

Can Social Networks Counter Support for Political Violence? Evidence from a Network Dyad Experiment*

ERIN ROSSITER[†], CARLY WAYNE[‡], TAYLOR CARLSON[§]

February 8, 2025

Abstract

Public support for political violence is increasing in democracies around the world, including the United States, threatening democratic stability and safety. While much research focuses on when citizens support violence, less explores interventions to reduce such support. This study examines whether conversations about political violence within real-world networks can shift attitudes. Using a sophisticated experimental design that allows us to sample hard-to-survey populations, engage real-world network dyads, and observe synchronous interpersonal conversations, we recruit people with real-world social ties who oppose (moderates) and support (extremists) political violence, randomly assigning these pairs to have a conversation about political violence or not. Strikingly, we find that conversations do not change attitudes among extremists, but do cause moderates to weaken their opposition to political violence. Our findings suggest that unlike cross-partisan conversations, encouraging political discussions among friends and family within the context of political violence is a risky strategy, as these conversations may foment rather than reduce extremism.

*We thank the Weidenbaum Center on the Economy, Government, and Public Policy and the National Science Foundation (2241834) for funding support. We thank Charles Jamerlan at CloudResearch for technical support in implementing the study on Connect. We thank participants at conferences and workshops at BYU, University of Chicago, Toronto Behavior Workshop, Peace Science, and PNOC for feedback on the pre-analysis plan.

[†]Nancy Reeves Dreux Assistant Professor, University of Notre Dame, erossite@nd.edu

[‡]Assistant Professor, Washington University in St. Louis, carlywayne@wustl.edu

[§]Associate Professor, Washington University in St. Louis, tncarlson@wustl.edu

1 Introduction

Support for — and acts of — political violence are increasing in democracies around the world, including the United States (Kalmoe and Mason, 2022*b*; Kleinfeld, 2021). From storming the United States Capitol Building to assassination attempts on presidential candidates, Americans across political divides are engaging in political violence. Beyond these highly salient moments of direct political violence, the United States Capitol Police reported that they investigated more than 8,000 threats against members of Congress in 2023 alone (Winston, 2024), and a survey of local and state elected officials found that one in six experienced a threat due to their job and 30% knew a colleague who had left their job due to concerns about safety (Edlin, 2022).

Critically, not only are acts of political violence increasing, but the broader public’s willingness to support violence is increasing, as well. A recent nationally representative survey found that 36% of Americans agreed that “the traditional American way of life is disappearing so fast that we may have to use force to save it” (Wintemute et al., 2023). While likely an overestimate due to vague descriptions of violence (Westwood et al., 2022*a*), other measures show that less than 10% of Americans support specific acts of violence toward political outgroup members, including vandalism (7.6%), assault (3.7%), and murder (2.1%) (Iyengar et al., 2024). These measures show that a majority of Americans reject violence toward political outgroup members. Nevertheless, even a small percentage of Americans still translates to millions of people who support these acts. Importantly, support for political violence is not a belief polarized along partisan lines. Republicans and Democrats report similar levels of support for these acts (Kalmoe and Mason, 2022*b*). The support we see today — and any increase in this support that may occur — represents a growing danger to American democracy.

How can we stall or even reverse support for political violence? Although people may support political violence for a wide range of reasons, we need to investigate paths to deradicalization. Some may point to top-down approaches, arguing that political elites need to lead the charge against political violence, but changing behavior among elites is challenging. Instead, bottom-up approaches, focusing on reducing support for political

violence among voters first, might be a path forward with more levers we can realistically pull to change behavior more broadly.

In this paper, we examine one potential bottom-up approach to reducing support for political violence: peer social networks. Social networks have been shown to be a powerful way to increase behaviors beneficial to a democracy, like voting (Bond et al., 2012; Nickerson, 2008; Sinclair, 2012), and have a tremendous influence on shaping political preferences (Huckfeldt and Sprague, 1995; Sinclair, 2012; Zuckerman, 2005). Networks wield this influence when members share political information, beliefs, and behaviors with other network members, and interpersonal communication is a key way this occurs (McClurg, 2006). Regardless of whether preferences change as a result of information provision and persuasion or social pressure, social influence within networks is a key component of political attitudes that could be leveraged to try to change support for political violence.

We argue that interpersonal communication between in-network social ties can change support for political violence, with those who oppose political violence (moderates) persuading those who support political violence (extremists) to be less supportive. Testing our argument in real social networks is challenging because most Americans do not support political violence in the first place and there is a high degree of homophily within social networks, which presents endogeneity challenges (Minozzi et al., 2020). To test these expectations, we field a sophisticated, multi-wave experiment that allows us to sample hard-to-survey populations, engage real-world network dyads, and observe synchronous interpersonal conversations.¹ Specifically, we first recruit people who oppose political violence (moderates), but know someone personally who supports political violence (extremists), who we also recruit into the study. After each individual completes a pre-survey, these dyads then participate in a third wave where they are randomly assigned to have a conversation with each other about political violence or not, before answering various questions about their political views, including support for political violence.

In contrast to our expectations, we find that moderates are not able to persuade ex-

¹Our pre-registration is available at <https://osf.io/wm3jg>

tremists to be less supportive of political violence via conversation. Conversations can indeed be persuasive, but the most persuasive network tie is the extremist. Extremists' support for political violence is unchanged after conversations with a moderate social tie, but moderates weaken their opposition to political violence after these same conversations. While moderates still broadly express opposition to political violence, this opposition is more ambivalent.

Altogether, we make three key contributions. First, we demonstrate that interpersonal communication between real network ties can dampen opposition to political violence. Although conversations across lines of difference has been shown to have positive effects, such as reducing affective polarization (Levendusky, 2023; Mutz, 2006; Rossiter, 2022; Rossiter and Carlson, 2024a), it may not be an elixir to other problems challenging contemporary American democracy. Consistent with findings that cross-partisan conversation does not reduce support for anti-democratic attitudes (Voelkel et al., 2022), cross-cutting conversations about political violence support might not push attitudes in the direction that strengthens democracy. This also suggests that future research on conversations could more carefully consider asymmetrical treatment effects between conversation partners. Second, we extend research focused on the depolarizing effects of inter-party conversation by examining conversations between people who might share the same partisan identity, but differ in other attitudes. Given extensive partisan homophily within social networks, it is important to consider contexts in which conversations occur between copartisans, but where disagreement can still be present. Third, we introduce a new approach to group experiments that allows researchers to identify causal relationships within existing social networks. Importantly, our findings also have key implications for understanding support for political violence in global context. Specifically, our results draw important parallels with extant research on radicalization and the potential dangers of real-world social ties drawing friends and loved ones into more extremist belief systems.

2 Political Violence in the United States

Political violence is defined as using physical force to produce or resist political change (Della Porta, 1995), and is typically used to describe violent acts committed by non-state actors. Political violence can take many forms, including destruction of property, threatening physical harm to others (Kalmoe and Mason, 2024), assault, and even murder. These violent actions can occur within the context of organized protests advocating for political change or less centralized actions by individuals. Political violence can be levied against those who work directly in government, such as political leaders, police officers, or election workers, but it can also be targeted at civilians and voters. The key that makes violent actions *political* is whether the purpose is to achieve a political goal.

Acts of political violence have increased in the United States over the past several decades. Kalmoe and Mason (2024) highlight that between 2016 and 2021, there was a ten-fold increase in the number of threats against Congress, which was met with a five-fold increase in spending on personal security among members of Congress. An ACLED report indicates that in 2020 and 2021 in the United States, there were 610 armed demonstrations, which were 6.5 times more likely to include violence compared to unarmed demonstrations (Jones, 2022). Although threatening political violence is less common than other forms of political speech (Kalmoe and Mason, 2024), violent political protests are less common than peaceful demonstrations, and physical violence against elected officials is rare, these threats and actions have been increasing in the United States and demand social science inquiry.

Following the call to better understand the rise in political violence in the United States, several scholars have taken a step back to examine public support for political violence. There is considerable debate over how best to measure support for political violence, given that individuals might define political violence differently and not pay adequate attention when responding to surveys, which can inflate estimates (Carey et al., 2021; Kalmoe and Mason, 2022a; Westwood et al., 2022a,b). Despite challenges in accurately measuring support for political violence, survey data using a variety of instruments generally converges on the notion that a small percentage of the American public supports

“physically aggressive political behavior” (Kalmoe and Mason, 2022a), such as murder or assault. Nationally representative survey data suggests that less than 2 percent of Americans support murdering prominent politicians and less than 6 percent of Americans support assaulting prominent politicians between 2022 and 2025 (Westwood and Lelkes, 2024).

Although public support for political violence appears to be relatively low in the United States, even a small percentage of citizens supporting violence nonetheless amounts to a large number of people.² Public support for political violence may beget actual acts of violence. Moreover acts of political violence, even by a small minority, can still have a tremendous impact on democratic functioning by diminishing trust in institutions and contributing to democratic backsliding, not to mention the immediate, direct effects on the victims of political violence. It is therefore important to understand *why* Americans support political violence so that we can investigate methods to dismantle radicalization processes underlying these views.

3 Deradicalization

With increased attention to political violence, it makes sense that scholars have also taken an interest in understanding (de)radicalization. Radicalization is defined as: “a process of escalation from nonviolent to increasingly violent repertoires of action that develops through a complex set of interactions unfolding over time” (Della Porta, 2018). Although there is considerable debate over the extent to which extreme attitudes are directly linked to behavioral outcomes, such as violence, understanding radicalization processes is central for understanding how to mitigate potential precursors to political violence.

Some scholars point to systemic factors, such as modernization or cultural traditions that support violence for driving radicalization (Della Porta, 2018). Here, individuals come to have more extreme preferences and view violence as an appropriate means to achieve their political goals in part because of these ingrained root causes. From a more

²The United States current population includes roughly 262 million adults. If 1% of them support political violence, this would be roughly 2.6 million people.

relational perspective, other scholars highlight the importance of repeated interactions between relevant actors, such as repeated encounters between social movements and the police, that push people toward justifying violence (Della Porta, 2018). On an even more micro-level, a crucial pathway to political violence operates from the bottom up through social networks. Individuals are most likely to get involved in violent political movements when someone they personally know is already a member of the organization and ushers them in via conversation (Della Porta, 1988; Jasko et al., 2017; Passy, 2003; Sageman, 2004). Thus, social ties are a potential double-edged sword: moderate social ties could pull people away from extremist spirals, but extremist ties could function in the opposite way, fomenting radicalization among a broader circle of friends and family members.

Much of the research on radicalization is rooted in understanding Jihadist ideology (Bloom et al., 2019; Mitts, 2022; Piazza and Guler, 2021; Ying, 2024), leaving us with limited understanding of how patterns of Jihadist radicalization might extend to support for political violence more broadly. For example, Jihadist radicalization is generally viewed as highly centralized, with clear leaders driving radicalization efforts, but much of the political violence observed in the United States in recent years comes from decentralized “coalition[s] of social networks” that orbit around “online propaganda hubs and forums” (SPLC). Targeting extremist leaders in the United States might not be as effective as other strategies, like countervailing messaging, due to the decentralized leadership structure that characterizes extremist movements in the United States (Berger, 2021). Instead, targeting deradicalization efforts at the grassroots level might be more effective.

In summary, the radicalization literature generally suggests that social networks can be a powerful source of radicalization, but it is not clear whether this interpersonal dynamic can be leveraged to reverse radicalization. Much of what we know about radicalization comes from studies focused on Jihadist ideology, which may provide a useful framework for understanding radicalization to political violence more broadly, but might not fully extend to broad, cross-partisan support for political violence in the United States today. As such, we seek to investigate the extent to which social networks can be leveraged to support deradicalization efforts via tempered support for political violence

in the contemporary United States context.

4 Why Social Networks Could Reduce Support for Political Violence

Social networks have long been shown to affect political attitudes and behavior, particularly in ways that support political participation and democratic functioning. For example, people are more likely to turn out to vote or donate to political causes (Sinclair, 2012), work on political campaigns (Kenny, 1992), or otherwise engage civically (Klofstad, 2007), when they are encouraged to do so from someone within their social networks. Aside from encouraging political participation, social networks have been shown to provide information about candidates and current events for people who otherwise pay little attention to the news (Eveland Jr and Hively, 2009; Eveland Jr et al., 2012; Lazarsfeld et al., 1968), though the information shared socially can be distorted, potentially contributing to polarization and misinformation (Carlson, 2019, 2024). Political conversations, particularly those between people who disagree, have recently been shown to decrease attitudinal (Klar, 2014) and affective (Kalla and Broockman, 2020; Levendusky, 2023; Rossiter, 2022; Rossiter and Carlson, 2024a) polarization, while increasing tolerance for the other side (Mutz, 2006; Mutz et al., 1996). Together, these findings suggest that social networks can be leveraged to uphold key pillars of democratic functioning, including political participation, an informed public, and tolerance.

There are several reasons why social networks are thought to influence political attitudes and behavior. First, some research suggests that social networks can shape political engagement through social pressure (Gerber et al., 2008). In this case, individuals conform to those in their social network, taking on similar candidate or policy preferences, or agreeing to participate in politics, as a result of pressures to fit in with the group and avoid social sanctioning (Sinclair, 2012). Specifically examining political participation, previous work has varied the amount of social pressure, revealing that the more social pressure that is exerted, the more likely people are to participate (Bond et al., 2012;

Gerber et al., 2008).

Second, social networks can provide information essential for people to learn about politics, form policy or candidate preferences, and engage in politics. [Eveland Jr and Hively \(2009\)](#) demonstrates that political discussion within social networks increases political knowledge. The information shared socially could also lead people to update their policy preferences. While nearly impossible to disentangle, the idea here is that peers not only pressure others into sharing their views, but they also provide information from which people can come to form their own preferences. ([Eveland Jr et al., 2012](#)). Similarly, [McClurg \(2003\)](#) further articulates how information shared within social networks can contribute to participation, writing: “...informal conversations between network partners expose people to political information from the surrounding social environment. . . . People also may be exposed to information about the mechanics of electoral politics and involvement. Information about which candidate to support, why to support that candidate, when the candidate is holding a rally, or even how to just get involved are all types of information that can be effectively exchanged by word of mouth” (p. 7).

Third, drawing on both social pressure and information provision, social networks can facilitate persuasion. A myriad of empirical approaches have documented persuasion within social networks, including panel data ([Baker et al., 2020](#)) and field experiments ([Broockman and Kalla, 2016](#)). Interpersonal communication tends to be affect-laden, which can be a key basis of political attitudes ([Taber and Lodge, 2006](#)) and persuasion more broadly ([Petty et al., 2015](#)). As such, interpersonal interactions can amplify emotional content while social pressure simultaneously activates the desire to belong ([Del Vicario et al., 2016](#)). Emotional content is often further increased by “previous affective ties” between recruiters (extremists who support political violence) and potential targets (moderates who do not support political violence) ([Della Porta, 1988](#)) and the development of new “affective bonds” between these network ties ([Sageman, 2004](#)). When it comes to political attitudes, persuasion is a function of the substantive and emotional content of the information shared, as well as the nature of the relationship within the social network.

How can interpersonal communication within social networks be leveraged to *deradicalize*? We argue that the very informational and emotional mechanisms that are effective for extremists to radicalize moderates could also be used by moderates to deradicalize extremists. Specifically, extremists may moderate their views if they are exposed to counter-attitudinal messages that debunk misinformation underlying their beliefs or place social and emotional pressure to disavow their beliefs that break societal norms, such as support for political violence. Critically, these informational and emotional pathways to deradicalization are most likely to be successful when they are used by individuals who are perceived to belong to an ideological ingroup (Nyhan and Reifler, 2012). Interpersonal conversations are well-suited for deradicalization via these pathways, relative to other contexts such as mass media, because interpersonal settings require attention and engagement to uphold social norms and maintain the communication. The narrative flow of information conversation lends itself to providing cohesive alternative explanations of incorrect information (Lewandowsky et al., 2012), and mixed social conversations can redirect “self-conscious” emotions such as empathy, guilt, or shame (Scheff, 1988) in service of broader social identities (Gaertner et al., 2015; Suhay, 2015) that walk individuals back from extremist beliefs.

While we have learned a lot about the impact social networks can have on attitudes and behavior, we still have much to learn. Substantively, we know little about how social networks can be used to reduce support for political violence. Methodologically, it is hard to generalize extensively from the current body of work because so many different approaches have been used. Studies utilizing observational data have the benefit of real world networks, but they lack any causal identification. However this doesn’t tell us anything about whether the network ties persuaded each other to share similar views. Yet, studies using experimental designs (more common), are limited in their external validity, typically focusing on cross-partisan conversations between strangers, which are less likely to occur in the real world.

5 Hypotheses

Together, we argue that interpersonal communication within social networks could be a central route to reducing support for political violence in the United States. The extant literature demonstrates that (1) political violence is on the rise in the United States, with Americans being increasingly concerned about it, rendering it worthy of scientific study; (2) social networks can play an essential role in increasing support for political violence in some contexts (e.g., Jihadist ideology), but we know less about how they contribute to *deradicalization* more broadly; and (3) social networks have been shown to affect a wide range of political attitudes and could be uniquely suited for helping moderates pull their more radical network ties away from political violence. Building on this body of research, we derive three pre-registered hypotheses.

First, we hypothesize that when extremists who support political violence discuss political violence with a moderate peer from their social network, they will become less supportive of political violence. The key idea here is that the moderate social tie will be able to use informational and emotional appeals uniquely tailored to their extremist social tie to persuade them to be less supportive of political violence. Moderates can use their preexisting relationship with extremists to craft arguments and approach the topics in ways that they think would be most effective for that unique person. Moreover, as a credible source from someone they trust and respect, the counter-attitudinal perspectives might be taken more seriously.

Primary Hypothesis: Among those who support political violence, conversation on the topic with an in-network source who does not support political violence will *decrease* their support for political violence, relative to no conversation.

Building on this primary hypothesis, we propose two secondary hypotheses. Our first secondary hypothesis investigates the effect of a conversation between moderates and extremists on the *moderate*. We expect that the moderate will further decrease their support for political violence as a result of this conversation. The conversation will give

the moderate the opportunity to solidify their views while trying to persuade their peer to become less supportive of political violence. As a result, the moderate should further entrench their opposition to political violence.

Secondary Hypothesis 1: Among those who do not support political violence, conversation on the topic with an in-network source who disagrees will *decrease* their support for political violence, relative to no conversation.

We introduce another secondary hypothesis in which we expect that the causal effect of the conversation on extremists' support for political violence will be stronger in conversations between copartisans. That is, when a moderate tries to persuade her copartisan extremist peer to oppose political violence, she will be more successful than when a moderate tries to persuade her outpartisan extremist peer. The idea here is that being copartisans creates a shared social identity, leading to more common ground between the network ties. The shared social identity and common ground can lead the moderate to be viewed as a more credible source on the issue where the network ties disagree (political violence), which could in turn make her more persuasive than scenarios in which she did not share such common ground.

Secondary Hypothesis 2: Among those who support political violence, the effect of conversation on decreasing support for political violence will be stronger if they perceive their in-network source to be a copartisan, relative to a non-copartisan.

6 Research Design

To test these expectations, we design and field a sophisticated experiment that allows us to sample hard-to-reach populations, engage real-world network dyads, and observe synchronous interpersonal conversations and their effects on downstream political attitudes. We use a three-wave survey design, with an interactive experiment in Wave 3. In Wave 1 we collect pre-treatment covariates for Americans who primarily do *not* support political

violence (“moderates”). These Wave 1 respondents then recruit individuals from within their own personal networks who they believe support political violence (“extremists”). In Wave 2, we collect the pre-treatment covariates for these referred individuals. The experimental portion of our study occurs in Wave 3 where each moderate-extremist dyad returns to the platform at a mutually agreeable time and is randomly assigned to participate in a synchronous online conversation about political violence or not. Following treatment (conversation or not), all subjects answer questions about their support for political violence and related political attitudes and behaviors.

6.1 Sampling: Waves 1 and 2

We use Cloud Research’s *Connect for Researchers* platform to field the study. Connect is a crowdsourcing platform for online research that is well-regarded among survey researchers and has recently been demonstrated to contain high-quality respondents (Kay, 2024). The primary reason we chose Connect is because of its promise for us to use its panel of participants who do not support political violence to recruit a real network dyad who does support political violence to participate in the study with them, as we will discuss next. We also chose Connect because it allows researchers to directly communicate with subjects while preserving anonymity (similar to platforms like MTurk), which is important for this multi-wave study to minimize attrition and avoid collecting personally identifiable information.

In this research, we are interested in a specific kind of network dyad in the American public. We are interested in people who do not support political violence (the vast majority of Americans) who have someone in their life who disagrees on these views, supporting political violence in some form (the minority of Americans). As such, we are not interested in people who have entirely homogeneous networks with respect to support for political violence.³ Among network dyads that feature disagreement on this topic, our

³In a nationally representative YouGov survey fielded in May 2024, we found that 27% of Americans who do not approve of “people reacting to political events with violence (e.g., driving cars through crowds of protesters, harassing voters outside of polling stations, starting fires, threatening harm against elected officials)” know someone who they think definitely (7%) or probably (20%) approves of such political violence. Similarly, only 30% of Americans who approve of such political violence say that they do not know anyone who disapproves. These trends were consistent among both Democrat and Republican

interest is in those who might reasonably have a conversation about these views. Our population of interest is Americans who a) have at least somewhat heterogeneous networks with respect to political violence support, and b) are willing to potentially discuss those views with each other.⁴ As such, a key scope condition of this project is that our results will have limited ability to generalize beyond that population. However, importantly, this also means that our experiment does not test a hypothetical intervention to reduce support for political violence where the real-world implementation and scalability of it is unclear. Instead, our experiment estimates the causal effect of conversations that are actually happening in network dyads in our population of interest. We believe that focusing on a narrower population of interest provides important external validity to our estimated treatment effects.⁵

There are, however, several key hurdles to recruiting even this more narrow population of interest. First, we need to locate respondents who do not support political violence, but, crucially, (1) know someone who does, and (2) are actually willing to reach out to invite this person to participate in the study. Based on pilot testing, we observed that fewer than 20% of moderate respondents likely satisfy these two conditions. As such, in our study, we first employ a pre-treatment screener, qualifying respondents for Wave 1 if they report that they “personally know someone who they think would ever approve of Americans using violence to achieve political goals.” 5,091 respondents who answered yes to the screener question then participated in Wave 1 of the study. Of these respondents, 769 (15%) self-reported inviting a partner to participate in Wave 2 of the study. Of these 769 invitations, 287 (37%) invited “extremists” completed Wave 2. This created

respondents and it is not the case that people assumed that out-partisans were necessarily supportive of violence. Together, this suggests that although the majority of Americans are in homogeneous social networks when it comes to support for political violence, more than a quarter of Americans believe that someone in their social networks disagrees with them on this issue. Thus, we are not describing networks composing a small minority of Americans.

⁴In the May 2024 YouGov survey referenced previously, we found that 39% of Americans who disapprove of political violence, but know someone who approves, would be willing to have a conversation with them about their views on contentious conflicts that can result in violence.

⁵Holliday et al. (2024) similarly stress the importance of external validity in conversation-based interventions, noting that studies aimed at reducing polarization through cross-partisan conversation do not typically evaluate treatment effects among groups representative of the population of interest. Rossiter and Carlson (2024b) make a similar point, evaluating the importance of selection into cross-partisan conversation, though they do not find heterogeneous treatment effects conditional on pre-treatment preferences for selection into political discussion.

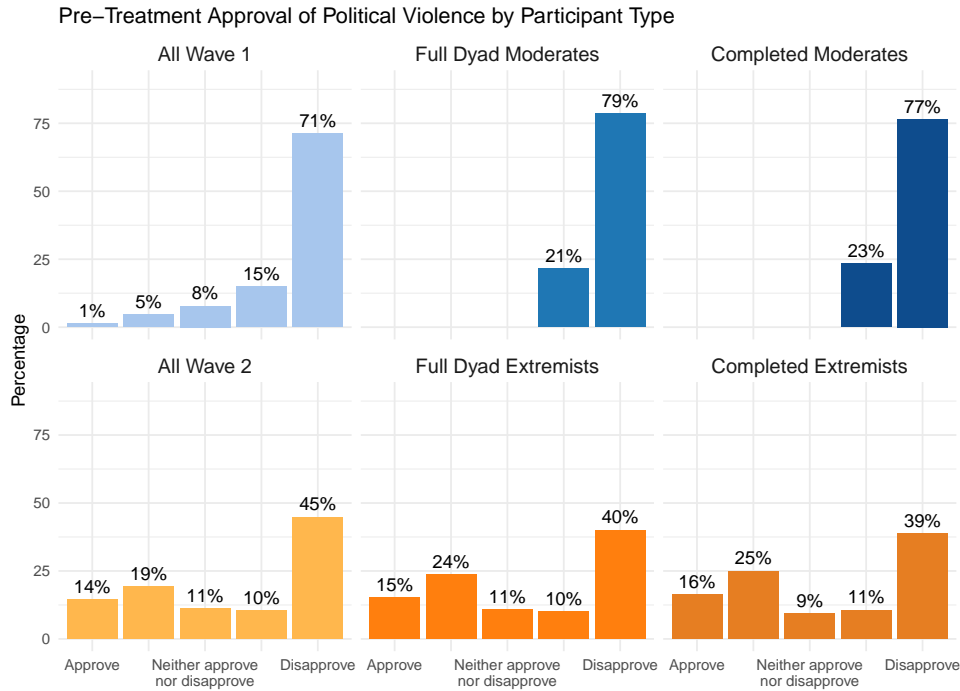


Figure 1: Respondent Support for Violence by Treatment Wave

287 potential dyads, which were then randomized to treatment or control. Of these 287 dyads (574 individuals), 216 dyads (432 people) (75%) completed Wave 3 of the study.

As we demonstrate in Figure 1, we do not observe differential attrition across moderates based on ex ante political violence support. 71% of moderates who participate in Wave 1 strongly oppose political violence, as compared to 78% of moderates who invited an extremist social tie to take the Wave 2 survey, and 76% of moderates who completed the full study. For extremists, two important attributes of political violence support should be noted. First, as anticipated, moderates incorrectly inferred their peer's support for political violence about half the time. Of the extremists who agreed to take Wave 2, 53% *also* disapprove of political violence (or, at least, *say* they do). However, political violence support is also significantly higher in the extremist sample than among our moderate sample. Moreover, we do not see patterns of differential attrition between Waves 2 and 3 among extremists, meaning that extremists who sign up to participate in the study and complete their individual survey (Wave 2) are not distinct from those who return to complete the interactive experimental portion of the study (Wave 3), at least based on observable characteristics. However, we cannot rule out the possibility

that more extreme individuals chose not to respond to the Wave 2 invitation in the first place, because we do not have data on the extremists who were invited, but declined to participate.

6.2 Treatment: Wave 3

Wave 3 contains the randomized experiment. Wave 3 starts with a few questions about the contact the dyad members had between previous Waves 1, 2 and 3. Dyads are then randomly assigned to treatment or control.⁶ To improve balance and power in our relatively small- n study, we implement block randomization to treatment or control. Specifically, we randomize dyads to treatment (conversation) or control (no conversation) within copartisan or non-copartisan blocks. Copartisan blocks are defined as dyads where both participants are (1) a Democrat or Independent who leans toward being a Democrat, (2) a Republican or Independent who leans toward being a Republican, or (3) are both true Independents. Otherwise, dyads are categorized as cross-partisan.

If assigned to the treatment condition, the network dyad will engage in a synchronous online discussion about political violence using an online chat software called Chatter (Rossiter, 2022), prior to answering questions about support for political violence. This app has an interface for participants like iMessage or WhatsApp. It facilitates text-only communication. It is a good fit for our experiment because researchers can pre-specify which participants are members of each chatroom.

Upon returning to participate in Wave 3, both discussion partners in the treatment condition were given a task to complete prior to the conversation. The *moderate* conversation-partner was asked to read a passage that shared several conversational tools they could use to have a productive, persuasive conversation with their social tie regarding political violence, equipping them with conversational strategies to more effectively communicate the negative consequences of political violence to their partner. Specifically, we highlighted sharing personal experiences with political violence, including emotional reactions to learning about political violence in the news, sharing powerful stories of ex-

⁶Importantly, dyad partners do not know prior to opening Wave 3 what their treatment status is, limiting the potential for differential attrition due to treatment status.

amples of political violence, providing information about social norms (i.e., the majority of Americans do not support political violence), and providing information about the effects of political violence, such as democratic backsliding.

This text is important substantively – as we can then analyze the actual conversation text to see if different conversational tools utilized by the moderate were more or less effective in persuading extremists to not support political violence. Ethically this background information provides the moderates with tools to steer the conversation *against* political violence. Briefly put, by preparing the moderate for the conversation, we aim to give the moderate a persuasive “advantage.” They will be prepared to start the conversation, they will have elaborated on their views beforehand, and they will have several discussion points at their disposal during the conversation.

The *extremist* instead receives a filler task that asks them to clear their mind by focusing on several different potential aspects of their daily life. The filler task the treated extremist completes is designed to be completely unrelated to politics, but give the extremist in the pair a reading task of similar length prior to conversation. This is important substantively (e.g. that both participants are entering the conversation with a similar level of survey effort and fatigue), but also technically, as it ensures respondents reach the chatroom at a similar time.

Both dyad members are then provided instructions on how to engage in the chat portion of the study (Figure 2). These instructions are designed to do a few important things. First, they provide respondents with specific examples that can provide a jumping-off point for the conversation, encouraging a more robust, substantive discussion. Second, the examples are purposefully designed to include violence across party lines, including prominent instances of political violence by both Democrats and Republicans. The prompt is also designed to minimize social desirability concerns regarding the discussion of political violence by including language about justice, core values, and fundamental rights. By encouraging respondents to think broadly about what political violence means and what it is used for, we hoped respondents would have a more nuanced discussion about this difficult topic. Finally, ending the prompt with specific questions

Your partner will be joining you in this chatroom. Please begin by writing a message saying hello so they know you are here. You may need to wait a few minutes for your partner to catch up and arrive in the chatroom.

When your partner arrives, we ask that you discuss the following topic: **the use of political violence in the United States today.**

Some Americans deeply disagree with their government and with other people about politics. For example, they may see political officials or institutions as corrupt, they may view the actions of other members of the public as going against their core values, or they may believe government policies violate fundamental rights. In these disagreements, Americans sometimes try to change things through a variety of actions, including voting or protesting.

Sometimes, however, Americans **also use violence**, including assaulting people with opposing views, harassing public officials, destroying or defacing property, taking up arms to fight the state, or harming police officers or soldiers.

Some recent events in the United States related to political violence include, but are not limited to:

- Assault or assassination attempts against politicians, including former President Donald Trump (July 13, 2024), Michigan Governor Gretchen Whitmer (October 8, 2020), Congressman Steve Scalise (June 14, 2017), former House Speaker Nancy Pelosi's spouse Paul Pelosi (October 28, 2022)
- Violence against government employees and their families including police officers, election officials, judges, public health officials and teachers
- Violent protests and riots surrounding issues such as the Israeli-Palestinian Conflict or the Black Lives Matter Movement, both of which have included arson, destruction of property, and injuries and death for protestors, counter-protestors, law enforcement officers, and civilian bystanders
- Violence that occurred on January 6, 2021, at the US Capitol

We would like you to have a conversation about the use of political violence in general to achieve political goals. For example, you might talk about the following topics:

- How often do you think the ends justify the means of political violence?
- What are some potential costs or problems that result from people resorting to political violence?
- What should people do instead of using violence? How can people advocate for or achieve political goals without violence?

Please chat for 15 minutes. The timer begins counting down when your partner arrives. When the timer reaches 0, please return to the survey for a few final questions.

Figure 2: Conversation Instructions for Dyads in Treatment Condition

they might address in the conversation is also designed to provide structure, giving respondents key questions they might try to answer in the course of their conversation. The dyad then proceeds to a 15-minute text-based conversation prior to answering questions independently after the conversation.

In the control condition, dyad members instead proceed directly to the dependent variable section of the study without having *any* conversation or interaction.⁷ We carefully considered the best control condition for this experiment. We are interested in the persuasive, deradicalizing effect of conversation with a real network contact who does not support political violence on the views of those who do. Our desired baseline is observing support for political violence when the network dyad did not have that conversation. The question then becomes—should they have no conversation or a different kind of conversation? If our control condition had the dyads talk, but about another topic

⁷As long as participants had returned to the study, we did have them engage in a conversation *after* measuring our political violence dependent variables of interest. Control condition participants were asked to have a conversation about the state of the US economy. They were given similarly structured filler tasks with information about the economy and various discussion topics. Because these conversations occurred after we measured the dependent variables, they do not affect causal inference at all. However, they do allow us to investigate whether the conversations about political violence were distinct from conversations about other political topics in terms of length and tone, for example.

(whether political or nonpolitical), there might still be persuasion on support for political violence because the participants know that is why they are included in the study. It was impossible to recruit network dyads with our snowball sampling approach who differ in attitudes on support for political violence without participants knowing that is why they were recruited. Therefore, the source cue and/or the act of interpersonal communication alone—even without discussing political violence—with someone who disagrees on support for political violence could lead extremists to moderate their views on the topic. Therefore, we decided on a *no-treatment* control, with control dyads *not* engaging in any conversation prior to answering the main attitudinal dependent variable questions. As such, our treatment is a bundled treatment of (1) conversation with a network dyad who has known disagreement on the topic of interest and (2) talking about the topic of interest. While future research could include additional experimental conditions to tease out how these different elements of a conversation affect the outcome, we could not include more than two experimental arms in our study for power reasons. Recruiting real network dyads with differing attitudes on support for political violence is a steep hurdle, so to maximize power, we opted for two conditions alone.

6.3 Measurement

Our primary measure of support for political violence re-asks respondents the question they received in Waves 1 or 2 of the study, immediately after their conversation about political violence (treatment), or immediately upon entering Wave 3 (control). Specifically, the question is worded as follows:

Some Americans deeply disagree with their government and with other people about politics. For example, they may see political officials or institutions as corrupt, they may view the actions of other members of the public as going against their core values, or they may believe government policies violate fundamental rights. In these disagreements, Americans sometimes try to change things through a variety of actions, including voting or protesting.

Sometimes, however, Americans **also use violence**, including assaulting people with opposing views, harassing public officials, destroying or defacing property, taking up arms to fight the state, or harming police officers or soldiers.

What do you think? Do you approve or disapprove of Americans using violence to achieve political goals?

Do you approve or disapprove of Americans using violence to achieve political goals?

- Approve
- Somewhat approve
- Neither approve nor disapprove
- Somewhat disapprove
- Disapprove

Because of the methodological debate over the most appropriate way to measure support for political violence (Kalmoe and Mason, 2022a; Westwood et al., 2022a,b), we also measure violence support with three additional metrics. First, we ask a follow-up question pertaining to the descriptive text above. This question assesses support through the *frequency* with which political violence may be justified (“When it comes to using violence to achieve political goals, how often do you think the ends justify the means?”). This phrasing may reduce the social desirability effects associated with supporting political violence, where people may be more willing to say violence is justified under certain circumstances rather than saying that they agree in general with the use of political violence.

We also ask respondents a 4-item agree-disagree scale of support for violence, adapted from Nivette et al. (2017), operationalizing support as an additive index of this scale, which reduces measurement noise compared to a single item measure. Finally, we also include the support for violence scale from Westwood et al. (2022a), asking about support for specific acts of political violence, including protesting without a permit, vandalism, assault, arson, assault with a deadly weapon, and murder. We adapt this scale in two ways. First, this scale originally included specific, randomized male names. For simplicity, we use the term “a man,” throughout. Second, this scale originally asked specifically about *partisan* violence. We altered items to be violence against those with opposing political beliefs, instead.⁸

⁸The scenarios were presented in vignettes similar to this arson example: Next, you will read several scenarios. Please indicate if you support or oppose the actions of the person in the scenario. A man was convicted of arson. He was arrested by police as he attempted to run from a fire he started at his local Board of Elections office. Although he waited for the building to close for the night, several adjacent buildings were still occupied. Do you support or oppose his actions?

6.4 Analysis Strategy

To test our primary hypothesis, we focus on **extremists** in the dyad recruited via the snowball sampling procedure. We have two inclusion criteria. First, the extremist’s dyad must have completed Wave 3—including randomization to treatment and control, completion of their assigned experimental condition, and measurement of the primary outcome. Second, the extremist must have answered the pre-treatment support for political violence item stating they either “Approve,” “Somewhat approve,” or “Neither approve nor disapprove” of political violence. This is because moderates often have *incorrect* perceptions about the extremists in their network and invite someone to the study who actually *disapproves* of political violence. In our data, approximately 50% of recruited “extremists” in fact report that they do *not* support political violence. We do not preemptively exclude these dyads from taking the full study and analyze these dyads as an additional exploratory test.⁹

As such, our main estimation approach is as follows. We regress extremists’ post-treatment support for political violence (Y_i) on their treatment status (X_i), and the pre-treatment measure of the outcome P_i to increase precision in our small sample (Clifford et al., 2021):

$$Y_i = \beta_1 X_i + \beta_2 P_i + \epsilon_i \quad (1)$$

For Secondary Hypothesis 1 (treatment effects of conversation among the moderates) we use the same model as with the extremists. Our first inclusion criteria is the same—moderates are only included in treatment effect estimation if they completed Wave 3 (randomized to treatment, engaged in treatment, and answered political violence outcome). Our second inclusion criteria is that moderates do not support political violence, but we are already screening moderates in Wave 1 for this criterion, so any moderate that completes Wave 3 is included in treatment effect estimation:

⁹It is important that including these dyads (where the extremist actually also disapproves of political violence) in Wave 3 does not jeopardize estimation of our primary treatment effects of interest. This is because we omit extremists from estimation based on a *pre-treatment* measure and omitted extremists are thus balanced across treatment and control groups.

$$Y_i = \beta_1 X_i + \beta_2 P_i + \epsilon_i \quad (2)$$

Finally, Secondary Hypothesis 2 predicted that treatment effects would be stronger for extremists in copartisan dyads. To estimate this treatment effect, we again include only extremists that meet the inclusion criteria outlined above, namely, those snowball sampled extremists who indeed display pre-treatment support for political violence. Our main interest is β_3 —the coefficient on the interaction between the treatment indicator (X_i) and the indicator for whether the extremist perceives the moderate to be a copartisan or not (C_i). To operationalize copartisanship, we use the perceived partisan similarity between the moderate and extremist, according to the extremist as they answered in Wave 2 of the survey. This means that if an extremist reports they are a Republican and that they perceive that their friend who recruited them shares their partisanship, this dyad will be labeled as co-partisan, even if the moderate actually self-reports their partisanship as Democrat. We again include the pre-treatment measure of the outcome to increase precision:

$$Y_i = \beta_1 X_i + \beta_2 C_i + \beta_3 (X_i * C_i) + \beta_4 P_i + \epsilon_i \quad (3)$$

After assessing these three main hypotheses, we then proceed to assess potential mechanisms underlying these relationships, downstream persuasion effects for other attitudes and behavior, and potential other heterogeneous treatment effects based on respondent characteristics, which we describe in more detail below.

7 Results

7.1 Extremists and Moderates are Demographically and Ideologically Distinct

We begin by evaluating the characteristics of the extremists and moderates in our data. Note that, because treatment is assigned at the *dyad-level*, observable differences between

Table 1: Demographic Differences Between Moderates and Extremists

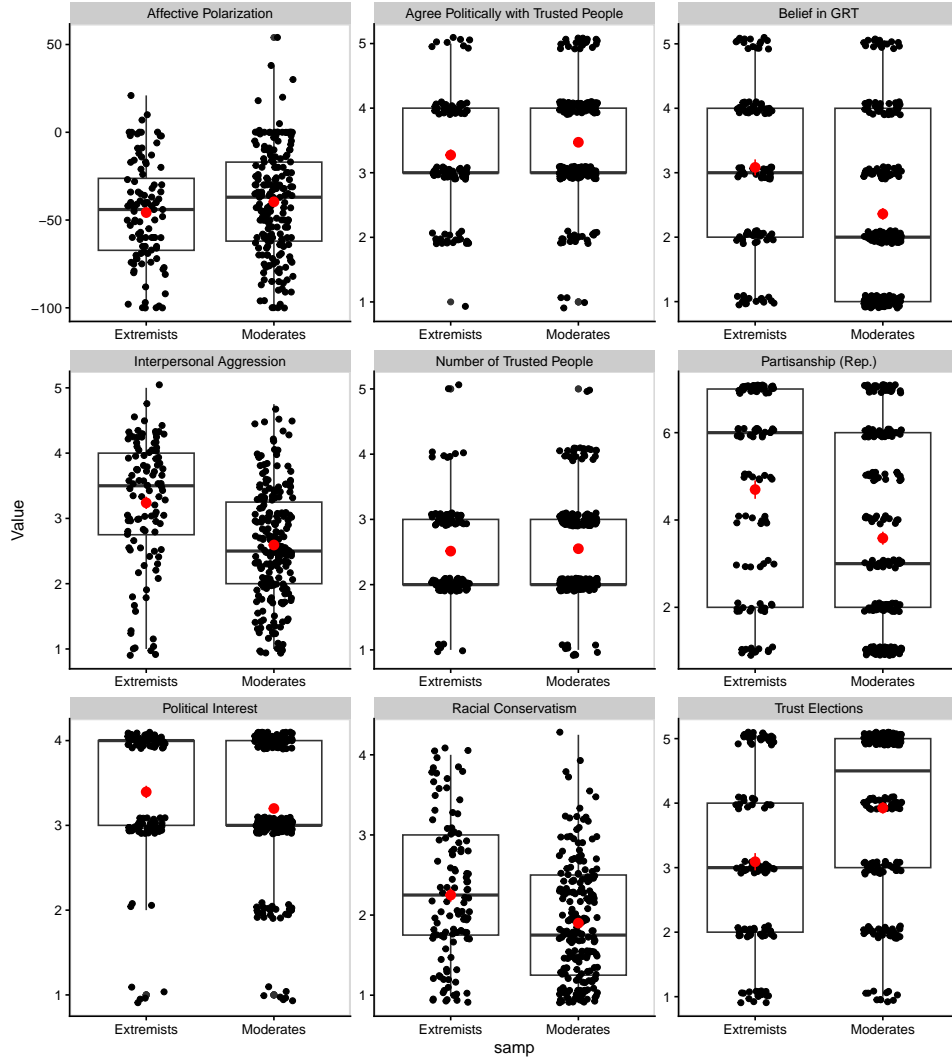
	Moderates mean	Extremists mean	<i>p</i> -value for difference
Democrat	0.578	0.321*	<0.001
White	0.704	0.615*	<0.001
Black	0.202	0.321*	=0.021
Men	0.432	0.661*	<0.001
Evangelical	0.400	0.371	=0.777
Unemployed	0.073	0.037	=0.126
Married	0.411	0.385	=0.639
N	287	109	

discussion partners are expected and do not affect causal inference. However, because we have ample pre-treatment data on extremists’ and moderates’ political and social attitudes, we are able to evaluate the extent to which they differ on these dimensions, in addition to disagreeing about support for political violence.

Consistent with recent work highlighting that Democrats and Republicans support political violence at similar levels ([Kalmoe and Mason, 2022b](#); [Westwood and Lelkes, 2024](#)), we recruited a similar amount of extremists who identified as Democrats (43.1 percent) and Republicans (48.1 percent). This means that our findings are unlikely to be driven (or dampened) by one partisan group being over-represented in the sample. That said, a greater proportion of extremists in our sample were Republican, compared to moderates, which were 54.6 percent Democrats and 37.5 percent Republicans. When it comes to strength of partisanship, however, *exactly* the same percentage of extremists and moderates were strong partisans (37 percent). It is, therefore, not the case that extremists, based on views about political violence, are simply stronger partisans than their moderate counterparts. However, the average levels of interest in politics in our sample overall are much higher than the US adult population, likely based on our recruitment strategy.

Beyond typical political characteristics, we investigated personality characteristics and social attitudes that have been associated with extremism in previous work. Figure 3 illustrates the average levels of each characteristic among moderates and extremists

Figure 3: Distribution of Dispositional & Attitudinal Differences among Moderates and Extremists



in our data. We find that, on average, respondents who support political violence have higher levels of interpersonal aggression, conspiratorial thinking, racially conservative views, belief in great replacement narratives, and antipathy towards outpartisans. They also have a smaller number of people they trust with whom they agree politically and are less likely to trust elections. Finally, extremists are significantly more interested in politics than their moderate social ties, =which is consistent with recent evidence that people who are more interested in politics also have more extreme policy preferences (Krupnikov and Ryan, 2022), but is an important point we return to further in the discussion.

7.2 Network Ties have Common Ground, Discuss Politics Frequently, and Disagree on Political Violence

Although we deliberately recruited network ties who disagree with each other on support for political violence, we observe substantial similarity on many political and social characteristics. Table 2 shows the percentage of dyads who considered their social tie to have each characteristic in common. We find high levels of homophily in social characteristics in our dyads, with 93 percent of dyads perceiving similarity on gender, and 92 percent on race. Politically, 88 percent of people within a dyad believed that their social tie voted for the same candidate for president in 2020, which dropped slightly to 86 percent in the 2024 election. There is somewhat less homophily when it comes to age and religion, but the people within a dyad are still overwhelmingly similar to one another, even if they disagree on their support for political violence.

Table 2: Perceived Similarity Between Network Ties

	Percent of Dyads
Gender	93
Race	92
2020 Vote Choice	88
2024 Vote Choice	86
Party	86
Age	83
Religion	82

Notes: Both moderates (Wave 1) and extremists (Wave 2) were asked “Do you think you and [initials of network tie] have the following things in common or not?” Followed by the seven characteristics listed in the table. The percentages in the table reflect cases in which both the moderate and extremist reported that they had this characteristic in common. Percentages are based on dyads who completed the experiment (Wave 3) entirely.

These dyads reflect various types of real-world relationships: 35.8 percent are spouses, 18.5 percent are other family members,¹⁰ 32.6 percent are close friends, and a small proportion are acquaintances (4.9 percent) or colleagues (7.3 percent). This is broadly reflective of political discussion networks more generally, where previous work finds that

¹⁰Broken down further, these family members are: siblings (32.9 percent), parent-child (51.9 percent), in-laws (5.1 percent), and other relatives, including aunts, uncles, cousins, nieces, nephews, and grandparents (10.1 percent).

political discussions commonly take place within families (Conover et al., 2002).¹¹

Across these distinct types of relationships, political discussion was relatively common among the network ties who participated in this study, further supporting the external validity of our design. Only 1.9 percent of moderates and 2.8 percent of extremists reported that politics never comes up when they talk to this social tie, with 38.4 percent of moderates and 27.8 percent of extremists reporting that politics comes up frequently. The majority of both moderates (73 percent) and extremists (63.9 percent) report that political violence, specifically, has come up in conversation with their social tie in the past.

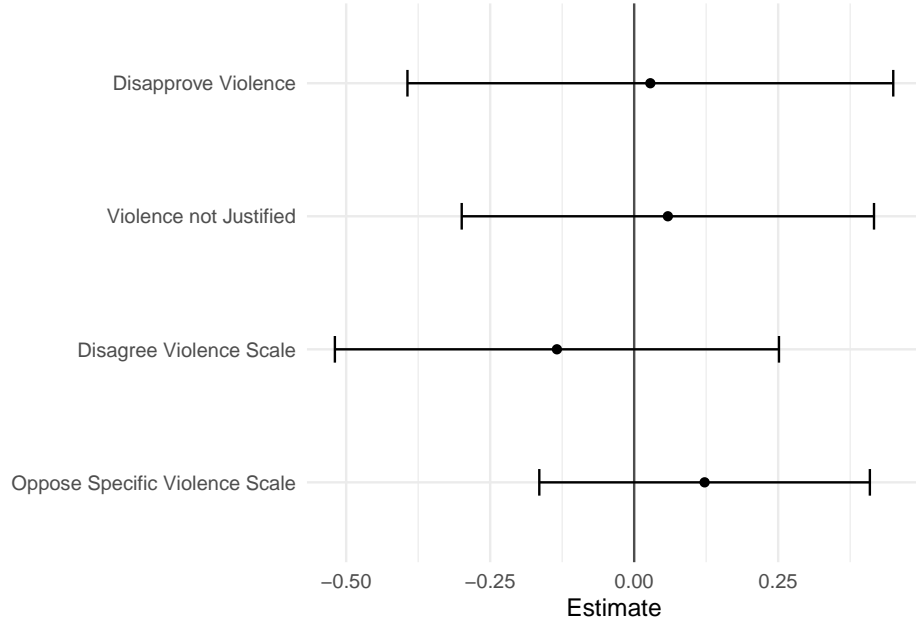
Participants' past experiences discussing politics with each other reveal some degree of disagreement, as expected. 58.5 percent of moderates and 45.7 percent of extremists reported that political conversations lead to arguments at least sometimes. These political conversations are most commonly viewed as neutral by both moderates (32.4 percent) and extremists (35.2 percent), but 32 percent of moderates and 24 percent of extremists considered political conversations with each other to be somewhat or very heated in the past. Despite occasional heated disagreements, the majority of both moderates (64 percent) and extremists (60.6 percent) report that they see eye-to-eye on most things or everything, aside from their views on political violence. This suggests that we are overwhelmingly capturing network dyads that agree with each other on other political issues, perhaps those more commonly studied in network research, but disagree on the topic at hand: political violence. The dyads, therefore, have substantial "common ground" on which to draw in shaping their conversations and persuasion strategies.

7.3 Conversations Did Not Affect Extremists' Support, but Reduced Opposition Among Moderates

Turning to our main results, we next test whether conversations about political violence with moderates would reduce support for political violence among extremists, our primary

¹¹Although this is consistent with previous research, we note that the logistics of this study might have made it easier to complete within a household, thereby potentially inflating the percentage of spouse dyads.


Figure 4: Opposition to political violence among extremists in treatment and control



hypothesis. As shown in Figure 4, there is no statistically significant difference in support for political violence (*DisapproveViolence*) among extremists who discussed political violence with a moderate and those who did not ($\beta = 0.03$, $SE = .22$). We present the results for analyses using only extremists who met our pre-registered criteria (e.g. were *actually* extremists), but in the appendix, we show that the results are substantively and statistically the same when we include all extremists. Importantly, the effect size here is quite close to zero, indicating that the lack of a significant effect is not due solely to a low statistical power and correspondingly large confidence interval.

Given the complexity of measuring support for political violence (Westwood et al., 2022a), we also pre-registered that we would examine treatment effects among additional measures of support for political violence, including how often they think political violence is justified, a scaled support for violence variable (Nivette et al., 2017), support for political violence across a range of specific scenarios (Westwood et al., 2022a), and an additive index of support in those scenarios. Specifically, presented participants with brief scenarios in which a man was convicted of protesting without a permit, vandalism, assault, arson, assault with a deadly weapon, and murder, then asked if they supported the man's actions. The *Oppose Specific Violence Scale* results in Figure 4 reflect the

Figure 5: Opposition to political violence among moderates in treatment and control




average support across these scenarios.¹² Indeed, Figure 4 shows consistently null results across *all* four of these measures of support for political violence. Across all of these measures, we do not find evidence that conversations with moderates reduced extremists' support for political violence.

Although our primary focus was on extremists, we also pre-registered a secondary hypothesis regarding the effect these conversations would have on the moderates, expecting that the conversation with an extremist would further reduce moderates' support for political violence, relative to no conversation (Secondary Hypothesis 1), due to the defensive bolstering that would take place prior to and during the conversation. As shown in Figure 5, however, we found the opposite: moderates who discussed political violence with an extremist from their social network reduced their opposition to political violence, compared to moderates who did not have a conversation ($\beta = -0.34$, $SE = 0.11$). Just as with our test of our Primary Hypothesis, we examine multiple measures of support for political violence and find consistent results: the belief that violence is justified, and support for specific violent scenarios all significantly increase.

¹²Breaking these out into individual items, we find that opposition significantly softens for: protesting without a permit (marginally), assault, arson, and murder. Vandalism and assault with a deadly weapon are not significantly different. See Appendix for tables.

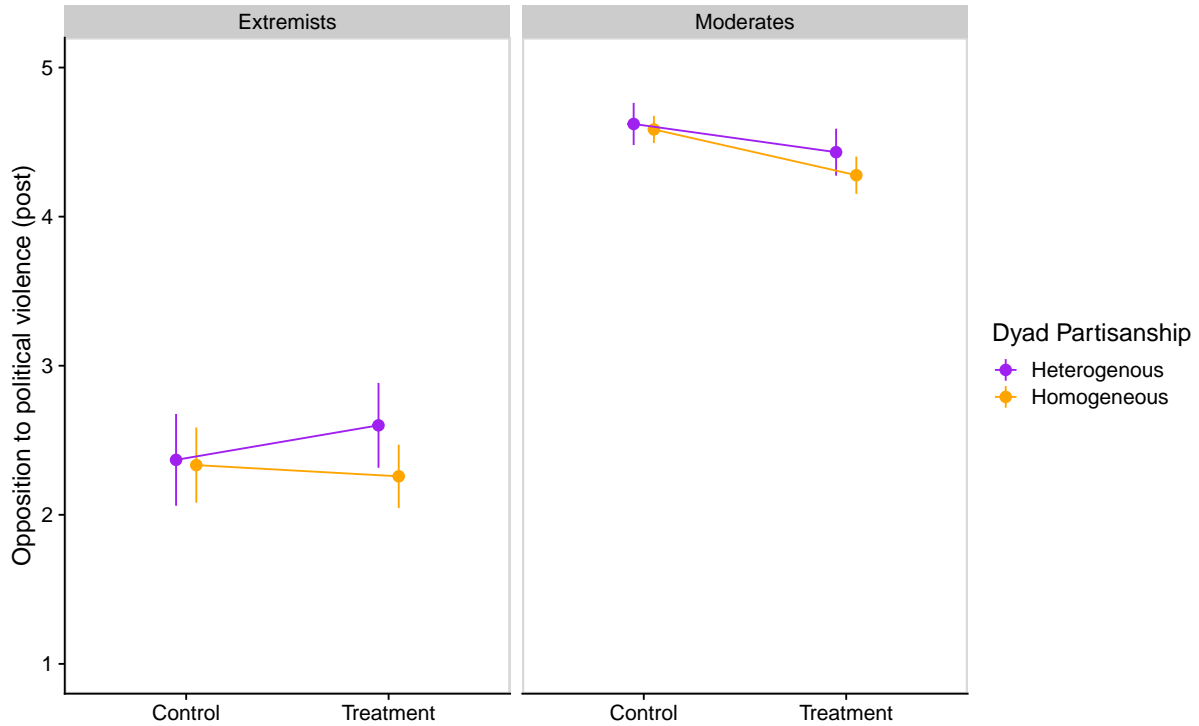
Figure 6: Mapping Changes in Support for Political Violence Post-treatment



Importantly, however, we also find that 94% of moderates did not cross the “midpoint” from opposing to supporting political violence. As we illustrate graphically in Figure 6, moderates in the treatment group were most frequently moving from “strongly opposing” political violence to “somewhat opposing” political violence after discussions with their extremist conversation partner. This makes sense, given that 76% of moderates start out reporting their “strong opposition” of political violence. Importantly, we do see some evidence of extremists moving away from political violence after their conversation, but not in enough numbers to be statistically significant. Moreover, movement tends to be concentrated among those whose support for violence was less firm prior to conversation (e.g those who had indicated they only somewhat supported violence).

To test our final pre-registered hypothesis, that support for political violence would decrease most among extremists in co-partisan dyads (Secondary Hypothesis 2), we use a

Figure 7: Effects of Conversation on Support for Political Violence Does Not Vary by Co- versus Cross-Partisan Dyads



model in which we interact treatment status with an indicator for whether the extremist perceives the moderate to be a copartisan (1) or not (0). As Figure 7 showcases, we do not find any significant interaction effect. Though we did not pre-register this expectation among moderates, we also run the same interaction model for this group, and find similarly null results. Respondents' views on political violence do not differ across homogeneous versus heterogeneous conversation pairings. However, we note that we are underpowered to definitively conclude that this is a true null effect given the high statistical power required for interaction models, particularly given that cross-partisan social dyads are quite rare in our study (as in the world).

8 What did they talk about?

A unique advantage of our research design is that we are also able to directly investigate potential mechanisms, using the full transcripts of the conversations available from Chat-

ter to critically evaluate various potential routes to persuasion. We begin by investigating the tenor of the conversations as a whole, where the unit of analysis is the conversation. We note that any analyses of these transcripts are post-treatment and focus only on those in the treatment group who had a conversation about political violence prior to measuring support for political violence. As a result, these analyses are *not* to be interpreted causally.

We hand-coded each conversation for the presence of a variety of characteristics, including the use of the persuasion tools discussed in the prompt given to moderates: personal experiences with political violence, powerful individual stories, negative effects of political violence, and social norms against violence. We also examined whether specific prominent examples of political violence were brought up, including: January 6, assassination attempts of officials, Israel-Palestine protests, BLM protests, and violence or threats against government employees. Finally, we also coded the tone of the conversation (hostile or not), whether the subjects' relationship was discussed, if values were mentioned, and if subjects' debated the definition/scope of violence.¹³

A central component of our design is that we recruited people who knew each other personally to have these conversations (or not), hypothesizing that their social relationship would allow the moderates to more successfully persuade their extremist social ties to be less supportive of political violence. We found that 24.8 percent of conversations specifically referenced their personal relationship. Sometimes, this was simply by using pet-names (buddy, sweetheart, mom), but other times, participants would reference past experiences they shared together. This gives us confidence that participants were indeed recruiting people they knew personally and that in at least about a quarter of the conversations, the relationship itself was relevant to how they discussed political violence.

We provided participants with a lot of material from which to spark their conversa-

¹³We also used this manual coding as an opportunity to review the quality of the conversations, coding conversations as bots or at least partially using chatGPT (7.3%) or off-topic (3.7%). 13.8% of conversations made some reference to being part of a study. Most of these were comments like "let's look at the questions they gave us in the prompt," but others were more direct like "god, I hope they never read these." None of the comments indicated that participants were aware of the hypotheses associated with the study, but there could still be Hawthorne effects changing respondent behavior because they know they are being observed.

tions, so we needed to examine the extent to which these topics populated the conversations themselves. We found that 31.2 percent of conversations discussed the assassination attempts against President Trump, 28.4 percent discussed January 6, 22.9 percent discussed Black Lives Matter, 13.8 percent discussed violence against government employees, including police, election workers, and military officers, and 11.9 percent discussed protests surrounding the Israel-Palestine conflict. Participants also added new topics we did not supply, such as the Civil and Revolutionary wars, and the Civil Rights Movement in the 1960s.

We also provided moderates with a list of strategies they could try to open these conversations with their extremist partners. We found that 55.1 percent of conversations discussed the effects of political violence, making it the dominant strategy. However, most participants did not use the information we provided and instead discussed effects of violence with respect to the victims of violence. For example, one moderate respondent argued, as follows in their conversation, with the following response from the extremist partner:

Moderate: I get where you're coming from. People are definitely frustrated, especially when it feels like their voices aren't being heard but at the same time, doesn't violence kind of make things worse? Like, it can hurt innocent people and sometimes it just causes more chaos instead of real change.

Extremist: Yeah, but if peaceful methods aren't getting results then what's left? Look at how long people have been fighting for certain rights. Sometimes it feels like the only way to be taken seriously is by pushing back with force.

14.7 percent of conversations included a discussion of personal experience with political violence, most commonly recounting how they felt emotionally when they learned about events like the January 6 insurrection, though some discussed interactions with protests, particularly for those about the Israel-Palestine Conflict. For example, one moderate recounted a bus driver asking them to close the windows due to excessive tear gas during a riot in Pioneer Square in Seattle.¹⁴ 11.9 percent of conversations discussed powerful stories, explaining specific moments of political violence. Although this might

¹⁴ "The protests during covid and the black lives matter was pretty long going and alot of ppl got hurt. Think that was justifiable? [...] Weird stuff like rioters in pioneer square and the bus driver asking us to close the windows due to the tear gas... TEARGAS!! My gosh."

have been the simplest narrative strategy to discuss from our suggestions, we found that social norms were the least commonly used, with only 7.3 percent of conversations discussing the fact that the majority of Americans do not support political violence.¹⁵ While it seems that our prompts were helpful for some conversations, giving moderates a jumping-off point, most discussion partners were able to tailor the conversation around their own social contexts. For example, many moderates tried to offer alternative approaches to political violence, as the below quote exemplifies:

I hear what you're saying. It's frustrating when it feels like our voices aren't being heard. But I still believe that violence is never the answer. There are always other ways to make a difference, even if they seem difficult or time-consuming. We can continue to organize peaceful protests, reach out to our representatives, and build alliances with other groups who share our goals. By working together and staying persistent, we can achieve real change without resorting to violence.

This ability to bring in unique ideas is a key feature of the design, allowing individuals to have unstructured conversations more similar to how they might approach them in reality.

Overall, these conversations were civil, with only 8.3 percent including language or comments reflecting a negative tone. For most conversations, participants expressed their disagreement respectfully and often ended conversations on a positive note, thanking their discussion partner. For example, one respondent ended their chat with “It has been a great conversation, thank you. I agree, more Americans agree either each other than they think and we should come together more.” One way participants discussed their differences in support for political violence was to actually discuss how they define political violence. 34.9 percent of conversations defined political violence. In many of these cases, this is how the moderate seemed to soften their opposition to political violence, coming to consider different contexts under which it could be justified, such as in self-defense or for oppressed groups. For example:

¹⁵Though, one moderate respondent did explicitly use the information provided in our prompt, saying: ““True, there is no doubt the frustration is real, but I’ve read that like over 96% don’t support political violence. It’s not just about hurting others but it can backfire and make things harder in the long run. What do you think about that?”. Their conversation partner responded: “I mean, I get that most people don’t like it but I feel like those people haven’t experienced real oppression. Sometimes you’ve got to break things to get people to actually notice and change things.”

Well, generally, I am of the opinion that political violence is never okay. However, after listening to your points, I think I don't necessarily disagree with you / if people are using political violence only on property, and it is for a cause I support that they have tried to solve peacefully already, I can see how political violence may be the only option / I still would struggle to ever think violent protests that harm another human would be okay though.

Similarly, about 40.4 percent of conversations discussed broader values, such as opposition to violence because it is immoral, regardless of whether it is politically motivated. Though often the conversation turned to situations under which violence would be morally permissible as well. For instance one extremist partner argued: "I'm not saying I'm going out and breaking knees for america. I am just of the mind that america is worth fighting for! Even if it is within our own damn democracy if our values and heritage are being whisked away from us!!!"

We suspected that some support for political violence might be rooted in misinformation and coded for whether conversations included conspiracy theories. We found that 12.8 percent of conversations referenced conspiracy theories. Most often, participants alluded to the assassination attempts against President Trump being staged by conservatives to boost support by making him a martyr, or liberals affiliated with ANTIFA instigating violence during the January 6 insurrection at the Capitol to frame Republicans. Instigators at protest events (both left- and right-leaning) were also sometimes cited as potential "paid agitators" to make protestors look bad.

Conversations rarely referenced racial animus (6.4 percent), but partisan animus was much more common (26.6 percent). Racial animus was typically expressed within conversations about Black Lives Matter protests, while negative comments about partisans came up in conversations about several different topics, most commonly the assassination attempts at President Trump and the January 6 insurrection.

Together, analyzing the transcripts reveals that our prompts were able to guide the conversations about half of the time, but participants otherwise found their own approaches to discussing political violence with their peer. The conversations were overwhelmingly respectful and tended to focus more on how to define political violence and the conditions under which it is acceptable. This might explain some of the treatment

effects we observed: it could be the case that moderates broadened their scope of what they consider to be political violence, leading them to soften their opposition to it.

9 Discussion and Conclusion

This project addresses a question of critical importance: how can known social ties help reduce support for political violence among their family and friends? As support for political violence – and acts of political violence – have spiked in recent years, there is growing concern that social acceptance around the use of political violence is contributing to its potential normalization, creating an environment that fails to sanction—and even supports—the few people who choose to act on these beliefs. As more individuals within a network endorse acts of political violence, those who oppose such actions may feel their voice diminish, potentially leading to self-censorship when other network members support political violence. With increasing tacit approval of political violence, a new norm favoring it may emerge. Support for political violence is an important phenomenon to explain and thwart. Even if supporters would never commit violent acts themselves, it is dangerous in and of itself to normalize violence for the few who then commit these acts.

However, this normalization process may be reversed if trusted individuals from real-world personal networks can effectively leverage their source credibility and exert ingroup social pressures to reinforce norms of non-violence and dissuade network ties away from political violence. Indeed, as [Kleinfeld \(2021\)](#) writes, “...people committing far-right violence — particularly planned violence rather than spontaneous hate crimes — are older and more established than typical terrorists and violent criminals. They often hold jobs, are married, and have children. Those who attend church or belong to community groups are more likely to hold violent, conspiratorial beliefs ([Pape et al., 2022](#)). These are not isolated ‘lone wolves’; they are part of a broad community that echoes their ideas” (p. 161). Exposing these individuals to members of their community who *disagree* with them about political violence may thus be a crucial intervention for changing views.

Our project thus offers four primary contributions. First, we explore the impact of discussions within real social networks on attitudes towards political violence. In doing so, we expand upon existing literature concerning support for political violence in the United States by incorporating how social relationships can abate political violence. Tested interventions to reduce political violence are often non-social, like correcting meta-perceptions with information (e.g., [Mernyk et al., 2022](#)). However, in the mega-study on interventions to reduce outparty animosity, support for undemocratic norms, and support for partisan political violence, the interventions that reduced support for partisan political violence the *most* involved observing an indirect or mock social experience ([Voelkel et al., 2022](#)). For example, the best performing intervention to reduce support for partisan violence corrected outparty misperceptions, but did so by showing a video of other people thinking through their answers and their reactions to learning they overestimated outpartisan responses ([Voelkel et al., 2022](#)). Our experiment advances this literature by estimating the effects of an intervention that is *social* and thus is likely *stronger* and more durable, given the pattern of evidence suggesting social interventions are stronger than non-social interventions ([Bond et al., 2012](#); [Gerber et al., 2008](#); [Voelkel et al., 2022](#)). The intervention we test is likely more *scalable* as well, as it uses existing network dyads and conversations that happen within them ([Edsall, 2024](#)).

A second key contribution is our focus on existing relationships. For example, research by [Rossiter and Carlson \(2024a\)](#) has shown the efficacy of interpersonal conversation in changing important attitudes such as affective polarization, however did not examine whether preexisting social ties enhance this effect. Similarly, research by [Wayne \(2022; 2024\)](#) has demonstrated the impact of group discussion on how individuals make political choices, but not with individuals with real-world network ties, a key limitation. While our main focus is on network dyads as an avenue to reduce political violence, we will also shed light on how relationships are being affected by dissent on democracy-threatening beliefs. Some people feel as though they are losing family members to extreme, usually conspiratorial, political beliefs. It appears many people in this situation wish to discuss their concerns with their family member in an effort to combat the harmful beliefs and the

potential that those beliefs continue to spiral into other harmful attitudes and behaviors (e.g., [Faye, 2023](#); [Muncaster, 2022](#)). Our results, coming from real network ties, will shed light on one strategy concerned family members may consider taking in these situations.

A final contribution of this research is to understand the role concerned citizens can play in thwarting rising levels of political violence and strengthening their democracy. Concerned citizens can vote for elected officials who are against political violence and uphold democratic norms ([Graham and Svolik, 2020](#)), citizens can continue to pressure elected official once in office, and citizens can also support the free press who hold the powerful to account. In this project, we look beyond these larger-scale actions of voting and advocacy, and we shed light on a more grassroots approach of citizens playing a role within their own networks to address the issue of growing support for political violence.

References

- Baker, Andy, Barry Ames, and Lúcio Rennó. 2020. *Persuasive peers: social communication and voting in Latin America*. Princeton University Press.
- Berger, John M. 2021. *Extremismo: La serie de conocimientos esenciales de MIT Press*. Ediciones UC.
- Bloom, Mia, Hicham Tiflati, and John Horgan. 2019. “Navigating ISIS’s preferred platform: Telegram1.” *Terrorism and Political Violence* 31(6): 1242–1254.
- Bond, Robert M, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle, and James H Fowler. 2012. “A 61-million-person experiment in social influence and political mobilization.” *Nature* 489(7415): 295–298.
- Broockman, David, and Joshua Kalla. 2016. “Durably reducing transphobia: A field experiment on door-to-door canvassing.” *Science* 352(6282): 220–224.
- Carey, John M., Gretchen Helmke, Brendan Nyhan, and Susan C. Stokes. 2021. “Tempered expectations and hardened divisions a year into the Biden presidency.” *Brightline Watch* <https://brightlinewatch.org/tempered-expectations-and-hardened-divisions-a-year-into-the-biden-presidency/>.
- Carlson, Taylor N. 2019. “Through the grapevine: Informational consequences of interpersonal political communication.” *American Political Science Review* 113(2): 325–339.
- Carlson, Taylor N. 2024. *Through the grapevine: Socially transmitted information and distorted democracy*. University of Chicago Press.
- Clifford, Scott, Geoffrey Sheagley, and Spencer Piston. 2021. “Increasing precision without altering treatment effects: Repeated measures designs in survey experiments.” *American Political Science Review* 115(3): 1048–1065.
- Conover, Pamela Johnston, Donald D Searing, and Ivor M Crewe. 2002. “The deliberative potential of political discussion.” *British journal of political science* 32(1): 21–62.

- Del Vicario, Michela, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. 2016. "The spreading of misinformation online." *Proceedings of the national academy of Sciences* 113(3): 554–559.
- Della Porta, Donatella. 1988. "Recruitment processes in clandestine political organizations." *International social movement research* 1: 155–169.
- Della Porta, Donatella. 2018. "Radicalization: A relational perspective." *Annual Review of Political Science* 21(1): 461–474.
- Della Porta, Donatella Alessandra. 1995. "Social movements, political violence and the state."
- Edlin, Ruby, Turquoise Baker. 2022. "Poll of Local Election Officials Finds Safety Fears for Colleagues — and Themselves." *Brennan Center for Justice* <https://www.brennancenter.org/our-work/analysis-opinion/poll-local-election-officials-finds-safety-fears-colleagues-and>.
- Edsall, Thomas. 2024. "Why Losing Political Power Now Feels Like 'Losing Your Country'." <https://www.nytimes.com/2024/04/24/opinion/partisan-polarization-republicans-democrats.html?smid=tw-share>.
- Eveland Jr, William P, and Myiah Hutchens Hively. 2009. "Political discussion frequency, network size, and "heterogeneity" of discussion as predictors of political knowledge and participation." *Journal of communication* 59(2): 205–224.
- Eveland Jr, William P, Myiah J Hutchens, and Alyssa C Morey. 2012. "Social networks and political knowledge." *The SAGE handbook of political communication* pp. 241–252.
- Faye, Shon. 2023. "I'm Worried About Losing Touch With My Conspiracy Theorist Parents." <https://www.vogue.com/article/conspiracy-theorist-parents>.
- Gaertner, Samuel L, John F Dovidio, Rita Guerra, Eric Hehman, and Tamar Saguy. 2015.

- “A common ingroup identity: Categorization, identity, and intergroup relations.” In *Handbook of prejudice, stereotyping, and discrimination*. Psychology Press pp. 433–454.
- Gerber, Alan S, Donald P Green, and Christopher W Larimer. 2008. “Social pressure and voter turnout: Evidence from a large-scale field experiment.” *American political Science review* 102(1): 33–48.
- Graham, Matthew H, and Milan W Svobik. 2020. “Democracy in America? Partisanship, polarization, and the robustness of support for democracy in the United States.” *American Political Science Review* 114(2): 392–409.
- Holliday, Derek E., Shanto Iyengar, Yphtach Lelkes, and Sean J. Westwood. 2024. “Meaningfully Reducing Partisan Animosity at Scale: Evaluating the Scope Conditions of Bridging Interventions for the Reduction of Polarization.” *Working Paper* <https://www.dropbox.com/scl/fi/ho16v4l6ojnlpi3tu3gux/HollidayIyengarLelkesWestwood2024bridge?rlkey=52y6qqbk53tgmtvs5a3gqvwh0&e=1&dl=0>.
- Huckfeldt, Robert, and John Sprague. 1995. “Citizens, politics, and social communication: Information and influence in an election campaign.” *Cambridge University Press*.
- Iyengar, Shanto, Yphtach Lelkes, and Sean Westwood. 2024. “America’s Political Pulse April 2024 Dashboard.” <https://polarizationresearchlab.org/americas-political-pulse/>.
- Jasko, Katarzyna, Gary LaFree, and Arie Kruglanski. 2017. “Quest for significance and violent extremism: The case of domestic radicalization.” *Political psychology* 38(5): 815–831.
- Jones, Sam. 2022. “Fact Sheet: Updated Armed Demonstration Data Released A Year After the 6 January Insurrection Show New Trends.” *ACLEDD* <https://acleddata.com/2022/01/05/updated-armed-demonstration-data-released-a-year-after-the-6-january-insurrection->

- Kalla, Joshua L, and David E Broockman. 2020. "Reducing exclusionary attitudes through interpersonal conversation: Evidence from three field experiments." *American Political Science Review* 114(2): 410–425.
- Kalmoe, Nathan P, and Lilliana Mason. 2022a. "A holistic view of conditional American support for political violence." *Proceedings of the National Academy of Sciences* 119(32): e2207237119.
- Kalmoe, Nathan P, and Lilliana Mason. 2022b. *Radical American partisanship: Mapping violent hostility, its causes, and the consequences for democracy*. University of Chicago Press.
- Kalmoe, Nathan P, and Lilliana Mason. 2024. "Threats as Political Communication." *Political Communication* 41(1): 162–171.
- Kay, Cameron Stuart. 2024. "Extraverted introverts, cautious risk-takers, and selfless narcissists: A demonstration of why you can't trust data collected on MTurk."
- Kenny, Christopher B. 1992. "Political participation and effects from the social environment." *American Journal of Political Science* pp. 259–267.
- Kertzer, Joshua D, Marcus Holmes, Brad L LeVeck, and Carly Wayne. 2022. "Hawkish biases and group decision making." *International Organization* 76(3): 513–548.
- Klar, Samara. 2014. "Partisanship in a social setting." *American journal of political science* 58(3): 687–704.
- Kleinfeld, Rachel. 2021. "The rise of political violence in the United States." *J. Democracy* 32: 160.
- Klofstad, Casey A. 2007. "Talk leads to recruitment: How discussions about politics and current events increase civic participation." *Political Research Quarterly* 60(2): 180–191.
- Krupnikov, Yanna, and John Barry Ryan. 2022. *The other divide*. Cambridge University Press.

- Lazarsfeld, Paul F, Bernard Berelson, and Hazel Gaudet. 1968. *The people's choice: How the voter makes up his mind in a presidential campaign*. Columbia University Press.
- Levendusky, Matthew. 2023. *Our common bonds: Using what Americans share to help bridge the partisan divide*. University of Chicago Press.
- Lewandowsky, Stephan, Ullrich KH Ecker, Colleen M Seifert, Norbert Schwarz, and John Cook. 2012. "Misinformation and its correction: Continued influence and successful debiasing." *Psychological science in the public interest* 13(3): 106–131.
- McClurg, Scott D. 2003. "Social networks and political participation: The role of social interaction in explaining political participation." *Political research quarterly* 56(4): 449–464.
- McClurg, Scott D. 2006. "The electoral relevance of political talk: Examining disagreement and expertise effects in social networks on political participation." *American Journal of Political Science* 50(3): 737–754.
- Mernyk, Joseph S, Sophia L Pink, James N Druckman, and Robb Willer. 2022. "Correcting inaccurate metaperceptions reduces Americans' support for partisan violence." *Proceedings of the National Academy of Sciences* 119(16): e2116851119.
- Minozzi, William, Hyunjin Song, David MJ Lazer, Michael A Neblo, and Katherine Ognyanova. 2020. "The incidental pundit: Who talks politics with whom, and why?" *American Journal of Political Science* 64(1): 135–151.
- Mitts, Tamar. 2022. "Countering violent extremism and radical rhetoric." *International Organization* 76(1): 251–272.
- Muncaster, Sue. 2022. "My Gentle, Intelligent Brother Is Now A Conspiracy Theorist And His Beliefs Are Shocking." <https://www.huffpost.com/entry/brother-conspiracy-theory>.
- Mutz, Diana C. 2006. "Hearing the other side: Deliberative versus participatory democracy." *Cambridge University*.

- Mutz, Diana Carole, Paul M Sniderman, and Richard A Brody. 1996. *Political persuasion and attitude change*. University of Michigan Press.
- Nickerson, David W. 2008. “Is voting contagious? Evidence from two field experiments.” *American political Science review* 102(1): 49–57.
- Nivette, Amy, Manuel Eisner, and Denis Ribeaud. 2017. “Developmental predictors of violent extremist attitudes: A test of general strain theory.” *Journal of research in crime and delinquency* 54(6): 755–790.
- Nyhan, Brendan, and Jason Reifler. 2012. “Misinformation and Fact-checking.” *Research Findings* .
- Pape, Robert, the Chicago Project on Security, and Threats. 2022. “UNDERSTANDING AMERICAN DOMESTIC TERRORISM: Mobilization Potential and Risk Factors of a New Threat Trajectory.” https://d3qi0qp55mx5f5.cloudfront.net/cpost/i/docs/americas_insurrectionists_online_2021_04_06.pdf?m.
- Passy, Florence. 2003. “Social networks matter. But how?” *Social movements and networks. Relational approaches to collective action* pp. 21–48.
- Petty, Richard E, John T Cacioppo, and Jeff A Kasmer. 2015. “The role of affect in the elaboration likelihood model of persuasion.” In *Communication, social cognition, and affect (PLE: Emotion)*. Psychology Press pp. 117–146.
- Piazza, James A, and Ahmet Guler. 2021. “The online caliphate: Internet usage and ISIS support in the Arab world.” *Terrorism and political violence* 33(6): 1256–1275.
- Rossiter, Erin L. 2022. “Measuring agenda setting in interactive political communication.” *American Journal of Political Science* 66(2): 337–351.
- Rossiter, Erin L, and Taylor N Carlson. 2024a. “Cross-partisan conversation reduced affective polarization for republicans and democrats even after the contentious 2020 election.” *The Journal of Politics* 86(4): 000–000.

- Rossiter, Erin L., and Taylor N. Carlson. 2024b. “To Discuss or Not to Discuss? How Selective Exposure to Political Discussion Conditions Experimental Findings on Polarization.” *Working Paper* .
- Sageman, Marc. 2004. “Understanding terror networks.” *University of Pennsylvania* .
- Scheff, Thomas J. 1988. “Shame and conformity: The deference-emotion system.” *American sociological review* pp. 395–406.
- Sinclair, Betsy. 2012. *The social citizen: Peer networks and political behavior*. University of Chicago Press.
- Suhay, Elizabeth. 2015. “Explaining group influence: The role of identity and emotion in political conformity and polarization.” *Political Behavior* 37: 221–251.
- Taber, Charles S, and Milton Lodge. 2006. “Motivated skepticism in the evaluation of political beliefs.” *American journal of political science* 50(3): 755–769.
- Voelkel, Jan G, Michael Stagnaro, James Chu, Sophia Pink, Joseph Mernyk, C Redekopp, M Cashman, Q Submitters, J Druckman, D Rand et al. 2022. “Megastudy identifying successful interventions to strengthen Americans’ democratic attitudes.” *Northwestern University: Evanston, IL, USA* .
- Wayne, Carly, Mitsuru Mukaigawara, Joshua D Kertzer, and Marcus Holmes. 2024. “Diplomacy by committee: Assessing resolve and costly signals in group settings.” *American Journal of Political Science* .
- Westwood, Sean, and Yphtach Lelkes. 2024. “America’s Political Pulse.” <https://polarizationresearchlab.org/americas-political-pulse/>.
- Westwood, Sean J, Justin Grimmer, Matthew Tyler, and Clayton Nall. 2022a. “Current research overstates American support for political violence.” *Proceedings of the National Academy of Sciences* 119(12): e2116870119.

- Westwood, Sean J, Justin Grimmer, Matthew Tyler, and Clayton Nall. 2022*b*. “Reply to Kalmoe and Mason: The pitfalls of using surveys to measure low-prevalence attitudes and behavior.” *Proceedings of the National Academy of Sciences* 119(32): e2207584119.
- Winston, Adrienne, Rashard Rose. 2024. “US Capitol Police investigated more than 8,000 threats against lawmakers in 2023, up from the previous year.” *CNN* <https://www.cnn.com/2024/01/19/politics/uscp-threat-assessment-lawmakers-2023/index.html>.
- Wintemute, Garen J, Sonia L Robinson, Andrew Crawford, Daniel Tancredi, Julia P Schleimer, Elizabeth A Tomsich, Paul M Reeping, Aaron B Shev, and Veronica A Pear. 2023. “Views of democracy and society and support for political violence in the USA: findings from a nationally representative survey.” *Injury epidemiology* 10(1): 45.
- Ying, Luwei. 2024. “Military Power and Ideological Appeals of Religious Extremists.” *Journal of Politics* .
- Zuckerman, Alan S. 2005. *The social logic of politics: Personal networks as contexts for political behavior*. Temple University Press.