

Original Research

The Political Influence of Non-Politicized Users: A Survey Experiment on Engagement Intentions toward Protest Information on Twitter

Sangyeon Kim^{id}

Department of
Communication & Media,
Ewha Womans University

Corresponding to

Sangyeon Kim

Department of
Communication & Media,
Ewha Womans University,
Seoul, Republic of Korea.

Email: sangyeonkim@ewha.ac.kr

Disclosure Statement

No potential conflict of interest
was reported by the authors.

Received

28 Dec 2025

Revised

8 Apr 2026

Accepted

14 Apr 2026

ABSTRACT

How do the political characteristics of Twitter accounts shape users' engagement intentions toward protest information? This article argues that the political characteristics of accounts sharing protest information affect how that information is interpreted and engaged with by other Twitter users. Specifically, I suggest that whether Twitter accounts are perceived as overtly political or nonpolitical can shape how users respond to signals about political protests. I hypothesize that nonpolitical accounts may generate higher engagement intentions toward protest messages than political accounts, as they are seen as less biased or more trustworthy. To test this theory, I conducted an online experiment using vignettes that simulate the Twitter environment. Participants were exposed to protest-related Tweets and were asked whether they would retweet or like them, with some accounts presenting political traits in their profiles and others appearing nonpolitical. Contrary to my expectations, the results did not reveal a statistically significant difference in participants' responses between political and nonpolitical profiles. However, the study revealed unexpected patterns, including the role of education in shaping retweet behavior differently across political groups and the influence of context-specific factors, such as protest types and images, on user engagement. These findings suggest that individual characteristics and content features may interact in complex ways, warranting further exploration.

KEYWORDS

social media, social network, political communication, network effects, survey experiment

People obtain information from various media sources, including traditional outlets like television, radio, and newspapers, as well as newer platforms like Twitter and Facebook. How individuals process information from these external sources has long been a focus of media effects research (Albertson & Lawrence, 2009; Arceneaux & Johnson,

2013; Cohen, 1963; Feldman, 2011; Holbert et al., 2010; Levendusky, 2013; Martin & Yurukoglu, 2017; McCombs & Shaw, 1972; S. Iyengar, 1991; Stroud, 2011). As social media increasingly becomes central to news consumption, understanding how users perceive and engage with information within these platforms has drawn considerable scholarly attention (Feezell, 2018; Munger, 2020). This shift is particularly important for understanding engagement with protest information on platforms like Twitter, which has been recognized for its role in shaping political discourse by both network and social movement scholars (Barberá et al., 2015; Steinert-Threlkeld, 2017).

Unlike traditional media, social media platforms come with built-in social networks, connecting users from the moment they join. These networks embed social cues, such as “likes” and “retweets,” within shared content, including political information (Munger, 2020). To understand engagement with protest information on social media, it is essential to account for the interdependence among users. Research on collective action highlights the importance of network structures—such as size (Centola, 2013), density (Macy, 1990), degree distribution (Watts, 2002), centralization (Marwell et al., 1988), and weak ties (Centola and Macy, 2007)—in influencing how this interdependence shapes the spread of protest-related information (Bakshy et al., 2012).

This article examines how node-level characteristics in social media networks—specifically, whether Twitter accounts are perceived as political—affect how individuals interpret and share protest information. I argue that political traits of accounts shape inter-user influence in Twitter protest hashtag networks. Users may be more likely to share protest content from accounts they perceive as nonpolitical, viewing them as less biased or more credible compared to overtly political accounts, which could make nonpolitical accounts more effec-

tive in generating engagement intentions toward protest content. To evaluate this hypothesis, I conducted an online experiment using simulated Twitter profiles with varying political traits. Contrary to expectations, the results showed no statistically significant differences in retweet or like behavior between participants exposed to political and nonpolitical profiles. Nonetheless, the study revealed unexpected patterns that merit further investigation.

Two unregistered exploratory patterns emerged from the data, reported here strictly as hypothesis-generating. First, education was associated with retweet intentions in opposite directions across partisan groups. Among liberals, higher education was associated with higher retweet intentions, consistent with existing literature (Larreguy & Marshall, 2017; Le & Nguyen, 2021; Perrin & Gillis, 2019; Sunshine Hillygus, 2005), whereas among conservatives the association was negative. This pattern runs counter to established findings and warrants replication before firm conclusions can be drawn.

Second, an exploratory analysis using protest-type-specific dependent variables found that in the liberal group, participants who viewed a gun control tweet paired with a nonpolitical profile reported higher retweet intentions than those who viewed it with a political profile. This directional pattern is consistent with the hypothesis but, given its unregistered status, should be understood as motivating future pre-registered research rather than as confirming the hypothesis under specific conditions.

The Role of Social Networks on Engagement with Protest Information in Social Media

Research indicates that social media users often encounter political content incidentally while seeking entertainment (Feezell, 2018; Lin and Lu, 2011; Quan-Haase & Young, 2010). Munger (2020) introduces the idea of a “credibility

cascade,” where social signals like “Likes” and “Retweets” enhance the credibility of news initially perceived as “fake,” especially among less digitally literate users. On platforms like Twitter and Facebook, diffusion mechanisms like retweets and shares help spread information. Opinion leaders—central, knowledgeable individuals—play a key role in amplifying content and influencing user engagement (Park, 2013; R. Iyengar et al., 2011; Schäfer & Taddicken, 2015; Van Eck et al., 2011).

Other studies highlight the importance of individuals’ perceptions in the diffusion of information within social networks. Content characteristics significantly influence perceptions and sharing likelihood, with Romero et al. (2011) noting that different hashtags show varied diffusion patterns on Twitter. Peer effects also matter, as users are more likely to share content shared by friends, especially among “weak ties” (Bakshy et al., 2012).

In contentious politics and collective action, the network position of social media accounts is essential for understanding the spread of protest information (Barberá et al., 2015) and user participation in protests (Steinert-Threlkeld, 2017). Key network components, such as size (Centola, 2013), density (Macy, 1990), degree distribution (Watts, 2002), and weak ties (Centola & Macy, 2007), influence protest evolution. The literature highlights the interdependence of social media users, recognizing both content-level characteristics (e.g., hashtags) and structural elements (e.g., tie types) in shaping perceptions and diffusion of political information.

Building on this literature, I aim to examine how a *node-level* characteristic—specifically, the political traits of users’ neighbors—affects their perceptions of the information they encounter and their subsequent engagement decisions. Bakshy et al. (2012) notes the importance of node connections; weak ties can enhance peer effects in information diffusion, consistent with Network Science classics (Granovetter, 1983). However,

I argue that *characteristics of nodes* should also be considered, as they influence the credibility of the signals regarding political events conveyed to other nodes.

This article focuses on the political characteristics of nodes within networks, specifically investigating how Twitter users respond to protest information shared by accounts with varying levels of political traits in their profiles. The political traits in Twitter profiles can indirectly signal the credibility of the protest information shared, affecting how other users perceive and choose to share that content. I will theorize how these political characteristics condition the effectiveness of user accounts in eliciting engagement intentions toward protest content among other users.

THEORY AND HYPOTHESIS

As demonstrated in the previous section, network structures are vital in shaping the dynamics of information diffusion in social media. However, the question of which nodes in a political network hold the most influence remains underexplored. While previous research identifies “opinion leaders” as central figures in information contagion (Meng et al., 2011; Park, 2013; R. Iyengar et al., 2011; Schäfer & Taddicken, 2015; Van Eck et al., 2011), it is less clear if they hold similar power in Twitter networks during protest events. I argue that when it comes to generating engagement intentions toward protest content, politicized accounts (often seen as opinion leaders) are less effective than non-politicized accounts. This is due to the unique dynamics of political signaling on Twitter, where signals from non-politicized users are perceived as more credible and inclusive.

Twitter accounts involved in protest-related discussions can generally be classified as either politicized or non-politicized, based on their profile information. Politicized accounts explicitly

express interest in political topics or reveal their partisanship, while non-politicized accounts do not feature such information. In the context of Twitter protest networks, politicized accounts can be viewed as opinion leaders on political issues because of their visible alignment with political matters. Users may regard these accounts as experts or knowledgeable about political protests, but I argue that their influence is more limited than suggested by existing literature.

In contrast, non-politicized accounts, which do not typically engage in political discourse, are perceived as more credible when they post about protests. This perception arises from the “informational cascade” framework used in the study of contentious politics, which explains how mass protest participation can grow through social observation (Kuran, 1991; Lohmann, 1994). When non-political accounts share protest-related content, it signals to others that the protest is relevant beyond the politically active and could garner broad support. This broad appeal is critical to encouraging widespread engagement and participation.

In the study of social media, Munger (2020) adapts the concept of “informational cascade” to what he terms a “credibility cascade.” Credibility cascade refers to the process by which internet news from platforms often labeled as “clickbait” can earn credibility through social validation. Social signals like “Likes” and “Retweets” encode recommendations from users, which lend legitimacy to stories initially perceived as dubious or fake. As this process repeats, the perceived credibility of the information increases, despite its origin (Munger, 2020).

Building on Munger (2020)’s concept of a “credibility cascade,” I argue that Tweets about protests from non-politicized accounts are more likely to generate engagement intentions because these accounts, which do not usually signal political engagement, lend greater credibility to the event. Conversely, signals from politicized accounts are likely to be discounted, as they may

make protests seem like niche, partisan efforts rather than inclusive movements. This resonates with prior studies that find most social media users are not primarily engaged with political content (Feezell, 2018; Lin and Lu, 2011; Quan-Haase & Young, 2010; Seattle, 2018). For these users, content from politicized accounts may seem irrelevant or even off-putting, reducing the likelihood of further engagement.

Additionally, politicized accounts may be seen as less credible by users who avoid partisan content due to perceived bias (Arceneaux & Johnson, 2013). Tweets from politicized accounts can lead others to perceive the protest as a partisan effort, which diminishes the likelihood of participation or further engagement. In contrast, when non-politicized accounts, which typically avoid political topics, share protest-related content, it signals that the protest has significance beyond the usual politically engaged audiences. These accounts have a more neutral image, and their involvement in political discourse stands out.

Moreover, the decision for a non-politicized account to post about a protest carries greater perceived weight because it involves potential risks. Non-politicized users might face backlash from their followers—who may expect non-political content—if they suddenly engage in political discussions. This social risk elevates the credibility of their messages, as followers interpret the act of posting about a protest as a deliberate and significant choice, further enhancing the content’s importance and perceived relevance to other users. Given this, I hypothesize the following:

Hypothesis. Users report lower retweet and like intentions toward protest-related Tweets when the Tweets are published by politicized accounts compared to non-politicized accounts.

Because this hypothesis is tested using a vignette experiment that elicits self-reported behavioral

responses, the dependent variables, retweet and like intentions, serve as proxies for engagement behavior rather than measures of actual diffusion in a networked environment.

To test this hypothesis, I control for the influence of partisanship, as literature on echo chambers shows that political information diffusion is often shaped by users' partisan alignment (Cinelli et al., 2021; Garimella et al., 2018; Gillani et al., 2018). While it would be ideal to simultaneously test both partisanship and account characteristics, practical constraints make it more feasible to focus on account characteristics. Therefore, I control for partisanship by grouping respondents according to their self-reported political affiliation, as further detailed in the methods section.

RESEARCH DESIGN

I recruited 2,509 U.S. respondents through Amazon Mechanical Turk (MTurk) on May 4th and 5th, 2023. As discussed further in the Results section, the sample proved to be disproportionately composed of politically active individuals, with 88% reporting protest participation in the past 12 months, a figure that far exceeds the 9% reported in the American National Election Studies (American National Election Studies, 2020). This has important implications for the generalizability of the findings, which are addressed in a dedicated validity discussion below. Table 1 outlines how participants were randomly assigned to four treatment groups and two control groups, based on their responses to partisanship questions. As mentioned earlier, in addition to the political profile treatment—the main theoretical focus of this paper—I introduced a partisanship

treatment, given the strong correlation between one's partisanship and support for protest events, especially when the protests have a partisan nature. For instance, it is highly unlikely that Republicans would support protests advocating for abortion rights, given the party's stance on the issue. Thus, participants were first assigned to either the liberal or conservative group, which matched protest events to their partisanship, based on their answers to the American National Election Studies (ANES) partisanship questions.¹ After this, participants were assigned to one of three groups: the Political group, the Non-political group, or the control group. The assignment primarily reflected the political characteristics of Twitter accounts, determined by whether or not the profile contained political information. All groups were shown Tweets related to political protests, but the distinction between groups lay in the profile page displayed alongside the Tweets. The Political group was exposed to Tweets from profiles that contained political traits, while the Non-political group was shown protest-related Tweets from non-political profiles. Control groups were shown the Tweets without any account profile information.

After providing informed consent, participants answered questions about their basic personal information and partisanship. They were then shown four Tweets, along with profile information, and asked whether they would like to retweet or like the Tweets they read. The profile information varied depending on their experimental group assignment. Participants in the Political group were shown profiles that contained political traits, while those in the Non-political group were shown profiles that did not include any political characteristics. Both groups were presented with four Tweets related to protest

¹ T Respondents who did not identify with either the Democratic or Republican parties could label themselves as "Independent" or members of "Some other party." In such cases, a follow-up question asked about their political leaning: "How would you describe your political views?" with options ranging from extremely liberal to extremely conservative. For instance, if a respondent identified as "Independent" but chose "slightly liberal" for their political leaning, they would be assigned to the liberal treatment group.

Table 1. *Distribution of Treatment Assignment*

	Partisanship Treatment			
	Liberal	Conservative	Sum	
Political Profile Treatment	Political	374	457	831
	Nonpolitical	375	458	833
	Control	382	463	845
	Sum	1,131	1,378	2,509

Note. Distribution of treatment assignment. Participants are first assigned to either the Liberal group or the Conservative group based on their self-reported partisanship. Then, the participants were randomly assigned to one of three groups – the Political group, the Nonpolitical group, or the Control group.

events. Importantly, the protest events were aligned with each participant's partisanship, as they were previously assigned to either the Liberal group or the Conservative group based on their responses to partisanship questions.

Treatments

To simulate exposure to protest information in a Twitter-like environment, I fabricated Tweets and account profiles for the four groups. I generated protest-related Tweets on two issues: abortion rights and gun control. These topics were chosen for two main reasons: they were receiving significant attention from both political parties at the time, and this attention had translated into actual protests and counter-protests. As of January 9, 2023, abortion rights-related protests were widespread, with a dozen states having implemented near-total abortion bans, allowing only minimal exceptions. This galvanized both sides: Democrats, who largely sought to protect abortion rights, and Republicans, who mostly supported the new bans. Similarly, gun control discussions were highly active due to a series of tragic incidents, including the Michigan State

University and Nashville elementary school shootings. Protests calling for stricter gun control erupted across the United States, while gun rights supporters were sharply critical of the movement. Therefore, I created two Tweets about abortion rights protests and two Tweets about gun control protests, reflecting the perspectives and positions of both Democrats and Republicans. To ensure a realistic experience of exposure to protest-related content on Twitter, the Tweets were modeled after actual protest-related posts and designed to resemble current events.

The Twitter profiles used in the experiment were carefully fabricated to ensure consistency and comparability among the simulated accounts. First, all profile images featured human photos, even though not all Twitter users use human photos as their profile pictures. I used AI-generated human images to avoid the potential issue of participants recognizing the individuals in the photos.² The images were evenly divided by race and gender to reflect the diversity of the user pool and to minimize any potential influence of race or gender on participants' reactions. Specifically, I included one white man, one white woman, one black woman, and one

² The AI-generated images were created using the site <https://this-person-does-not-exist.com/en>. I generated four images by selecting specific options. For a white male, I selected "Male, No selection for age, White." For a white female, I selected "Female, No selection for age, White." For a black female, I selected "Female, No selection for age, Black." For a Hispanic male, I selected "Male, No selection for age, Latino Hispanic." I refreshed the page multiple times until satisfactory images appeared, as some images were not convincingly human.

Hispanic/Latino man. Each account name was also carefully selected to match the ethnicity of the profile image. The number of followers and following accounts for each profile was randomly drawn from a range of 500 to 2,000, ensuring the profiles appeared realistic while avoiding extreme follower counts that might influence participants. Importantly, within each experimental group, all profiles varied only in their political or nonpolitical traits, while the overall makeup of the profiles remained consistent across groups.

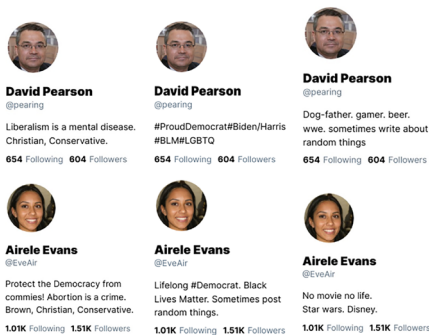
Figure 1 presents examples of the simulated profiles. The first column shows right-wing political profiles, the second column displays liberal profiles, and the final column features nonpolitical profiles with information unrelated to politics. It is important to note that all profile characteristics—such as race, gender, and follower count—were identical across the different types of profiles, except for the biography section, which highlighted each account’s political orientation or personal details. The political profiles were modeled after actual Twitter accounts that posted about the protest events examined in this study.

One potential concern with the conservative

political profiles is that they contain language that could be construed as hostile or uncivil—for instance, phrases such as “Liberalism is a mental disease”—which could confound political identity signaling with incivility. Ideally, direct measures of perceived incivility would have been collected to isolate these two elements. However, no such items were included in the survey. Instead, I draw on a feature of the design that provides indirect evidence. The degree of hostility varies across the stimuli: the conservative political profiles and tweets associated with the abortion rights scenario contain comparatively more aggressive language than those used in the gun control scenario (see Appendix for the full set of profile images and tweets). If the incivility present in some stimuli were a meaningful confound driving engagement responses independently of the political-versus-nonpolitical distinction, one would expect the treatment effect patterns to differ between the abortion and gun control conditions. The disaggregated regression results in Figures A1-A4 show no such differential pattern as the results are consistent across both protest-type conditions. I therefore treat this as indirect evidence that incivility did not serve as a meaningful confound, while acknowledging that future work should construct stimuli that vary political signaling and incivility independently to allow a cleaner test.

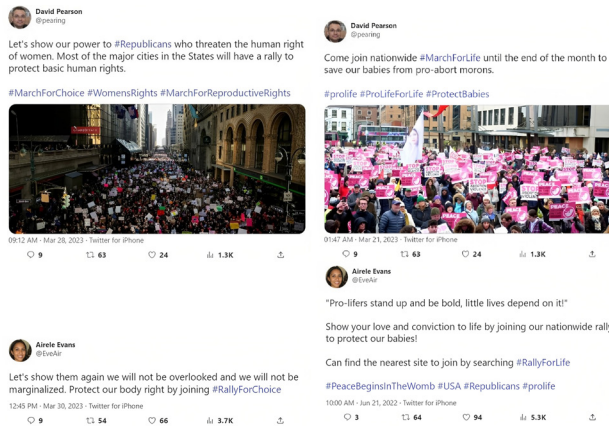
The Tweets used in the study were developed based on actual posts related to protest events concerning abortion rights and gun control. I collected approximately 40 Tweets discussing these events and crafted simulated Tweets for use in the vignette. The images included in the simulated Tweets were sourced from real Tweets that addressed the protest issues. To ensure comparability and realism, I also created popularity matrices for the Tweets. The number of quoted retweets was randomly drawn from a range of 0 to 10, while the number of retweets and likes was selected from a range of 10 to 100. This approach allowed the Tweets to appear as realistic as possible while minimizing the potential effects

Figure 1. Example sets of Simulated Twitter Profiles (White Man / Black Woman)



Note. While all the other features are held constant, the profile texts show differences – whether they show political content or not.

Figure 2. Example Sets of Simulated Tweets (White Man / Black Woman – Abortion Rights)



Note. The left column shows those from the Liberal group while the right column displays those from the Conservative group.

of excessively high or low engagement numbers. For the view count, I calculated it by multiplying 57 by the number of retweets. Although the popularity matrices differed within each group, their overall composition remained consistent across all groups.

Figure 2 presents examples of the simulated Tweets featuring either a white male or a black female. The first column displays Tweets encouraging participation in Pro-Choice protest events, while the second column showcases Tweets supporting Pro-Life protest events. Participants in the liberal group were shown images from the first column, whereas those in the conservative group were presented with images from the second column. All popularity metrics, except for the main content, were held constant.

Figures 3 and 4 display example sets of treatments that participants assigned to either the conservative group or the liberal group encountered during the experiment. The first and third images illustrate the political profile treatment, which features a profile containing political information or interests alongside

the protest-related Tweet. Participants in the political profile treatment group viewed this type of vignette. In contrast, the second and fourth images depict the non-political profile treatment, showcasing profiles that are entirely irrelevant to politics. Respondents in the nonpolitical profile treatment group were presented with these images. All participants were asked, "What will be your reactions when you see the Tweet? Please choose all choices that apply." with three options: Retweet, Like, and None. Control groups were shown only Tweets that aligned with their partisanship, without account information, as illustrated in Figure 3.

Additionally, I incorporated two screeners to measure respondents' attention. It should be noted that the survey did not include a formal manipulation check since no items were collected to assess whether participants perceived the profiles as political versus nonpolitical, or whether the profiles were judged as more or less credible. As a result, it is not possible to confirm that the treatment activated the intended perceptual mechanism. This is an important limitation for

Figure 3. Example of a Set of Treatments (White Man / Black Woman – Conservative)

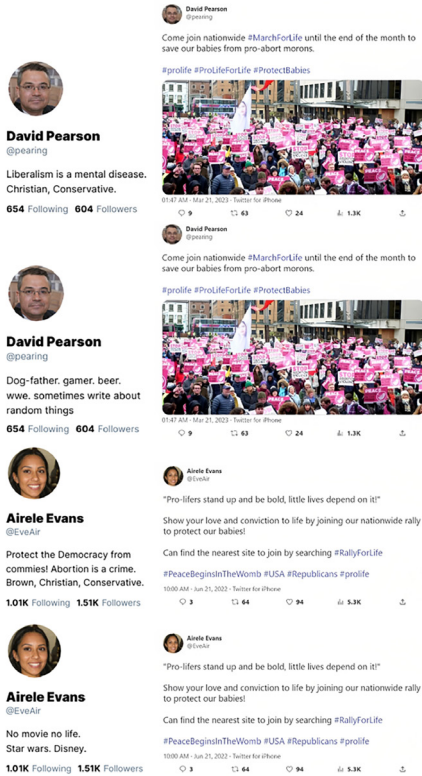
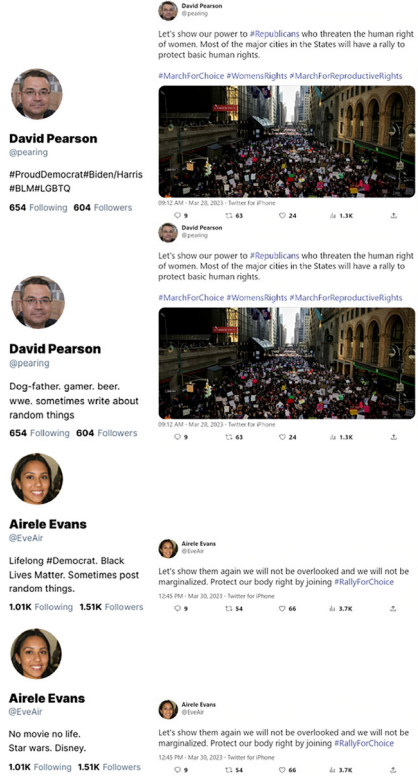


Figure 4. Example of a Set of Treatments (White Man / Black Woman – Liberal)



Note. The first and third rows show the treatments for the Political group while the second and last rows show those for the Nonpolitical group.

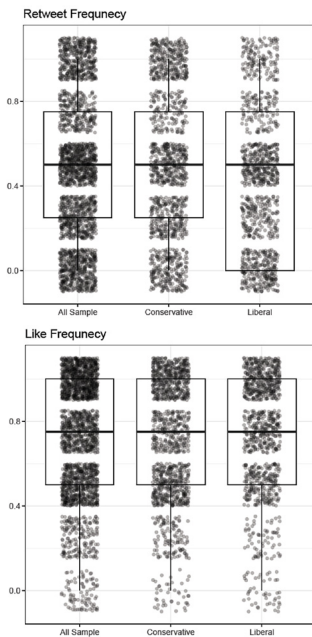
Note. The first and third rows show the treatments for the Political group while the second and last rows show those for the Nonpolitical group.

interpreting the results, which is addressed further in the Discussion. A multi-item scale that includes screeners with both high and low passage rates can be both economical and effective (Berinsky et al., 2021). Thus, I included one stand-alone screener with a low pass rate and one grid screener with a high pass rate.³

RESULTS

This section presents the findings from the statistical analyses. Figure 5 displays the distribution of two key dependent variables: normalized retweet frequency and like frequency, based on respondents' interactions. Normalized

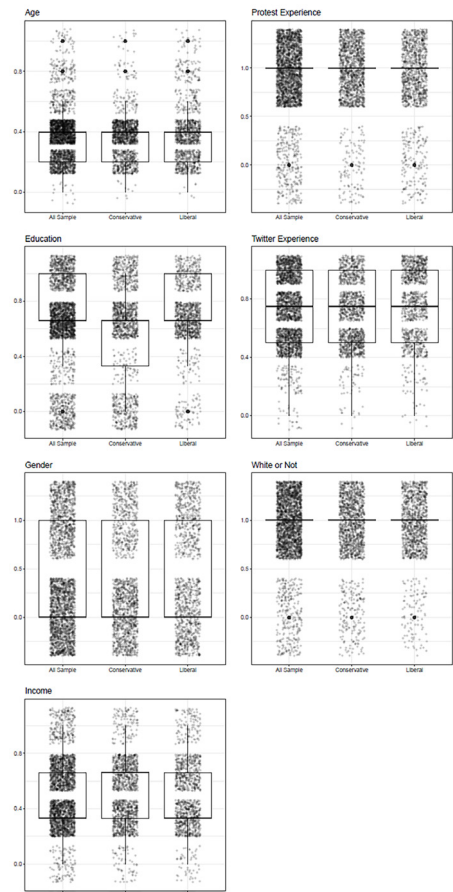
³ I used a “website” stand-alone screener and an “Obama was the first president” grid screener. See the full list of screeners by Berinsky et al. (2021) in their online appendix: <https://doi.org/10.1017/psrm.2019.53>.

Figure 5. *Distribution of Dependent Variables*

Note. The box plots depict the distribution of Retweet and Like frequency both at an aggregate level and across two groups. While the bottom line of the box shows the 25th percentile, the top line of the box shows the 75th percentile. The line inside the box shows the mean value.

scores for both metrics were calculated by summing responses to four key questions and dividing by the total number of questions in the vignette. Overall engagement, particularly the frequency of likes, was higher than anticipated. A substantial portion of respondents chose to like all the Tweets they viewed, indicating strong engagement with the presented content. Similarly, the mean retweet frequency was approximately 0.5, suggesting that respondents were selective, retweeting a significant portion of the Tweets they encountered.

Figure 6 illustrates the distribution of

Figure 6. *Distribution of Data on Respondents*

Note. The box plots depict the distribution of various population statistics both at an aggregate level and across two groups. While the bottom line of the box shows the 25th percentile, the top line of the box shows the 75th percentile. The line inside the box shows the mean value. Most of the variables are normalized to 1, meaning higher values indicate a high level in the variables (Age, Education, Twitter Experience, Income). The other three variables indicate a binary variable (Protest Experience, Gender, White).

respondents' personal characteristics, which generally align with the expected demographic patterns for Twitter users. The age distribution shows a right skew, while the gender distribution appears balanced. Education trends are slightly left-skewed, and income follows a normal distribution. The racial composition reflects a majority of white respondents (over 85%), which, while slightly higher than the 75% reported in American National Election Studies (2020), remains within a reasonable range. Additionally, a high proportion of respondents reported frequent Twitter usage, confirming the platform's relevance for this experiment.

A notable feature of the sample is the extremely high rate of protest participation: 88% of respondents indicated having participated in at least one protest or demonstration in the past 12 months, compared to 9% in the American National Election Studies (American National Election Studies, 2020). This discrepancy indicates that the sample is composed disproportionately of politically active individuals rather than the broader population of Twitter users that the paper's theoretical framework assumes.

This has direct consequences for interpreting the results. The paper's core argument assumes that ordinary, non-activist Twitter users update their engagement intentions based on credibility signals conveyed by account profiles, consistent with the credibility cascade framework (Munger, 2020). Highly committed activists, however, may engage with protest content primarily on the basis of pre-existing ideological commitment and collective identity rather than profile-level credibility assessments. For such individuals, whether the account sharing protest information is political or nonpolitical may be largely irrelevant

to their engagement decision, as their threshold for sharing is already low. Moreover, the high baseline engagement rates visible in Figure 5, with a substantial portion of respondents liking all tweets they viewed, are consistent with ceiling effects that would suppress any treatment effect regardless of its true size.

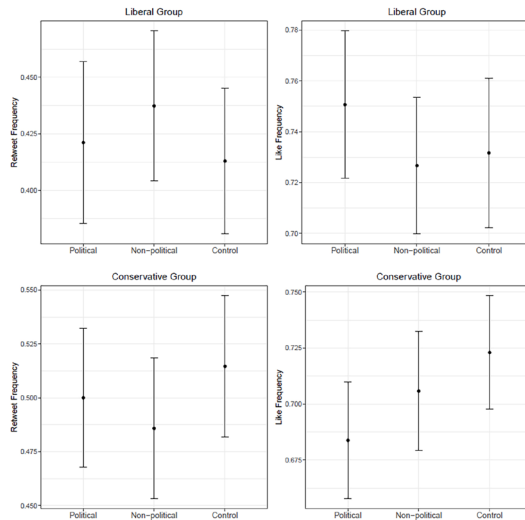
These considerations suggest that the null finding may partly reflect the composition of this particular sample rather than the complete absence of the hypothesized mechanism in the broader population. A sensitivity analysis comparing treatment effects among respondents who reported protest participation versus those who did not would help assess the extent to which sample composition drives the results. These results are reported in the Appendix (Figures A19–A22). The subgroup results show that treatment effects remain statistically non-significant in both the protest-experienced and non-experienced subgroups, suggesting that the null finding is not driven solely by the activist-heavy composition of the sample.⁴

Figure 7 displays the normalized mean values of retweet and like frequency across all groups. Contrary to the hypothesis, I did not find statistically significant differences between the political and nonpolitical groups regarding retweet and like frequency. In the liberal groups, the retweet frequency of the nonpolitical group was higher than that of the political group; however, this difference lacked statistical significance (p -value: 0.515).⁵ The like frequency exhibited the opposite trend, with more likes coming from the political group, but again, there was no statistical significance (p -value: 0.261). In the conservative group, the retweet frequency from the political group was higher than that from the nonpolitical group, contrary to the hypothesis;

⁴ Given that only approximately 12% of respondents reported no protest participation, the non-participant subgroup is small, and subgroup estimates should be interpreted with appropriate caution.

⁵ All p -values are calculated using randomization inference, which computes p -values for regression and difference-of-mean tests by considering fluctuations in randomized controlled trial data resulting from the randomization process itself.

Figure 7. Average Retweet and Like Frequency (Normalized to Zero)



Note. Difference-in-mean comparison tests are conducted by using the liberal treatment group and the conservative treatment group each as a population for different dependent variables – normalized retweet and like frequency (ranging from 0 to 1). Each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All *p*-values are calculated by using randomization inference.


however, this difference was also not statistically significant (*p*-value: 0.564). The like frequency aligned with the predicted direction suggested by the hypothesis, with more likes coming from the nonpolitical group, but once again, there was no statistical significance (*p*-value: 0.261).

Figures 8 and 9 depict the estimated effects obtained from multivariate regression models for both liberal and conservative groups, with retweet and like frequency as the dependent variables. The results align with the findings of the difference-of-mean comparison tests presented in Figure 7, revealing that the political profile treatment does not exhibit statistical significance in either the liberal or conservative groups for either dependent variable. Additionally, no significant effects of treatments were found in the disaggregated models that utilized the frequency of retweets and likes for each type of protest event (pro-life/

pro-choice protests and anti-gun control/pro-gun control protests) as their dependent variables, as demonstrated in Figures A1 to A4.

An important caveat for interpreting this null result is the absence of a manipulation check. The null finding is consistent with two distinct explanations: (1) the treatment was received as intended and the hypothesized credibility mechanism genuinely did not operate, or (2) participants did not perceive the profiles as meaningfully political versus nonpolitical in the first place, meaning the treatment failed to activate the intended perception. These two possibilities cannot be distinguished with the current data. Future studies should include direct perceptual measures, such as ratings of perceived politicalness and perceived credibility of the profile, to cleanly separate manipulation failure from genuine mechanism absence.

Figure 8. Multivariate Regression Model Results (DV: Retweet Probability)




Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

The following patterns were not pre-registered and emerged from exploratory analyses across a large number of model specifications. Given the corresponding risk of inflated false-positive rates, these findings should be treated as hypothesis-generating rather than confirmatory, and they warrant replication in future pre-registered studies.

One exploratory pattern in the results is a contrasting association between education and retweet intentions across the two partisan groups, as indicated in Figure 8. The education variable shows a positive association in the liberal group and a negative association in the conservative group. Existing literature documents a generally positive relationship between education and political engagement across countries (Larreguy

Figure 9. Multivariate Regression Model Results (DV: Like Probability)



Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

& Marshall, 2017; Le & Nguyen, 2021; Perrin & Gillis, 2019; Sunshine Hillygus, 2005); the negative pattern observed in the conservative group runs counter to this expectation and may reflect the partisan sorting of higher-educated conservatives away from protest activity. This pattern is reported descriptively and should not be interpreted as establishing a causal relationship.

A second exploratory pattern emerges from a model using retweet intentions for gun control protest-related questions with a relevant image as the dependent variable (Figure A9). In the liberal group, respondents who viewed the Tweet with a nonpolitical profile reported a 5 percentage point higher retweet intention compared to those who viewed it with a political profile ($p = 0.067$).

This directional pattern is consistent with the hypothesis, but given its unregistered status and the marginal p-value, it should be understood strictly as a pattern warranting future investigation rather than as supporting evidence for the hypothesis.

DISCUSSION

In this project, I developed a theory and hypothesis to examine the impact of the political characteristics of protest information producers on users' engagement intentions toward protest information on Twitter. I posited that users who view Tweets from accounts with political affiliations are less inclined to share protest-related information compared to those who observe Tweets from accounts without political affiliations. To test this hypothesis, I conducted an online survey experiment using a vignette that simulates actual Tweets users would encounter on the platform.

The hypothesis was not supported as all p-values exceeded .25 across both difference-of-means tests and multivariate regression models. This clear null warrants direct theoretical engagement rather than treatment as an incidental result. The paper's hypothesis extended the credibility cascade framework (Munger, 2020), which describes how social endorsements can confer legitimacy on information from initially low-credibility sources, to the domain of protest mobilization. The expectation was that non-politicized accounts would lend greater credibility to protest content, thereby encouraging wider engagement. However, protest mobilization may be governed by motivational logics that are fundamentally distinct from news credibility assessment. The collective action literature emphasizes that the decision to endorse or participate in protest activity involves moral commitment, collective identity, and strategic calculations about a movement's likely success (Kuran, 1991; Lichbach, 1995; Olson,

1971). In threshold models of collective action, individuals act when the perceived momentum of a movement exceeds their personal threshold (Kuran, 1991; Lohmann, 1994). All that said, theorized profile-level cues about a single account's political orientation may be too peripheral to shift these thresholds meaningfully, particularly when the protest content itself already provides the substantively relevant information for the engagement decision.

Furthermore, the experimental design deliberately matched protest events to respondents' partisan leanings, ensuring ideological alignment between the content and the participant. This means the primary motivational driver for engagement was held constant across conditions. Under these conditions, profile-level credibility cues may have lacked the capacity to generate meaningful additional variation in engagement intentions. As discussed in the Results section, the activist-heavy sample composition further reduces the likelihood of detecting a treatment effect, as highly committed individuals likely engage on the basis of pre-existing ideological commitments rather than source credibility assessments. Taken together, these considerations suggest that the credibility cascade logic, originally theorized in the context of news media credibility (Munger, 2020), does not straightforwardly extend to protest mobilization, where moral and ideological commitments may appear to dominate the engagement decision.

Beyond the null finding, two unregistered exploratory patterns emerged and are reported here as hypothesis-generating. In exploratory analyses, education was associated with retweet intentions in opposite directions across the two partisan groups—positively among liberals and negatively among conservatives. If replicated, this pattern would challenge existing literature documenting a generally positive relationship between education and political engagement (Larreguy & Marshall, 2017; Le & Nguyen, 2021;

Perrin & Gillis, 2019; Sunshine Hillygus, 2005), and it opens avenues for future investigation into how education and partisan identity jointly shape online engagement. A second exploratory pattern, from a model using protest-type-specific dependent variables, showed a directional association consistent with the hypothesis in the liberal group for gun control content paired with images. Given the unregistered nature of both patterns and the elevated false-positive risk across a large analytic space, these should be treated as motivating future pre-registered research rather than as confirmatory evidence.

Several limitations of this study warrant explicit acknowledgment. First, as with all vignette-based survey experiments, the dependent variables capture stated engagement intentions rather than observed behavior. Self-reported intentions are subject to social desirability pressures and may not map directly onto actual retweet or like decisions in a live platform environment. Second, participants in this study were exposed to four sequential simulated profile pages, which introduces the possibility of ordering and habituation effects. Responses to later stimuli may reflect diminished sensitivity to the profile-level manipulation rather than purely the experimental condition, a concern that is particularly relevant for interpreting the null finding. Third, since the study was conducted, Twitter has been rebranded as X under new ownership, accompanied by significant changes to the platform's features, user base, and content moderation practices. It remains an open question whether the dynamics of political signaling and credibility that motivated this study operate similarly on X, and future research should examine whether these findings generalize to the current platform context.

In summary, this study contributes to our understanding of how political characteristics of source accounts influence users' engagement intentions toward protest information on social media, while also highlighting the complexities inherent in user engagement. The exploratory

patterns regarding education levels, if replicated in future pre-registered studies, would suggest the need for a nuanced exploration of how demographic factors shape political behavior in online environments. Future research should focus on refining methodologies to better capture the dynamics of Twitter interactions, exploring the role of content style, and considering a broader array of user characteristics. By doing so, we can gain deeper insights into the mechanisms of online protest mobilization and the varying influences of political affiliations on user engagement.

REFERENCES

- Albertson, B., & Lawrence, A. (2009). After the credits roll: The long-term effects of educational television on public knowledge and attitudes. *American Politics Research*, 37(2), 275–300.
- American National Election Studies. (2020). *ANES 2020 time series study full release*. <https://electionstudies.org/data-tools/anes-question-search/>
- Arceneaux, K., & Johnson, M. (2013). *Changing minds or changing channels?: Partisan news in an age of choice*. University of Chicago Press.
- Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. (2012). The role of social networks in information diffusion. In *Proceedings of the 21st International Conference on World Wide Web* (pp. 519–528).
- Barberá, P., Wang, N., Bonneau, R., Jost, J. T., Nagler, J., Tucker, J., & González-Bailón, S. (2015). The critical periphery in the growth of social protests. *PLoS One*, 10(11), Article e0143611.
- Berinsky, A. J., Margolis, M. F., Sances, M. W., & Warshaw, C. (2021). Using screeners to measure respondent attention on self-administered surveys: Which items and how many? *Political Science Research and Methods*, 9(2), 430–437.


- Centola, D. M. (2013). Homophily, networks, and critical mass: Solving the start-up problem in large group collective action. *Rationality and Society*, 25(1), 3–40.
- Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3), 702–734.
- Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., & Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9), Article e2023301118.
- Cohen, B. C. (1963). *The press and foreign policy*. Princeton University Press.
- Feezell, J. T. (2018). Agenda setting through social media: The importance of incidental news exposure and social filtering in the digital era. *Political Research Quarterly*, 71(2), 482–494.
- Feldman, L. (2011). Partisan differences in opinionated news perceptions: A test of the hostile media effect. *Political Behavior*, 33(3), 407–432.
- Garimella, K., De Francisci Morales, G., Gionis, A., & Mathioudakis, M. (2018). Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In *Proceedings of the 2018 World Wide Web Conference* (pp. 913–922).
- Gillani, N., Yuan, A., Saveski, M., Vosoughi, S., & Roy, D. (2018). Me, my echo chamber, and I: Introspection on social media polarization. In *Proceedings of the 2018 World Wide Web Conference* (pp. 823–831).
- Granovetter, M. (1983). The strength of weak ties: A network theory revisited. *Sociological Theory*, 1, 201–233.
- Holbert, R. L., Garrett, R. K., & Gleason, L. S. (2010). A new era of minimal effects? A response to Bennett and Iyengar. *Political Behavior*, 60(1), 15–34.
- Iyengar, R., Van den Bulte, C., & Valente, T. W. (2011). Opinion leadership and social contagion in new product diffusion. *Marketing Science*, 30(2), 195–212.
- Iyengar, S. (1991). *Is anyone responsible? How television frames political issues*. University of Chicago Press.
- Kuran, T. (1991). Now out of never: The element of surprise in the East European revolution of 1989. *World Politics*, 44(1), 7–48.
- Larreguy, H., & Marshall, J. (2017). The effect of education on civic and political engagement in nonconsolidated democracies: Evidence from Nigeria. *Review of Economics and Statistics*, 99(3), 387–401.
- Le, K., & Nguyen, M. (2021). Education and political engagement. *International Journal of Educational Development*, 85, Article 102441.
- Levendusky, M. S. (2013). Why do partisan media polarize viewers? *American Journal of Political Science*, 57(3), 611–623.
- Lichbach, M. I. (1995). *The rebel's dilemma*. University of Michigan Press.
- Lin, K. Y., & Lu, H. P. (2011). Why people use social networking sites: An empirical study integrating network externalities and motivation theory. *Computers in Human Behavior*, 27(3), 1152–1161.
- Lohmann, S. (1994). The dynamics of informational cascades: The Monday demonstrations in Leipzig, East Germany, 1989–91. *World Politics*, 47(1), 42–101.
- Macy, M. W. (1990). Learning theory and the logic of critical mass. *American Sociological Review*, 55(6), 809–826.
- Martin, G. J., & Yurukoglu, A. (2017). Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9), 2565–2599.
- Marwell, G., Oliver, P. E., & Prah, R. (1988). Social networks and collective action: A theory of the critical mass. III. *American Journal of Sociology*, 94(3), 502–534.
- McCombs, M. E., & Shaw, D. L. (1972). The agenda-setting function of mass media. *Public Opinion Quarterly*, 36(2), 176–187.
- Meng, F., Wei, J., & Zhu, Q. (2011). Study on the impacts of opinion leader in online consuming decision. In *Proceedings of International Joint*

- Conference on Service Sciences (pp. 140–144).
- Munger, K. (2020). All the news that's fit to click: The economics of clickbait media. *Political Communication*, 37(3), 376–397.
- Olson, M. (1971). *Logic of collective action*. Harvard University Press.
- Park, C. S. (2013). Does Twitter motivate involvement in politics? Tweeting, opinion leadership, and political engagement. *Computers in Human Behavior*, 29(4), 1641–1648.
- Perrin, A. J., & Gillis, A. (2019). How college makes citizens: Higher education experiences and political engagement. *Socius*, 5, Article 2378023119859708.
- Quan-Haase, A., & Young, A. L. (2010). Uses and gratifications of social media: A comparison of Facebook and instant messaging. *Bulletin of Science, Technology Society*, 30(5), 350–361.
- Romero, D. M., Meeder, B., & Kleinberg, J. (2011). Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th International Conference on World Wide Web* (pp. 695–704).
- Schäfer, M. S., & Taddicken, M. (2015). Opinion leadership—Mediatized opinion leaders: New patterns of opinion leadership in new media environments? *International Journal of Communication*, 9, 195–212.
- Seattle, J. E. (2018). *Frenemies: How social media polarizes America*. Cambridge University Press.
- Steinert-Threlkeld, Z. C. (2017). Spontaneous collective action: Peripheral mobilization during the Arab Spring. *American Political Science Review*, 111(2), 379–403.
- Stroud, N. J. (2011). *Niche news: The politics of news choice*. Oxford University Press.
- Sunshine Hillygus, D. (2005). The missing link: Exploring the relationship between higher education and political engagement. *Political Behavior*, 27, 25–47.
- Van Eck, P. S., Jager, W., & Leeflang, P. S. (2011). Opinion leaders' role in innovation diffusion: A simulation study. *Journal of Product Innovation Management*, 28(2), 187–203.
- Watts, D. J. (2002). A simple model of global cascades on random networks. In *Proceedings of the National Academy of Sciences* (pp. 5766–5771).

Appendix


A1. Regression Results for Specific Types of Questions

Figure A1. Multivariate Regression Model Results (DV: Retweet Probability for Abortion-Related Questions)




Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

Figure A2. Multivariate Regression Model Results (DV: Like Probability for Abortion-Related Questions)




Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

Figure A3. Multivariate Regression Model Results (DV: Retweet Probability for Gun-Related Questions)



Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.


Figure A4. Multivariate Regression Model Results (DV: Like Probability for Gun-Related Questions)



Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.


A2. Regression Results for Single Questions

Figure A5. Multivariate Regression Model Results (DV: Retweet Probability for an Abortion-Related Question with a Relevant Image)




Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

Figure A6. Multivariate Regression Model Results (DV: Like Probability for an Abortion-Related Question with a Relevant Image)




Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

Figure A7. *Multivariate Regression Model Results (DV: Retweet Probability for an Abortion-Related Question without a Relevant Image)*




Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

Figure A8. *Multivariate Regression Model Results (DV: Like Probability for an Abortion-Related Question without a Relevant Image)*




Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

Figure A9. Multivariate Regression Model Results (DV: Retweet Probability for a Gun-Related Question with a Relevant Image)




Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

Figure A10. Multivariate Regression Model Results (DV: Like Probability for a Gun-Related Question with a Relevant Image)




Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

Figure A11. *Multivariate Regression Model Results (DV: Retweet Probability for a Gun-Related Question without a Relevant Image)*



Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

Figure A12. *Multivariate Regression Model Results (DV: Like Probability for a Gun-Related Question without a Relevant Image)*



Note. Two models are fitted by using the liberal treatment group and the conservative treatment group each as population. The x-axis depicts the coefficients in regression analysis. The y-axis shows the variables included in the regression analysis. For both, each dot and the corresponding line indicate the point estimate and the 95% confidence interval. All p-values are calculated by using randomization inference.

A3. Balance Tables

Table A1. Balance Table for Conservative Groups

	Conservative / Political	Conservative / Nonpolitical	Conservative / Control
Education	0.59	0.60	0.58
White	0.83	0.89	0.88
Income	0.51	0.52	0.50
Gender	0.33	0.30	0.34
Age	0.37	0.36	0.36
Twitter Frequency	0.70	0.69	0.70
Protest Experience	0.91	0.91	0.93

Note. The numbers in the table are the mean values of all observations within the group. Most of the variables are normalized to 1, meaning higher values indicate a high level in the variables (Age, Education, Twitter Experience, Income). The other three variables are binary variables (Protest Experience, Gender, White).

Table A2. Balance Table for Liberal Groups

	Liberal / Political	Liberal / Nonpolitical	Liberal / Control
Education	0.69	0.69	0.68
White	0.92	0.90	0.91
Income	0.50	0.51	0.52
Gender	0.41	0.39	0.44
Age	0.40	0.40	0.40
Twitter Frequency	0.77	0.76	0.78
Protest Experience	0.84	0.85	0.85

Note. The numbers in the table are the mean values of all observations within the group. Most of the variables are normalized to 1, meaning higher values indicate a high level in the variables (Age, Education, Twitter Experience, Income). The other three variables are binary variables (Protest Experience, Gender, White).

A4. Simulated Account Profiles and Tweets

This section presents the full set of simulated Twitter account profiles and tweets used in the experiment. Profiles are organized by treatment condition. Within each condition, profiles 1-2

correspond to the abortion rights scenario and profiles 3-4 correspond to the gun control scenario.

Figure A13. Liberal Political Profiles (Profiles 1-2: Abortion Rights Scenario, Profiles 3-4: Gun Control Scenario)



Figure A14. Liberal Nonpolitical Profiles (Profiles 1-2: Abortion Rights Scenario, Profiles 3-4: Gun Control Scenario)



Figure A15. Conservative Political Profiles (Profiles 1-2: Abortion Rights Scenario, Profiles 3-4: Gun Control Scenario)







Note. Profiles 1-2 contain comparatively more hostile language (e.g., “Liberalism is a mental disease”; “Protect the Democracy from commies!”) than profiles 3-4, which are assertive but not hostile in tone.

Figure A16. Conservative Nonpolitical Profiles (Profiles 1-2: Abortion Rights Scenario, Profiles 3-4: Gun Control Scenario)

 David Pearson @pearing Dog-father. gamer. beer. wwe. sometimes write about random things 654 Following 604 Followers	 Airele Evans @EveAir No movie no life. Star wars. Disney. 1.01K Following 1.51K Followers	 Raina Lopez @lopran cook for fun, listen to live. mostly tweet about things I care. 1.43K Following 1.8K Followers	 Mario Chapa @mario10 Full-time gamer. sometimes coding. 916 Following 764 Followers
---	--	---	--

Figure A17. Simulated Tweets Shown to Participants in the Liberal Condition (Tweets 1-2: Abortion Rights; Tweets 3-4: Gun Control)

Note. Identical tweets were shown across the political and nonpolitical profile conditions.

Figure A18. Simulated Tweets Shown to Participants in the Conservative Condition (Tweets 1-2: Abortion Rights; Tweets 3-4: Gun Control)

The figure displays four simulated tweets arranged in a 2x2 grid. Each tweet includes a profile picture, name, handle, text, an image, and engagement metrics (replies, retweets, likes).

- Top-Left:** David Pearson (@pearing) promotes a nationwide #MarchForLife rally to save babies. Includes a photo of a large crowd at a rally.
- Top-Right:** Airele Evans (@EveAie) encourages pro-lifers to stand up and protect babies. Includes a photo of a person holding a sign.
- Bottom-Left:** Raina Lopez (@topran) asks for support to protect constitutional rights. Includes a photo of a person holding a 'DON'T TREAD ON ME' sign with a crossed gun and '2nd AMENDMENT' text.
- Bottom-Right:** Mario Chapa (@mariorio) demands no more restrictions on firearms. Includes a photo of a person holding a 'NO MORE GUNS' sign.


Note. Identical tweets were shown across the political and nonpolitical profile conditions.

A5. Subgroup Sensitivity Analysis (Protest Experience)

The following figures present regression results disaggregated by respondents' self-reported protest experience. Figures A19 and A20 show retweet and like intention models estimated among respondents who reported having participated in at least one protest (approximately


88% of the sample). Figures A21 and A22 show the same models among respondents who reported no protest participation (approximately 12% of the sample). Given the small size of the non-participant subgroup, estimates in Figures A21-A22 should be interpreted with caution.

Figure A19. Subgroup Analysis: Multivariate Regression Results (DV: Retweet Probability; Respondents with Protest Experience)




Note. Two models are fitted using the liberal and conservative treatment groups as separate populations. The x-axis depicts the coefficients; the y-axis shows the predictors. Each dot and line represent the point estimate and 95% confidence interval. All p-values calculated via randomization inference.

Figure A20. Subgroup Analysis: Multivariate Regression Results (DV: Like Probability; Respondents with Protest Experience)




Note. Two models are fitted using the liberal and conservative treatment groups as separate populations. The x-axis depicts the coefficients; the y-axis shows the predictors. Each dot and line represent the point estimate and 95% confidence interval. All p-values calculated via randomization inference.

Figure A21. Subgroup Analysis: Multivariate Regression Results (DV: Retweet Probability; Respondents without Protest Experience)



Note. Two models are fitted using the liberal and conservative treatment groups as separate populations. The x-axis depicts the coefficients; the y-axis shows the predictors. Each dot and line represent the point estimate and 95% confidence interval. All p-values calculated via randomization inference.

Figure A22. Subgroup Analysis: Multivariate Regression Results (DV: Like Probability; Respondents without Protest Experience)



Note. Two models are fitted using the liberal and conservative treatment groups as separate populations. The x-axis depicts the coefficients; the y-axis shows the predictors. Each dot and line represent the point estimate and 95% confidence interval. All p-values calculated via randomization inference.