are especially interested in opportunities that might lead respondents to disconnect their state and local partisan choices from their national partisan choices. This pre-analysis plan describes a survey that embeds four experiments. Three provide different types of information: one on voters' views, one on legislators' actions, and one on the policy status quo. The fourth increases the salience of state and local politics for policing-related policy.

## Research strategy

The authors purchased a module in the 2022 Cooperative Election Survey (CES). The CES is a well-respected nationally representative election survey in operation since 2006. Modules are available for research teams to purchase. A module allows the team to design their own survey to be conducted among 1,000 respondents, situated within the broader common content survey administered across all (roughly 60,000) respondents.

The module includes a longer survey in October 2022 (the pre-election wave) and a shorter survey in November 2022 (the post-election wave) conducted among the same respondents (i.e., it is a short panel survey). After the election, the CES team (and not our own research team) also work with Catalist to use administrative data to validate whether respondents turned out in the November 2022 election. This is called "validated turnout" and we plan to use it as an outcome.

Prior to finalizing the CES module, we conducted a pilot experiment using Lucid. The results of this pilot informed our experimental design and our pre-analysis plan, but will not be used in the final analysis. Those results are available upon request.

# Design and empirical strategy

The survey will include four experimental assignments, which we will cross randomize. Three experiments will be conducted during the pre-election survey, and a fourth during the post-election survey. For two of our experiments (Experiments 1 and 2, discussed below), we will randomize whether they appear in the pre-election or the post-election survey.

All CES modules are conducted after the common content. The CES 2022 common content has eight questions about policing which will be asked of all respondents before any of our questions and before any of our experiments. We will make use of these baseline views. We will also make use of respondents' partisan leanings asked before our module.

All analyses will use CES-provided sample weights.

## Making use of baseline information

Many of the hypotheses we lay out below are conditional on individuals' pre-experimental characteristics. We will use three sets of pre-experimental characteristics.

First, we will use respondents' partisan lean, specifically their 7-point party identification collected in the CES Common Content (pid7). We will define "Democrats" as those who identify as a "Strong Democrat", "Not very strong Democrat", or independent who tends to lean towards the Democratic Party. Similarly, we will define "Republicans" as "Strong Republican", "Not very strong Republican", or those who tend to lean towards the Republican Party. We will use this classification (based fully on common content questions collected prior to our module being administered) to analyze heterogeneity for every experiment we conduct, as well as for the randomization of experiment 3. In this pre-analysis plan, we will sometimes use the terms "liberal" and "Democrat" interchangeably, and likewise for "conservative" and "Republican". Classification, however, will always be based on party identification and not ideological identification. Finally, some respondents identify as independents with no partisan lean. We consider these respondents neither Democrats nor Republicans. They will not be used in estimation for testing any of the hypotheses (listed below) that depend on party.

Second, we will use baseline characteristics that we collect ourselves. These include the following three sets of questions:

**PRE002**
**Police performance evaluation**
Thinking about the police today, how would you rate the following aspects of police performance?
Rows:

| | |
|---|---|
| PRE002a | Police response times |
| PRE002b | Solving crimes |
| PRE002c | Police-community relations |
| PRE002d | The way the police treat citizens |
| PRE002e | Keeping communities safe |
| PRE002f | Using the right amount of force for each situation |
| PRE002g | Treating different racial and ethnic groups equally |

Columns:

1. Very good

2. Somewhat good
3. Neither good nor bad
4. Somewhat bad
5. Very bad

**PRE003**
**Police approval**
As a whole, do you approve or disapprove of the way the police in the United States are doing their job?
1. Approve
2. Disapprove

**PRE004**
**Police confidence**
How much confidence do you, yourself, have in the police?
1. A great deal
2. A lot
3. A moderate amount
4. A little
5. None

We will define respondents who strongly oppose the police as those who satisfy all three of the following criteria: 1) They answered "Somewhat bad" or "Very bad" to at least 5 of the questions in PRE002, 2) They disapprove of police performance in question PRE003, and 3) They have little or no confidence in the police in question PRE004. Below, we will refer to these respondents as having negative affect towards the police.

We will define respondents who strongly support the police as those who satisfy all three of the following criteria: 1) They answered "Somewhat good" or "Very good" to at least 5 of the questions in PRE002, 2) They approve of police performance in question PRE003, and 3) They have a great deal or a lot of confidence in the police in question PRE004. Below, we will refer to these respondents as having positive affect towards the police.

Third, the common content asks eight questions about policing. The first is:

[CC22_307] {single} Do the police make you feel...?
(Allows one selection)
◯ [1] Mostly safe
◯ [2] Somewhat safe
◯ [3] Somewhat unsafe
◯ [4] Mostly unsafe

The other seven relate to specific policy preferences:

[CC22_334grid] {grid} Do you support or oppose each of the following proposals?

| | **[1]** Support | **[2]** Oppose |
|---|---|---|
| **[CC22_334a]** Eliminate mandatory minimum sentences for non-violent drug offenders. | ◯ | ◯ |

| | | |
|---|---|---|
| **[CC22_334b]** Require police officers to wear body cameras that record all of their activities while on duty. | ◯ | ◯ |
| **[CC22_334c]** Increase the number of police on the street by 10 percent, even if it means fewer funds for other public services. | ◯ | ◯ |
| **[CC22_334d]** Decrease the number of police on the street by 10 percent, and increase funding for other public services | ◯ | ◯ |
| **[CC22_334e]** Ban the use of choke holds by police | ◯ | ◯ |
| **[CC22_334f]** Create a national registry of police who have been investigated for or disciplined for misconduct. | ◯ | ◯ |
| **[CC22_334g]** End the Department of Defense program that sends surplus military weapons and equipment to police departments. | ◯ | ◯ |
| **[CC22_334h]** Allow individuals or their families to sue a police officer for damages if the officer is found to have “recklessly disregarded” the individual’s rights. | ◯ | ◯ |

Note that CC22_334a is related to sentencing but not policing. Thus, we will only consider the seven questions about policing (CC22_334b-CC22_334h). Using only respondents who were NOT assigned to our sample (roughly 59,000 out of the full 60,000 person sample), we will restrict to those classified (by 7-point party identification) as either Democrats or Republicans. We will create a dummy variable for whether the Respondent is a Democrat, and we will estimate linear probability LASSO regression predicting whether a respondent is a Democrat based on their answers to the 7 policy questions listed above (CC22_334b-CC22_334h), as well as all two-way interactions. We will use the fitted value of this regression as our respondents’ “Police reform liberalism”: Higher values means that the respondent holds views on policing policies which are more similar to Democratic voters than Republican voters. Voters with very low scores have views much more like Republicans than Democrats, and voters with intermediate scores have views that do not cleanly belong to either party’s voters.

## Analyzing binary outcomes

Most of the outcomes we propose to analyze are binary. We have conducted a series of power analyses (available upon request) and concluded that linear probability models and logistic regressions have (nearly) identical statistical power for the sample sizes and base rates that we are considering. Thus, all analyses will be based on linear probability models, which (1) are easier to interpret, (2) facilitate much easier comparison across different regressions, samples, and outcomes (since the coefficients do not mechanically depend on the base rates), (3) lend themselves immediately to instrumental variables regressions with no further assumptions (which we will use in one analysis; see Hypothesis 4.4), and (4) avoid concerns about incidental parameter bias, since we do not know how many controls we will use (see the below discussion of our data-driven approach to selecting optimal controls).

## Experiment 1: Information about policing policies in place in large cities

### Motivation

We suspect that the heated political rhetoric around policing has led many respondents to form beliefs which are false. We expect that a better understanding of the prevalence of certain practices can change respondents' support for specific policies. For example, finding that a particular practice is more widespread than one expected might increase support for the practice becoming mandatory (either because it must not be as disruptive to police operations as the respondent expected, or because it must actually matter for police departments if they voluntarily adopted it), or could reduce support for the practice becoming mandatory (because discovering that it is widespread might reduce confidence in its efficacy).

Ultimately, we view policy preferences as motivated by the expected effects of the policy being considered, and knowledge about the prevalence of the policy indirectly informs its expected effects. For example, we expect many respondents to hold relatively broad views on policing, such as "Policing in the United States is not working," for instance. In the presence of such broad views, we expect that informing respondents that a particular practice is widespread will lead them to conclude that the practice is ineffective, since this is the best way to reconcile the idea that policing is not working with the idea that a practice is widespread.[1] An alternative is to provide information directly about the effects of these practices, but as we noted above, existing research suggests that efforts at overt persuasion tend to be ineffective or even backfire. Since we expect respondents to already hold fairly firm views about the state of policing in the United States, we opted instead to provide information about the prevalence of the practices.

### Design

We ask respondents to guess the prevalence of four different practices among large police departments. The treatment groups is then informed of the actual prevalence of the practice according to the most recent Bureau of Justice Statistics' Law Enforcement Management and Administrative Statistics (LEMAS) Survey (conducted 2016). We then ask respondents whether they support a particular reform proposal that is closely connected to the practice. Finally, we ask a series of questions about perceptions of interest groups' and parties' policing agendas, as well as a broad question about reform orientation.

[1] It is worth mentioning that our design will allow us to test these sorts of predictions. For example, we will be able to determine whether, in response to treatment, respondents who believed these practices to be rare become more likely to "oppose reforms like the ones we asked about above" because "reforms like these are not enough."

The experimental protocol is as follows. Note that all questions are asked of all respondents, and the only difference between treatment and control is whether true information is revealed after we collect their beliefs.

We want to ask you some questions about the policies in place in large police departments.

We're going to ask you about the 50 largest police departments in the country. All of these departments have more than 1,000 officers.

The smallest of these departments are the St. Louis Police Department, the Oklahoma City Police Department, and the Cincinnati Police Department. Police Departments in cities larger than these are also included in the list.

**EXP4_001a (Background checks and misconduct registry)**
Among these departments: What fraction do you believe conduct background checks before hiring an officer?
**TREATMENT ONLY, ONLY AFTER ANSWERING**: It's actually 100% of departments.
**EXP4_001b**
Do you support creating a national registry of police who have been investigated for or disciplined for misconduct?

1. Support
2. Oppose

**EXP4_002a (Racial diversity and affirmative action)**
What fraction of police officers do you think are Black?
**TREATMENT ONLY, ONLY AFTER ANSWERING**: It's actually 18% of officers.
**EXP4_002b**
Do you support requiring that new police hires are racially representative of their communities (affirmative action)?

1. Support
2. Oppose

**EXP4_003a (Community policing)**
What fraction of departments require training in being responsive to community needs?
**TREATMENT ONLY, ONLY AFTER ANSWERING:** It's actually 90% of departments.
**EXP4_003b**
Should all departments be required to train officers in being responsive to community needs?

1. Support
2. Oppose

**EXP4_004a (Choke holds)**
In what fraction of departments are choke holds and neck restraints banned from use?
**TREATMENT ONLY, ONLY AFTER ANSWERING:** It's actually 58% of departments.
**EXP4_004b**
Do you support banning the use of choke holds by police?

1. Support
2. Oppose

**EXP4_005: Group agenda for policing**
Please rate whether you agree or disagree with the following statements
Rows:

EXP4_005a The Democratic Party's policy agenda for policing is dangerous
EXP4_005b The Republican Party's policy agenda for policing is dangerous
EXP4_005c The Democratic Party's policy agenda for policing is likely to improve policing in important ways
EXP4_005d The Republican Party's policy agenda for policing is likely to improve policing in important ways
EXP4_005e Black Lives Matter's policy agenda for policing is dangerous
EXP4_005f Police unions' policy agenda for policing is dangerous
EXP4_005g Black Lives Matter's policy agenda for policing is likely to improve policing in important ways
EXP4_005h Police unions' policy agenda for policing is likely to improve policing in important ways

Columns
1. Strongly agree
2. Somewhat agree
3. Neither agree nor disagree
4. Somewhat disagree
5. Strongly disagree

**EXP4_006: Reform preference**
Some people oppose reforms like the ones we asked about above, and they offer various reasons. Do you agree with any of the following statements?
1. Reforms like these are not necessary.
2. Reforms like these will make it too difficult for police to do their jobs.
3. Reforms like these are not enough.
4. I don't agree with any of these statements.

**EXP4_007: Police reform attitude**
Overall, which of the following statements do you most agree with?
1. Because of fundamental problems, policing as an institution needs to be completely rebuilt.
2. Although there are problems with policing, necessary changes can be made through reforms within the current system.
3. Little or nothing needs to be done to reform policing.

**EXP4_008: Decrease police on the street**
Should the United States decrease the number of police on the street by at least 10% and shift the funding towards other public services?
1. Yes
2. No

**EXP4_009: Increase police on the street**
Should the United States increase the number of police on the street by at least 10%, shifting the money from other public services?
1. Yes
2. No

## Substantive hypotheses

***Hypothesis 1.1A***: Information increases conservatives', but not liberals', support for reforms.

Liberals frame interventions like these as "common sense reforms." Conservatives typically argue either that they are unnecessary because the police do a good job without them (the "reforms are unnecessary” argument), or that they will interfere with police operations (the “reforms are disruptive” argument). We think that informing conservatives that these reforms are already fairly widespread will make them appear more important than previously believed (since police departments have determined a need for them already) or less threatening to police operations (since many departments already function with them in place). Thus, we expect information to make conservatives more supportive of making these reforms mandatory. For liberals, we believe support will already be sufficiently strong that information has no effects.

***Hypothesis 1.1B***: Treatment effects among conservatives will be driven by those with low baseline beliefs about how widespread these practices are.

Conservatives opposed to these reforms typically view them either as unnecessary or disruptive. Both views imply that these practices should be rare. We believe information should be most effective among those who are most “surprised” by the high prevalence of these practices. (Note: We will use question EXP4_006 to distinguish between these two conservative arguments; see mechanism test 1.1 below.)

***Hypothesis 1.2A***: Information will make conservatives more supportive of police reform, but will make liberals more supportive of complete overhaul of policing.

As discussed above, we expect treatment to increase conservative support for the specific policies we ask about. However, we also expect it to shift them towards support for reform more broadly. For liberals, on the other hand, we expect that information showing that these practices are already very widespread will reduce their support for incrementalism or the “reform” position advocating for “common sense” policies. Instead, we expect this information to push them towards the view that “policing as an institution needs to be completely rebuilt”.

***Hypothesis 1.2B***: Treatment effects will be driven by those with low baseline beliefs about how widespread these practices are.

We expect the responses above to be stronger among those who were more surprised to find that these practices are as widespread as they are.

***Hypothesis 1.3A***: Information will make conservatives friendlier towards reform-oriented groups and more skeptical of anti-reform groups, while it will make liberals more skeptical of the more moderate reform-oriented groups.

Above, we argued that we expect information to increase conservatives’ openness to reform, and to make liberals prefer complete overhaul rather than incremental reform. Here, we note that these positions are attached to specific parties and interest groups in American politics: The Democratic Party typically supports moderate reforms while the Republican Party typically supports only a very limited reform package. Black Lives Matter typically adopts a stronger and more radical pro-reform stance than the Democratic Party, while police unions are typically even more strongly opposed to reform than the Republican Party is. We expect changes in support for parties and interest groups to mirror changes in stances on reform.

***Hypothesis 1.3B***: Treatment effects will be driven by those with low baseline beliefs about how widespread these practices are.

We expect the responses above to be stronger among those who were more surprised to find that these practices are as widespread as they are.

***Hypothesis 1.4***: Treatment will not have significant effects on respondents’ views on police funding.

As we argued above, we view policing is a complex issue that has been simplified to fit into broader cultural conflicts, and where information is filtered through nationalized media and partisan frames. As such, we see many substantively important policy issues as low salience (because of the simplifying nature of cultural conflict) where beliefs are wrong (because national media abstracts from respondents' local context). However, police funding is a very high salience issue which is typically discussed as funding *cuts* or *increases* (i.e., always *relative* to the current status quo, regardless of what the status quo is). Thus, we expect respondents' minds to be made up on issues of policing funding changes, and expect our information to have no significant effect on preferences.

## Controls

Our main interest is in the effects of treatment, which we randomly assign. Thus, treatment is uncorrelated with all other variables (asymptotically), and controls are not necessary for causal inference. However, as is well-known, inclusion of controls can improve precision and reduce standard errors in an experimental setting. Here, we discuss three types of controls that we will use in the experiments below. For simplicity of exposition, all estimating equations presented below suppress the control variables, although we plan to include them in the final analysis.

First, as noted above, our protocol randomizes whether Experiment 1 is conducted pre-election or post-election. We will always control for a dummy for this randomized characteristic. Note that when Experiment 1 is conducted post-election, we will not control for any terms describing the treatment conditions of the pre-election assignment (see Experiment 3 for a discussion of this topic). Our experiment also randomizes the order in which EXP4_001-EXP4_004 are asked, and all analyses will include dummy variables for order to account for any order effects.

Second, four of our key dependent variables are asked in the CES Common Content: "Do you support creating a national registry of police who have been investigated for or disciplined for misconduct?" (our EXP4_001b, CES' CC22_334f), "Do you support banning the use of choke holds by police?" (our EXP4_004b, CES' CC22_334e), "Should the United States decrease the number of police on the street by at least 10% and shift the funding towards other public services?" (our EXP4_008, CES' CC22_334d), and "Should the United States increase the number of police on the street by at least 10%, shifting the money from other public services?" (our EXP4_009, CES' CC22_334c). In all cases, we will control for the CES response in analyzing our treatment. This makes this a within-subjects design. Since the wording is identical, controlling for pre-experimental views is more efficient than controlling for baseline demographics. When these questions are the dependent variable, we will control for the pre-experimental response but not baseline demographics (discussed below). When these questions are included in the construction of the dependent variable (e.g., counts of how many reforms the respondent supports), we will control for the pre-experimental response and baseline demographics. Both sets of regressions will continue to control for whether the experiment was conducted pre- or post-election.

Third, we will include a set of demographic controls that are correlated with support for policing reforms. As noted above, this is to improve precision of our estimated treatment effects, not to reduce bias or bolster causal inference. To select demographic controls, we will use the CES sample which was *not* assigned our experimental module (most of the sample). We will regress our continuous measure of "police reform liberalism" (discussed above) on a large set of demographic characteristics from the CES Common Content using a LASSO regression.[2] This will determine a

[2] Specifically, we will use a quadratic in age (2022-birthyr), a dummy variable for gender (gender), a dummy variable for having a 4-year college degree or a postgraduate degree (educ), dummy variables for all 7 race categories asked about in race (which includes "Hispanic or Latino"), a dummy variable for being registered to

parsimonious set of key predictors of support for police reform, and we will include these demographic variables in all regressions below.

Note that our application of this control variable selection approach must account for the fact that the training data (the CES sample that was not assigned our module) is roughly 59,000 individuals, while our final analysis sample will only be roughly 1,000 individuals. To account for this, we will first run the LASSO regression using all 59,000 non-module respondents using Stata's lasso2 command with all parameters set to their defaults. We will then draw 50 independent random samples (with replacement) from the non-module respondents. Each of these 50 samples will include 1,000 respondents only. For each of these 50 samples, we will estimate the same LASSO regression, and determine the number of variables selected. We will calculate the median number of variables selected across each of these 50 samples, and we will choose that number of variables from the 59,000-respondent LASSO (obviously, choosing the most important variables).

## Empirical implementation

As noted above, all binary outcome variables will be analyzed with a linear probability model. Let $Supp^J$ be a dummy variable equaling 1 if the respondent reports supporting the policy described in EXP4_00Jb for $J \in \{1{:}\, Misconduct\ database, 2{:}\, Affirmative\ action, 3{:}\, Community\ policing, 4{:}\, Choke\ holds\}$. Let T1 be a dummy variable equaling 1 if the respondent was assigned to treatment in experiment 1.

Separately for Democratic and Republican respondents,[3] we will estimate a linear regression of the form:

$$Supp^J = \alpha^J + \beta^J T1 + \varepsilon \qquad (1)$$

In addition to running four regressions of the form described in equation (1), we will also estimate a single regression for a composite of all four reforms that we ask about. Define $SuppCount$ as a count of the number of policies that the respondent supports, which can range from 0 to 4.[4] We will estimate a Negative binomial regression to determine whether treatment moves total support. This regression will be of the form:

$$SuppCount \sim NegBin(\exp(\alpha + \beta T1)) \qquad (2)$$

We will test whether $\beta > 0$.

***Hypothesis 1.1A***: Information increases conservatives' support for reforms but not liberals': For conservatives, $\beta^J, \beta > 0$ while for liberals $\beta^J, \beta = 0$.

vote, dummy variables for the 9 Census Divisions based on state of residence (inputstate), dummy variables for 7-point ideological self-identification (CC22_340grid), dummy variables for 4 responses on urban/rural (urbancity), dummy variables for 6 responses on November 2022 voting intentions (CC22_363), dummy variables for 7-point partisan identification (pid7), dummy variables for 4-point self-reported importance of religion (pew_religimp), a dummy variable for being married (marstat), a dummy variable for being a citizen (cit1), dummy variables for the five possible responses about news interest (newsint), a quadratic in income (where we impute income as the midpoint of the ranges reported in faminc_new), the presence of children (child18), union membership (union), and 2020 Presidential vote choice (presvote20post).

[3] We will also estimate this for independent respondents who report not leaning towards either party. However, that will be a small sample, and that will be an under-powered test that we do not expect to be statistically significant (and for which we do not have a hypothesis).

[4] $SuppCount \equiv \sum_j Supp^j$

To test for heterogeneous effects, we will define $Low^j$ to record the gap between respondents' believed prevalence and the true prevalence, as a fraction of the true prevalence:

$$Low^J = \begin{cases} 1 - Belief^J/Truth^J \ if \ Belief^J < Truth^J \\ 0 \ otherwise \end{cases}$$

In other words, $Low^j$ varies from 0 to 1, with zero indicating the respondent was either correct about the true prevalence or over-estimated the prevalence, 1 indicating that the respondent erroneously thought no departments enacted the practice, 0.5 indicating that the respondent thought the practice was half as common as it actually is, etc.

Separately for Democratic and Republican respondents, we will estimate:

$$Supp^J = \alpha^J + \beta_1^J Low^J + \beta_2^J T1 + \beta_3^J (T1 \times Low^J) + \varepsilon$$

Similarly, we wish to test whether total support (rather than only question by question) is affected by treatment. To do so, we will define $MeanLow$ as the average of the four $Low^J$ variables. We will then estimate:

$$SuppCount \sim NegBin(\exp(\alpha + \beta_1 MeanLow + \beta_2 T1 + \beta_3 T1 \times MeanLow))$$

***Hypothesis 1.1B***: Treatment effects among conservatives will be driven by those with low baseline beliefs about how widespread these practices are: For conservatives, we expect that $\beta_2^J, \beta_2$ is small or zero, while $\beta_3^J, \beta_3 > 0$.

To test for effects on general orientation towards reform, overhaul, or the status quo, we will define three dummy variables:

- $Reform = 1$ if the respondent chose "Although there are problems with policing, necessary changes can be made through reforms within the current system." in response to Question EXP4_007, and equals zero otherwise.
- $Overhaul = 1$ if the respondent chose "Because of fundamental problems, policing as an institution needs to be completely rebuilt." in response to Question EXP4_007, and equals zero otherwise.
- $StatusQuo = 1$ if the respondent chose "Little or nothing needs to be done to reform policing." in response to Question EXP4_007, and equals zero otherwise.

We will then estimate:

$$\mathrm{Y} = \alpha^Y + \beta^Y T1 + \varepsilon \tag{3}$$

for $Y \in \{Reform, Overhaul, StatusQuo\}$.

***Hypothesis 1.2A***: Information will make conservatives more supportive of police reform, but will make liberals more supportive of complete overhaul of policing: For conservatives, $\beta^{Reform} > 0$ and $\beta^{StatusQuo} < 0$, while for liberals, $\beta^{Reform} < 0$ and $\beta^{Overhaul} > 0$.

We will again test for heterogeneous effects among those with persistently low beliefs.

$$\mathrm{Y} = \alpha^Y + \beta_1^Y MeanLow + \beta_2^Y T1 + \beta_3^Y T1 \times MeanLow + \varepsilon$$

***Hypothesis 1.2B***: Treatment effects will be driven by those with low baseline beliefs about how widespread these practices are: For conservatives, $\beta_2^{Reform} + \beta_3^{Reform} > 0$, $\beta_3^{Reform} > 0$,

$\beta_2^{StatusQuo} + \beta_3^{StatusQuo} < 0$, and $\beta_3^{StatusQuo} < 0$. For liberals, $\beta_2^{Reform} + \beta_3^{Reform} < 0$, $\beta_3^{Reform} < 0$, $\beta_2^{Overhaul} + \beta_3^{Overhaul} > 0$, and $\beta_3^{Overhaul} > 0$.

Finally, will estimate a series of linear regressions to see whether treatment affects perspectives on (1) whether the respondent believes that group *G's* policies on policing are dangerous, for $G \in \{BLM, Dem, Rep, Unions\}$, and (2) whether the respondent believes that group *G's* policies on policing are likely to improve policing in important ways.

We will estimate the regressions in two different ways, and expect both to yield substantively similar results.

First, we will define eight dummy variables for respondents' views on the eight groups we ask about in Question EXP4_005. $Danger^G$ will indicate that the respondent either Strongly agrees or Somewhat agrees that group $G$'s policies on policing are dangerous. $Help^G$ will indicate that the respondent either Strongly agrees or Somewhat agrees that group $G$'s policies on policing are likely to improve policing in important ways.

We will estimate a series of eight linear regressions to see whether treatment moves these views:

$$Danger^G = \delta_0^G + \delta_1^G T1 + \varepsilon \tag{4}$$

$$Help^G = \theta_0^G + \theta_1^G T1 + \varepsilon \tag{5}$$

Second, we will estimate linear regressions which use responses to the 4-point ordinal scale questions about the danger and helpfulness of these groups' policies:

$$DangerOrdinal^G = \delta_0^G + \delta_1^G T1 + \varepsilon \tag{6}$$

$$HelpOrdinal^G = \theta_0^G + \theta_1^G T1 + \varepsilon \tag{7}$$

***Hypothesis 1.3A***: Information will make conservatives friendlier towards reform-oriented groups and more skeptical of anti-report groups, while it will make liberals more skeptical of the more moderate reform-oriented groups: For conservatives, $\delta_1^G < 0, \theta_1^G > 0$ for $G \in \{BLM, Dem\}$ and $\delta_1^G > 0, \theta_1^G < 0$ for $G \in \{Rep, Unions\}$. For liberals $\delta_1^G < 0, \theta_1^G > 0$ for $G = BLM$ and $\delta_1^G > 0, \theta_1^G < 0$ for $G \in \{Dem, Rep\}$.

We will again test for heterogeneous effects depending on baseline beliefs by estimating:

$$Danger^G = \delta_0^G + \delta_1^G T1 + \delta_2^G MeanLow + \delta_3^G T1 \times MeanLow + \varepsilon$$

$$Help^G = \theta_0^G + \theta_1^G T1 + \theta_2^G MeanLow + \theta_3^G T1 \times MeanLow + \varepsilon$$

$$DangerOrdinal^G = \delta_0^G + \delta_1^G T1 + \delta_2^G MeanLow + \delta_3^G T1 \times MeanLow + \varepsilon$$

$$HelpOrdinal^G = \theta_0^G + \theta_1^G T1 + \theta_2^G MeanLow + \theta_3^G T1 \times MeanLow + \varepsilon$$

***Hypothesis 1.3B***: Treatment effects will be driven by those with low baseline beliefs about how widespread these practices are: For conservatives, $\delta_3^G < 0, \theta_3^G > 0$ for $G \in \{BLM, Dem\}$ and $\delta_3^G > 0, \theta_3^G < 0$ for $G \in \{Rep, Unions\}$. For liberals $\delta_3^G < 0, \theta_3^G > 0$ for $G = BLM$ and $\delta_3^G > 0, \theta_3^G < 0$ for $G \in \{Dem, Rep\}$.

Finally, we will estimate effects on support for increasing and decreasing police funding based by estimating the following two regressions separately for conservative and liberal respondents (as noted above, we will control for baseline views expressed in the CES Common Core, making this a within-subjects design).

$$Supp^{inc} = \alpha^{inc} + \beta^{inc}T1 + \varepsilon$$

$$Supp^{dec} = \alpha^{dec} + \beta^{dec}T1 + \varepsilon$$

where $Supp^{inc}$ and $Supp^{dec}$ are dummy variables for supporting increasing and decreasing police funding, respectively.

***Hypothesis 1.4***: Treatment will not have significant effects on respondents' views on police funding. We expect that neither $\beta^{inc}$ nor $\beta^{dec}$ will be significant for Republicans or Democrats.

## Additional heterogeneity and tests for the mechanism

Here we describe several tests we will conduct which are built into the design of our experiment but where we do not have a specific hypothesis.

***Mechanism test 1.1***: We will test whether information increases support for reform by addressing the "reforms are unnecessary" argument or by addressing the "reforms are disruptive" argument.

Based on Question EXP4_006, we will code four dummy variables:

- $Unnecessary = 1$ if the respondent chose "Reforms like these are not necessary." in response to Question EXP4_006, and equals zero otherwise.
- $Disruptive = 1$ if the respondent chose "Reforms like these will make it too difficult for police to do their jobs." in response to Question EXP4_006, and equals zero otherwise.
- $Inadequate = 1$ if the respondent chose "Reforms like these are not enough." in response to Question EXP4_006, and equals zero otherwise.
- $None = 1$ if the respondent chose "I don't agree with any of these statements." in response to Question EXP4_006, and equals zero otherwise.

Our main interest is in focusing on conservative respondents and regressing these four dummy variables on treatment:

$$Y = \alpha^{Y} + \beta^{Y}T1 + \varepsilon$$

We expect that $\beta^{None} > 0$. If $\beta^{Unnecessary} < 0$ then we will conclude that information increases conservative support for reform because the fact that these practices are widespread suggests that they must have some value. If $\beta^{Disruptive} < 0$ then we will conclude that information increases conservative support for reform because it reduces the belief that implementing these practices would be excessively disruptive to police operations.

***Mechanism test 1.2***: As argued above, we expect that partisanship leads to sharply different responses to new information. Specifically, we argued that information would increase Republicans' support reform, but would not affect Democrats' support (and may even reduce support by increasing the belief that these "common sense" reforms are ineffective). As is well-documented, Democrats and Republicans hold different views on the police and different support for reform measures. We will use baseline information to test whether partisanship matters for generating differential responses to treatment, above and beyond any role for "affect" towards the police or for baseline openness to reform.

As described above, we will create dummy variables for whether a respondent strongly opposes ("negative affect") or strongly supports ("positive affect") the police. Note that it is possible for respondents to have neither positive nor negative affect towards the police.

We also described above how we will generate a continuous measure of "police reform liberalism". We will convert this continuous measure to be two dummy variables by creating one dummy variable for respondents in the bottom quartile ("anti-reform") and one dummy variable for respondents in the top quartile ("pro-reform"). Thus, roughly half of respondents will be classified as neither pro-reform nor anti-reform.

To test whether partisanship matters beyond affect towards the police or pro-/anti-reform attitudes, we will estimate a series of regressions based on the full sample:

$$SuppCount \sim NegBin[\alpha_0 + \alpha_R Rep + \alpha_D Dem + \alpha_T T1 + \alpha_{RT} Rep \times T1 + \alpha_{DT} Dem \times T1]$$

$$\begin{aligned} SuppCount \sim NegBin[&\beta_0 + \beta_R Rep + \beta_D Dem + \beta_T T1 + \beta_{RT} Rep \times T1 + \beta_{DT} Dem \times T1 \\ &+ \gamma_N NegAffect + \gamma_P PosAffect + \gamma_{NT} NegAffect \times T1 + \gamma_{PT} PosAffect \times T1] \end{aligned}$$

$$\begin{aligned} SuppCount \sim NegBin[&\theta_0 + \theta_R Rep + \theta_D Dem + \theta_T T1 + \theta_{RT} Rep \times T1 + \theta_{DT} Dem \times T1 \\ &+ \lambda_P ProRef + \lambda_A AntiRef + \lambda_{PT} ProRef \times T1 + \lambda_{AT} AntiRef \times T1] \end{aligned}$$

As we argued above, we expect Democrats and Republicans to respond differently to our information treatment. This can be tested by testing the null that $\alpha_{RT} - \alpha_{DT} = 0$. If that null is rejected, then we conclude that Democrats and Republicans respond differently to the same information.

We can test whether those with positive and negative affect towards the police respond differently by testing the null that $\gamma_{NT} - \gamma_{PT} = 0$. If partisanship matters above and beyond any differences in affect, then we would continue to reject the null $\beta_{RT} - \beta_{DT} = 0$ (i.e., partisanship would be important even after accounting for heterogeneous treatment effects by respondent affect).

We can test whether those who support or oppose reforms at baseline respond differently by testing the null that $\lambda_{PT} - \lambda_{AT} = 0$. If partisanship matters above and beyond any differences in baseline reform support, then we would continue to reject the null $\theta_{RT} - \theta_{DT} = 0$ (i.e., partisanship would be important even after accounting for heterogeneous treatment effects by respondent support for reforms).

***Mechanism test 1.3***: It is plausible that our treatment effects are different among those randomly assigned to receive experiment 1 during the pre-election survey (October) or the post-election survey (November). We will interact a dummy variable for pre-/post-election status with our main treatments (equations 2, 3, 4, and 5, above), although we have no hypotheses about how the effects will differ.

## Experiment 2: Information about voters' views

### Motivation

In light of heightened affective party polarization, political homogeneity in social networks, and increased consumption of traditional and social media that is aligned with ones existing views, we hypothesize that people have various inaccurate beliefs about the actual views of voters of the opposite party.[5] In many cultural narratives, we believe that Democrats are portrayed as being "anti-cop" while Republicans are portrayed as defending the police without exception. Both are oversimplifications.

In this experiment, we will first collect descriptive evidence from a nationally representative sample about the actual prevalence of a view, as well as the beliefs of same-partisans and opposite-partisans about that view's prevalence. We will then experimentally test whether accurate information about the beliefs of the parties induces individuals to change their views on their other party, their own party, or the importance of police reform.

In a pilot experiment, we collected respondents' views and beliefs on several policing-related questions. For a variety of reasons, we concluded that this was too cumbersome for respondents. Thus, in our CES experiment, we ask only about views and beliefs for a single question: Do you think the police are doing a good job holding officers accountable when misconduct occurs? We selected this specific question because accountability has been at the heart of most recent policing debates, and most reforms being advocated for hinge crucially on whether people believe that police officers are already held accountable or not.

Moreover, this is an issue with a meaningful but not infinite partisan divide. According to Pew in June 2020, only 13% of Democrats and 51% of Republicans answered "Yes". Thus, on the one hand, focusing on differences in views on accountability does illustrate some of the real partisan divides on police-related views and does not attempt to hide the important differences. On the other hand, the fact that 48% of Republicans feel that the police *do not* do a good job holding officers accountable illustrates a clear gap between actual voters' actual views and the oversimplified narrative of partisan identity. Thus, we felt that this was a good choice of question to fairly reflect true partisan gaps while also challenging expectations.

After collecting respondents' baseline beliefs about both parties' views, we inform them of the true prevalence of these views. We then ask a series of questions to explore respondents' views on the parties, their policy agendas, broad reform orientation, and views on allied interest groups.

## Design

As noted above, our design randomizes respondents to receive either Experiment 1 OR Experiment 4 in the pre-election survey, and to receive the other in the post-election survey.

**EXP2_001: Police accountability**

Do you think the police are doing a good job holding officers accountable when misconduct occurs?

1. Yes
2. No

**EXP2_002: Democratic voters agree level**

How many **<u>Democratic</u>** voters do you think agree with you about that?

1. Very few (0-20%)
2. Some, but not most (20-40%)
3. About half (40-60%)
4. Most, but not all (60-80%)
5. Nearly all (80-100%)

**EXP2_003: Republican voters agree level**

[5] Similarly, Ahler and Sood (2018) show that voters hold inaccurate beliefs about the demographic composition of the opposite party.

How many **Republican** voters do you think agree with you about that?

1. Very few (0-20%)
2. Some, but not most (20-40%)
3. About half (40-60%)
4. Most, but not all (60-80%)
5. Nearly all (80-100%)

**EXP2_004: Information**

**AFTER ANSWERING EXP2_004, IF TREATMENT, AND IF EXP2_001 = = YES**:
Actually, 13% of Democrats and 51% of Republicans agree with you.

**AFTER ANSWERING EXP2_004, IF TREATMENT, AND IF EXP2_001 = = NO**:
Actually, 87% of Democrats and 47% of Republicans agree with you.

**EXP2_005: Party evaluation**

Please rate whether you agree or disagree with the following statements

Rows:

EXP2_005a The Democratic Party's efforts to increase police accountability go too far.
EXP2_005b The Democratic Party's efforts to increase police accountability don't go far enough.
EXP2_005c The Republican Party's efforts to increase police accountability go too far.
EXP2_005d The Republican Party's efforts to increase police accountability don't go far enough.
EXP2_005e The Democratic Party does a good job representing its voters' views.
EXP2_005f The Republican Party does a good job representing its voters' views.

Columns

1. Strongly agree
2. Somewhat agree
3. Neither agree nor disagree
4. Somewhat disagree
5. Strongly disagree

**EXP2_006: Police reform attitude 3pts**

Overall, which of the following statements do you most agree with?

1. Because of fundamental problems, policing as an institution needs to be completely rebuilt.
2. Although there are problems with policing, necessary changes can be made through reforms within the current system.
3. Little or nothing needs to be done to reform policing.

**EXP2_007: Group agenda for policing**

Please rate whether you agree or disagree with the following statements

Rows:

EXP2_007a The Democratic Party's policy agenda for policing is dangerous
EXP2_007b The Republican Party's policy agenda for policing is dangerous
EXP2_007c The Democratic Party's policy agenda for policing is likely to improve policing in important ways
EXP2_007d The Republican Party's policy agenda for policing is likely to improve policing in important ways
EXP2_007e Black Lives Matter's policy agenda for policing is dangerous
EXP2_007f Police unions' policy agenda for policing is dangerous
EXP2_007g Black Lives Matter's policy agenda for policing is likely to improve policing in important ways
EXP2_007h Police unions' policy agenda for policing is likely to improve policing in important ways

Columns

1. Strongly agree
2. Somewhat agree
3. Neither agree nor disagree
4. Somewhat disagree
5. Strongly disagree

**EXP2_008: Decrease police on the street**
Should the United States decrease the number of police on the street by at least 10% and shift the funding towards other public services?
1. Yes
2. No

**EXP2_009: Increase police on the street**
Should the United States increase the number of police on the street by at least 10%, shifting the money from other public services?
1. Yes
2. No

## Substantive hypotheses

***Hypothesis 2.1***: In response to our information treatment, *Democrats* who *agree* with their party (i.e., believe the police *are not* doing a good job holding officers accountable) will
***A***: become more frustrated by the Democratic Party
***B***: become more frustrated by the Republican Party
***C***: become more supportive of Black Lives Matter
***D***: shift towards a posture of complete overhaul rather than incremental reform

***Hypothesis 2.2***: In response to our information treatment, *Democrats* who *disagree* with their party (i.e., believe the police *are* doing a good job holding officers accountable) will
***A***: become more concerned about the Democratic Party's agenda
***B***: become more open to the Republican Party's agenda
***C***: become more concerned about Black Lives Matter

***Hypothesis 2.3***: In response to our information treatment, *Republicans* who *agree* with their party (i.e., believe the police *are* doing a good job holding officers accountable) will
***A***: become more concerned about the Democratic Party's agenda
***B***: shift towards a posture of no change in policing

***Hypothesis 2.4***: In response to our information treatment, *Republicans* who *disagree* with their party (i.e., believe the police are not doing a good job holding officers accountable) will
***A***: become more open to the Democratic Party's agenda
***B***: become more concerned about the Republican Party's agenda
***C***: become more frustrated that the Republican Party does not represent its voters
***D***: become more concerned about police unions

We think voters intuitively understand that policy platforms are compromises of constituents' views. We also expect most voters to hold inaccurate beliefs about the prevalence of accountability-related views. In the nationally representative survey data that we based our information treatment on (from Pew), 13% of Democrats believe the police are doing a good job holding officers accountable, and 51% of Republicans believe this. We expect that most voters would expect both numbers to be higher, both for Democrats and for Republicans. We think that voters will be surprised by the unanimity among Democrats, and the fact that this is a roughly 50/50 question for Republicans. All of our hypotheses are based on this expectation, but it is important to acknowledge that our experimental design will allow us to confirm or reject this hypothesis because we collect data on beliefs prior to the information treatment.

With that expectation in mind, we think that Democrats who support increased accountability (87% of Democrats) will be frustrated that their party has not achieved more, despite near complete unanimity on the topic among Democrats, and that Republicans are nearly evenly divided (Hypothesis 2.1A). We also expect this to increase frustration with the Republican Party by leading Democrats to conclude that the Republican Party's obstructionist position is not even strongly supported by their own voters (2.1B). Overall, we expect this frustration to lead to a broader skepticism of the traditional political system, leading to more support for an "outsider" reform group like Black Lives Matter (2.1C) and a broader posture of complete overhaul rather than incremental reform (2.1D).

Some Democrats, however, do believe that the police are doing a good job holding officers accountable. As noted above, we think that these voters will be surprised to learn about near-unanimity among their co-partisans. Since we think they understand that parties aim to be representative of their voters, we think they will update their beliefs to increasingly feel that the Democratic Party is too radical on reform and accountability (2.2A). We also think that they will be surprised that Republicans are less united and more divided than they expected. Thus, we expect them to find the Republican position less extreme than they previously believed (2.2B). We suspect all voters view Black Lives Matter as being "to the left" of (or more strongly pro-reform than) the Democratic Party. We think that as these Democratic voters become more concerned about their own party, they will similarly increasingly feel that Black Lives Matter is too extreme (2.2C).

As noted above, we expect Republicans to be surprised by the unanimity of the Democratic Party's consensus on this issue and, given their understanding of the party's position as being a compromise among their voters, will come to believe the party is more radical than previously expected. We think this will lead to more concerns about the Democratic Party's policy agenda (2.3A) and even some backlash in which they increasingly oppose any reforms (2.3B).

Finally, we suspect that the Republicans who *do not* believe the police are doing a good job holding officers accountable will 1) see themselves as a substantial minority of their party and 2) see their party as broadly opposed to reform. In light of this, we expect that informing them that they are a much larger share of the party than they expected will lead them to become more frustrated by their party's agenda (2.4B) and more interested in the Democratic Party's policy agenda as an alternative (2.4A). We expect that they will become more frustrated that their party does not more actively support reform, despite having a reasonable amount of support among Republican voters (2.4C), and that they will attribute this to the influence of police unions on the party (2.4D).

## Empirical implementation

Our design (which collects respondents' perceptions prior to providing them with information) directly allows us to identify Republicans and Democrats who are out of step with their party. It also measures all the outcomes that we describe above.

All regressions presented below will control for variables selected by the LASSO regression described above (see "Experiment 1: Controls").

Let $Agree_5^J$ denote whether the respondent either strongly agrees or agrees with the $J^{th}$ of 6 statements (lettered a-f) presented in EXP2_005 (the six statements about whether each party goes too far, not far enough, and represents its voters). We will regress each of these six statements on an indicator for treatment, separately for each of the four samples described above (Democrats and Republicans who do and do not agree with their party). We denote the samples as *PG* where $P \in \{R: Repub, D: Dem\}$ and $G \in \{A: Agree, D: Disagree\}$. We will estimate:

$$Agree_5^J = \alpha_{PG}^J + \beta_{PG}^J T2 + \varepsilon$$

The second main (and admittedly very similar) regression is one based on the 8 statements (lettered a-g) presented in EXP2_007 (the eight statements about whether each party, Black Lives Matter, and police unions are dangerous and whether they are likely to improve policing). Let $Agree_7^J$ denote whether the respondent agrees with statement $J \in \{a, b, \dots, g\}$ listed in EXP2_007. Using the same definition of groups and the same notation as above, we will estimate:

$$Agree_7^J = \theta_{PG}^J + \gamma_{PG}^J T2 + \varepsilon$$

These two regressions (estimated for each of the four subsamples) will allow us to test many of the hypotheses described above:

***Hypothesis 2.1A***: Respondents become more frustrated by the Democratic Party: $\beta_{DA}^b > 0, \gamma_{DA}^c < 0$ (more likely to say Democratic Party doesn't go far enough, less likely to say Democratic Party likely to improve policing)

***Hypothesis 2.1B***: Respondents become more frustrated by the Republican Party: $\beta_{DA}^d > 0, \gamma_{DA}^d < 0$ (more likely to say Republican Party doesn't go far enough, less likely to say Republican Party likely to improve policing)

***Hypothesis 2.1C***: Respondents become more supportive of Black Lives Matter: $\gamma_{DA}^e < 0, \gamma_{DA}^G > 0$ (less likely to say Black Lives Matter agenda is dangerous, more likely to say Black Lives Matter likely to improve policing)

***Hypothesis 2.2A***: Respondents become more concerned about Democratic Party: $\beta_{DD}^a > 0, \gamma_{DD}^a > 0$ (more likely to say Democratic Party goes too far, more likely to say Democratic Party agenda is dangerous)

***Hypothesis 2.2B***: Respondents become more open to Republican Party: $\beta_{DD}^d < 0, \gamma_{DD}^b < 0$ (less likely to say Republican Party doesn't go far enough, less likely to say Republican Party agenda is dangerous)

***Hypothesis 2.2C***: Respondents become more concerned about Black Lives Matter: $\gamma_{DD}^e > 0$ (more likely to say Black Lives Matter agenda is dangerous)

***Hypothesis 2.3A***: Respondents become more concerned about Democratic Party: $\beta_{RA}^a > 0, \gamma_{RA}^a > 0$ (more likely to say Democratic Party goes too far, more likely to say Democratic Party agenda is dangerous)

***Hypothesis 2.4A***: Respondents become more open to Democratic Party: $\beta_{RD}^a < 0, \gamma_{RD}^a < 0$ (less likely to say Democratic Party goes too far, less likely to say Democratic Party agenda is dangerous)

***Hypothesis 2.4B***: Respondents become more concerned about Republican Party: $\beta_{RD}^{d} > 0$, $\gamma_{RD}^{b} > 0$ (more likely to say Republican Party doesn't go far enough, more likely to say Republican Party agenda is dangerous)

***Hypothesis 2.4C***: Respondents become more likely to say Republican Party doesn't represent its voters: $\beta_{RD}^{f} < 0$

***Hypothesis 2.4D***: Respondents become more concerned about police unions: $\gamma_{RD}^{f} > 0$ (more likely to say police unions' agenda is dangerous)

As above, we will also create non-binary dependent variables $OrdinalAgree_5^J$ and $OrdinalAgree_7^J$ which will be 5-point scales mirroring the 5-point scales we used when we asked the question. We will test the above hypotheses using these continuous outcomes as well.

Note that many hypotheses include two parameters. In these cases, we will report the significance of each individual coefficient, as well as the results of an F test for the joint significance of the coefficients.

To test the final two hypotheses about the general orientation towards reform, overhaul, or the status quo, we will define three dummy variables:

- $Reform = 1$ if the respondent chose "Although there are problems with policing, necessary changes can be made through reforms within the current system." in response to Question EXP2_006, and equals zero otherwise.
- $Overhaul = 1$ if the respondent chose "Because of fundamental problems, policing as an institution needs to be completely rebuilt." in response to Question EXP2_006, and equals zero otherwise.
- $StatusQuo = 1$ if the respondent chose "Little or nothing needs to be done to reform policing." in response to Question EXP2_006, and equals zero otherwise.

We will then estimate (again, separately for subgroups of respondents depending on party and agreement with party, denoted by $PG$):

$$\mathrm{Y} = \alpha_{PG}^{Y} + \beta_{PG}^{Y} T2 + \varepsilon \tag{3}$$

for $Y \in \{Reform, Overhaul, StatusQuo\}$.

***Hypothesis 2.1D***: Respondents will shift from incremental reform to complete overhaul: $\beta_{DA}^{Reform} < 0$ and $\beta_{RA}^{Overhaul} > 0$.

***Hypothesis 2.3B***: Respondents will shift from incremental reform to complete overhaul: $\beta_{RA}^{Reform} < 0$ and $\beta_{RA}^{StatusQuo} > 0$.

Finally, it is important to note that we expect only a small share of Democrats to disagree with their party (because there appears to be near unanimous agreement), and therefore we expect Hypotheses 2.2A-C (based on the $DD$ sample) to be under-powered.

## Experiment 3: Information about legislative behavior

### Motivation

Increases in mass party polarization are well-documented. This polarization obviously undermines legislative compromise. Above, we have already explored the effects of providing facts that undermine simplistic partisan narratives and the effects of efforts to shift respondents' perceptions of the opposite party. In this experiment, we test whether we can broker compromise using "interest group" endorsements of policy positions. Our particular focus is on the fact that actual legislation tends to be more balanced than the extreme positions one would typically infer from polarized rhetoric. Thus, we focus on actual legislation and experimentally vary information about true features of the person who proposed the legislation.

We focus on a specific piece of legislation which includes features that are objectionable to both Democrats and Republicans, and where the proposer represents "groups" which are held in high esteem by both Democrats and Republicans. Specifically, we focus on a bill by Representative Val Demings, who is a Democrat and a Black woman (both of which are signs of credibility for Democrats), as well as a former police officer and police chief (both of which are signs of credibility for Republicans). The bill focuses on low clearance rates for violent crimes (something the left frequently emphasizes as a criticism of the police and a reason to "defund") and its implications are fairly critical of the police (something we use in our framing and summary). However, the bill proposes to improve clearance rates with more funding and higher police officer salaries (a "pro-police" measure the right often calls for and strongly opposed to goals associated with the left).

Rather than asking about views on this specific legislation, we aim to separate respondents' from broader debates in national politics by asking about similar state-level legislation. We think this is policy-relevant, since the most important decisions governing policing are made by state and local elected officials, and because we suspect many respondents have few other grounds for choosing state legislative candidates.

### Design

All respondents will be shown the following prompt and associated questions:

We want to ask you about policies towards the police. Please read the following information carefully. We may ask you some questions about the proposed policy or the people who proposed it on the next page.

[NAME AND DESCRIPTION DEPENDS ON TREATMENT] has a bill to increase funding for police departments, while also increasing accountability. She worries that police departments are not doing a good job solving violent crime, and wants to increase hiring and pay of detectives to address that limitation.

**EXP1_001: Experiment 1 manipulation check**
Which of the following do you remember from the previous question about the person proposing the new policies towards the police?

1. Black
2. Democrat
3. Former police chief
4. None of the above

The treatment varies the name and description provided about who introduced the bill to increase police funding. This randomization depends on respondents' pre-experiment partisan loyalties (discussed above in Section "Making use of baseline information"). Respondents will be assigned to one of four conditions:

Control (both parties): Member of Congress Val Demings

Treatment 1 (Democrats only): Member of Congress Val Demings *who is Black*

Treatment 2 (Democrats only): Member of Congress Val Demings *a leading House Democrat*

Treatment 3 (Republicans only): Member of Congress Val Demings *a former police officer and police chief*

After the manipulation check (EXP1_001: "Which of the following do you remember…"), respondents will be shown a short vignettes that depends on their treatment assignment:

In many states, [DESCRIPTION DEPENDS ON TREATMENT] are introducing bills similar to the one described before.

Their treatment status determines the description as such:

Control (both parties): members of the state legislature

Treatment Democrats (whether treatment are 1 or 2): *Republican* members of the state legislature

Treatment Republicans: *Democrat* members of the state legislature

We then ask about support for the bill:

**EXP1_003: Support for bill**
Would you want your own State Representative to support these bills or oppose them?

1. Definitely support
2. Probably support
3. Neither support nor oppose
4. Probably oppose
5. Definitely oppose

Next, we ask whether this would affect vote choice. Our interest is in whether this issue can induce *party-switching*, as well as whether our experimentally manipulated endorsements drive party-switching. Thus, we will ask the vote choice question differently for Democrats and Republicans.

**EXP1_004D (wording for Democrats only): Change vote choice**
If the Democratic candidate in your State District opposed this bill, but the Republican candidate supported it, how might it affect your voting decision in that election?

1. Would definitely vote Republican
2. Would make me lean towards the Republican but might still vote Democrat
3. Would not affect my vote
4. Would make me lean towards the Democrat but might still vote Republican
5. Would definitely vote Democrat

**EXP1_004R (wording for Republicans only): Change vote choice**
If the Republican candidate in your State District opposed this bill, but the Democratic candidate supported it, how might it affect your voting decision in that election?

1. Would definitely vote Republican
2. Would make me lean towards the Republican but might still vote Democrat

3. Would not affect my vote
4. Would make me lean towards the Democrat but might still vote Republican
5. Would definitely vote Democrat

## Hypotheses

Our main interest is in whether treatment increases support for "compromise" reforms like the one proposed by Representative Demings, and whether any increase in support could plausibly translate into vote choice.

***Hypothesis 3.1***: Treatment increases respondents' support for bills like these.

As argued above, we see this bill as having elements that both Democrats and Republicans would object to. In light of this, we do not expect either group of voters to show strong support for the bill. However, given mixed overall impressions of the bill (given its "conflicting" content), we expect respondents to respond strongly to an endorsement from the interest group that they trust, respect, and see as credible on policing issues.

***Hypothesis 3.2***: Treatment will increase support for "opposite party" politicians.

Overall, we expect respondents to have a pretty limited understanding of the responsibilities of state government. Thus, we propose that information about policing-related issues (after priming them that state officials are particularly important for this) will be able to move respondents' vote choice.

***Hypothesis 3.3***: Treatment effects on vote choice will be smaller than effects on support for the bill.

At the same time, we recognize that other issues also play a role in vote choice (as well as partisan identity). Thus, we expect treatment effects to be smaller.

## Empirical implementation

All regressions presented below will control for variables selected by the LASSO regression described above (see "Experiment 1: Controls").

Additionally, however, it is important to note that we have a powerful control for regressions studying support for the legislation. The CES Common Content asks about support for increasing and decreasing police funding. Since increasing police funding is a component of the described legislation, we will control for whether the respondent supports increasing police funding by 10% in the CES Common Content (question **CC22_334c**).[6]

Let $T_3 = 1$ if the respondent was assigned to treatment in Experiment 3. To analyze the effects of treatment on support for the bill and vote choice, we will use the full continuous variation in our 5-point scales.

Specifically, we will define $Supp$ as an outcome variable based on the respondent's answer to "Would you want your own State Representative to support these bills or oppose them?" We will define $Supp = 2$ if the respondent answered "Definitely support", 1 for "Probably support", 0 for "Neither support nor oppose", -1 for "Probably oppose", and -2 for "Definitely oppose".

[6] Technically, CC22_334c asks about "Increas[ing] the number of *police on the street*" (emphasis added) which is subtly different than the Demings legislation, which we describe as "increas[ing] hiring and pay *of detectives*" (emphasis added).

We will define $Vote$ as an outcome variable based on the respondent's answer to "how might it affect your voting decision in that election?" We will define $RepVote = 2$ if the respondent answered "Would definitely vote Republican", 1 for "Would make me lean towards the Republican but might still vote Democrat", 0 for "Would not affect my vote", -1 for "Would make me lean towards the Democrat but might still vote Republican", and -2 for "Would definitely vote Democrat".

Since our core interest is in party-switching, we will construct a variable called $Switch$ is a variable to record the propensity for switching votes to the opposite party, which is obviously a combination of $RepVote$ and respondents' baseline party preferences:

$$Switch = \begin{cases} RepVote \text{ } if \text{ } respondent \text{ } is \text{ } Democrat \\ -RepVote \text{ } if \text{ } respondent \text{ } is \text{ } Republican \end{cases}$$

We will estimate the effects of treatment using basic OLS regressions (including the controls discussed above). We will estimate:

$$Supp = \alpha + \beta T_3 + \varepsilon$$

$$Switch = \theta + \gamma T_3 + \varepsilon$$

As noted above, our main substantive hypotheses focus on $\beta$ and $\gamma$. As we have done throughout the paper, we will estimate these regressions separately for Democrats and Republicans.

***Hypothesis 3.1***: Treatment increases respondents' support for bills like these: $\beta > 0$.

***Hypothesis 3.2***: Treatment will increase support for "opposite party" politicians: $\gamma > 0$.

Our final hypothesis (3.3) is that treatment effects will be larger for support $Supp$ than for vote choice $RepVote$. One challenge in testing this hypothesis is that the units are different: $Supp$ ranges from "Definitely support" to "Definitely oppose" while $RepVote$ ranges from "Would definitely vote Republican" to "Would definitely vote Democrat." These are fundamentally different scales.

Thus, we will normalize both coefficients by the partisan difference in responses among the control group. Specifically, using only the control group and only respondents that we classify as either Democrats or Republicans, we will estimate:

$$Supp = \eta_0 + \eta_1 Dem + \varepsilon$$

$$RepVote = \nu_0 + \nu_1 Dem + \varepsilon$$

where $Dem = 1$ if the respondent is a Democrat and zero otherwise. This normalization will allow us to test hypothesis 3.3 in meaningful terms. But first, we need to estimate a single regression for both hypotheses 3.1 and 3.2. To do so, we will estimate (re-using notation from above) the following two regressions using all respondents classified either as Democrats or Republicans:

$$Supp = \alpha + \beta T_3 + \varepsilon$$

$$Switch = \theta + \gamma T_3 + \varepsilon$$

***Hypothesis 3.3***: Treatment effects on vote choice will be smaller than effects on support for the bill: $0 < \frac{\beta}{\eta_1} < \frac{\gamma}{\nu_1}$

In testing this, we will simply use our estimates of $\widehat{\eta_1}$ and $\widehat{\nu_1}$ without accounting for the fact that they are estimated (i.e., we will not use the delta method). This is because our core interest is not inherently in the ratios of these parameters, but rather in the relative size of $\beta$ and $\gamma$, after normalizing them to have sensible and comparable units.

## Robustness: Excluding Florida

Our goal in this experiment is to experimentally vary the respondents' information about Val Demings, the Representative who proposed the bill. In 2022, Val Demings was running for Senate in Florida. Thus, it is likely that Florida respondents already had substantial information about Ms. Demings from her campaign, and our experiment probably does not control (and plausibly does not even influence) their perception of the proposer.

Thus, we will conduct a robustness check in which we drop Florida respondents from all analyses of Experiment 3. We expect this to matter little for the results, and in this case, we will include these results in the appendix only. If the results are substantially different without the Florida respondents, however, we will rely on the non-Florida results as the main results.

## Additional heterogeneity and tests for the mechanism

Here we describe several tests we will conduct which are built into the design of our experiment but where we do not have a specific hypothesis.

***Mechanism test 3.1***: We will test whether Democrats' response to a funding-increasing bill are larger when the endorser is a leading House Democrat or a Black woman.

As argued above, we believe that Democrats will shift their support for this "mixed" legislation when given additional details about the proposer. Our primary treatment effect will be based on the composite treatment (i.e., regardless of whether Democrats are assigned to treatment 1 or treatment 2). However, we intentionally designed the experiment to test whether Democrats are more responsive to partisan signals about the endorser (treatment 2: "a leading House Democrat") or identity-based signals about the endorser (treatment 1: "who is Black"). We will test which type of endorsement Democrats respond to more strongly by estimating:

$$Supp = \alpha + \beta_B T_3^B + \beta_D T_3^D + \varepsilon$$

$$RepVote = \theta + \gamma_B T_3^B + \gamma_D T_3^D + \varepsilon$$

where $T_3^B = 1$ if the respondent was assigned treatment 1 in Experiment 3 ("Val Demings *who is Black*") and zero otherwise, $T_3^D = 1$ if the respondent was assigned treatment 2 in Experiment 3 ("Val Demings *a leading House Democrat*") and zero otherwise, and both regressions are estimated only for Democrats.

We will test whether $\beta_B, \beta_D, \gamma_B, \gamma_D$ are individually significantly different from zero, but we will also test the hypotheses that $\beta_B = \beta_D$ and $\gamma_B = \gamma_D$. We will present p-values for both sets of hypotheses.

***Mechanism test 3.2***: We will test whether the effects of endorsements differ for those with higher or lower levels of pre-experiment support for police reform.

To implement this test, we will use our measure of "police reform liberalism", discussed above in the section "Making use of baseline information". Recall that this is a continuous measure based respondents pre-experiment answers to seven reform-related questions asked in the CES Common Content. Higher values imply that the respondents' stated views on the seven questions are more strongly associated with Democrats, while lower values imply that they are more strongly associated with Republicans. We denote respondents' police reform liberalism as $PRLib$.

To test for differential responses, we will estimate the following regression separately for Democrats and Republicans:

$$Supp = \alpha + \beta_1 T_3 + \beta_2 T_3 \times PRLib + \varepsilon$$

We will test whether $\beta_2$ is statistically significantly different from zero, separately for Democrats and Republicans.

In addition, we will test (again for each party) whether each of the following additional quantities are statistically significantly different from zero:

$$\beta_1 + \beta_2 \min(PRLib|Party = P)$$

$$\beta_1 + \beta_2 \mathrm{E}(PRLib|Party = P)$$

$$\beta_1 + \beta_2 \max(PRLib|Party = P)$$

The first quantity represents the effect of treatment on the respondent from Party P who is least supportive of reform, the second quantity represents the effect of treatment on the average respondent from Party P, and the third quantity represents the effect of treatment on the respondent from Party P who is most supportive of reform. (We list only the quantities based on the estimated $\beta$'s (i.e., only from the $Supp$ regression) but we will also test the same quantities based on the $\gamma$'s (i.e., from the $Switch$ regression).) Our test of whether $\beta_2 = 0$ already tests whether these quantities differ from one another, however it is of independent interest to know whether the treatment effects for these different types of respondents are statistically significant.

It is important to note that the results from Mechanism test 3.2 can be very different for both parties, but can also be differently signed for the $Supp$ regression and the $Switch$ regression. For example, it is possible that an endorsement from a former police officer and police chief will have the largest effects on support for the bill among Republicans with the most strong anti-reform views on policing, because these individuals may have more allegiance to former police officers. But that the same endorsement might have the largest effects on *vote switching* among those with the *least* strong anti-reform views on policing, because these individuals may be more open to voting for Democrats (i.e., maybe be more likely to be on the margin of switching their vote). In this case, $\beta_1 > 0$ and $\beta_2 < 0$ (i.e., they will be opposite signed), but $\gamma_1 > 0$ and $\gamma_2 > 0$ (same signed). In this case, support for the bill and support for party switching would show the same sort of average effects ($\beta$ and $\gamma$ from earlier would be of the same sign), but the heterogeneity would show the opposite set of patterns. Some reviewers may be inclined to see these heterogeneity results as unreliable noise, but we consider this a perfectly plausible result, and will offer a substantive interpretation of any statistically significant results from Mechanism test 3.2, rather than seeing them as a statistical noise or in contradiction with the results from Hypotheses 3.1 and 3.2.

## Adjusting for earlier experiments

It is important to note that this experiment will be run *after* either Experiment 1 or Experiment 2 (depending on the respondents' randomization). Our cross randomization ensures that we can still use treatment assignment in Experiment 3 to estimate the causal effects on our outcomes of interest. However, this causal effect must be interpreted as the effect amidst a weighted average of other types of experimentally manipulated information (i.e., amidst the information provided in

Experiment 1 or 2). As noted in the recent econometric literature (Muralidharan, Romero, and Wuthrich, 2022; henceforth MRW), this is a different estimand than the effects of Experiment 3's information being provided in a pure control "business as usual" sample. If our Experiment 3 treatment has no interaction effects with our other types of treatment, then the two estimands are the same. However, if there are interaction effects between the types of treatment (which we consider likely), then simply analyzing the effects of Experiment 3 without accounting for these interactions will fail to control the size of a test which tests the hypothesis that the effects amidst business as usual would be zero (see MRW).

MRW discuss several solutions to control size. One is to include all interaction terms between the treatments, and to present this in the appendix. We will do so. They note that in this regression, the main effect of being assigned treatment in Experiment 3 (without being interacted with treatment assignment in the other experiments) is a consistent estimate of the regular estimand of interest (the effects of treatment in Experiment 3 amidst business as usual), since this coefficient is only identified from the sample assigned to control in Experiments 1 and 2. MRW call this the "long regression." However, as MRW note, this test based on the long regression is very low-powered.

One solution proposed by MRW is discussed in relation to audit studies. They recommend that experimentally varying some characteristic (say race) is a consistent estimate of the average causal effect of race in the population, if all other experimentally-varied characteristics are set to match the population distribution. Another solution proposed by MRW is to make an ex ante assumption about the magnitude of the interaction coefficients which, if correct, can yield a much more powerful test of the hypothesis that the main effect is equal to zero.[7]

Our approach is to combine insights from these two different approaches. Like the second, we will invoke assumptions on the structure of interactions. Like the first, we will invoke assumptions that allow us to relate the "short model" (i.e., regressions using only the Experiment 3 treatment assignment without including any interactions) to a meaningful estimand. Specifically, we assume that any effects of Experiments 1 and 2 that affect the effects of Experiment 3 (i.e., any interactions between earlier treatment and later treatment) operate through respondents' overall orientation towards reform (questions **EXP4_007** in Experiment 1 and **EXP2_006** in Experiment 2, which are identical questions). Under this assumption, we can reweight the sample so that the group treated in Experiment 1 or 2 has the same distribution of reform orientation as the control group in those experiments, and our "short model" regressions will yield consistent estimates of the population average causal effect of our Experiment 3 treatment.

In short then, we will present three sets of results for all Experiment 3 analyses:

1. The "long model" regression results including interactions with earlier treatment assignments.[8] These test whether our Experiment 3 treatment has heterogeneous effects

[7] It is worth reiterating the central advice of MRW: "there is no free lunch." Making fewer assumptions (e.g., about the magnitude of interactions) reduces the power of size-controlling hypothesis tests. Having more flexible tests reduces power, while having less flexible tests fails to control size. And basing decisions on empirical estimates (rather than ex ante data-free assumptions) introduces severe data-dependent model selection biases.

[8] Specifically, let $T_1 = 1$ if the respondent was assigned to receive Experiment 1 before Experiment 3 and was assigned to treatment in Experiment 1, and zero otherwise. Let $T_2 = 1$ if the respondent was assigned to receive Experiment 2 before Experiment 3 and was assigned to treatment in Experiment 2. Let $T_3 = 1$ if the respondent was assigned treatment in Experiment 3. If we are interested in the effects of $T_3$ on some outcome $Y$, then the "long model" regression is $Y = \alpha_0 + \beta T_3 + \alpha_1 T_1 + \alpha_2 T_2 + \alpha_3 T_1 \times T_3 + \alpha_4 T_2 \times T_3 + \varepsilon$ , where we

depending on treatment in Experiments 1 and 2, but these are well-known to be under-powered tests.

2. The unadjusted "short model" that comes from regressing outcomes on Experiment 3 assignment (and the LASSO-selected controls) using CES-provided sample weights. This is the most straightforward way to analyze the effects of Experiment 3, but the estimated effects that it yields are not consistent estimates for the effects of Experiment 3 amidst a business as usual status quo, and such tests often over-reject the null that those effects are zero (see MRW). Note that these are the regressions described above in "Empirical implementation".
3. An adjusted "short model" that additionally adjusts the CES-provided weights of respondents assigned to treatment in either Experiment 1 or Experiment 2 so that the distribution of responses on reform orientation matches the distribution among respondents assigned to control in Experiments 1 and 2.[9]

All three sets of results will be included in the paper, but we reserve the right to decide which will be in the main body and which will be in the appendix, depending on the results (for example, if the second and third sets of analyses produce the same results, then we will avoid the details of this discussion in the main text, and only include the third set of results in the appendix).

## Experiment 4: Priming the salience of state and local politics

### Motivation

As mentioned above, we see policing as occupying a strange place in issue space: It is highly salient, but since politics has largely nationalized and most policing responsibilities fall to states and localities, political energy and attention is largely directed towards elected officials with little actual control over the issue. Our aim in this experiment is to highlight this wedge and remind respondents of the importance of state and local elections for policing issues, along with the fact that these elections are largely disregarded by voters.

### Design

Our aim is to emphasize low-salience state and local elections to respondents, by (1) telling them that state and local elections are often neglected but important for policing, and (2) by asking them whether their state is electing its Attorney General this year. We suspect that most voters will not be aware of whether their state is electing its Attorney General, and that pointing this out will highlight the neglected nature of these races.

To this end, the treatment group <u>only</u> will be given the following text and questions:

Many people worry that voters spend too much time following national politics while important issues like policing are mainly decided by state and local elected officials. For example, this November, most states will elect their Attorney General (the state's chief law enforcement officer). Cities may elect their mayor or city council members, who have a lot of influence over local policing.

---

interpret $\beta$ as the effects of $T_3$ among "business as usual" (i.e., among respondents assigned to control in Experiment 1 or 2).

[9] Let $w_j$ be the weight of some individual assigned to treatment in either Experiment 1 or 2 who chose $Y = j$ where $Y$ is reform orientation and $j \in \{1,2,3\}$ denotes the respondent's choice among the three options. Then the adjusted weight $\widetilde{w}_j$ is given by $\widetilde{w}_j = w_j \times \frac{\Pr(Y = j|T_1 + T_2 = 0)}{\Pr(Y = j|T_1 + T_2 = 1)}$.

**TURNOUT_002**
Is your state electing its Attorney General this year?

1. Yes
2. No
3. I don't know

**TURNOUT_003**
Is your state electing its Mayor this year?

1. Yes
2. No
3. I don't know

**TURNOUT_004**
**Descriptive Text/ Reading Section**
State and local elections typically don't receive as much attention as federal elections, even though they have more impact on policing. If your state is electing its Attorney General this year, you can learn about the candidates here:

https://www.stateside.com/election/2022-attorneys-general-elections

After this text, all respondents will be given the following question:

**TURNOUT_005: Multiple Choice**
In which types of elections do you take policing policies into account when choosing who to vote for? (Select all that apply)

1. Mayor
2. Governor
3. Attorney General
4. State legislative races
5. President
6. US Senate (federal)
7. US House of Representatives (federal)

## Hypotheses

It is, of course, relevant to know whether our treatment affects whether voters take policing into account when choosing candidates for mayoral, gubernatorial, state attorney general, or state legislative races. However, our design is not focused on these outcomes because we see no way to disentangle true effects from experimenter demand effects.[10] Thus, our core empirical approach and hypotheses focus on the CES-collected validated turnout outcome. This outcome is available well after the election, and is based on working with Catalist to validate (using administrative records) the actual turnout status of each CES respondent.

***Hypothesis 4.1***: Priming respondents to think about state and local elections will increase turnout.

[10] By this, we mean a situation in which respondents only *report* thinking about policing when choosing the Mayor, for instance. But they report this because we told them that it is important to do so, while in reality, they do not actually change their voting behavior.

Turnout in mid-term elections (like 2022) is low for a variety of reasons. We expect that under these conditions, reminding respondents that other non-Presidential elections are important for policing will increase their electoral participation.

***Hypothesis 4.2***: Turnout effects will be driven by marginal voters (i.e., those who are less likely to vote).

We expect our prime to be most effective among the voters who are less attached. We expect that many respondents are already committed to voting, and among these respondents, we do not expect our treatment to have any effects.

***Hypothesis 4.3***: Turnout effects will be driven by those in states with an Attorney General election.

Since our prime explicitly focuses on the Attorney General, we expect respondents to be particularly responsive in the states where there is an Attorney General elections.

***Hypothesis 4.4***: Priming respondents to think about state and local politics will be more effective when combined with information treatments that increase pro-reform sentiment.

We expect several of our treatments above to increase support for policing reforms. Here, we prime respondents to recognize that low-salience elections are important for policing. We expect these two treatments to interact. Specifically, we expect that when our information interventions increased respondents' support for police reform, they will be more responsive to reminders that state and local elections are important.

### Empirical implementation

All regressions presented below will control for variables selected by the LASSO regression described above (see "Experiment 1: Controls").

Additionally, however, it is important to note that we have a powerful control for our turnout regressions. The CES Common Content asks about turnout intent (prior to our survey module):

[CC22_363] {single} Do you intend to vote in the 2022 general election on November 8th?
(Allows one selection)

◯ **[1]** Yes, definitely
◯ **[2]** Probably
◯ **[3]** I already voted (early or absentee)
◯ **[4]** I plan to vote before November 3rd
◯ **[5]** No
◯ **[6]** Undecided

For all turnout regressions, we will exclude respondents who report having already voted as well as non-citizens, and we will control for fixed effects for the remaining five choices respondents choose from in CC22_363 (Yes, definitely; Probably; I plan to vote before November 3rd; No; Undecided). We suppress these fixed effects below for notational simplicity. Unlike earlier regressions, the bulk of these regressions will be run for the full sample, regardless of partisan identification (see Hypothesis 4.4 for an exception, and Mechanism test 4.1 for tests of heterogeneous effects by partisan identity).

Finally, for our turnout regressions, we will only use the sample that Catalist was able to successfully validate. In past CES surveys, this has been a large share of the sample, but we obviously do not currently have this variable and do not know exactly how Catalist and the CES team will implement validation this year. Thus, we cannot pre-specify exactly the coding that will be used or the share of the sample that will be successfully validated.

All turnout regressions will be estimated using a dummy variable $Turnout$ equaling one if the voter did cast a vote and zero otherwise and will be based on linear probability models. Let $T_4 = 1$ if the respondent was assigned to treatment in Experiment 4, and zero otherwise.

Our main test of whether our priming treatment increased respondents' turnout will be based on the regression:

$$Turnout = \alpha + \beta T_4 + \varepsilon$$

***Hypothesis 4.1***: Priming respondents to think about state and local elections will increase turnout: $\beta > 0$.

As noted above, we will always include for fixed effects for respondents' vote intention. To test Hypothesis 4.2, we will interact our treatment with some of these vote intentions. One challenge is defining this ex ante, without access to the data. For example, it is obvious that "undecided" voters are marginal voters to which our hypothesis will apply. But in the 2020 CES, this was only 4% of the sample. Choosing a category that is too small will undoubtedly lead to a low-powered test, but expanding the classification of marginal voters to include less obvious responses to the turnout intent question will yield a less precise test of our hypothesis. We aim to balance these concerns by testing Hypothesis 4.2 using three different definitions of "marginal." For all three definitions, we will only count voters as "marginal" if they report already being registered to vote.

Next, let $VI(j_1, \dots j_n)$ be a dummy variable such that $VI(j_1, \dots, j_n) = 1$ if the respondent reports being registered to vote <u>and</u> chose vote intent equal to $j_1$, …, or $j_n$. Let $J = \{1{:}\, Yes, definitely; 2{:}\, Probably; 4{:}\, I\ plan\ to\ vote\ before\ November\ 3rd; 5{:}\, No; 6{:}\, Undecided\}$ be the vector of possible responses.

We will first estimate whether treatment has larger effects on turnout among undecided voters only; then estimate whether there are also effects on voters choosing "probably" or "no"; and finally estimate whether it has larger effects on turnout among voters who chose any of the three (undecided, probably, or no). These regressions are given by:

$$Turnout = \alpha_0 + \alpha_1 T_4 + \alpha_2 T_4 \times VI(6) + \varepsilon$$

$$Turnout = \beta_0 + \beta_1 T_4 + \beta_2 T_4 \times VI(6) + \beta_3 T_4 \times \mathrm{VI}(2{,}5) + \varepsilon$$

$$Turnout = \gamma_0 + \gamma_1 T_4 + \gamma_2 T_4 \times VI(2{,}5{,}6) + \varepsilon$$

where (as above) we suppress the notation that we control for fixed effects for the different choices of voter intent, but we always include the main effect for respondents' choice in the vote intent question.

***Hypothesis 4.2***: Turnout effects will be driven by marginal voters: $\alpha_2, \beta_2, \beta_3, \gamma_2 > 0$ and $\beta_2 > \beta_3$.

We will use only t-tests for the significance of individual coefficients (not F-tests for joint significance). We acknowledge now that many of these tests will be under-powered due to few respondents selecting the given option, but it is impossible to tell how important this is until we have seen the data.

Define a dummy $AG_s$ that equals one if the state in which the respondent lives is electing its Attorney General is 2022 and zero otherwise. We will estimate:

$$\text{Turnout} = \beta_0 + \beta_1 T_4 + \beta_2 \text{AG}_s + \beta_3 \text{T}_4 \times \text{AG}_s + \varepsilon$$

***Hypothesis 4.3***: Turnout effects will be driven by those in states with an Attorney General election: $\beta_3 > 0$.

Finally, we will test whether exogenous increases in support for policing reform increase the effects of our turnout prime. To do so, we will use an instrumental variables (IV) strategy. We will estimate our IV regression using two-stage least squares.

The first stage will be based on our results from Experiments 1 and 2. Specifically, as before, we will define three dummy variables:

- $Reform = 1$ if the respondent chose "Although there are problems with policing, necessary changes can be made through reforms within the current system." in response to Question EXP2_006, and equals zero otherwise.
- $Overhaul = 1$ if the respondent chose "Because of fundamental problems, policing as an institution needs to be completely rebuilt." in response to Question EXP2_006, and equals zero otherwise.
- $StatusQuo = 1$ if the respondent chose "Little or nothing needs to be done to reform policing." in response to Question EXP2_006, and equals zero otherwise.

As before, let $T_1$ be a dummy indicating that the respondent was assigned Experiment 1 in the pre-election survey and was assigned to treatment in Experiment 1; let $T_2$ be a dummy indicating that the respondent was assigned Experiment 2 in the pre-election survey and was assigned to treatment in Experiment 2.

Hypotheses 1.2, 2.1, and 2.3 all imply that either $T_1$ or $T_2$ will affect $Reform$, $Overhaul$, or $StatusQuo$, at least for some subset of respondents. This implies that those (randomly assigned) variables can be used as instruments for respondents' reform-orientation. However, those hypotheses also lay out some hypotheses about heterogeneous treatment effects (depending on respondents' beliefs reported in Experiment 1 or views reported in Experiment 2, and depending on respondents' party affiliation). If those hypotheses are correct, then including those interaction terms can increase the strength of our instruments without invoking additional identification assumptions because the standard IV assumptions allow for interactions when treatment is randomly assigned. However, if those hypotheses are incorrect, then including the interactions will reduce the strength of our instruments, raising standard weak instrument concerns about over-rejection of the null and low coverage confidence intervals.

For this reason, we are choosing not to specify a first stage of our IV regressions in this pre-analysis plan. Instead, we commit that we will use any of the interactions discussed in Hypotheses 1.2, 2.1, and 2.3 that have statistically significant effects on any of the three outcomes ($Reform$, $Overhaul$, and $StatusQuo$) in a linear probability model (where we will use 10% as the cutoff for significance). We believe that this balances the importance of pre-registering our statistical approach with the reality that including non-meaningful instruments will exacerbate precision problems in our IV regression of interest.

Hypotheses 1.2, 2.1, and 2.3 all differ depending on respondent's partisan identity. We will estimate these IV regression separately by pre-experiment party identification – which allows for the first stage *and* the second stage to differ depending on party but will be less powerful – as well as pooled for both parties (excluding independents) where we will include partisanship as an interaction alongside the other interactions in the first stage (and second stage, without being interacted with treatment).

Letting $E_1$ and $E_2$ denote dummy variables indicating that the respondent was assigned Experiment 1 or Experiment 2 in the pre-election survey, respectively, our first stage regression will be a subset of:

$$Y = \alpha_0 + \alpha_1 T_1 + \alpha_2 T_2 + \alpha_3 MeanLow \times E_1 + \alpha_4 T_1 \times MeanLow + \alpha_5 Agree_2 \times E_2 + \alpha_6 Agree_2 \times T_2 + \varepsilon$$

where $Agree_2$ is a dummy variable denoting that the respondent agreed with their party in Experiment 2, all other variables are defined as above, the final first stage will exclude treatment interactions which are not individually significant, we interpret an interaction between zero and a missing value as being equal to zero (because we do not observe $MeanLow$ or $Agree_2$ for respondents not assigned to Experiment 1 or 2, respectively), and we will estimate these first stage regressions for $Y \in \{Reform, Overhaul, StatusQuo\}$.

The second stage, then, will be a subset of:

$$Turnout = \beta_0 + \beta_1 T_4 + \beta_2 \hat{Y} + \beta_3 T_4 \times \hat{Y} + \beta_4 MeanLow \times E_1 + \beta_5 Agree_2 \times E_2 + \beta_6 MeanLow \times E_1 \times T_4 + \beta_7 Agree_2 \times E_2 \times T_4 + \varepsilon$$

where we note that the terms in the first stage which are not randomly assigned instruments must be included in the second stage (including with an interaction with $T_4$, in our case) so as not to drive identification. Again, we will estimate this regression for $Y \in \{Reform, Overhaul, StatusQuo\}$.

***Hypothesis 4.4***: Priming respondents to think about state and local politics will be more effective when combined with information treatments that increase pro-reform sentiment: $\beta_3 > 0$ when $Y = Reform$.

Although our primary hypothesis is about $\beta_3$ when $Y = Reform$, we will also report and interpret the statistical significance of $\beta_1$ and $\beta_2$ because these terms are of inherent interest. For example, when $Y = Reform$, then $\beta_1$ corresponds to the effects of our Experiment 4 priming treatment on turnout when we *did not* exogenously shift respondents' support for reform, and $\beta_2$ corresponds to the effects on turnout of exogenously shifting respondents' support for reform when we *did not* also prime them to think about state and local politics. We will also report and interpret the statistical significance of results when $Y = Overhaul$ and $Y = StatusQuo$, which are also of inherent interest, although we are not pre-specifying specific hypotheses.

## Additional heterogeneity and tests of the mechanism

Here we describe several tests we will conduct which are built into the design of our experiment but where we do not have a specific hypothesis.

***Mechanism test 4.1***: We will test whether priming has differential effects on turnout depending on respondents' partisan affiliation.

Our main approach is to estimate

$$Turnout = \beta_0 + \beta_1 \mathrm{Dem} + \beta_2 \mathrm{Rep} + \beta_3 \mathrm{T}_4 + \beta_4 \mathrm{Dem} \times \mathrm{T}_4 + \beta_5 \mathrm{Rep} \times \mathrm{T}_4 + \varepsilon$$

where $Dem$ and $Rep$ are dummy variables for Democrats and Republicans (as defined above) and independents with no stated partisan lean are the omitted category. We will test not only whether Democratic and Republican partisans respond differently than non-partisan independents (the individual significance of $\beta_4$ and $\beta_5$) and whether there are effects for Democrats and Republicans (the null hypotheses of $\beta_3 + \beta_4 = 0$ and $\beta_3 + \beta_5 = 0$), but also whether Democrats and Republicans respond differently than one another (the null hypothesis of $\beta_3 + \beta_4 = \beta_3 + \beta_5$, which is equivalent to $\beta_4 - \beta_5 = 0$).

However, unlike other questions, we will also use a more flexible notion of partisan identity to study heterogeneous effects on turnout. This is because turning out to vote (1) is the most fundamental form of political participation and (2) is necessarily linked to a broad range of issues, whereas the other questions we've analyzed were focused on policing. Thus, we will classify respondents into five groups on the basis of the CES Common Content question pid7: Strong Democrats; Not very strong Democrats combined with independents who lean towards the Democrats; Independents who claim not to lean towards either party; Not very strong Republicans combined with independents who lean towards the Republicans; and Strong Republicans. In the 2020 CES, these groups are all roughly equal sized (15-25% of voters fall into each group).

We will estimate:

$$\begin{aligned} Turnout = \beta_0 + \beta_1 \mathrm{StrongDem} + \beta_2 \mathrm{WeakDem} + \beta_3 \mathrm{StrongRep} + \beta_4 \mathrm{WeakRep} + \beta_5 \mathrm{T}_4 + \beta_6 \mathrm{T}_4 \\ \times \mathrm{StrongDem} + \beta_7 \mathrm{T}_4 \times \mathrm{WeakDem} + \beta_8 \mathrm{T}_4 \times \mathrm{StrongRep} + \beta_9 \mathrm{T}_4 \times \mathrm{WeakRep} + \varepsilon \end{aligned}$$

In this specification, $\beta_5$ is the treatment effect for independents without any partisan lean, and $\beta_6$-$\beta_9$ are differential effects depending on partisan identification. We will present two sets of hypothesis tests. First, we will test whether each group shows a significant effect of treatment (e.g., $\beta_5 + \beta_6 = 0$). There are five such hypotheses. Second, we will test whether each group shows a significantly different treatment effect from another group (e.g., $\beta_5 + \beta_6 = \beta_5 + \beta_7$). There are 10 such hypotheses.

***Mechanism test 4.2***: We will test whether priming has differential effects on turnout depending on respondents' pre-prime stance on policing.

We will test for heterogeneous effects of our priming treatment depending on a variety of individual characteristics collected before Experiment 4. Each of these regressions will be of the form:

$$Turnout = \alpha + X'\beta_1 + \beta_2 T_4 + T_4 X'\beta_3 + \varepsilon$$

We will use many characteristics, and we are registering no hypotheses on any of these.

- Respondents' police reform liberalism
- Respondents' affect towards the police
- Dummy variables for all possible values of respondents choice in response to the Common Content question about whether the police make respondents feel safe (CC22_307)
- Dummy variables for whether the respondent agrees or strongly agrees that each of the four groups we ask about (the Democratic Party, the Republican Party, Black Lives Matter, and police unions) are likely to improve policing
- Dummy variables for whether the respondent is Black or Hispanic (two dummy variables)
- Respondents who could correctly identify the party of their governor

***Mechanism test 4.3***: We will test whether priming has differential effects on turnout depending on respondents' local area increase in crime rates.

Our prime increases the salience of local elections for policing, but it also plausibly provides respondents with *information* about the importance of state and local elections for policing. That is, it is possible that respondents simply did not know how much influence state and local elections have. An obvious question, then, is whether any treatment effects on turnout are more consistent with the salience or the information explanation.

To answer this question, we propose to use independent variation in the salience of policing, unrelated to our experiment. Our main insight is that the recent increase in violent crime has (1) affected different communities differently, and (2) been very salient in political discussions and debates. As a result, for some respondents (those who live in MSA's with large increases in violent crime), policing is already a highly salient issue. In this case, our priming experiment should have relative modest effects on the salience of policing issues. However, our priming effect could still have large effects on the information respondents have about policing, particularly if they do not understand the importance of state and local politics for policing.

Put differently, if the main effects of our prime come from increasing the *salience* of policing, then our treatment effects should be smaller in the presence of a large crime spike, since salience will already be high. If the main effects of our prime result come from increasing *information* about electoral politics and policing, then our treatment effects should be larger in the presence of a large crime spike, since these are the environments where respondents will be most interested in exercising democratic accountability. Thus, we will separate between these two mechanisms by interacting our treatment with local increases in crime rates.

One challenging in pre-specifying this hypothesis is that crime data is notoriously problematic. It is released with a very long lag, and the FBI's switch from the UCR to the NIBRS system in 2021 has lead to an even higher non-response rate than in the past. In response to these issues, many teams of journalists and researchers are aiming to collect systematic crime data in "real time" (or with a less extreme delay than the FBI's data). At the time of writing this pre-analysis plan, it is impossible to know what data sources will be available and which will be most reliable or up-to-date (in part because we cannot know when Catalist will validate voter turnout with the administrative records). Thus, we will reach out to experts and collect recommendations. Nonetheless, it is important to note that this is an important "degree of freedom" in our statistical analysis that we cannot fully avoid with a pre-analysis plan.

With some recommended measure of the change in crime during the time leading up to the 2022 election, we will estimate:

$$Turnout = \beta_0 + \beta_1 T_4 + \beta_2 \text{CrimeSpike} + \beta_3 \text{T}_4 \times \text{CrimeSpike} + \varepsilon$$

for all respondents who have not yet voted (regardless of partisanship). If $\beta_3 > 0$ then we will conclude that our priming treatment is primarily effective because of information, if $\beta_3 < 0$ then we will conclude that our priming treatment is primarily effective because of salience.

***Mechanism test 4.4***: We will test whether priming has differential effects on turnout depending on election closeness.

Our priming treatment reminds respondents of the importance of state and local elections, and in particular emphasizes the Attorney General. As is common in studies of political participation, any effects on voter turnout are consistent with both expressive motives (in which respondents derive utility simply from expressing their political preferences) and instrumental motives (in which respondents are motivated by election outcomes and the possibility that their vote will affect those outcomes). Even though there is a very small probability that the respondent's own vote will affect the election outcome, we expect that by reminding respondents that relative few voters engage with these elections, it is plausible that respondents will feel instrumental incentives to vote.

To separate between these two incentives, we will focus on the respondents in the 30 states which *are* electing their Attorney General this year. We will collect data on the winner's share of the top-2 candidate vote. While this information is obviously not available when respondents are deciding whether to vote, we expect respondents to be able to form rough expectations of whether or not the election will be close. Let $WinVoteShare$ be the share of the top-2 candidate vote received by the election winner, minus 0.5 (so that $WinVoteShare = 0$ corresponds to an asymptotically close election).

The key prediction of instrumental voting that we aim to exploit is that the incentives to vote are stronger when the election is closer. If our priming treatment increases instrumental motives to vote, then these effects should be stronger in closer elections. To test this we will estimate:

$$Turnout = \beta_0 + \beta_1 T_4 + \beta_2 \mathrm{WinVoteShare} + \beta_3 \mathrm{T}_4 \times \mathrm{WinVoteShare} + \varepsilon$$

If $\beta_1 > 0, \beta_3 < 0$ then our treatment increases turnout when the election is very close ($\beta_1$ is the treatment effect when $WinVoteShare = 0$ at a 50% win) and these effects become smaller as the winner's vote share becomes larger. In this case, we will interpret it as evidence in favor of instrumental motives. Otherwise, we will interpret it as evidence consistent with expressive motives. If the evidence is consistent with expressive motives, we will estimate the regression above separately for Republicans and Democrats, using own-party vote share in place of winner's vote share, and present estimates non-parametrically. This will tell us whether people prefer expressing themselves more when their candidate will win safely or when their candidate will lose by a large margin.

### Adjusting for inattention

In many experiments, respondents do not pay attention to all questions. For example, Fowler, et al. (2022) show that a small but meaningful share of apparent moderates in past CES waves have actually been respondents randomizing their answers to questions. Since inattention likely changes in response to survey length and content, it is difficult to know inattention rates for any particular survey module.

For this reason, our survey is designed such that in the middle of our survey module we intentionally ask respondents about increasing police funding and decreasing police funding. Since it is impossible for someone to simultaneously support increasing *and* decreasing police funding, respondents who choose both options must be randomizing answers (evidence of inattention). If some fraction *q* of respondents are randomizing, then *q/4* of them should choose the contradictory answers of increasing and decreasing police funding. Moreover, our treatment effects should be attenuated by *q*, since our treatments cannot have effects on respondents whose answers are randomized.[11]

Thus, in the paper we will report inattention-adjusted estimates of our effects in experiments 1 and 2. Namely, if the regressions above estimate that the treatment effects are $\hat{\beta}$ then the inattention-adjusted estimate is $\hat{\beta}/(1-4z)$, where $z$ is the fraction of respondents who report wishing to increase and decrease police funding. Note that adjusting the estimate by a fixed scalar does not affect the p-value or statistical significance. We will, of course, also include unadjusted estimate in the paper, as these are the most straightforward way to analyze our data.

[11] If the true treatment effect for the attentive sample is $\beta$ and the inattentive sample is 0 (by construction, since one cannot affect randomized answers), then the sample treatment effect is $(1-q)(\beta) + q(0)$, where $q$ is the fraction which randomizes. This is a standard weighted average of coefficients, weighted by the relative sample size.

## Prioritizing hypotheses

Our main and most central hypotheses are 1.2A, 1.3A, 2.1A, 2.1B, 2.4A, 2.4B, 3.1, and 4.1.

## References

Ahler, Douglas J., and Gaurav Sood. "The parties in our heads: Misperceptions about party composition and their consequences." *The Journal of Politics* 80.3 (2018): 964-981.

Fowler, Anthony, Seth J. Hill, Jeffrey B. Lewis, Chris Tausanovitch, Lynn Vavreck, and Christopher Warshaw. 2022. "Moderates." *American Political Science Review* Forthcoming.

Muralidharan, Karthik, Mauricio Romero, and Kaspar Wuthrich. 2022. “Factorial designs, model selection, and (incorrect) inference in randomized experiments.” *Working Paper*.