



Current research overstates American support for political violence

Sean J. Westwood^{a,1}, Justin Grimmer^{b,c,d}, Matthew Tyler^{b,e}, and Clayton Nall^e

Edited by Anthony Fowler, University of Chicago, Chicago, IL; received September 15, 2021; accepted February 7, 2022 by Editorial Board Member Margaret Levi

Political scientists, pundits, and citizens worry that America is entering a new period of violent partisan conflict. Provocative survey data show that a large share of Americans (between 8% and 40%) support politically motivated violence. Yet, despite media attention, political violence is rare, amounting to a little more than 1% of violent hate crimes in the United States. We reconcile these seemingly conflicting facts with four large survey experiments ($n = 4,904$), demonstrating that self-reported attitudes on political violence are biased upward because of respondent disengagement and survey questions that allow multiple interpretations of political violence. Addressing question wording and respondent disengagement, we find that the median of existing estimates of support for partisan violence is nearly 6 times larger than the median of our estimates (18.5% versus 2.9%). Critically, we show the prior estimates overstate support for political violence because of random responding by disengaged respondents. Respondent disengagement also inflates the relationship between support for violence and previously identified correlates by a factor of 4. Partial identification bounds imply that, under generous assumptions, support for violence among engaged and disengaged respondents is, at most, 6.86%. Finally, nearly all respondents support criminally charging suspects who commit acts of political violence. These findings suggest that, although recent acts of political violence dominate the news, they do not portend a new era of violent conflict.

political violence | affective polarization | democratic norms

Provocative recent work (1–3)—cited in PNAS (4, 5), *The American Journal of Political Science* (6), 60 other articles and books, and 40 news articles that, together, have garnered over 2,281,133 Twitter engagements—asserts that large segments of the American population now support politically motivated violence. These studies report that up to 44% of Americans would endorse hypothetical violence in some undetermined future event (1–3, 7). This survey work fits within a media landscape that regularly raises the specter of political violence. Since 2016, we counted 2,863 mentions of political violence on news television, more than 630 news stories about political violence, and over 10 million tweets on the topic of the January 6 riot alone (see *SI Appendix, section 1* for details for all counts in this paragraph). Political violence, however, remains exceedingly rare in the United States, amounting to 48 incidents (8) in 2019 (the most recent year for which data are available) compared to 4,526 incidents of nonpolitical violent hate crimes (9) and 1,203,808 total violent crimes (10) documented by the Department of Justice.

In this paper, we reconcile supposedly significant public support for political violence with minimal actual instances of violent political action. To do this, we use four survey experiments that assess respondents' reactions to specific acts of violence, where we experimentally manipulate whether partisanship motivated the activity and the severity of the violence. Using these studies, we identify two reasons why current survey data overestimate support for political violence in the United States.

First, ambiguous survey questions cause overestimates of support for violence. Prior studies ask about general support for violence without offering context, leaving the respondent to infer what “violence” means. Using detailed treatments and precisely worded survey questions, we resolve this ambiguity and reveal that support for violence varies substantially depending on the severity of the specific violent act. With our measures, assault and murder attract minimal support, while low-level property crimes gain higher (although still low) support. Moreover, even though segments of the public may support violence or report that it is justified in the abstract, nearly all respondents still believe that perpetrators of well-defined instances of severe political violence should be criminally charged.

Second, disengaged survey respondents cause an upward bias in reported support for violence. Prior survey questions force respondents to select a response without providing

Significance

Recent political events show that members of extreme political groups support partisan violence, and survey evidence supposedly shows widespread public support. We show, however, that, after accounting for survey-based measurement error, support for partisan violence is far more limited. Prior estimates overstate support for political violence because of random responding by disengaged respondents and because of a reliance on hypothetical questions about violence in general instead of questions on specific acts of political violence. These same issues also cause the magnitude of the relationship between previously identified correlates and partisan violence to be overstated. As policy makers consider interventions designed to dampen support for violence, our results provide critical information about the magnitude of the problem.

Author contributions: S.J.W. designed research; S.J.W., J.G., and M.T. performed research; M.T. and C.N. contributed new reagents/analytic tools; S.J.W. analyzed data; and S.J.W. and J.G. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. A.F. is a guest editor invited by the Editorial Board.

Copyright © 2022 the Author(s). Published by PNAS. This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: sean.j.westwood@dartmouth.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2116870119/-DCSupplemental>.

Published March 18, 2022.

a neutral midpoint or a “don’t know” option. This causes disengaged respondents—satisficers (11)—to select an arbitrary or random response (12). Current violence support scales are coded such that four of five choices indicate acceptance of violence. In the presence of arbitrary responding, such a scale will overstate support for violence. Across all four studies, we show that disengaged respondents report higher support for violence.

Accounting for these sources of error, our four studies show that American support for political violence is less intense than prior work asserts (Fig. 1) and is contingent on the severity of the violent act. Depending on how the question is asked, we show that the median of existing estimates of the public’s support for partisan violence is nearly 6 times larger than the median of our estimates (18.5% versus 2.9%). While recent political events show that extreme political groups are willing to engage in violence, these groups are likely to overlap with the narrow segment of the population who already support political violence. As policy makers consider interventions designed to dampen support for violence, our results demonstrate that support for violence is not a mass phenomenon, indicating that antiviolence measures should be appropriately tailored to match the scale of the problem.

Support for Partisan Violence Is Lower than Previously Reported

Partisan animosity, often referred to as affective polarization (13), has increased significantly over the last 30 y. While Americans are arguably no more ideologically polarized than in the recent past, they hold more-negative views toward the political opposition and more-positive views toward members of their own party. This pattern has been documented across several measures of animosity and has raised alarm among scholars across disciplines about the potential consequences of growing partisan discord (e.g., ref. 14).

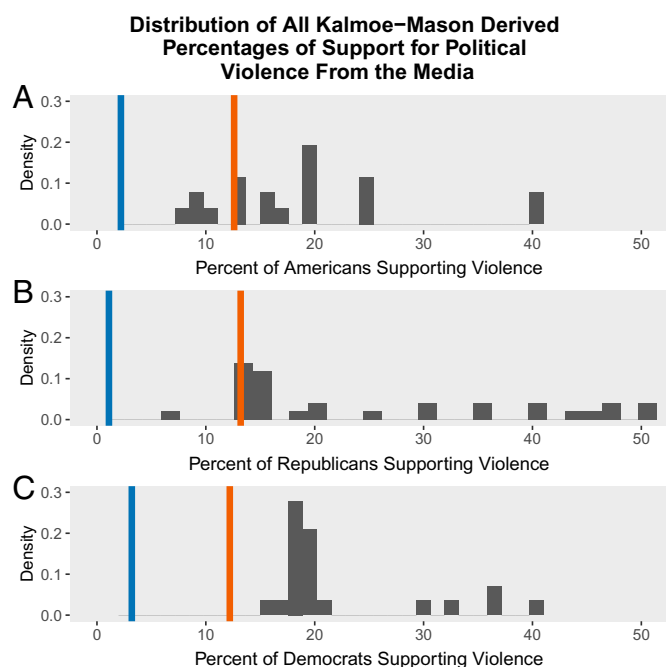


Fig. 1. We created a census of all reported estimates of support for violence using the Kalmoe–Mason measure in the media (this includes work done by authors other than Kalmoe and Mason). This figure shows their distribution. We report this in the full sample (A), for Republicans (B), and for Democrats (C). To contextualize the problems, in these estimates, we overlay the largest estimates (orange line) and smallest estimates (blue line) from the studies that follow. There is large variation in the reported values, but all are significantly larger than our preferred estimate. See *SI Appendix* for additional details.

Numerous studies have documented the negative interpersonal, “apolitical” (15) consequences of affective polarization, including politically based discrimination against job applicants (16), prospective romantic partners (17), workers (18), and even scholarship recipients (for a review, see ref. 13). These findings have created substantial concerns over partisan animosity’s pervasive effects on American social life (19).

Yet, evidence suggests that affective polarization is not related to and does not cause increases in support for political violence (20, 21) and is generally unrelated to political outcomes (21, 22). Moreover, partisan violence appears to be unrelated to many other political variables (2). We are therefore left with a phenomenon that is not explained by the current literature on partisan animosity, that is rarely observed in the world, but that is based on prior work supported by a near majority of the American population (1–3).

We show that documented support for political violence is illusory, a product of ambiguous questions and disengaged respondents. We now explain how each causes political violence to appear more popular than it is in the public.

Ambiguous Questions Create Upward Bias in Estimates of Support for Violence.

Even if respondents truthfully report their views on political violence, vague questions make it impossible to compare responses across individuals, and render sample averages uninterpretable. For example, a measure from Kalmoe and Mason (hereafter, Kalmoe–Mason) (2, 3) asks about perceived justification for partisan violence generally: “How much do you feel it is justified for [respondent’s own party] to use violence in advancing their political goals these days?” But the estimand measured by this survey item is unclear because it leaves ambiguous what “violence” refers to. Another question from Robert Pape (23), “The use of force is justified to restore Donald Trump to the presidency,” offers a specific motivation, but, like the Kalmoe–Mason measures, leaves definition of “violence” to the respondent to fill in. As a simplistic example, suppose that respondents interpret the question as asking about either partisan-motivated assault or partisan-motivated murder (both acts of violence). If one individual interprets violence as “assault” while another interprets violence as “murder,” then these responses are not comparable, and therefore we cannot make an inference about which respondent expresses more support for political violence (24). This also affects mean expressed support for violence. The quantity $P(\text{support partisan violence})$ is an average of respondents who interpret the question as asking about assault and others interpreting the question as asking about murder. The conditional average support for partisan violence and the relative prevalence of the components of the mixture are unknown, $P(\text{support partisan violence}) = P(\text{support partisan violence}|\text{assault})P(\text{assault}) + P(\text{support partisan violence}|\text{murder})P(\text{murder})$.

It is impossible to know, from existing responses to vague questions, whether respondents support severe, moderate, or minor forms of violence, which could range from support for violent overthrow of the government to minor supporting assault at a local protest. We address this concern in two ways across our four survey experiments. First, we use two different levels of violence for study 1, study 2, and study 3: assault and murder. Second, in study 4, we vary the underlying violent act along a taxonomy of severity.

Disengaged Respondents Upwardly Bias Measures of Support for Political Violence.

The goal of all surveys is to capture genuine opinions from a sample. However, it is well known that not all respondents engage in the thought, consideration, and reflection

necessary to provide reasoned responses to all questions (25), and some may even overreport rare and negative traits/opinions to troll researchers (26). As the complexity of the work needed to answer a question increases (i.e., thinking about meaning, filling in details in ambiguous questions, forming opinions on a question a respondent has never previously considered, etc.) and motivation to deeply engage decreases, respondents are more likely to satisfice (12). When satisficing, respondents may simply select a neutral midpoint (11), randomly select a response (27), or even leave a survey (25). We suspect that the vague and ambiguous nature of current survey measures of political violence are especially likely to cause respondents to satisfice.

Two features of the current survey designs cause the problem. First, existing measures of support for partisan violence collapse response categories to indicate support (1, 2). For example, one survey question asks respondents, “How much do you feel it is justified for Democrats to use violence in advancing their political goals these days?” and uses a five-point Likert-like scale with options “Not at all,” “A little,” “A moderate amount,” “A lot,” and “A great deal.” Ref. 2 then recodes the responses “A little” to “A great deal” as indicating support for partisan violence and “Not at all” as opposing partisan violence. Second, such survey questions fail to offer a neutral midpoint or a “don’t know” option. If these imperfect options or frustration from the ambiguous nature of the actual question cause a respondent to disengage from the survey task and satisfice (11), they are likely to arbitrarily pick from the set of imperfect options. But, in this example, satisficers picking a random response would end up indicating support for violence four times out of five.

To formalize this example, the goal is to measure the true preferences for partisan violence in the population, which we will call $\mathbb{E}[Y]$. This quantity is estimated from a representative survey of the population by taking a mean of a survey question, $\mathbb{E}[Y^{\text{survey}}]$. If some disengaged respondents satisfice, then the estimated support for partisan violence will be

$$\mathbb{E}[Y^{\text{survey}}] = \mathbb{E}[Y \mid \text{Engaged}]P(\text{Engaged}) + \mathbb{E}[Y^{\text{satisfice}} \mid \text{Disengaged}]P(\text{Disengaged}),$$

where reported support when satisficing, $Y^{\text{satisfice}}$, might be different from the true support Y depending on the survey respondent’s behavior when satisficing. If $\mathbb{E}[Y^{\text{satisfice}} \mid \text{Disengaged}] > \mathbb{E}[Y \mid \text{Disengaged}]$, then the survey-based estimate will be larger than the true level of support for violence. This condition is likely to hold under current survey-based approaches to measuring preferences for partisan violence where four of five response options indicate support for violence (80% of possible responses). If respondents choose their response at random with a uniform probability, then the chance that they would appear to support partisan violence is 0.8. If true $\mathbb{E}[Y \mid \text{Disengaged}] < 0.8$, then the presence of disengaged respondents will cause an upward bias. In an extreme example, if no one actually supports partisan violence, but 31% of respondents—the proportion who fail our engagement test in study 1—in a survey answer at random, a survey would find that $0.31 \times 0.8 = 24.8\%$ of respondents support partisan violence. This is very close to the amount of inflation we see in partisan violence in our following studies.*

We take explicit steps to address disengaged respondents who satisfice. We offer satisficers response options that are less likely to upwardly bias estimates: a balanced five-point scale with a neutral

midpoint. This brings the measure in line with standard and methodologically robust approaches to measurement, and reduces the chances that a satisficer will randomly select a response indicating support for violence. We also report our estimates based on individuals who are engaged—passing a comprehension check—and individuals who are disengaged, or fail a comprehension check.

Assessing Partisan Differences in Who Commits Political Violence. Concern about political violence in the United States is often associated with increasing levels of affective polarization between Democrats and Republicans (14). But existing measures of support for partisan violence tend to not assess whether providing information about the partisanship of who committed the act of violence affects support or opposition for the act of violence. Providing this information is important, because there are two potential interpretations of a positive effect. If the response is sincere, it could be that copartisans give additional leeway for acts committed by copartisans. But, if the response is insincere, it could be that partisans, in general, are merely offering support for their party—a version of partisan cheerleading. While randomizing information about partisanship alone is insufficient to distinguish between these two possibilities, failure to find a difference in a well-powered study provides strong evidence that neither leeway nor cheerleading occur.

To assess how partisanship affects support for violence, in our study 1 and study 2, we explicitly vary information about the partisanship of who committed the acts of violence. As we show below, we fail to find a consistent partisan difference—implying that there is little evidence for a general leeway or cheerleading effect.

While we find little evidence of partisan cheerleading among all partisans, we might worry that a specific subset of partisans engage in explicit partisan cheerleading. To make this assessment, in study 3, we use existing survey questions to identify partisan cheerleading (28) and find that partisan cheerleaders inflate support for violence, but those cheerleaders comprise only a small share of respondents and therefore do not appear to meaningfully affect results.

Methods

To uncover how these sources of error affect perceptions of partisan violence, we conducted four survey experiments. We fielded our first survey (which contained study 1 and study 4) via Qualtrics panels in January 2021—starting 2 d after the violence of January 6. This allows us to test our predictions during a period when partisan discord and violence dominated news coverage. Our second survey (study 2) was fielded in April 2021, also on Qualtrics panels. Our final survey (study 3) was fielded in November of 2021 on the YouGov panel. This allows us to verify that our results are not dependent on proximity to the Capitol riots or on a specific survey panel.

The Qualtrics data were collected from Qualtrics panels and utilized quota sampling. Respondents were recruited from panel members by email. All surveys were restricted to Democrats and Republicans. Leaners were coded as partisans. For Qualtrics data, we quota sampled on age, sex, and race/ethnicity to match Census targets. The sample is generally very representative of the population (*SI Appendix, Tables S1, S19, S30, and S39*). These data were analyzed without survey weights per our preanalysis plan. The YouGov data were sampled with the standard YouGov matching procedure. YouGov data were analyzed with provided weights.†

The survey flow was as follows: consent, attention check, demographics, covariates (including the measure from refs. 1, 2), randomized treatment, engagement test, and then outcome questions. Our experiments were approved by

*We note that, while not observed here, if true support for violence were above 0.8, the bias would be negative. Also, if the true prevalence rate among the disengaged were 0.8, then the bias for the population parameter would be zero.

†By necessity, weights were not used when estimating partial bounds.

institutional review boards at Stanford, Dartmouth, and University of California, Santa Barbara. Participants were asked to give consent after reading an information sheet.

All four experiments were preregistered. For our Qualtrics data, following our preanalysis plan, we excluded participants who failed a preregistration attention check (a question asking respondents to make two specific response choices) and those who completed the survey in less than one-third of the median complete time. Neither of these choices altered the demographic composition of the sample, as purged respondents were not counted toward quotas and were replaced. Both of these choices work against us by removing disengaged subjects, which means that our estimates are conservative, as these design choices remove respondents who are most likely to respond to all survey questions at random. Those who remain and satisfice are likely doing so because of flawed, ambiguous, or insufficiently contextualized questions and not because of general inattention.

Measuring Engagement. To assess how satisficing affects responses, we group participants based on their cognitive engagement with our survey, measured as the retention of information from vignettes. Respondents who can identify the state where the event in the vignette they just read occurred (information that was repeated in the headline and up to two times in the text) are coded as engaged, and those who cannot are coded as disengaged. Reading short passages is not a cognitively intense task, but committing content to memory is (29), and reveals engagement with cognitive work. Across our three studies, we group our sample into “engaged” respondents—those who are sufficiently motivated to carefully read and retain information—and “disengaged” respondents—those who can follow instructions but are less likely to retain facts or carefully evaluate questions.[‡]

Outcome Questions. We use three core questions for study 1, study 2, and study 3 and deploy two strategies to identify satisficing’s upward bias on estimates of support for violence. First, in our justification outcome question, we use a dichotomous scale which should prompt disengaged respondents to randomly respond, which we can detect by observing the response distribution. Second, we offer a balanced scale with a neutral midpoint when measuring support, building on prior work that suggests disengaged respondents will respond with a “don’t know” or neutral midpoint response (11). Our three questions are as follows:

- 1) “Was the suspect justified or unjustified?” using a dichotomous outcome scale. This differs from refs. 1–3, where the authors ask a hypothetical question with an unbalanced five-point Likert scale that is then recoded in their analysis to a dichotomous measure.
- 2) “Do you support or oppose the actions of [suspect]?” This is measured with a balanced Likert with a neutral midpoint and is intended to separately capture support for a suspect when a violent act is committed.
- 3) “Should the suspect face criminal charges?” We capture responses with a dichotomous yes/no scale, and the question is intended to determine whether the respondent thinks that someone who commits an act of violence should or should not face charges.

Respondents Reject Extreme Violence, Whether It Is Political or Not

Study 1, study 2, and study 3 show that, as preregistered, respondents overwhelmingly reject both political and nonpolitical violence, and disengaged survey respondents show higher measured support for political violence. We find no evidence of partisan effects, as partisans from both sides express similar tolerance for political violence. We also find higher (although still low) levels of support for the less violent act in study 1 relative to the more violent act in study 2 and study 3.

[‡]SI Appendix, Table S64 shows that removing disengaged respondents does not meaningfully change the demographics of our sample (age, gender, race, partisanship, income, and education). Another concern is that we are conditioning on a posttreatment outcome. However, our goal is not to measure the causal effect of engagement (30), but to merely show that responses differ based on engagement.

To avoid the problem of ambiguous question wording, our design presents a detailed act of violence, which prevents respondents from substituting their own definition of “violence” when answering our outcome questions.

In study 1 ($n = 1,002$), we randomly assigned participants to read one of two stories based on real acts of political violence. In the first story, a Democratic driver was charged with hitting a group of Republicans in Florida who were registering citizens to vote. In the second story, a Republican driver was charged with assault for driving his car through Democratic protesters in Oregon. Respondents were also randomized to see the original version of the story that included partisan details or a version of the story that was altered to remove any reference to partisan motivation.

In this study, we focused on reporting details from real events. This means that, while comparable, the Democratic and Republican stories varied in several ways. To ensure that any effects we identify are not the result of those differences, we conducted a second version of this experiment. Study 2 ($n = 1,023$) used a single contrived story of violence in Iowa. To test the bounds of support for political violence, this story reported an extreme form of violence: murder. Similar to study 1, participants were randomly assigned to see a story with a Republican or Democratic shooter engaging in politically motivated violence or an apolitical act of murder. This story was necessarily fabricated to limit the differences across treatment conditions.

Study 3 ($n = 1,863$) is a replication of study 2 using the YouGov panel with the following alterations: 1) We removed the apolitical condition to focus on attitudes toward partisan violence, and 2) we removed covariate questions to reduce survey time.

Disengaged Responses Lead to Higher Estimates of Support for Political Violence.

At first glance, the results of this experiment appear to align with prior surveys. Across conditions where the driver’s actions are presented as political violence, we find that 21.07% of respondents in study 1 say the attack was justified. We find a similarly high level of support for the apolitical versions, where 20.56% of respondents in study 1 say the driver’s action is justified. The overall support for violence is lower in study 2 and study 3, reflecting the greater severity of the violence, with 10.02% of respondents in study 2 describing the political homicide as justified and 6.70% of the respondents describing the apolitical homicide as justified. In study 3, 10.11% described the homicide as justified.

But this is biased upward by respondents who fail the engagement test (~31% of respondents in study 1, 19% of respondents in study 2, and 19% of weighted respondents in study 3). For the political treatments, 37.87% of respondents who fail the engagement test say the driver’s actions were justified, while only 12.06% of respondents who passed the engagement test agree that the driver’s actions are justified. For the nonpolitical treatment, we find that 44.06% of respondents who failed the engagement test say the driver’s actions were justified, but only 10.93% of respondents who passed the engagement test say the driver’s actions are justified. Similarly, for study 2, in the political treatments, we find that 33.82% of the respondents who fail the engagement test say the shooter’s actions were justified, but only 4.37% of individuals who passed the engagement test say the action was justified. In the nonpolitical treatments, we find a similar large gap: 25.93% of respondents who fail the engagement test say the action was justified, but 2.68% of those who passed say the action was justified. The same pattern is found in study 3 (YouGov data), with 28.25% of disengaged respondents saying the shooting

Support for Violence Among Engaged and Disengaged Respondents

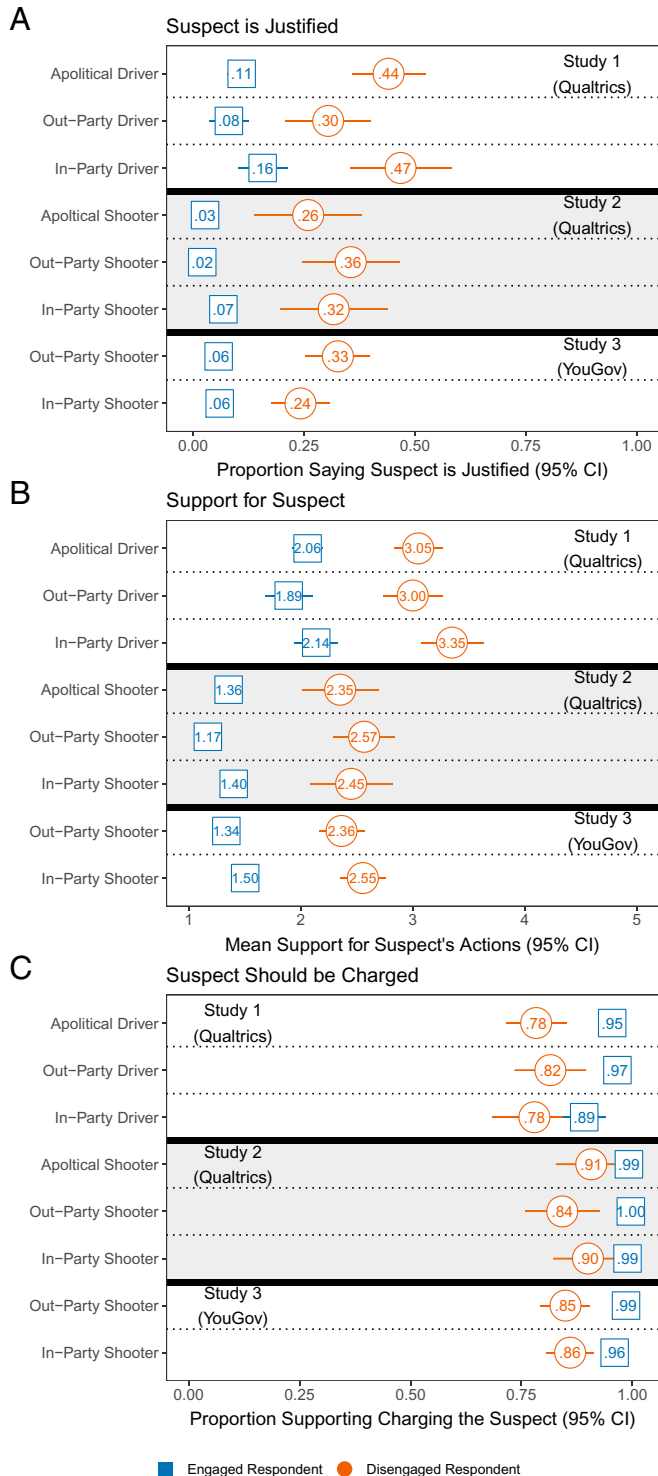


Fig. 2. This figure shows attitudes toward violence for each of our three measures: justification (A), support (B), and should the subject be charged (C). We plot group means and 95% CIs. For the YouGov data (study 3), we utilize survey weights. Providing partisan motivations has no effect on support for violence relative to identical, but apolitical, violence.

was justified, while only 5.86% of engaged respondents say the shooting was justified.

Fig. 2 shows that this overall pattern is found across all treatment conditions in both studies. The red circles and lines in Fig. 2 show disengaged respondents, while teal squares and lines show

engaged respondents. In all cases, disengaged responses indicate significantly greater justification and support for political violence relative to engaged responses.

When it comes to our third outcome question, support for charging the accused, we see a different pattern. Unlike the first two outcome questions, which are abstract moral judgments, this question is concrete: Should those who commit a crime face legal consequences? Consistent with the specificity of this question, we find much higher overall agreement. In study 1, pooling across our conditions, 92.70% of respondents who passed the engagement test want the suspect in the politically motivated violent crime charged, while 79.88% of disengaged respondents want the suspect in the politically motivated violent crime charged. For study 2, 99.30% of engaged respondents and 86.76% of disengaged respondents want the suspect charged. Similarly, in study 3, 97.30% of engaged respondents and 85.48% of disengaged respondents want the suspect charged

Abstract Questions and Disengaged Respondents Inflate Support for Violence. Respondents who fail our engagement test express much higher rates of support for the hypothetical political violence measure used in extant observational studies (which we included in all our studies pretreatment). We show problems with disengaged respondents with two sets of analyses. First, we show, in Table 1, that the current hypothetical question developed by refs. 1 and 2 (measured here with a balanced Likert with a neutral midpoint) generates overestimates of public support for partisan violence, because of disengaged respondents. Across our three studies, we find that support for violence on this measure is nearly twice as large in the disengaged group as in the engaged group.

Second, we look for evidence of satisficing on our three outcome measures. Our preregistered expectation is that disengaged respondents provide upwardly biased responses to abstract questions. We find substantial support for this hypothesis in the data. As detailed earlier, our questions vary in the extent to which they demand a well-considered response. Questions of justification and support require reflection on the criminal act, a personal moral code, and social norms, whereas asking if a person who committed a violent act should be charged requires no such introspection. Assuming respondents are cognitive misers who satisfice to escape considered thought where possible, we should then expect more satisficing on the first two questions than on the third (11).

This is borne out in our data. Fig. 3A shows that, when presented with a dichotomous question and no “don’t know” option, disengaged respondents essentially randomly split their responses between the two choices, while engaged respondents overwhelmingly report that the driver is not justified. Fig. 3B shows that, when disengaged respondents are presented with five choices that include a neutral midpoint, the modal response is the midpoint, with the remaining respondents splitting their responses between the remaining four categories. Both response strategies are consistent with satisficing. A plurality of engaged respondents report strongly opposing violence.

Table 1. Kalmoe–Mason support for violence measure by engagement

	Support for violence		
	Kalmoe–Mason measure % (N)		
	Study 1	Study 2	Study 3
Disengaged respondents	55 (312)	43 (190)	41 (354)
Engaged respondents	21 (690)	26 (833)	19 (1,509)
Combined estimate	32 (1,002)	29 (1,023)	23 (1,863)

Distributions of Responses Among Engaged and Disengaged Respondents

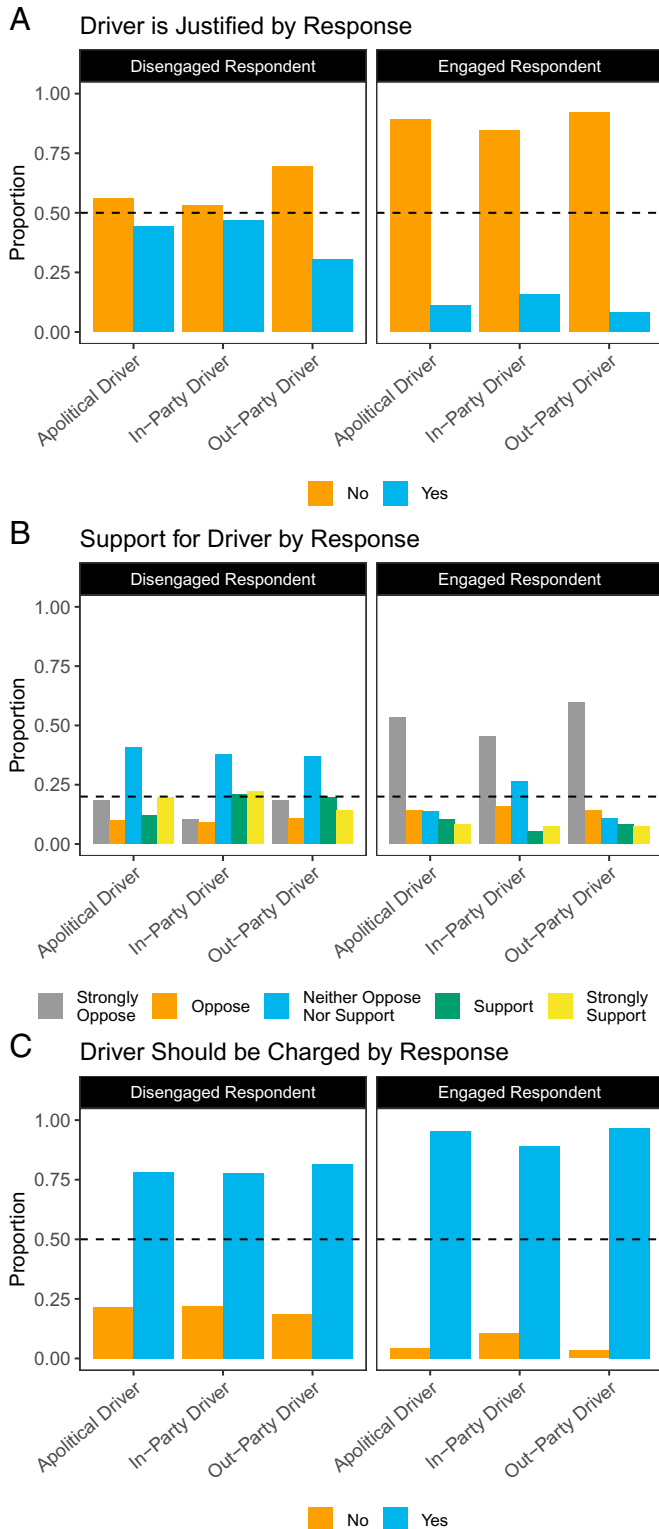


Fig. 3. The response distribution for justification (A), support (B), and charging preferences (C) by engagement for study 1. High levels of support for political violence can be partially attributed to random responding by disengaged respondents, especially when questions are vague.

Fig. 3C shows that, when answering a simpler question with clear normative expectations—charging criminals for crimes—disengaged and engaged respondents are much more comparable. It is also possible that respondents deemed the information in

the newspaper articles we provided insufficient to establish moral justification but sufficient to determine a preference for criminal charges.

Results from study 2, where the reported crime was murder, show a more dramatic difference between the engaged and the disengaged. For engaged respondents, justification peaks at 6.81%, support peaks at 2.15%, and willingness to excuse the suspect from criminal charges peaks at 1.08%. This compares to disengaged respondents where justification peaks at 35.50%, support peaks at 17.10%, and willingness to excuse the suspect from criminal charges peaks at 15.8%. Disengaged respondents report support that is ~8 times greater than engaged respondents.

Study 3, our YouGov replication of study 2, produces very similar results. Justification is ~5.5 times larger for disengaged (32.60%) versus engaged (5.98%) respondents, support is ~9 times larger for disengaged (26.10%) versus engaged (2.89%) respondents, and willingness to excuse the suspect from criminal charges is ~4 times larger for disengaged (15.10%) versus engaged (3.98%) respondents.

These results suggest that overestimates of support for political violence on surveys are partially explained by satisficing and random responses because of flawed questions.

Support for Political Violence Is Lowest for the Most Severe Crimes

We have, so far, demonstrated that disengaged respondents create upward bias in support for political violence and that this is a function of the amount of thought questions require of respondents. Our expectation is that offering additional information—that a suspect has been convicted of a specific crime—reduces question ambiguity enough to attenuate differences between disengaged and engaged respondents. By reporting an exact crime, we are also able to bound what support for political violence exists by crime severity.

Study 4 ($n = 1,009$) captures support for nullifying convictions for a set of politically motivated crimes (some violent and some not) that vary in severity from protesting without a permit to murder. To administer the survey, we first asked standard demographic and covariate batteries and administered a neutral vignette that mentioned a state. We coded engagement by asking respondents to identify the state where a news event occurred in a pretreatment and unrelated vignette (31). Each respondent then read a short prompt informing them that a man, “Jon James Fishnick,” had been convicted of a crime and faces sentencing in the coming week. We then randomly selected a single crime (protesting without a permit, vandalism, petty assault, arson, assault with a deadly weapon, and murder) along with details specifying that the crime was partisan and committed against a member of the opposing party. Participants were then asked to suggest a sentence for Fishnick that ranged from community service to more than 20 y in prison.

Fig. 4 shows the frequency of each suggested sentence by crime and by respondent engagement. When the crime is non-violent (protesting without a permit, vandalism) a near majority of both engaged and disengaged respondents support the minimal penalty of community service. A minimally violent crime (assault—throwing rocks leading to an injury) sees most respondents suggest a term in jail, although about 20 to 25% of respondents still support community service. However, a clear inflection point arrives when the crimes become violent and serious. For the remaining three crimes, respondents overwhelmingly support lengthy prison terms. Almost no engaged respondents favor community service as punishment for severe crimes: arson

Distributions of Proposed Sentences Among Engaged and Disengaged Respondents




Fig. 4. In this study, we remove as much ambiguity as possible by identifying a specific crime for which someone has been convicted. This additional context makes differences between engaged and disengaged respondents largely vanish. Furthermore, respondents, especially engaged ones, punish more-severe violent crimes with longer prison sentences. This suggests that, although support for political violence exists in the electorate, it is primarily constrained to support for minor crimes.

(3.8% of engaged respondents), assault with a deadly weapon (4.6%), and murder (2.6%). Indeed, the majority of engaged respondents believe more than 20 y in prison is the appropriate punishment for murder.

In addition to asking about the appropriate punishment, we asked whether the governor should pardon Fishnick. *SI Appendix, Fig. S2* shows that, on average, respondents only support a pardon for minor crimes. Engaged respondents are, however, much more likely than disengaged respondents to oppose a pardon for serious acts of violence.

Disengaged Respondents Bias Estimates of the Correlates of Political Violence

Our primary goal thus far has been to precisely estimate the levels of support for partisan violence in the public. However, others focus on a second goal: finding the characteristics of individuals that predict support for violence (2, 32, 33). But the same issues that create bias in estimates of support for violence also cause bias in estimates of the relationship between supporting violence and other variables. This is because the usual rules of vanilla measurement error are not applicable with disengaged survey respondents, who are likely to remain disengaged across several questions and therefore cause nonrandom measurement error. The consequence is that disengaged survey respondents can create measurement error that causes bias in an unknown direction and, in some cases, can make the relationships between variables appear stronger rather than weaker.

To get intuition for how this can occur, consider a simple example. Suppose our goal is to measure how much support for violence differs across a dichotomous attribute, X . As in our analyses above, we suppose that our respondents are divided into engaged and disengaged individuals. We will further suppose that being disengaged affects both the reported support for violence and the measured value of X , biasing both upward. As a hypothetical example, suppose that $P(\text{Violence} \mid \text{Engaged}, X = 1) = 0.15$, $P(\text{Violence} \mid \text{Engaged}, X = 0) = 0.05$, that $P(X = 1 \mid \text{Engaged}) = P(X = 0 \mid \text{Engaged}) = 0.5$, and that $P(\text{Engaged}) = 0.8$. But, for disengaged respondents, we suppose that $P(\text{Violence} \mid \text{Disengaged}, X = 1) = P(\text{Violence} \mid \text{Disengaged}, X = 0) = 0.8$, and that $P(X = 1 \mid \text{Disengaged}) = 0.8$. The true difference among the engaged respondents is $P(\text{Violence} \mid \text{Engaged}, X = 1) - P(\text{Violence} \mid \text{Engaged}, X = 0) = 0.1$. But, because of the nonrandom measurement error among the disengaged respondents, the estimated difference using the overall data is $P(\text{Violence} \mid X = 1) - P(\text{Violence} \mid X = 0) = 0.217$. Nonrandom measurement error from disengaged respondents causes the relationship between X and support for violence (measured as the difference in average support for violence at levels of X) to be more than twice as large as the true relationship.

We find evidence that this bias occurs when assessing predictors of political violence. The literature has identified three significant predictors of support for violence: partisan social identity, aggression, and hostile sexism (2, 32, 33). Here we focus on the largest predictor: aggression [as measured in our work with the Buss-Perry Short Form (34) from study 2]. As we show in Fig. 5A, the proportion of respondents who are engaged decreases rapidly at high levels of reported aggressive personality. Fig. 5B shows that, as a result, disengaged respondents are disproportionately represented among those with the highest levels of reported aggressive personality.

The higher reported levels of aggressive personality are coupled with the higher levels of support for violence among disengaged respondents that we documented above, resulting in disengaged respondents creating a stronger relationship between aggressive personality and support for violence. Fig. 6 shows that, if we use all respondents and the original measure of violence support from ref. 2, moving from the least to most aggressive personalities is associated with an 82 percentage point increase in support

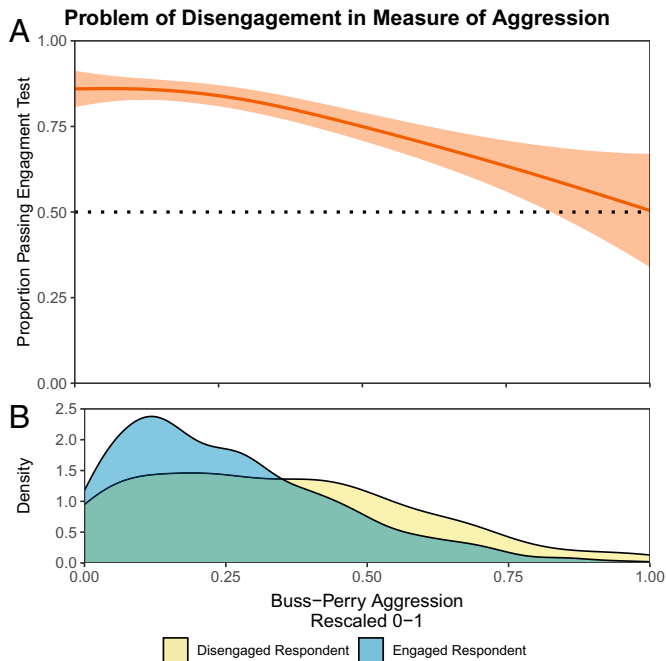


Fig. 5. This figure shows the problems with estimating correlates of support for violence when measures are biased. (A) The proportion of respondents who are disengaged by scored level of aggression on the Buss-Perry scale. (B) The distribution of aggression by engagement.

for violence. That same shift goes down to 67 percentage points among just the engaged respondents with the original measure. But, if we focus on only the engaged respondents, using our more precise measure, that same large shift from least to most aggressive is associated with a 20 percentage point increase in support for violence. Taken together, using imprecise survey questions and failing to account for disengaged respondents produces a relationship between aggressive personality and support for violence that is ~ 4 times too large.

Finally, Fig. 5A suggests that the assumption of a linear relationship obscures a nonlinear relationship (35). We continue to assume a linear relationship to provide an apples-to-apples comparison to the analysis in ref. 2. In *SI Appendix, Table S66*, we provide binned estimates of the relationship between aggression and violence.

Recommendations

Our goal is not to argue that there is no support for political violence in America. Recent events demonstrate that groups of American extremists will violate the law and engage in violence




Fig. 6. This plot shows that the relationship between aggression and support for political violence—as measured as the regression coefficient from a linear regression of support for violence on aggressive personality—is biased upward by disengaged respondents. Moreover, the relationship is much smaller when using a more precise measure of support for political violence.

to advance their political goals. Instead, our purpose is to show that, when attempting to estimate support for political violence among the public, care and precision is required. Generic and hypothetical questions offer respondents too many degrees of freedom and require greater cognition than a sizable portion of the population will engage in. We suggest that future attempts to measure support for political violence 1) utilize specific examples with sufficient details to remove the need for respondents to speculate, 2) benchmark results against general support for all violence, and 3) capture support for crimes that vary in severity.

Conclusion: Limited Support for Political Violence

Our results show support for political violence is not broad based and is, on average, ~ 13 times lower than the average estimate previously reported by Kalmoe-Mason (1, 2) and 6 times lower than the estimate provided by Pape (23). The public overwhelmingly rejects acts of violence, whether they are political or not. Our evidence suggests that extant studies have reached a different conclusion because of design and measurement flaws. When disengaged respondents are not excluded from analysis, measured support for violence is biased upward. Our evidence suggests that this is because disengaged respondents are satisficing in response to ambiguous questions. Vague questions about acceptance of partisan violence demand too much interpretation from respondents, yielding incorrect inferences about support for severe political violence. Not only is support for violence low overall, but support drops considerably as political violence becomes more severe. The most serious form of political violence—murder in service of a political cause—is widely condemned.

Importantly, our results are not conditional on partisanship (*SI Appendix, Tables S2, S20, and S33*). Our results are robust to several other predicted causes of political violence. We find that several standard political measures (i.e., affective polarization and political engagement) are less predictive of support for political violence than are general measures of aggression [measured using the Buss-Perry scale (34); *SI Appendix, Tables S10 and S26*], suggesting that tolerance for violence is a general human preference and not a specifically political preference.⁵ We also find that social desirability [measured with the Marlowe-Crowne scale (36)] does not temper support for political violence on surveys, suggesting that social desirability is not responsible for our lower estimates of support.

In study 3, we address two alternative mechanisms: partisan cheerleading and respondent trolling. We find that both significantly inflate support for violence, but do so for both engaged and disengaged respondents, suggesting that these mechanisms offer additional reasons to be skeptical of prior estimates. To test for partisan cheerleading (37), we use the design from ref. 28. Partisan cheerleaders are significantly more likely to support partisan violence across all three of our measures (*SI Appendix, Table S34*), but this is unlikely to drive our results, as this represents 3.6% of the sample, and cheerleaders are nearly evenly split between disengaged respondents ($n = 33$) and engaged respondents ($n = 38$). Secondly, we test for trolling using a shark bite question (26) as deployed on the American National Election Studies (the expectation is that responses above the known rate indicate trolling behavior). Trolling respondents inflate support for violence on two of our three measures (*SI Appendix, Table S33*), but, again, they represent a small portion of the sample (2.7%) and are

⁵We do, however, find that strong partisans are more likely to support violence.

split between engaged ($n = 17$) and disengaged respondents ($n = 34$). Removing cheerleaders and trolls decreases mean support for political violence from 1.42 to 1.39 (a change of 0.03 points).

Another concern is that focusing on engaged respondents is misleading because true support for violence might be correlated (positively or negatively) with disengaged survey responding. To address this, we derive partial identification bounds assuming that the true support for violence among disengaged respondents is not observable from the survey question (see *SI Appendix, section S9* for details on the methods used below). For example, in the study 3 outcomes asking about murder, if we assume that true support for violence among disengaged respondents is anywhere between 0% and 100%, then the 95% CI expands from [1.76%, 4.22%] to [0.94%, 24.95%]. However, if we cap true support among the disengaged at a more plausible yet still alarming number, such as 20% (approximately the median value reported in prior work), then the partial identification CI shrinks considerably to [0.94%, 6.86%].[¶] We note that 6.3% support is less than the minimum support for violence reported in Fig. 1. Overall, these bounds suggest that, unless disengaged respondents are orders of magnitude more pro-violence than engaged respondents, the population average support for violence is still much lower than previous estimates have implied.

Of course, it is important to understand that, while we show that support for political violence is lower than expected, it is not precisely measured as zero. An important next step is identifying why remaining support exists and where, specifically, violent political action is likely to emerge. Future work could randomize attention and identify what crimes people default to when asked generic violence questions.

[¶]For completeness, we note the other outcomes from study 3. For the justification outcome, the average support among just engaged respondents is [3.83%, 7.11%], with 0 to 100 disengaged support the full population average is [2.63%, 26.97%], and with 0 to 20 disengaged support the full population average is [2.63%, 10.14%]. For charging the attacker, the average charge rate among just engaged respondents is [95.00%, 97.71%], with 0 to 100 disengaged support the full population charge rate is [74.24%, 98.36%], and with 0 to 20 disengaged support the full population charge rate is [92.29%, 98.36%]. Note that all of these estimates are for respondents assigned to the in-party shooter condition, and no survey weights were used.

1. L. Diamond, L. Drutman, T. Lindberg, N. P. Kalmoe, L. Mason, Americans increasingly believe violence is justified if the other side wins. *Politico*, 1 October 2020. <https://www.politico.com/news/magazine/2020/10/01/political-violence-424157>. Accessed 2 March 2022.
2. N. P. Kalmoe, L. Mason, *Radical American Partisanship: Mapping Violent Hostility, Its Causes, & What It Means for Democracy* (University of Chicago Press, 2022).
3. J. M. Carey, G. Helmke, B. Nyhan, S. C. Stokes, American democracy on the eve of the 2020 election. *Bright Line Watch* (2020). <https://brightlinewatch.org/american-democracy-on-the-eve-of-the-2020-election/>. Accessed 2 March 2022.
4. L. M. Bartels, Ethnic antagonism erodes Republicans' commitment to democracy. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 22752–22759 (2020).
5. K. Clayton *et al.*, Elite rhetoric can undermine democratic norms. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2024125118 (2021).
6. J. E. Uscinski *et al.*, American politics in two dimensions: Partisan and ideological identities versus anti-establishment orientations. *Am. J. Polit. Sci.* **65**, 877–895 (2021).
7. D. A. Cox, Support for political violence among Americans is on the rise. It's a grim warning about America's political future. *Business Insider*, 26 March 2021. <https://www.aei.org/opa-eds/support-for-political-violence-among-americans-is-on-the-rise-its-a-grim-warning-about-americas-political-future/>. Accessed 2 March 2022.
8. S. G. Jones, C. Dooze, N. Harrington, The escalating terrorism problem in the United States. <https://www.csis.org/analysis/escalating-terrorism-problem-united-states>. Accessed 2 March 2022.
9. Federal Bureau of Investigation, Hate crime statistics, 2019. <https://ucr.fbi.gov/hat-crime/2019>. Accessed 2 March 2022.
10. Federal Bureau of Investigation, FBI releases 2019 crime statistics. <https://www.fbi.gov/news/pressrel/press-releases/fbi-releases-2019-hate-crime-statistics>. Accessed 2 March 2022.
11. J. A. Krosnick, S. Narayan, W. R. Smith, Satisficing in surveys: Initial evidence. *New Dir. Eval.* **1996**, 29–44 (1996).
12. J. A. Krosnick, Response strategies for coping with the cognitive demands of attitude measures in surveys. *Appl. Cogn. Psychol.* **5**, 213–236 (1991).
13. S. Iyengar, Y. Lelkes, M. Levendusky, N. Malhotra, S. J. Westwood, The origins and consequences of affective polarization in the United States. *Annu. Rev. Polit. Sci.* **22**, 129–146 (2019).
14. E. J. Finkel *et al.*, Political sectarianism in America. *Science* **370**, 533–536 (2020).
15. J. N. Druckman, S. Klar, Y. Krupnikov, M. Levendusky, J. B. Ryan, (Mis)estimating affective polarization. *J. Polit.*, 10.1086/715603 (2020).

Our results offer critical context to stakeholders, citizens, and politicians on the nation's response to political protests in Portland and the events following the 2020 presidential election. A small share of Americans support political violence, but most of this support comes from a troubling segment of the public who support violence in general. Even among this group, support is further contingent on the severity of the violent act and is generally limited to relatively minor crimes. Political violence is a problem in every public, but, as our results show, it is important to carefully and accurately measure such support before raising alarm that might not be warranted. This is especially true when these alarms direct attention, funding, and concern away from other critical policy debates (38).

Violence of the sort seen on January 6 is, at most, concentrated at the extremes of the parties, and, despite the massive news coverage of political violence, the underlying acts are very rare in comparison to general crime trends. Nevertheless, any amount of support for political violence is troubling and worthy of exploration. Researchers should set their sights on these pockets of extremism and organized violent activity—not the casual and frequently underconsidered opinions of everyday voters. Mainstream Americans of both parties have little appetite for violence—political or not.

Data Availability. Experimental data and all replication code have been deposited in Harvard Dataverse (<https://doi.org/10.7910/DVN/ZEH08E>). We have also prepared a Code Ocean replication capsule (<https://doi.org/10.24433/CO.4651754.v1>).

ACKNOWLEDGMENTS. For helpful comments, we thank Alex Coppock, Jamie Druckman, Matt Gentzkow, Andrew Hall, Shanto Iyengar, Jennifer Jerrit, Samara Klar, Yphtach Lelkes, Neil Malhotra, Jonathan Mummolo, Erik Peterson, and Brandon Stewart.

Author affiliations: ^aDepartment of Government, Polarization Research Lab, Dartmouth College, Hanover, NH 03755; ^bDepartment of Political Science, Stanford University, Stanford, CA 94305; ^cDemocracy and Polarization Lab, Stanford University, Stanford, CA 94305; ^dHoover Institution, Stanford, CA 94305; and ^eDepartment of Political Science, University of California, Santa Barbara, CA 93106-9420

16. K. Gift, T. Gift, Does politics influence hiring? Evidence from a randomized experiment. *Polit. Behav.* **37**, 653–675 (2015).
17. G. A. Huber, N. Malhotra, Political homophily in social relationships: Evidence from online dating behavior. *J. Polit.* **79**, 269–283 (2017).
18. C. McConnell, Y. Margalit, N. Malhotra, M. Levendusky, The economic consequences of partisanship in a polarized era. *Am. J. Pol. Sci.* **62**, 5–18 (2018).
19. J. N. Druckman, S. Klar, Y. Krupnikov, M. Levendusky, J. B. Ryan, How affective polarization shapes Americans' political beliefs: A study of response to the COVID-19 pandemic. *J. Exp. Polit. Sci.* **8**, 223–234 (2020).
20. Y. Lelkes, S. J. Westwood, The limits of partisan prejudice. *J. Polit.* **79**, 485–501 (2017).
21. D. Brookman, J. Kalla, S. Westwood, Does affective polarization undermine democratic norms or accountability? Maybe not. *Am. J. Polit. Sci.*, in press.
22. J. G. Voelkel *et al.*, Interventions reducing affective polarization do not improve anti-democratic attitudes. OSF Preprints [Preprint] (2021). <https://doi.org/10.31219/osf.io/7evmp> (Accessed 2 March 2022).
23. R. Pape, Understanding the American insurrectionist movement: A nationally representative survey. *Chicago Project on Security and Threats* (2021). https://d3q10qp55mx5f5.cloudfront.net/cpost/i/docs/CPOST-NORC-UnderstandingInsurrectionSurvey_JUN2021_Topline.pdf. Accessed 2 March 2022.
24. G. King, C. J. Murray, J. A. Salomon, A. Tandon, Enhancing the validity and cross-cultural comparability of measurement in survey research. *Am. Polit. Sci. Rev.* **98**, 191–207 (2004).
25. R. Tourangeau, L. J. Rips, K. Rasinski, *The Psychology of Survey Response* (Cambridge University Press, 2000).
26. J. Lopez, D. S. Hillygus, Why so serious?: Survey trolls and misinformation. SSRN [Preprint] (2018). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3131087 (Accessed 2 March 2022).
27. J. A. Krosnick, Survey research. *Annu. Rev. Psychol.* **50**, 537–567 (1999).
28. B. F. Schaffner, S. Luks, Misinformation or expressive responding? What an inauguration crowd can tell us about the source of political misinformation in surveys. *Public Opin. Q.* **82**, 135–147 (2018).
29. C. C. Marella, "The effects of an integrated reading and writing curriculum on academic performance, motivation, and retention rates of underprepared college students," PhD thesis, University of Maryland, College Park, MD (1999).
30. K. Peyton, G. A. Huber, A. Coppock, The generalizability of online experiments conducted during the covid-19 pandemic. *J. Exp. Polit. Sci.*, 10.1017/XPS.2021.17 (2021).

31. J. V. Kane, Y. R. Velez, J. Barabas, Analyze the attentive & bypass bias: Mock vignette checks in survey experiments. *APSA Preprints* [Preprint] (2020). <https://preprints.apsanet.org/engage/apsa/article-details/5f825740efc0c2001974f0f7> (Accessed 2 March 2022).
32. T. Edsall, No hate left behind. *NY Times*, 13 March 2019. <https://www.nytimes.com/2019.03.13/opinion/hate-politics.html>. Accessed 2 March 2022.
33. L. Mason, K. Mason, What you need to know about how many Americans condone political violence—and why. *Washington Post*, 11 January 2021. <https://www.washingtonpost.com/politics/2021/01/11/what-you-need-know-about-how-many-americans-condone-political-violence-why/>. Accessed 2 March 2022.
34. P. M. Diamond, P. R. Magaletta, The short-form Buss-Perry Aggression Questionnaire (BPAQ-SF): A validation study with federal offenders. *Assessment* **13**, 227–240 (2006).
35. J. Hainmueller, J. Mummolo, Y. Xu, How much should we trust estimates from multiplicative interaction models? Simple tools to improve empirical practice. *Polit. Anal.* **27**, 163–192 (2019).
36. W. M. Reynolds, Development of reliable and valid short forms of the Marlowe-Crowne social desirability scale. *J. Clin. Psychol.* **38**, 119–125 (1982).
37. J. G. Bullock, G. Lenz, Partisan bias in surveys. *Annu. Rev. Polit. Sci.* **22**, 325–342 (2019).
38. A. E. Boydston, *Making the News: Politics, the Media, and Agenda Setting* (University of Chicago Press, 2013).