Reassessing Support for Political Aggression and Violence in the United States

Scott Clifford[1]
Associate Professor
Texas A&M University

Lucia Lopez
PhD Candidate
University of Houston

Lucas Lothamer
PhD Candidate
University of Houston

Abstract. Recent events have driven a surge in scholarly attention to public support for political violence in the United States. Yet, research paints a conflicting picture about the levels and correlates of support for violence. We argue these disagreements are partly due to researchers' measurement choices. After reviewing common practices and identifying measurement challenges, we introduce a measure designed to overcome these problems that allows respondents to choose their target of aggression. Across multiple studies, we compare our measure to two common alternatives. While we find similarities, our measure uncovers substantially more support for aggression and violence, particularly among weak partisans, holding implications for the levels and correlates of support for aggression. Further, by design, our measure provides information about the *type* of aggression that is endorsed and the most common targets. We conclude with recommendations for researchers studying support for political aggression.

Word count: 6,496

[1] Corresponding author.

Political violence has a long history in the United States, but recent events have renewed scholarly interest in the topic. The past several years have brought an increase in aggression against politicians, including death threats, vandalism, and "swatting" (e.g., Herrick and Thomas 2023). More extreme cases of political aggression have even gone so far as physical violence, such as the attack on Senator Nancy Pelosi's husband, the plan to kidnap Governor Gretchen Whitmer, and assassination attempts on former President Donald Trump. In response, many researchers have turned to investigating public support for political aggression (e.g., Armaly and Enders 2022; Holliday et al. 2024; Kacholia and Neuner 2022; Kalmoe and Mason 2022a; Uscinski et al. 2021).

While progress has been made in understanding public support for aggression, there is considerable conflict within the literature. For example, some research finds that the strength of partisan identity is the strongest predictor of support for violence (Kalmoe and Mason 2022a). Other research finds little or no relationship (Armaly and Enders 2022; Uscinski et al. 2021). There is also a significant debate over the prevalence of support for violence (Kalmoe and Mason 2022b; Westwood et al. 2022). In short, basic questions about the extent and correlates of support for political aggression and violence are contested.

In this manuscript, we argue that these conflicts stem, in part, from problems in the conceptualization and measurement of support for political aggression. For example, some prominent work primarily relies on survey questions that explicitly narrow the scope of political violence to *partisan* violence (e.g., Kalmoe and Mason 2019; Holliday et al. 2024). As a result, a strong partisan identity may be a prerequisite for supporting these actions, meaning these results capture only a specific form of political violence. Many survey questions also require

respondents to interpret vague terms like “violence” (for discussion, see Westwood et al. 2022). These problems are common in the literature, highlighting ongoing challenges in measuring support for political violence.

We move this literature forward in several ways. We begin by reviewing prominent ways of defining and measuring support for political aggression, as well as the shortcomings of each. Next, we develop a new measure that is designed to overcome these shortcomings. This measure first asks respondents to name the person doing the most harm in politics, then asks about support for both aggressive and violent actions toward that person. Then, we evaluate the validity and correlates of each measure. Descriptively, we find much lower support for aggression when it is explicitly limited to *partisan* aggression, and the partisan scale seems to be more strongly affected by satisficing and insincere responding. The three measures yield mostly similar stories about the predictors of support for aggression, though only the partisan violence scale finds a strong association with the strength of partisan identity. Finally, in a pre-registered experiment, we show that our measure yields much higher rates of support for aggression and violence than a partisan measure, even when specific actions (e.g., punching a politician) are held constant. We conclude with advice on how to study support for political violence and offer recommendations for future research.

Challenges in Measuring Support for Aggression and Violence

While there are many challenges to constructing valid survey measures, several specific challenges stand out when trying to capture support for political violence. We review these challenges below and discuss the implications for previous research.

There is considerable disagreement about the definition of both “political” and “aggression.” Moreover, as we discuss below, measures often do not match these definitions. In the discipline of psychology, scholars have more than 200 definitions of “violence” (Parrott and Giancola 2007). Within this literature, scholars generally define 'aggression' as behavior intended to harm an unwilling person, with 'violence' being an aggressive act aimed at causing extreme physical harm, such as injury or death (DeWall, Anderson, and Bushman 2011, pp. 246). (DeWall, Anderson, and Bushman 2011, pp. 246). Following this literature, we recognize that harm “can take many forms… such as physical injury, hurt feelings, or damaged social relationships” (Allen and Anderson 2017, p. 1-2). Thus, much like past work in political science, we include both aggressive and violent actions. Throughout the paper, we use “aggressive” as a broader term that includes violence, while we reserve the term “violence” for actions that cause physical harm.

While aggressive actions that do not give rise to physical harm may seem quite distinct from violent actions, they share common goals and consequences that make it important to study them together. Both types of behavior serve to influence political actors through fear and intimidation (Kalmoe and Mason 2024), rather than through conventional means, such as persuasion or electoral pressure. For example, some Republican lawmakers privately admitted to wanting to vote for Trump’s second impeachment but did not out of fear for their safety (Bump 2023). As this case illustrates, threats and harassment can create a fear of harm that affects political behavior.

Our definition includes many behaviors that have been classified as “political violence” in recent years, though we would describe them as aggression. For example, common survey

items have asked about whether it is okay to "destroy or deface property" or "harass or intimidate opposing elected officials" (States United Democracy Center Report, Kalmoe and Mason 2019, Uscinski et al 2021, Vegetti and Litvay 2022). While these behaviors do not meet the psychological definition of violence, as they don't clearly end with injury or death, they may instill the fear of violence and thus influence the behavior of political leaders. Thus, we also include similar behaviors in our measure, though we describe these as aggression rather than violence.

To define "political," we borrow from Kalmoe's (2014, 548) definition of political violence as "violence directed at political leaders and government by individuals or groups with the intention to achieve political aims." However, we prefer aggression over violence to reflect the broader class of behaviors studied in this literature. For these reasons, we define "political *aggression*" as needing to meet the following elements: 1) acts that cause harm or fear of harm, 2) directed at political actors, 3) by individuals or groups (but not the state), and 4) with the intention of achieving political aims.[2]

With a clearer definition of political aggression and violence, we now turn to challenges in measuring the concept and potential problems with past work.

Stipulated Motivations

While sharing similar conceptual starting points, scholars sometimes design their measures in ways that constrain the scope of the study. In particular, some scholars stipulate

[2] Political aggression, through intimidation or harm, also rises above incivility, which involves breaking norms of respect or reciprocity (e.g., Muddiman 2017).

the motivation for supporting political aggression in the survey question. For example, Kalmoe and Mason (2022a) rely heavily on a measure of *partisan* violence[3] that specifically asks about support for aggressive and violent actions against out-partisans. For instance, one item asks about support for violence against the out-party if that party wins the next election. By focusing on *partisan* violence, these measures likely narrow the scope of possible triggers for political violence to those who are motivated by partisan reasons.

By stipulating the reason or purpose for violence, these measures likely shape the predictors. Kalmoe and Mason (2022a) consistently find that the strength of partisan identity is one of the strongest predictors of support for partisan violence. Other work finding the same pattern relies almost exclusively on the same partisan violence measure (Kacholia and Neuner 2022; Mernyk et al. 2021) or a related scale that references ingroups and outgroups without explicitly mentioning partisan groups (Gøtzsche-Astrup 2021). In contrast, all of the studies that *do not* find a relationship with partisan strength use measures that do not explicitly stipulate any partisan or group-focused goals (Armaly and Enders 2022; Uscinski et al. 2021). Thus, these conflicting findings may be due to question design. Of course, if researchers aim to explain *partisan* violence specifically and make no claim to be measuring or studying other types of political violence, then these measures may be valid for that purpose.

Ambiguous and Generic Targets

[3] We follow the authors of this measure in describing it as "partisan violence," though it captures both aggression and violence.

Support for political aggression may vary depending on the target, such as whether they are a politician or a member of the mass public. Many common scales include both types of targets or are ambiguous. For example, a commonly used set of questions (e.g., Armaly and Enders 2022; Uscinski et al. 2021), which we refer to as the "abstract violence" scale,[4] has an item that references the use of violence against "members of the other side" and an item that references disagreement with the "government." These scale items seem to include mass and elite targets, respectively. The partisan violence scale includes items asking about sending threatening messages to partisan leaders, as well as harassing "ordinary" partisans. Other recent work includes items targeting both mass protestors and partisan officials (Holliday et al. 2024). Just as citizens have more negative feelings toward out-party elites than out-party voters (Druckman and Levendusky 2019), people may be more willing to support violence against elites.

Ambiguous Context

A closely related issue is the context provided for the survey questions. For example, consider the question "Is violence ever appropriate when citizens believe something is wrong with their government?" (Munis, Memovic, and Christley 2023). A respondent may reasonably answer "yes" because they believe it's justifiable in an authoritarian regime even though they

[4] This measure is also described by previous authors as measuring support for "violence," specifically. While all items in the scale use the term "violence," it is likely that many respondents are imagining aggressive rather than violent actions (Kalmoe and Mason 2022a; Westwood et al. 2022). See below for discussion.

would never condone it in the US (for evidence of how support varies by context, see Kalmoe and Mason 2022a). Thus, while researchers should avoid stipulating specific motivations for aggressive action, they also need to provide appropriate context.

Abstract Language

Many measures of violence require respondents to interpret terms like "violence" or "force" rather than specifying concrete actions. However, people may interpret these terms in different ways, creating measurement error (Westwood et al. 2022). To illustrate, Kalmoe and Mason (2022a) asked respondents who had endorsed political violence what they meant by the term. Some respondents interpreted the term to mean *lethal* violence. Many respondents interpreted the term more narrowly in terms of threats and insults, which do not constitute violence and some of which may not even constitute aggression. Thus, there is clear heterogeneity in what respondents have in mind when they endorse general statements about "violence," and this variation may inflate estimates of support.

In contrast to the concerns raised above, the use of abstract language and general scenarios may actually *underestimate* support for violence. A long literature on political values has demonstrated a "principle-application gap," in which people support general principles, but are willing to violate those principles in specific applications. For example, people may support democratic values in principle, but be unwilling to extend basic rights to disliked groups (Chong, McClosky, and Zaller 1983). Similarly, one might disagree that "violence is sometimes an acceptable way for Americans to express their disagreement with the government," while also believing that violence against a particular politician is justified because of their actions. Thus, the use of general measures may understate the breadth of support for political violence.

Acquiescence Bias

Acquiescence bias is a well-known phenomenon in which respondents are more likely to agree with a statement in a survey than to disagree with it, regardless of the substance of the statement (Pasek and Krosnick 2010; Saris et al. 2010). This effect particularly applies to questions using an agree-disagree response scale. Besides inflating endorsement, acquiescence can also bias correlations between constructs that are both measured on agree-disagree scales. Acquiescence bias increases the likelihood that some respondents agree with every statement offered in the survey. When those statements are all coded such that agreement indicates higher values of the construct, acquiescence bias will inflate the correlations between those two constructs (e.g., Archer and Clifford 2021; Kuru and Pasek 2016). Of course, including reversed items in agree-disagree scales can help reduce, though not eliminate bias (Leiton 2021).

There are two implications of acquiescence bias for the study of support for political violence. First, measures using an agree-disagree format will overstate support for political violence. Second, these measures may overstate correlations with other constructs measured with agree-disagree scales. For example, several studies evaluate associations between support for violence, measured using an agree-disagree scale, with measures such as conspiratorial thinking and populism – all measured with agree-disagree scales (Armaly and Enders 2022; Uscinski et al. 2021).

The Challenge of Measuring Rare Attributes

In addition to the problems described above, survey-based measures of support for violence face the challenge of sincerity in two ways. The first is the well-known problem of

satisficing, or minimizing the effort expended in taking a survey (Krosnick 1991). In its strongest form, respondents may provide meaningless responses without reading the questions. This is particularly problematic for measuring rare attitudes and behaviors because a small number of strong satisficers can make up a disproportionate share of those reporting that attitude. This problem of measurement error and rare attributes has been discussed in a variety of applications (Ansolabehere, Luks, and Schaffner 2015; Chandler and Paolacci 2017), including support for violence, and can inflate estimates of support (Westwood et al. 2022).

A related problem is insincere responding for humorous or deceptive motivations (e.g., to gain access to future surveys). This problem has been demonstrated with self-reports of belief in conspiracy theories and partisan misperceptions (Lopez and Hillygus 2018) and is most likely to occur for rare attributes. In our studies below, we address satisficing and insincere responding by measuring both response styles with multiple measures.

Partisan Cheerleading

As a final problem, questions about support for violence may be prone to partisan cheerleading. Cheerleading involves the insincere expression of beliefs or attitudes to convey affect towards the parties (for a review, see Bullock and Lenz 2019). Though there is little evidence for these effects on support for partisan violence, we consider it nonetheless (Westwood et al. 2022).

In summary, all existing measures of support for political aggression face shortcomings that are likely to affect conclusions about its prevalence and correlates. While these shortcomings have been documented in past work to varying degrees, scholars have yet to

develop and adopt a general measure of support for political aggression that attempts to address all of them. We turn to that task below.

Introducing a "Most Harmful Actor" Measure of Support for Political Violence

While our focus is on political aggression, the literature on political tolerance is instructive. A classic study on the topic assessed tolerance through attitudes toward communists, atheists, and socialists (Stouffer 1955), but was criticized for conflating the concept of political tolerance with attitudes toward those specific groups. To deconflict, Sullivan, Pierson, and Marcus (1979) developed a "content-controlled" measure that first asks about the group a respondent likes the least, then assesses tolerance for that group. The content-controlled measure is needed because political intolerance is predicated on a "willingness to permit the expression of those ideas or interests that one opposes" (Sullivan, Piereson, and Marcus 1979, 784). Similarly, political aggression is predicated on aggressing against a political actor that one opposes.

Building on the tolerance literature, we develop a parallel measure that we refer to as the "Most Harmful Actor" (MHA) measure, which aims to address the measurement challenges discussed above.[5] The first question reads: "Thinking about American politics today, what one person do you think is causing the most harm to America?" By asking respondents to name

[5] We ask about the actor causing the most harm, rather than the least-liked actor, to reflect our argument that political aggression is meant to influence others' actions. Targeting the actor having the most negative impact offers the greatest influence in politics.

their own target, we avoid constraining the target and thus the relevant motivations, which may be partisan, instrumental, or conspiratorial. However, the question sets the appropriate context – current US politics. Additionally, by focusing on a nominated actor, we avoid the principle-implementation gap in which respondents oppose violence in principle but support it against particular actors.

After nominating a most harmful actor, respondents are asked about support for a series of concrete aggressive and violent actions against that actor on a five-point favor/oppose scale. This design feature eliminates ambiguity in how a respondent interprets the term "violence" and allows the researcher to characterize the extent of aggression and violence a respondent is willing to support. The measure also avoids both agree/disagree scales and imbalanced scales in favor of a five-point favor/oppose scale. This choice reduces concerns about acquiescence bias and satisficing. Finally, the initial open-ended question about who is doing the "most harm" in American politics may serve as an expressive outlet, and thus may reduce partisan cheerleading (Yair and Huber 2020).

In the table below, we summarize the measurement challenges described above. Each column corresponds with one of three measures of support for political violence. The first two are the most common measures: support for abstract violence (e.g., Armaly and Enders 2022) and partisan violence (Kalmoe and Mason 2022a). The last column displays the Most Harmful Actor (MHA) measure. Each row corresponds with a measurement challenge and each cell evaluates how that measure fares on that challenge. We omit challenges that cannot be easily addressed by standard question design, such as measuring rare attributes and partisan cheerleading.

Table 1. Comparing Measures of Support for Political Violence

	Abstract Violence Measure (Uscinski and Parent 2014; Armaly and Enders 2022)	**Partisan Violence Measure (Kalmoe and Mason 2019, 2022a)**	**Most Harmful Actor Measure**
Defined	"general support for and the perceived acceptability of violence as a way of achieving one's political goals."	Ordinary partisans physically threatening or causing harm to opponents (including leaders and citizens). (p. 42)	"1) acts that cause harm or fear of harm 2) directed at political actors, 3) by individuals or groups (but not the state), and 4) with the intention of achieving political aims"
Measure Wording	On a strongly agree to strongly disagree scale: (1) It is acceptable to use violence in advancing political goals these days. (2) Violence is sometimes an acceptable way for Americans to express their disagreement with the government. (3) Violence is justified if the members of the other side act violently first.	Scaled from never to always: (1) When (if ever) is it OK for [in-partisans] to send physical threats and intimidating messages to [out-party] leaders? (2) When (if ever) is it OK for an ordinary [in-partisan] in the public to harass an ordinary [out-partisan] on the Internet in a way that makes them feel unsafe? From "not at all" to "a great deal": (3) How much is it justified for the [in-party] to use violence to advance their goals these days? (4) How justified would violence be if the [out-party] wins the 2020 presidential election?	Thinking about American politics today, what one person do you think is causing the most harm to America? How strongly would you support or oppose others taking the following actions against <actor>? Respondent evaluates specific actions, such as punching, threatening, or killing.
Issue #1: Stipulating Motivations	Does not stipulate motivations for violence.	Stipulates that the violence is against out-partisans.	Does not stipulate motivations for violence.
Issue #2: Ambiguous Targets	Includes items asking about elite (item 2) and mass partisan (item 3) targets.	Includes items asking about elite (item 1) and mass partisan (item 2) targets.	Holds target constant across questions because respondent selects the target.
Issue #3: Ambiguous Context	Item 1 is not anchored in a particular place and item 3 is not anchored in a particular place or time.	Time and place implied in items 3 and 4. Place implied in items 1 and 2 while time is left ambiguous.	Keeps time and place clear – "American politics today."
Issue #4: Abstract Language	All items rely on shared understanding of the term "violence."	Two items (1, 2) gauge support for specific, concrete actions. Two items (3, 4) rely on shared understanding of the term "violence."	All items gauge support for specific, concrete actions.
Issue #5: Acquiescence	All items use agree/disagree scales; no reversed items.	All items use construct-specific scales.	All items use construct-specific scales.

Potential Shortcomings

One possible shortcoming of the MHA measure is that it requires respondents to identify an actor who is causing harm. This could be problematic for two reasons. First, some respondents may not have the knowledge needed to name a relevant political actor. We expect that this is rare and that if a person cannot name an actor, then it is unlikely they would be motivated to support political aggression. Nonetheless, we address this in our studies below.

Second, a respondent may be unwilling to write out the name of an actor, as open-ended responses tend to require more effort from survey respondents. However, responses to open-ended questions are highly effective at identifying satisficing and fraudulent respondents (Kennedy et al. 2016; Kennedy et al. 2020). If a respondent gives an insufficient response to the most harmful actor question, it suggests that the respondent is giving other insincere or insufficient responses. We examine this below.

Item Validation Study

As an initial step, we conducted a study to validate specific actions to be used as outcome measures. Specifically, we sought to identify sets of actions that were perceived to constitute either violence or aggression and to contrast these with conventional forms of political participation.

We recruited 150 respondents from Mechanical Turk to complete the study on 12/18/2023 using the CloudResearch quality filter. Respondents were asked to evaluate 22 actions. We included five violence items that involved physical harm, such as punching a politician. Another eight aggression items involved intimidation without direct physical harm, such as sending threats or cursing in a politician's face. To contrast these actions with

conventional forms of political behavior, we included seven actions, such as writing a letter or donating money to a campaign. Finally, we included two "placebo" items that we had used in a pilot study.

According to our argument, actions that are aggressive but not physically harmful may still create fear that can influence political behavior. To evaluate this possibility, respondents evaluated all 22 actions on several dimensions. Most crucially, respondents evaluated four ways that each action could affect a politician: by causing physical harm, causing the politician to fear for their safety, causing the politician to change their beliefs, or causing the politician to fear being voted out of office. The latter two dimensions capture conventional forms of political influence, while the former capture political influence through aggression. Additionally, respondents rated the acceptability and legality of all actions.

We report the full details in the Appendix but summarize the main findings here. Using these ratings, we selected three sets of items corresponding with our a priori classifications: five physical harm items, six aggressive items, and seven conventional items. The average ratings for each set of items are shown for the four primary dimensions in Figure 1, below. As is clear, the violent actions are perceived as causing substantially more harm than either aggressive or conventional actions (ps < .001). While the violent actions were perceived as causing modestly more fear than the aggressive items (p < .001), both violent and aggressive actions were seen as causing substantially more fear than the conventional actions (ps < .001). Finally, while the conventional actions were perceived as causing little harm or fear, they were perceived as significantly more likely to change a politician's beliefs or to cause perceptions of *electoral* threat than either violent or aggressive actions (ps < .001). This confirms our

expectation that conventional forms of political behavior are perceived as influencing politicians through different mechanisms. Overall, our item validation study provides evidence for two sets of items that are seen as causing harm and fear or causing only fear.

Figure 1. Validation of Violent and Aggressive Actions

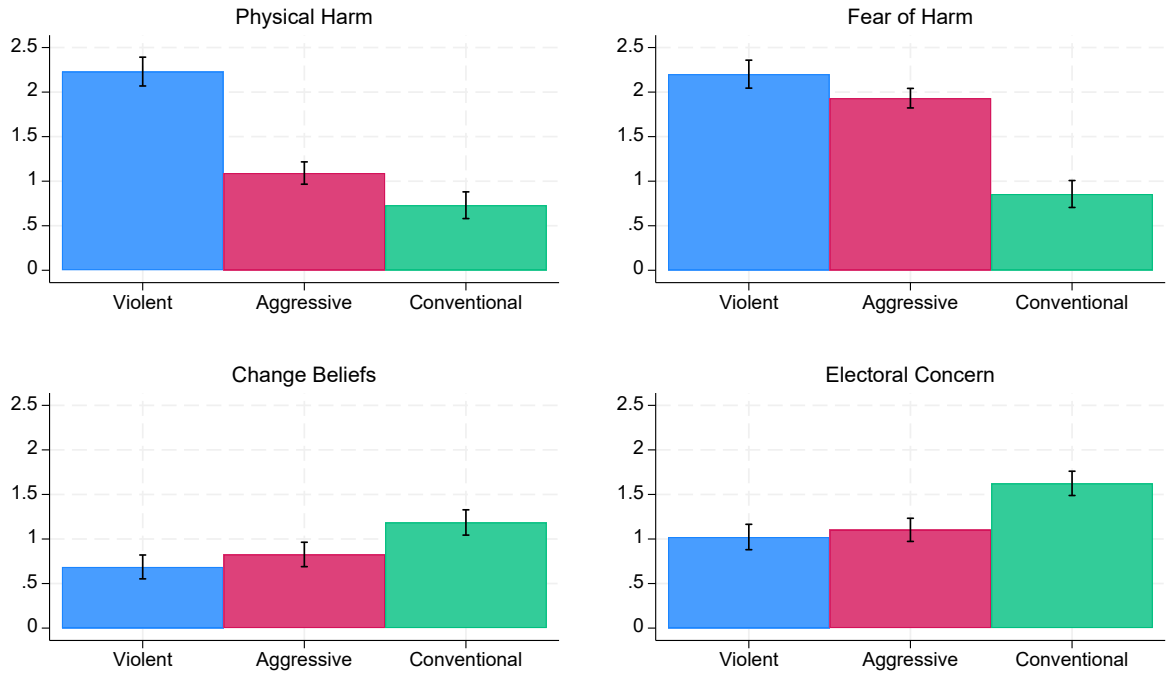

Note: Blue bars represent the average rating for five violent actions (e.g., punch in the face). Red bars represent six aggressive actions (e.g., sent threats). Green bars represent seven conventional actions (e.g., donating money). Panels show the mean rating for perceptions that the named action will cause physical harm (top left), cause them to fear for their safety (top right), cause them to change their beliefs (bottom left), or cause them to worry about being voted out of office (bottom right).

Study 1

To evaluate our new Most Harmful Actor (MHA) measure and to compare it to two popular alternatives, we recruited 1,013 respondents through Bovitz Forthright on Feb 1-6, 2024.[6] The sample was balanced to census demographics on age, gender, race, and region. Prior to fielding this study, we fielded a similar pilot study on MTurk. The results are substantively similar, but we focus our attention here on the more representative Forthright sample while presenting the MTurk results in the Appendix.

At the beginning of the survey, respondents reported their partisan identity and four-item measure of partisan social identity (Huddy, Mason, and Aarøe 2015),[7] conspiratorial predispositions (Uscinski, Klofstad, and Atkinson 2016), and trait aggression (Kalmoe 2014). Respondents then completed all three measures of political violence in random order: Most Harmful Actor, Partisan Violence, and Abstract Violence.

Measuring Satisficing

The survey included multiple items to measure response quality. Open-ended questions are particularly useful for identifying satisficing and fraudulent responses (e.g., Kennedy et al. 2020) so the opening section asked respondents to describe their political views in a single

[6] Invitations were sent to 3,161 respondents and 1,005 provided a complete interview, for a cooperation rate of 32%.

[7] Pure independents were randomly assigned to either the Democratic or Republican version of the partisan social identity scale.

word or phrase. Two authors coded these responses as either valid or invalid responses.[8] Additionally, the trait aggression grid included an instructed response item. We classify respondents who failed either attention check as satisficers.

Measuring Insincere Responding

The survey included two measures of insincere responding (Lopez and Hillygus 2018). In addition to a self-report of insincere responding behavior, respondents were asked to indicate which of five (if any) rare characteristics described them (e.g., more than seven feet tall). Following Lopez and Hillygus (2018), we categorized respondents as providing an insincere response if they self-report insincere responding at least "some of the time" or claimed more than one rare attribute.

Results

We begin with a discussion of data quality given that it has important implications for measuring rare characteristics. As is common in non-probability samples (Kennedy et al. 2016), there is evidence of low-quality responses. Ten percent of the sample failed at least one of the two attention checks and nine percent were flagged as providing insincere responses. Combining these measures, 16% showed some evidence of low-quality responses (see Appendix for further detail). Overall, rates of problematic responding are low in absolute terms, but may have implications for analyzing rare attitudes and behaviors.

The Most Harmful Actor (MHA) Measure

[8] Respondents reached a 94% agreement rate and a Krippendorf's alpha of 0.67.

We now turn to an analysis of the structure of the MHA scale. The first concern is that many respondents may not identify a most harmful actor. However, 89% of respondents named a single actor. Among the valid responses, the most common selections are Donald Trump (46%) and Joe Biden (46%), but respondents named a variety of others, including George Soros, Jerome Powell, and Clarence Thomas. Of those who did not name a single actor, some named a political group or stated they weren't sure (7%), while the rest gave a response indicative of satisficing (3%).[9] These responses are, to a large extent, an indicator of broader patterns of satisficing. The only consistent predictor of invalid responses is failing an earlier attention check (see Appendix for details). Thus, we find little evidence that the question is too demanding for respondents.

We now turn to analyzing the frequency of support for each action among respondents who provided a valid actor.[10] Consistent with past work, support for aggression or violence is low. Figure 2, below, plots the proportion of respondents who at least "somewhat" support each of the 10 actions. Support for violence generally falls below 10 percent, though 12-13% support punching the target in the face. Approximately 20% support publicly sharing the target's whereabouts and more than 25% support cursing in their face. Thus, while support for violence is low, support for aggression is more common. Nonetheless, all 10 items tend to load

[9] Coders reached a 98% agreement rate and a Krippendorf's alpha of 0.99.

[10] In the pilot study, we randomized the inclusion of some non-aggressive "placebo" items but found no evidence that these items affected the endorsement of aggression, so we omitted this feature from our current design.

well onto a single factor, and we find little evidence of a meaningful second factor (see Appendix for details).

Figure 2. Support for Aggression Against Most Harmful Act

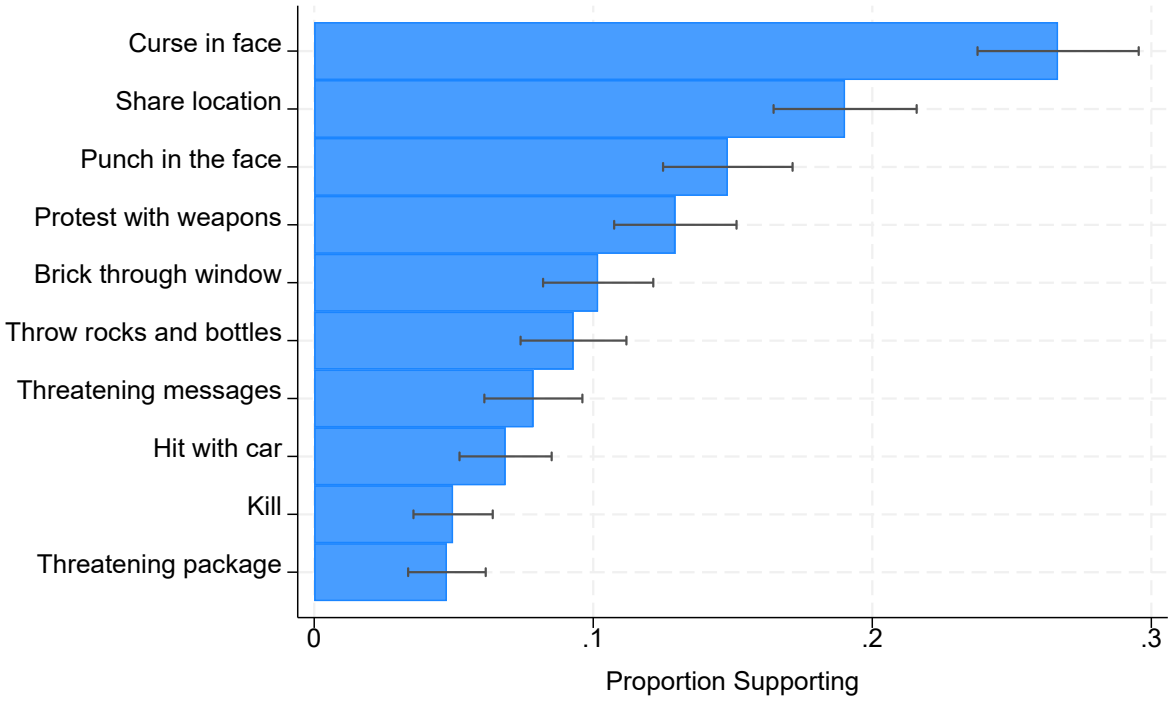

Note: Figure 2 shows the proportion of respondents who "somewhat" or "strongly" support each action against the person they selected as doing the most harm in American politics. Data is restricted to respondents who provided a valid response to the most harmful actor question.

Comparing Measures of Support for Political Aggression

We begin our comparison of the three scales with simple descriptive statistics that include all respondents, including satisficers. Figure 3 displays the distribution of each measure. For all three measures, support for political aggression is low, but this varies considerably by scale. About 29% of respondents scored at the minimum value of the MHA measure, while 30%

did so for the abstract measure. In contrast, 71% scored at the minimum value of the partisan violence scale, suggesting that respondents are less likely to endorse explicitly partisan aggression.

Of course, comparisons between scales are difficult to make due to item content, especially given variance in how respondents interpret the term "violence" in the abstract or partisan scales (Kalmoe and Mason 2022a). Both of these scales include threats and harassment, as well as physical violence, though details differ. In Study 2, however, we conduct a controlled experiment and show that these differences persist.

Figure 3. Distribution of Support for Political Aggression

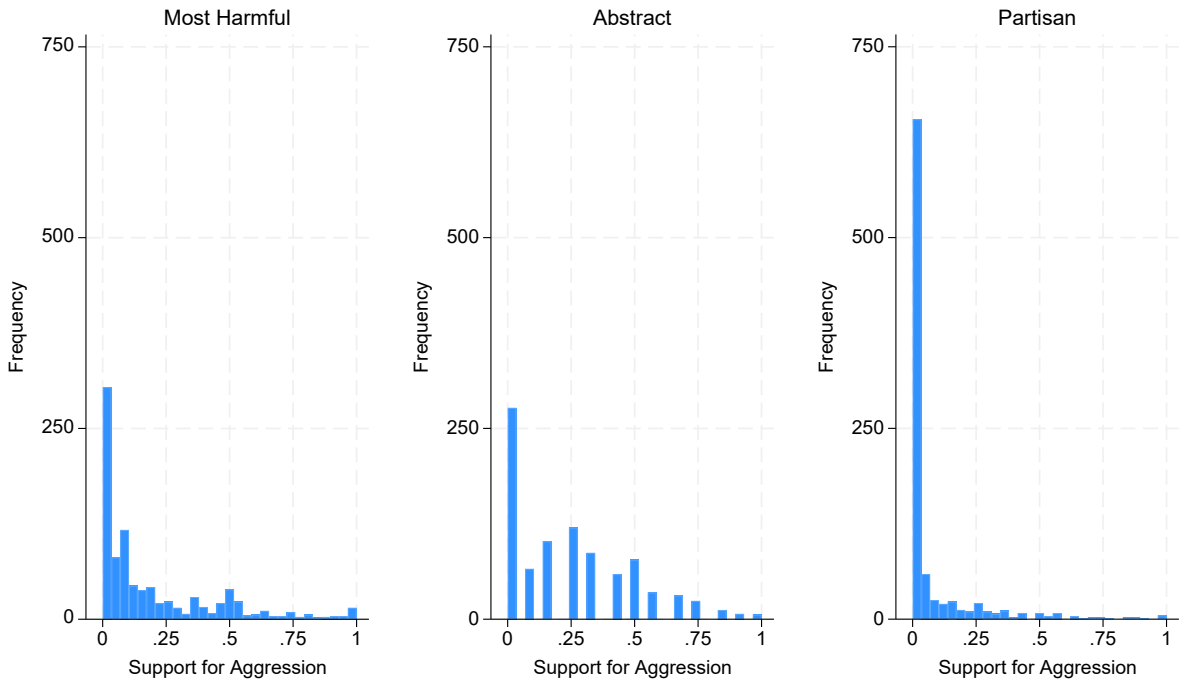

Note: Figure 3 shows the distribution of an additive index of the Most Harmful Actor, Abstract Violence, and Partisan Violence scales. All respondents are included in this analysis.

Of course, satisficing and insincere responding affect these numbers. The problem of measuring rare attributes suggests that respondents engaging in satisficing or insincere responding should make up a disproportionate share of those who endorse some level of aggression, as compared to those who don't. This should be particularly problematic for the partisan violence measure, since fewer endorse it. To illustrate, Figure 4 plots the proportion of respondents showing some evidence of problematic responding based on whether or not they scored at the minimum value of each measure. For all three measures, problematic respondents make up about 8-10% of those completely rejecting aggression. Among those endorsing some level of aggression, this figure climbs to 17-18% on the MHA and abstract scales, but to a troubling 28% of responses to the partisan violence scale. Thus, satisficing and insincere responding inflate estimates of endorsement of political aggression, but this is particularly problematic for the partisan measure, likely because it is less frequently endorsed.

More abstract measures may also be susceptible to order effects, as the survey shapes how respondents interpret the questions (e.g., the target and action). To test this possibility, we randomized the order of the scales. Question order significantly affected responses to the abstract measure, but not the MHA or partisan measures. Specifically, respondents endorsed significantly less abstract aggression when the questions came after the MHA measure (p = .016) or after the partisan measure (p = .046). These effects represent a 19% and 15% reduction in endorsement of abstract violence, respectively. These findings suggest that responses to the abstract measure may be uniquely susceptible to the context of the survey.

Figure 4. Respondents Who Endorse Aggression Are More Likely to Satisfice or Provide Insincere Responses

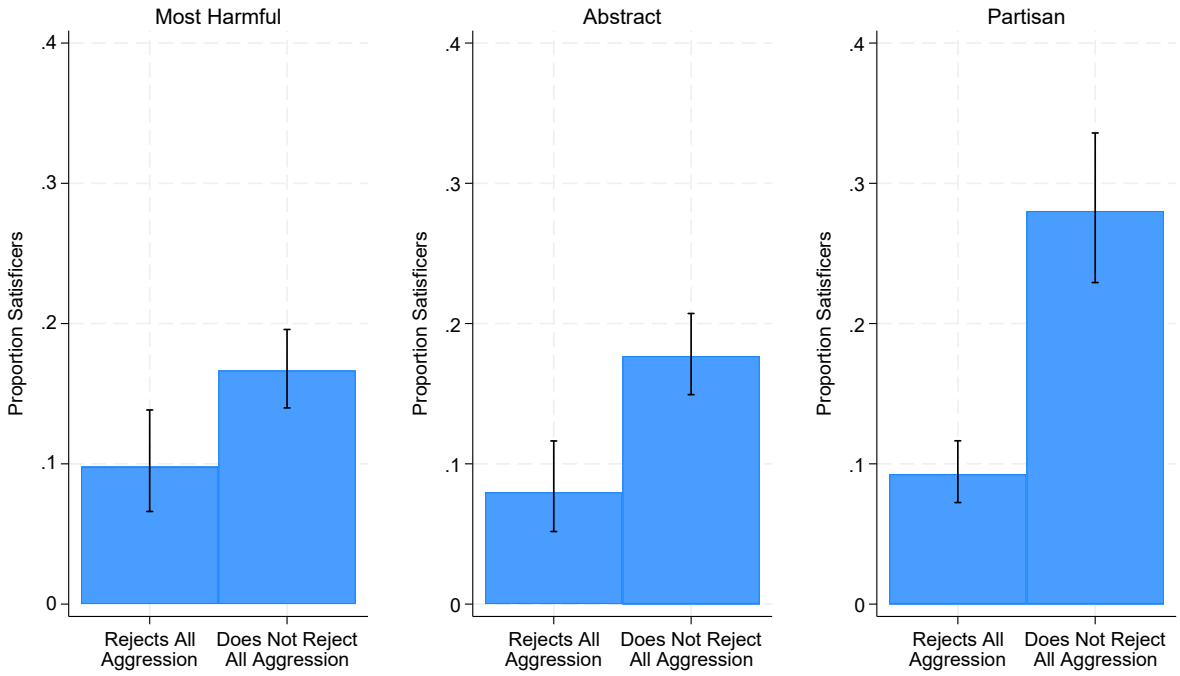

Note: Figure 4 shows the proportion of respondents who showed any evidence of satisficing or insincere responding for those who scored at the minimum value of each scale and those who scored above the minimum value.

We now turn to the relationships between measures. Figure 5 plots each pair of measures against each other with a loess curve. For this analysis, we only include respondents who gave a valid response to the most harmful actor question (n=904), as an analysis of the MHA scale would be unlikely to include invalid responses. As shown in the Appendix, this exclusion does not affect the representativeness of the sample. The MHA measure is moderately related to both the abstract (r = .51) and the partisan scale (r = .47), and the latter two are also moderately correlated with each other (r = .53). As the scatterplots make clear,

many respondents completely reject partisan violence but endorse some other form of violence. Respondents falling along the y-axis of the left and middle panels correspond with a principle-implementation gap – people who reject general statements about violence but are willing to endorse specific acts applied to their nominated actor. Meanwhile, respondents falling along the x-axis in these panels show the reverse pattern – they endorse some amount of violence in general, but reject it when faced with concrete actions applied to a specific person. This pattern is consistent with expressive motivations that are tempered by concrete scenarios. Overall, the measures are moderately related to each other but are far from redundant.

Figure 5. Relationships Between Measures of Support for Aggression

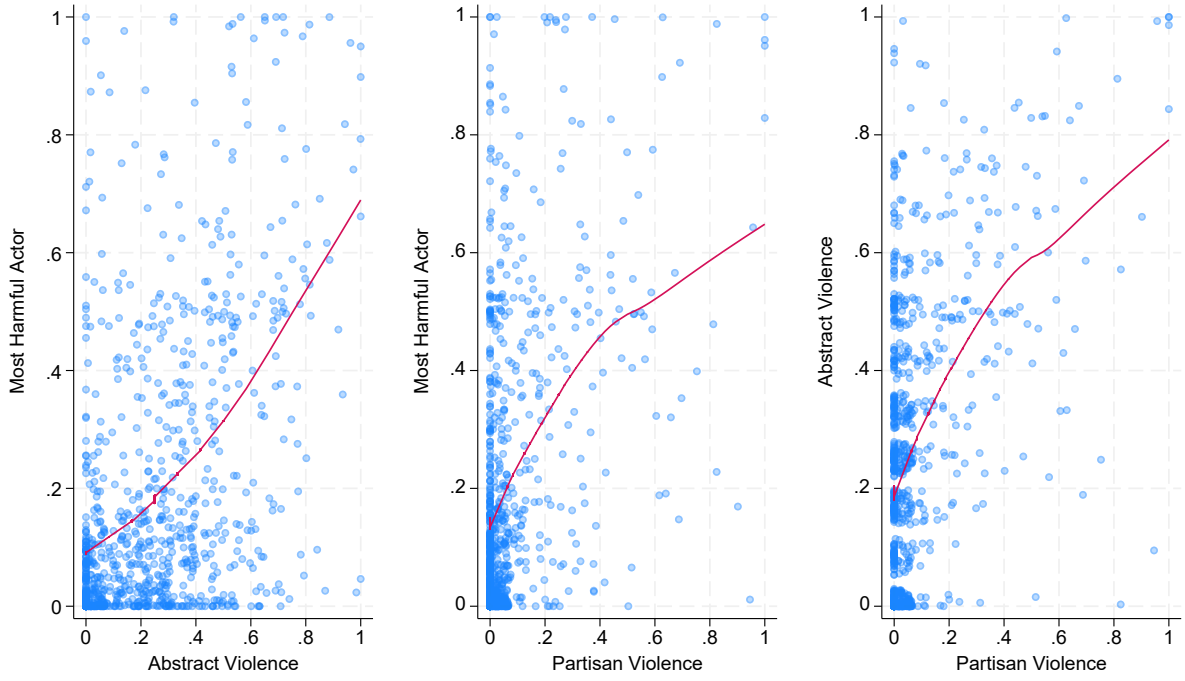

Note: Figure 5 shows a jittered scatterplot of the relationship between each pair of measures and a loess curve. Data is restricted to respondents who provided a valid response to the most harmful actor question.

The Correlates of Support for Violence

We now examine the correlates of support for political aggression and how they vary across measures. To do so, we estimate a series of two-limit tobit models that account for censoring at both ends of the distribution. As predictors, we include the social identity measure of partisanship and partisan strength, ideological extremity, conspiratorial predispositions, trait aggression, dichotomous indicators of satisficing and insincere responding, age, race, and gender. The coefficients are plotted in Figure 6.

Overall, the three measures tend to yield similar stories about the predictors of support for political aggression. However, one major disparity stands out – while the strength of partisan identity is one of the dominant predictors of support for partisan violence, it plays a minor or even inconsequential role for the MHA and abstract measures.

Several other differences stand out. Consistent with our bivariate descriptives, satisficing seems to play a more pronounced role in responses to partisan violence. There is also a discrepancy in the role of the direction of partisan identity. For the MHA, Democrats are significantly more likely to endorse aggression, while the opposite effect shows up in the abstract model. This result seems to be driven by respondents who select Donald Trump as their most harmful actor being significantly more supportive of aggression against him (see Appendix for further detail).

Figure 6. Correlates of Support for Political Aggression and Violence

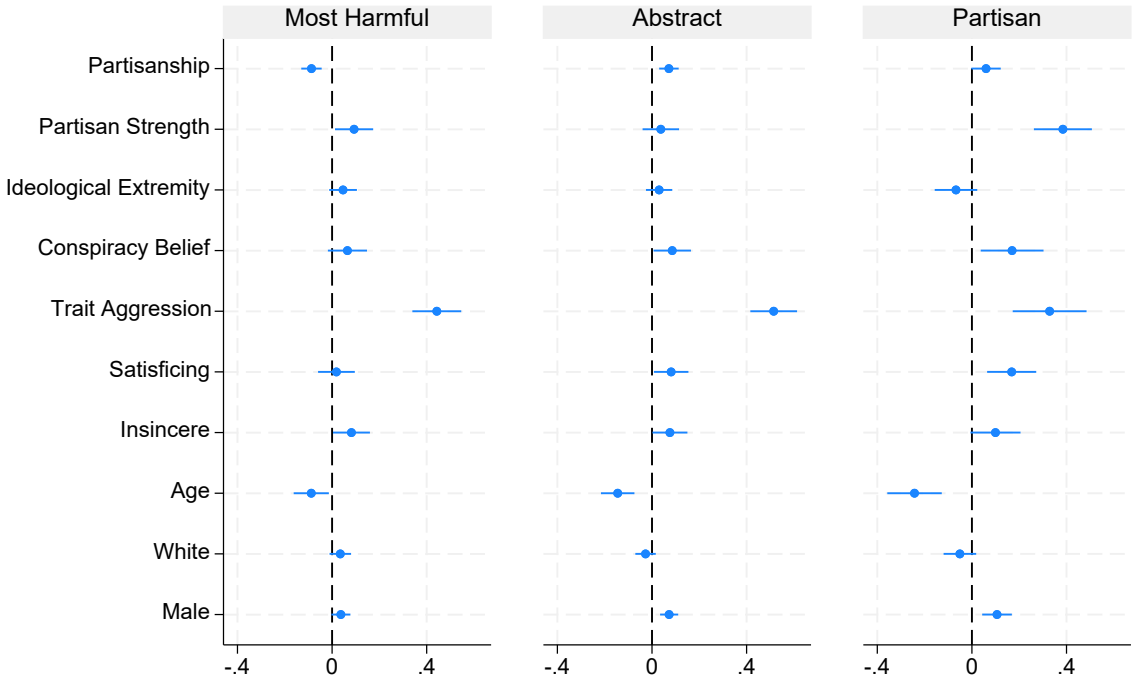

Note: Figure 6 plots the coefficients from tobit models predicting each measure. Data is restricted to respondents who provided a valid response to the most harmful actor question.

Study 2

Study 1 provided evidence that people are more likely to support aggression and violence against a nominated political actor than against out-partisans generally. However, direct comparisons between scales are difficult due to differing item content and scale construction. To provide a stricter test of how the target affects support for aggression and violence, we conducted a third study on Mechanical Turk on 4/5/2024 using the CloudResearch

quality filter. We aimed to recruit 1,000 respondents and 1,001 completed the study. Our study was pre-registered and we distinguish between confirmatory and exploratory tests below.[11]

All respondents were asked about their support for another person carrying out six aggressive and violent actions that range in severity from cursing in a person's face to hitting them with a car. All actions were drawn from Study 1. Respondents all assessed the same six behaviors, but we randomly assigned the target between subjects. Respondents were asked about actions against an out-party voter (*mass partisan*), an unnamed out-party politician (*elite partisan*), or the person they viewed as causing the most harm in American politics (*most harmful*).[12]

As in Study 1, we exclude from analysis respondents without a valid political actor (*n* = 28), following our pre-registration. To enable this exclusion without inducing post-treatment bias (Montgomery, Nyhan, and Torres 2018), all respondents received the open-ended most harmful actor question at the beginning of the survey, regardless of the experimental condition.[13] The rate of non-compliance was notably lower in this study, possibly due to the inclusion of text in the question stem mentioning that we would be asking further questions about their nominated actor.

---

[11] The anonymous registration can be viewed at: https://osf.io/szvqw/?view_only=868b86c0794f41e0adf40b0ca34b55ec

[12] Pure independents were randomly assigned to a partisan condition.

[13] Following the coding procedures from the previous studies, one of the authors, who was blind to the experimental conditions, coded the responses to the most harmful actor question.

Results

We begin with an exploratory description of the data. The left side of Figure 7, below, shows the mean level of support by experimental condition. Overall, support for aggression is low. On a scale ranging from 0 to 4, average support is only 0.38 when targeting partisan voters and 0.47 when targeting partisan politicians. However, support increases to 0.93 for the most harmful actor. Following our pre-registration, we formally test for differences using a two-limit tobit model and include dummy indicators of the most harmful actor and elite partisan conditions. We also control for the strength of partisan identity, trait aggression, and whether they passed an attention check. Supporting our main hypothesis, support for aggression is substantially higher in the most harmful condition than either the elite partisan (b = .78, p < .001) or mass partisan conditions (b = 1.19, p < .001). These findings are robust to alternative modeling strategies and exclusions (see Appendix for details). Consistent with expectations, respondents were also slightly more supportive of aggression toward politicians than voters (b = .40, p < .01).[14]

As an exploratory method of illustrating the size of these effects, we plot the percentage of respondents supporting each action by experimental condition in Figure 7. Averaging across items, respondents were more than twice as likely to endorse aggression against their nominated actor (17%) than against an elite (7%;) or mass partisan (5%). We can also make a

[14] Notably, this difference is modest compared to past work comparing polarized views towards party leaders and supporters (Druckman and Levendusky 2019), though this may be due to the low levels of endorsement.

rough comparison to a recent study (Holliday et al. 2024), which found that only 3% of the public supported the actions of a person throwing rocks at peaceful out-party protestors. Similarly, we find that 4% support throwing rocks at an out-party voter and 5% at an out-party politician. However, this figure increases to 12% for the most harmful actor. Overall, people make a small distinction between aggression against voters and politicians generally, but are dramatically *more likely to support aggression against the person they perceive to be doing the most harm in politics.*

Figure 7. Support for Political Aggression by Target

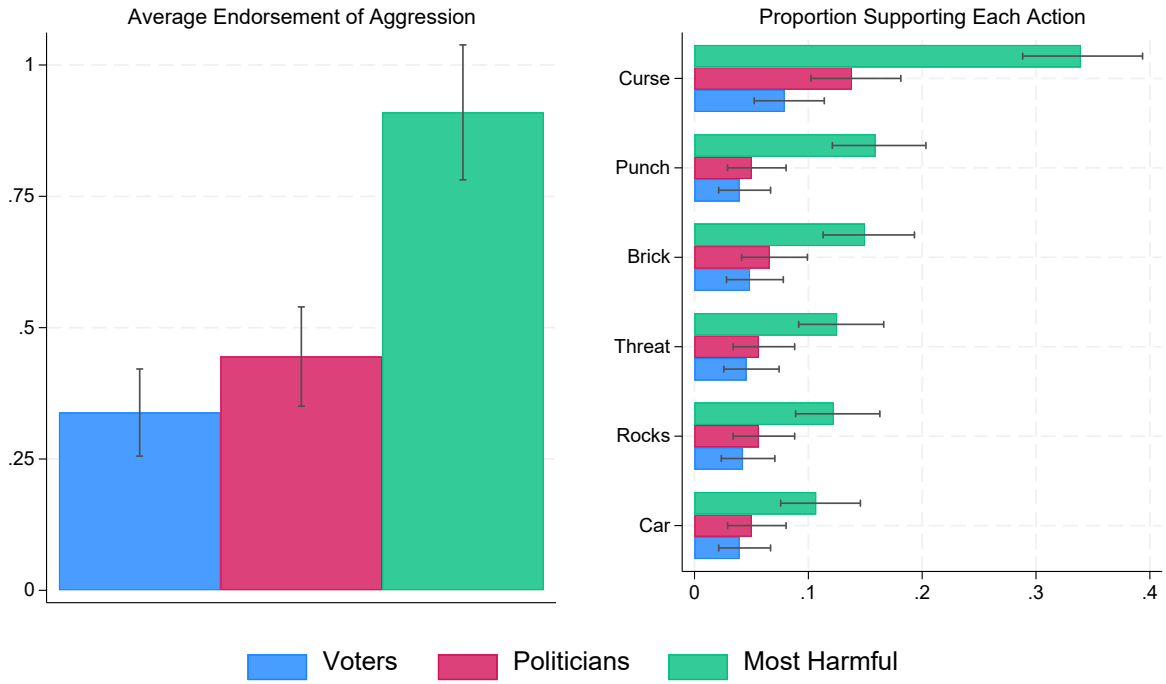

Note: The left-hand panel of Figure 7 shows the average endorsement, on a scale ranging from 0 to 4, of all six aggressive and violent actions. The right-hand panel shows the proportion of respondents who "somewhat" or "strongly" support each action against the target.

Conclusion

Overall, our paper makes several contributions. First, descriptively, we contribute to a debate over the prevalence of support for political aggression, finding that common measures may substantially underestimate support. This is because, while people overwhelmingly reject generic partisan aggression, they are more than twice as likely to endorse aggression against the person that they view as causing the most harm in politics. In other words, few people support indiscriminate partisan aggression, but two to three times as many support aggression against particular political actors. Of course, support for *others* engaging in aggressive behavior is clearly different from willingness to personally engage in it. However, the perceived acceptability of violence tends to be strongly related to actual violent behavior (e.g., Gendron, Williams, and Guerra 2011; Henry et al. 2000) and expectations of others' violent behavior can influence willingness to support and engage in violence (Arms and Russell 1997; Henry et al. 2000; McDoom 2013; Russell and Arms 1995; Werner and Hill 2010). Thus, support for political aggression may influence others' behavior as well as indicate a latent propensity to respond to elite calls to action.

Like past work (Westwood et al. 2022), we also find evidence that satisficing and insincere responding inflate estimates of support for aggression. However, this effect seems to be particularly strong for the common partisan violence measure, presumably due to its lower rates of endorsement. The extent of this problem will, of course, depend on the quality of the sample. Researchers must take steps to minimize and quantify the impact of inattentive and insincere respondents.

Second, we help reconcile conflicting evidence on the role of partisan identity by directly comparing multiple measures within a sample. Some scholars have argued that support for political aggression is primarily a consequence of strong partisan identities (Kalmoe and Mason 2022a) and several independent studies have upheld this finding (Kacholia and Neuner 2022; Mernyk et al. 2021). However, studies using different measures have found negligible relationships with the strength of partisan identity (Armaly and Enders 2022; Uscinski et al. 2021). By comparing measures within a study, our results suggest that the strength of partisan identity is strongly related to support for aggression only when the measure explicitly assumes partisan motivations. This suggests that support for political aggression is not unique to partisans, and that the partisan violence measure only captures a specific subset of supporters.

While we believe there are uses for each measure (for discussion on the value of multiple measurement approaches, see Kalmoe and Mason 2022b), the Most Harmful Actor measure introduced here fares well on several key dimensions. Besides being less affected by satisficing and capturing broader motivations for aggression, the MHA measure has the advantage of providing concrete evidence as to the specific type of aggression or violence that respondents are willing to endorse.[15] Thus, the MHA measure offers greater insight into both the breadth and depth of support for political aggression and we encourage researchers to carefully report how support varies across action type. Nonetheless, further work should be

[15] Other scholars have reported survey vendors refusing to field questions about political assassination. However, we have found that high-quality survey vendors are willing to field the MHA measure, perhaps because it is tailored to respondents' views.

done to develop the measure. In particular, it would be valuable to further examine the optimal items for inclusion in the scale as well as further investigating its factor structure.

Of course, responses to the MHA scale may vary depending on which political actors are most salient at the time. We view this as a feature rather than a bug. For example, support for aggression may be conditional on a politician's gender, which suggests that support for aggression is conditional on the political context. These features make the MHA measure particularly interesting for time series data, which would allow tracking changes in the breadth and depth of support for political aggression, as well as the most common targets of aggression. Overall, given the advantages of the Most Harmful Actor measure, we encourage scholars to adopt it in their study of support for political aggression and violence.

References

Allen, Johnie J. and Craig A. Anderson. "Aggression and Violence: Definitions and Distinctions." In *The Wiley Handbook of Violence and Aggression*, ed. Peter Sturmey (Hoboken: Wiley, 2017).

Ansolabehere, Stephen, Samantha Luks, and Brian F. Schaffner. 2015. "The Perils of Cherry Picking Low Frequency Events in Large Sample Surveys." *Electoral Studies* 40: 409–10.

Archer, Allison M. N. and Scott Clifford. 2021. "Improving the Measurement of Hostile Sexism." *Public Opinion Quarterly*.

Armaly, Miles T. and Adam M. Enders. 2022. "Who Supports Political Violence?" *Perspectives on Politics*: 1–18.

Arms, Robert L., and Gordon W. Russell. 1997. "Impulsivity, Fight History, and Camaraderie as Predictors of a Willingness to Escalate a Disturbance." *Current Psychology 1997 15:4* 15(4): 279–85.

Bullock, John G., and Gabriel Lenz. 2019. "Partisan Bias in Surveys." *Annual Review of Political Science* 22: 325-342.

Bump, Philip. "The Terrorizing Style in American Politics." *The Washington Post*, September 14, 2023. https://www.washingtonpost.com/politics/2023/09/14/romney-political-violence-republicans/

Chandler, Jesse, and Gabriele Paolacci. 2017. "Lie for a Dime: When Most Prescreening Responses Are Honest but Most Study Participants Are Impostors." *Social Psychological and Personality Science* 8(5).

Chong, Dennis, Herbert McClosky, and John Zaller. 1983. "Patterns of Support for Democratic

and Capitalist Values in the United States." *British Journal of Political Science* 13(4): 401–40.

DeWall, Nathan C., Craig A. Anderson and Brad J. Bushman. 2011. "The General Aggression Model: Theoretical Extensions to Violence. *Psychology of Violence* 1(3): 245-258.

Druckman, James N. and Matthew S Levendusky. 2019. "What Do We Measure When We Measure Affective Polarization?" *Public Opinion Quarterly* 83(1):114-22.

Gendron, Brian P., Kirk R. Williams, and Nancy G. Guerra. 2011. "An Analysis of Bullying Among Students Within Schools: Estimating the Effects of Individual Normative Beliefs, Self-Esteem, and School Climate." *https://doi.org/10.1080/15388220.2010.539166* 10(2): 150–64.

Gøtzsche-Astrup, Oluf. 2021. "Dark Triad, Partisanship and Violence Intentions in the United States." *Personality and Individual Differences* 173.

Henry, David, Nancy G. Guerra, Rowell L. Huesmann, Patrick H. Tolan, Richard VanAcker and Leonard Eron. 2000. "Normative Influences on Aggression in Urban Elementary School Classrooms." *American Journal of Community Psychology 2000 28:1* 28(1): 59–81.

Herrick, Rebekah and Sue Thomas. 2023. "Research Note: Rise in Violence Against U.S. Mayors: 2017 to 2021." *Social Science Quarterly* 104(2): 81-91.

Holliday, Derek E., Shanto Iyengar, Yphtach Lelkes, and Sean J. Westwood. 2024. "Uncommon and Nonpartisan: Antidemocratic Attitudes in the American Public." *Proceedings of the National Academy of Sciences* 121(13).

Huddy, Leonie, Lilliana Mason, and Lene Aarøe. 2015. "Expressive Partisanship: Campaign Involvement, Political Emotion, and Partisan Identity." *American Political Science Review*

109(1): 1–17.

Huff, Connor and Joshua D. Kertzer. 2018. “How the Public Defines Terrorism.” *American Journal of Political Science.* 62(1): 55-71.

Kacholia, Suhan, and Fabian G. Neuner. 2022. “Priming Partisan Identities and Support for Political Violence.” *Frontiers in Political Science* 4.

Kalmoe, Nathan P. 2014. “Fueling the Fire: Violent Metaphors, Trait Aggression, and Support for Political Violence.” *Political Communication* 31(4): 545–63.

Kalmoe, Nathan P. and Lilliana Mason. 2019. “Lethal Mass Partisanship: Prevalence, Correlates, and Electoral Contingencies.” In *National Capital Area Political Science Association* <https://www.dannyhayes.org/uploads/6/9/8/5/69858539/kalmoe___mason_ncapsa_2019_-_lethal_partisanship_-_final_lmedit.pdf>.

Kalmoe, Nathan P. and Lilliana Mason. 2022a. *Radical American Partisanship: Mapping Violent Hostility, Its Causes, and the Consequences for Democracy*. University of Chicago Press.

Kalmoe, Nathan P. and Lilliana Mason. 2022b. “A Holistic View of Conditional American Support for Political Violence.” *Proceedings of the National Academy of Sciences.* 119(32).

Kalmoe, Nathan P. and Lilliana Mason. 2024. “Threats as Political Communication.” *Political Communication* 41(1): 162-71.

Kennedy, Courtney, Andrew Mercer, Scott Keeter, Nick Hatley, Kyley McGeeney, and Alejandro Gimenez. 2016. “Evaluating Online Nonprobability Surveys.” *Pew Research Center*. http://www.pewresearch.org/2016/05/02/evaluating-online-nonprobability-surveys/.

Kennedy, Ryan, Scott Clifford, Tyler Burleigh, Phillip D. Waggoner, Ryan Jewell, and Nicholas J. G. Winter. 2020. “The Shape of and Solutions to the MTurk Quality Crisis.” *Political Science*

Research and Methods 8(4): 614–29.

Krosnick, Jon A. 1991. “Response Strategies for Coping with the Cognitive Demands of Attitude Measures in Surveys.” *Applied Cognitive Psychology* 5(3): 213–36.

Kuru, Ozan, and Josh Pasek. 2016. “Improving Social Media Measurement in Surveys: Avoiding Acquiescence Bias in Facebook Research.” *Computers in Human Behavior* 57: 82–92.

Lopez, Jesse, and D. Sunshine Hillygus. 2018. “Why So Serious?: Survey Trolls and Misinformation.” <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3131087>

McDoom, Omar Shahabudin. 2013. “Who Killed in Rwanda’s Genocide? Micro-Space, Social Influence and Individual Participation in Intergroup Violence:” *http://dx.doi.org/10.1177/0022343313478958* 50(4): 453–67.

Mernyk, Joseph S., Sophia L. Pink, James N. Druckman, and Robb Willer. 2021. *Correcting Inaccurate Metaperceptions Reduces Americans’ Support for Partisan Violence*.

Montgomery, Jacob M., Brendan Nyhan, and Michelle Torres. 2018. “How COnditioning on Posttreatment Variables Can Ruin Your Experiment and What to Do about it.” *American Journal of Political Science* 62(3): 760-75.

Muddiman, Ashley. 2017. “Personal and Public Levels of Political Incivility.” *International Journal of Communication* 11: 3182-3202.

Munis, B. Kal, Arif Memovic, and Olyvia R. Christley. 2023. “Of Rural Resentment and Storming Capitols: An Investigation of the Geographic Contours of Support for Political Violence in the United States.” *Political Behavior*.

Parrott, Dominic J. and Peter R. Giancola. 2007. “Addressing ‘The Criterion Problem’ in the Assessment of Aggression Behavior: Development of a New Taxonomic System.”

Aggression and Violent Behavior 12(3): 280-89.

Pasek, Josh, and Jon A. Krosnick. 2010. "Optimizing Survey Questionnaire Design in Political Science: Insights from Psychology." In *The Oxford Handbook of American Elections and Political Behavior*, ed. Jan E. Leighley. Oxford University Press, 27–49.

Russell, Gordon W., and Robert L. Arms. 1995. "False Consensus Effect, Physical Aggression, Anger, and Willingness to Escalate a Disturbance." *Aggressive Behavior* 21(5): 381–86.

Saris, Willem E, Melanie Revilla, Jon A Krosnick, and Eric M Shaeffer. 2010. "Comparing Questions with Agree/Disagree Response Options to Questions with Item-Specific Response Options." *Survey Research Methods* 4(1): 61–79.

Staerkle, Christian, and Alain Clemence. 2004. "Why People Are Committed to Human Rights and Still Tolerate Their Violation: A Contextual Analysis of the Principle-Application Gap." *Social Justice Research* 17(4): 389–406.

Stouffer, Samuel A. 1955. *Communism, Conformity, and Civil Liberties*. New York: Doublebay.

Sullivan, John L., James Piereson, and George E. Marcus. 1979. "An Alternative Conceptualization of Political Tolerance: Illusory Increases 1950s–1970s." *American Political Science Review* 73(3): 781–94.

Uscinski, Joseph E., Casey Klofstad, and Matthew D. Atkinson. 2016. "What Drives Conspiratorial Beliefs? The Role of Informational Cues and Predispositions." *Political Research Quarterly* 69(1): 57–71.

Uscinski, Joseph E., and Joseph M. Parent. 2014. *American Conspiracy Theories*. New York: Oxford Academic.

Uscinski, Joseph E., Adam M. Enders, Michele I. Seelig, Casey A. Klofstad, John R. Funchion,

Caleb Everett, Stefan Wuchty, Kamal Premaratne and Manohar N. Murthi. 2021. "American Politics in Two Dimensions: Partisan and Ideological Identities versus Anti-Establishment Orientations." *American Journal of Political Science*.

Vegetti, Federico and Levente Littvay. 2022. "Belief in Conspiracy Theories and Attitudes Toward Political Violence." *Italian Political Science Review* 52(1): 18-32.

Werner, Nicole E., and Laura G. Hill. 2010. "Individual and Peer Group Normative Beliefs About Relational Aggression." *Child Development* 81(3): 826–36.

Westwood, Sean J, Justin Grimmer, Matthew Tyler, and Clayton Nall. 2022. "Current Research Overstates American Support for Political Violence." *Proceedings of the National Academy of Sciences*.

Yair, Omer, and Gregory A. Huber. 2020. "How robust is evidence of partisan perceptual bias in survey responses? A new approach for studying expressive responding." *Public Opinion Quarterly* 84(2): 469-492.