Sympathy over Polarization: A Computational Discourse Analysis of Social Media Posts about the July 2024 Trump Assassination Attempt

Qingcheng Zeng[*1], **Guanhong Liu**[*2], **Zhaoqian Xue**[3], **Diego Ford**[4],
Rob Voigt[5], **Loni Hagen**[4], **Lingyao Li**[4]

[1]Northwestern University, Evanston, IL, USA
[2]University of Chicago, Chicago, IL, USA
[3]Georgetown University, Washington, DC, USA
[4]University of South Florida, Tampa, FL, USA
[5]University of California, Davis, Davis, CA, USA

qingchengzeng2027@u.northwestern.edu, lingyaol@usf.edu

Abstract

On July 13, 2024, an assassination attempt was made on Republican presidential candidate Donald Trump during a rally in Pennsylvania. This event triggered widespread discourses on social media platforms. In this study, we analyze posts from X (formerly Twitter) collected during the week preceding and following the incident to examine the short-term impact of this political shock on public opinion and discourse. Our investigation is guided by three central research questions. First (RQ1), we assess how public stance toward Donald Trump evolved over time and varied across geographic regions. Second (RQ2), we apply causal inference methods to determine whether the assassination attempt itself significantly influenced public attitudes, independent of pre-existing political alignments. Third (RQ3), we conduct topic modeling to identify shifts in dominant themes of online discussions before and after the event. Integrating large language model-based stance detection, difference-in-differences estimation, and topic modeling, our findings reveal a marked surge in sympathetic responses toward Trump in the immediate aftermath of the attempt, suggesting a unifying effect that temporarily transcended ideological and regional divides.

1 Introduction

Shock events, such as health crises or political violence, often catalyze abrupt shifts in public opinion toward political figures and partisan alignment (Mackintosh, 2014; Johansson et al., 2021). Among these, assassination attempts are especially extreme, drawing intense media scrutiny and provoking strong emotional responses with the potential to reshape the socio-political landscape (Pillemer, 2000; Atkeson and Maestas, 2012). The attempted assassination of Donald Trump during a campaign event in Pennsylvania in July 2024 exemplifies this type of political shock.

Existing literature offers two contrasting expectations toward the effect of assassination attempts on public attitudes: **sympathy** and **polarization**. The sympathy hypothesis suggests that dramatic events elicit cross-partisan unity and a temporary surge in public support, driven by affective responses rather than political alignment, as observed in Ronald Reagan's approval increase following the 1981 assassination attempt (Ostrom Jr. and Simon, 1985; Gilbert, 2013). In contrast, the polarization hypothesis anticipates intensified ideological divisions, as occurred after Martin Luther King Jr.'s assassination, which deepened racial and political fractures in American society (Putnam, 2007; Sokol, 2018; Whitlinger and Fretwell, 2019). These divergent frameworks point to fundamentally different outcomes in the wake of traumatic political violence.

Examining the survey data, Holliday et al. (2024) find that the 2024 Trump assassination event did not intensify partisan tensions. Democrats' attitudes remained unchanged, while Republicans exhibited a decreased support for partisan violence and showed no heightened hostility toward the opposing party. While survey research provides structured and representative insights into public opinion, it remains time- and resource-intensive (Groves et al., 2009). Social media, by contrast, offers immediate, large-scale reflections of public stances (González-Bailón et al., 2014; Pandarachalil et al., 2015), enabling real-time analysis of discourse dynamics. Metadata such as timestamps and geolocations further allow for fine-grained temporal and spatial mapping (Bollen et al., 2011; Pang and Lee, 2008). In a context of deep partisan polarization, social media has proven valuable for capturing ideological and geographic variation in responses to political shocks (Bennett et al., 2021; Rueda et al., 2023). In this study, we leverage social media to examine the impact of the July

[*]Equal contribution.

Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, pages 56–68

2024 assassination attempt on public discourse and investigate the following research questions:

RQ1: Descriptively, how do public stances toward Donald Trump vary over time and across regions?

RQ2: Does the assassination attempt significantly influence stances toward Trump, after accounting for political divisions?

RQ3: What are the major discussion themes before and after the event, and how do their associated stances evolve?

To answer these questions, we conduct the first large-scale descriptive study on X (formerly Twitter) examining discourse before and after the July 2024 attempted assassination of Donald Trump. We begin by designing a prompt for aspect-based stance detection toward Trump, evaluating it across several contemporary models, and selecting the best-performing model to annotate the full dataset. We then implement a Difference-in-Differences (DiD) framework to assess changes in stances across different state groupings (e.g., Red vs. Blue, Swing vs. non-Swing), measuring the event's impact on public stance. To capture thematic shifts, we apply LLM-based topic modeling (Pham et al., 2024), enabling a fine-grained analysis of discourse evolution. Our results indicate a general increase in favorable stance toward Trump, largely independent of state-level partisanship. We also identify substantial shifts in discussion topics, most notably a sharp decline in tweets referencing Trump scandals post-event. These patterns suggest a sympathy-driven public response, reflecting broad positive shifts in stance without accompanying signs of increased polarization.

2 Two Contrasting Hypotheses

Existing research on the effects of political shocks, such as assassination attempts, on public opinion typically centers on two competing perspectives: sympathy and polarization. The **sympathy** hypothesis posits that public stances can shift rapidly and broadly in favor of a political figure following a personal ordeal, independent of policy evaluations or partisan alignment (Ostrom Jr. and Simon, 1985). This mechanism is closely tied to the broader concept of the *Rally 'Round the Flag Effect* (Brody, 1991), which describes sharp increases in presidential approval during times of national crisis, including wars, terrorist attacks, and other emergencies (Mueller, 1970; Chanley, 2002; Hetherington and Nelson, 2003). A frequently cited example is President Ronald Reagan's approval surge following the 1981 assassination attempt, widely attributed to public sympathy rather than a reassessment of his political record (Ostrom Jr. and Simon, 1985; Brody, 1991; Gilbert, 2013). Applied to the Trump case, the sympathy hypothesis would predict a broad increase in favorable stances, potentially transcending entrenched partisan boundaries.

In contrast, the **polarization** hypothesis argues that major crises can deepen existing ideological and partisan divides, as individuals interpret events through divergent social and political lenses (Putnam, 2007). Rooted in Social Identity Theory, this perspective suggests that during crises, group identities become more salient, reinforcing in-group favoritism and out-group hostility (Tajfel and Turner, 2004; Huddy, 2003). For example, Sokol (2018) documents how Martin Luther King Jr.'s assassination intensified racial polarization, radicalizing both civil rights advocates and opponents. If polarization dominates the public response, the Trump assassination attempt may deepen partisan cleavages, manifesting, for instance, in regionally differentiated reactions between "red" and "blue" states.

3 Data and Methods

3.1 Data collection

We collected social media data from X using Brandwatch, a platform for historical data retrieval and large-scale analytics. The keyword "Trump" was chosen to broadly capture relevant discourse related to the event. Data spans July 7–20, 2024, covering one week before and after the July 13 assassination attempt to capture short-term public reaction while minimizing unrelated noise. To focus on the U.S. context, we included only English-language posts geotagged within the United States. Posts lacking state-level location metadata were excluded, resulting in a final dataset of 122,526 posts. Descriptive statistics for this dataset are provided in Appendix A.

3.2 Stance detection and validation

We employed LLMs to perform aspect-based stance detection on X posts, focusing on public stance toward Donald Trump. Each post was classified as favor, against, or neutral based on the stance it expressed. To evaluate model performance, we manually annotated a random sample of 300 posts. Three annotators independently labeled each post, with final labels assigned by majority vote. The average inter-annotator agreement, measured by Cohen's Kappa, was 0.67, indicating substantial consistency. This high-quality, human-annotated dataset served as a benchmark for assessing the accuracy and reliability of the LLM classifiers.

We evaluated five models: four LLMs, `Deepseek-V3-0324` (DeepSeek-AI et al., 2025), `Claude-3.5-Sonnet-20240620` (Anthropic, 2024), `GPT-4.1` (OpenAI, 2024), and `GPT-4.1-mini` (OpenAI, 2024), and a baseline sentiment classifier, RoBERTa sentiment (Barbieri et al., 2020). A task-specific prompt was developed to guide the models, including clear definitions of each stance category and explicit instructions for aspect-based classification. The full prompt is provided in Appendix B. Models were evaluated using four performance metrics—precision, recall, F1-score, and overall accuracy—across all three stance categories. As shown in Figure 3, GPT-4.1 consistently outperformed the other models across all metrics. Based on its superior performance, we selected GPT-4.1 to annotate the full dataset. Three examples of 'favor', 'neutral', and 'against' are provided in Appendix A.

3.3 Difference-in-differences (DiD) modeling

We then employed the DiD approach to estimate the assassination attempt's impact on public stance. The DiD framework is well-suited for this study due to the clearly defined intervention (assassination attempt), measurable outcome (public stance score), and natural division of states into treatment and control groups based on partisan alignment. The DiD implementation followed three steps (Card and Krueger, 2000; Besley and Case, 2002; Ladd and Lenz, 2009; Angrist and Pischke, 2009): (1) confirming the parallel trends assumption for treatment and control groups pre-intervention, (2) building regression models to estimate the treatment effect, and (3) performing robustness checks to ensure result reliability.

Intervention point specification. In this study, the intervention refers to the assassination attempt on Donald Trump on July 13, 2024, at 6:11 PM EDT (O'Donoghue et al., 2024). To ensure temporal consistency nationwide, all timestamps were converted to UTC (EDT+4). Data was segmented into pre-intervention (July 7–July 13, up to 10:11 PM UTC) and post-intervention (July 13, 10:11 PM UTC–July 20) periods. Public stance scores toward Donald Trump were computed using GPT-4.1 annotations, aggregated by state and date as a weighted average stance score for X posts, $S_{s,d}$:

$$S_{s,d} = \frac{\sum_{i=1}^{N_{s,d}} S_i}{N_{s,d}} \tag{1}$$

where: $N_{s,d}$ is the total number of tweets posted in state s on date d; S_i is the stance score for the i-th tweet in state s on date d.

Treatment-Control group division. The sympathy and polarization hypotheses differ in whether the assassination attempt drives distinct trends or magnitudes in public attitudes across social groups. Given the U.S.'s deep political polarization (Bennett et al., 2021; Rueda et al., 2023), we defined three treatment-control group divisions based on state-level political preferences: (1) Republican-leaning ("red") states as treatment, Democrat-leaning ("blue") states as control; (2) swing states as treatment, blue states as control; and (3) swing states as treatment, red states as control. This reflects the theoretical expectation that red states will show the strongest net stance increase compared to non-red states while swing states may show a shift in stance, potentially leaning toward one side. State classifications (red, blue, swing) are detailed in Table 1, based on Politico's public polling data (POLITICO, 2024).

Parallel trend assumptrion verification. The DiD framework requires that treatment and control groups would exhibit similar trends over time absent the intervention. Divergent trends render the DiD model unsuitable, necessitating parallel trend verification to ensure its applicability. We tested this parallel trend using linear regression for each state:

$$y_i = \beta_0 + \beta_1 x_i + \epsilon_i \tag{2}$$

where: y_i is the public stance score for a given day; x_i is the number of days since the earliest observation in that state; β_1 is the slope, representing the rate of change in public stance over time.

To compare mean slopes between groups, we conducted an independent samples t-test:

$$t = \frac{\bar{\beta}_T - \bar{\beta}_C}{\sqrt{\frac{s_T^2}{n_T} + \frac{s_C^2}{n_C}}} \tag{3}$$

where: $\bar{\beta}_T$ and $\bar{\beta}_C$ are mean slopes for the treatment group and the control group, individually; s_T^2 and s_C^2 are the variance of slopes in the treatment group and the control group, individually; n_T and

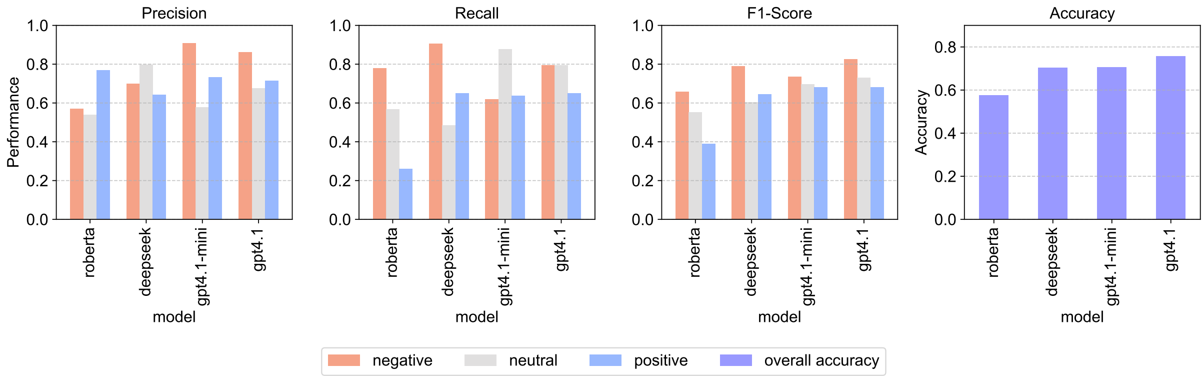

Figure 1: Performance evaluation of stance detection across different models.

Category	States
Swing States	Arizona, Georgia, Michigan, Nevada, North Carolina, Pennsylvania, Wisconsin
Blue States	California, Colorado, Connecticut, Delaware, Hawaii, Illinois, Maine, Maryland, Massachusetts, Minnesota, New Jersey, New Mexico, New York, Oregon, Rhode Island, Vermont, Washington, Washington D.C., Virginia
Red States	Alabama, Alaska, Arkansas, Idaho, Indiana, Iowa, Kansas, Kentucky, Louisiana, Mississippi, Missouri, Montana, Nebraska, North Dakota, Oklahoma, South Carolina, South Dakota, Tennessee, Texas, Utah, West Virginia, Wyoming, Florida, Ohio, New Hampshire

Table 1: Political Demarcation based on the Politico

n_C are the number of states in the treatment group and the control group, individually.

The t-tests for parallel trends (Table 2) yield t-statistics with p-values > 0.05, indicating no significant difference in pre-intervention trends across all treatment-control group pairs. This supports the parallel trends assumption, justifying the use of the DiD model.

Treatment v.s. Control Groups	T-Statistic
Red v.s. Blue States	0.274
Swing v.s. Red States	0.456
Swing v.s. Blue States	0.740

Table 2: T-Test for Parallel Trends Verification, significance codes: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Estimation frameworks. To estimate the causal effect of the assassination attempt on public stance toward Donald Trump, the DiD model is specified as:

$$\begin{aligned} Y_{s,t} &= \beta_0 + \beta_1 \text{Treatment}_s + \beta_2 \text{Post}_t \\ &+ \beta_3 (\text{Treatment}_s \times \text{Post}_t) \\ &+ \gamma X_{s,t} + \epsilon_{s,t} \end{aligned} \tag{4}$$

where $Y_{s,t}$ is the outcome variable (public stance score) for state s at time t; Treatment_s indicates the treatment group; Post_t indicates the post-intervention period; β_3 is the DiD effect from the interaction term; and $X_{s,t}$ represents control variables that may also influence the public stance. Notably, Cherlin (2021) argue that the white working class, grappling with job losses in an increasingly "open" America, forms a steadfast base of Trump support. Similarly, Winter (2023) notes that male voters are shifting toward the conservative Trump. Therefore, we include control variables—namely, the percentage of residents below poverty, female percentage, employment percentage, bachelor's degree percentage, and white percentage—obtained from the United States Census Bureau (U.S. Census Bureau, 2025), as well as time as an additional control variable.

To validate the robustness of the causal inference framework, we conducted a placebo test by setting the intervention break-point to July 11, 2024, around two days prior to the actual assassination attempt. This can ensure that spurious relationships or unaccounted temporal trends do not drive any observed effects in the primary analysis.

3.4 Topic modeling

Topic modeling is a widely used technique for identifying latent semantic structures within large text corpora. It has been effectively applied across a range of domains, including social media platforms (Li et al., 2024) and news media (Xian et al., 2024), to extract coherent themes and track discourse dynamics over time. Traditional topic modeling methods, such as Latent Dirichlet Allocation (Blei et al., 2003) or BERTopic (Grootendorst, 2022), often struggle with short and noisy text common to social media. To address these limitations, we adopted the TopicGPT framework (Pham et al., 2024), which integrates LLMs (GPT-4.1 in our context) into the topic modeling pipeline. TopicGPT generates candidate topics by prompting an LLM with representative document clusters, then refines these topics based on semantic consistency and distinctiveness. This approach enables the extraction of more coherent and human-interpretable topics from social media discourse.

4 Results

4.1 RQ1: Public Stance Changes

4.1.1 General Trend

We first showed the results and statistics in our final stance classification results in Table 3. Although after the assassination attempt, the number of posts is surging, the relative proportion indicates that after assassination, positive and neutral posts significantly rise and negative posts significantly drop, indicating that overall, the public tends to be much more positive after the political violence event.

Time Period	Favor	Against	Neutral
Before	6974 (14.9%)	28883 (61.8%)	10857 (23.2%)
After	21423 (28.3%)	23149 (30.5%)	31240 (41.2%)

Table 3: Stance Counts and Percentages Before and After

Further, we visualized the stance changes centering around Donald Trump in Figure 2 by taking *favor* as 1, *neutral* as 0, and *against* as -1. Analysis of X's stance toward Donald Trump also reveals a distinct temporal shift surrounding the July 13, 2024, assassination attempt (Figure 2a). Before July 13, stance scores averaged around -0.5. On the day of the event, these scores rise sharply to approximately -0.1, indicating a significant but temporary reduction in negativity. In the days following, the stance scores declines but stabilizes at about -0.3, suggesting a modest and persistent improvement compared to pre-event levels in a short time.

Disaggregating stance by political alignment reveals consistent trends across all groups, with no pronounced regional outliers (Figure 2b). Red States exhibited a slightly higher stance (around -0.45) compared to Blue States (approximately -0.60) before the event. Following the incident, both groups experienced noticeable positive shifts, with Red States peaking at 0 and Blue States reaching -0.15. Swing States followed a similar trajectory, reflecting a broadly consistent response across regions, regardless of ideological leanings. State-level stance trends (Figure 2c) for representative blue or red states like Texas, Florida, California, and New York further confirm this pattern, with all four experiencing comparable shifts, albeit with minor variations in magnitude.

4.1.2 User Level Analysis

To better understand the drivers behind the increase in favorable stance, we conduct a user-level analysis of two groups: (1) users who posted both before and after the assassination attempt, and (2) users who only posted after. This approach allows us to distinguish between individual-level attitude shifts and the influence of new participants. As shown in Table 4, among users active in both periods, 36.7% became more positive, compared to 26.4% who became more negative, and 36.9% who remained neutral or unchanged, indicating a net positive shift among continuing users. Among users who joined the conversation only after the event, a similar pattern emerges: 33.9% expressed a positive stance, outpacing the 26.5% with a negative stance. These findings suggest that the overall increase in positive stance is driven by both attitudinal shifts among existing users and the engagement of new users expressing supportive or empathetic views.

User sets	(Becoming) Favor	(Becoming) Against	Neutral/Unchanged
Retained Users	4216 (36.7%)	3038 (26.4%)	4241 (36.9%)
New Influx	13161 (33.9%)	10237 (26.5%)	15372 (39.6%)

Table 4: Stance Counts and Percentages Before and After

4.2 RQ2: Text Causal Analysis

To examine the assassination attempt's impact on public stance, we extend RQ1's descriptive analysis by applying the DiD model to isolate the event's causal effect, accounting for temporal fluctuations and concurrent factors.

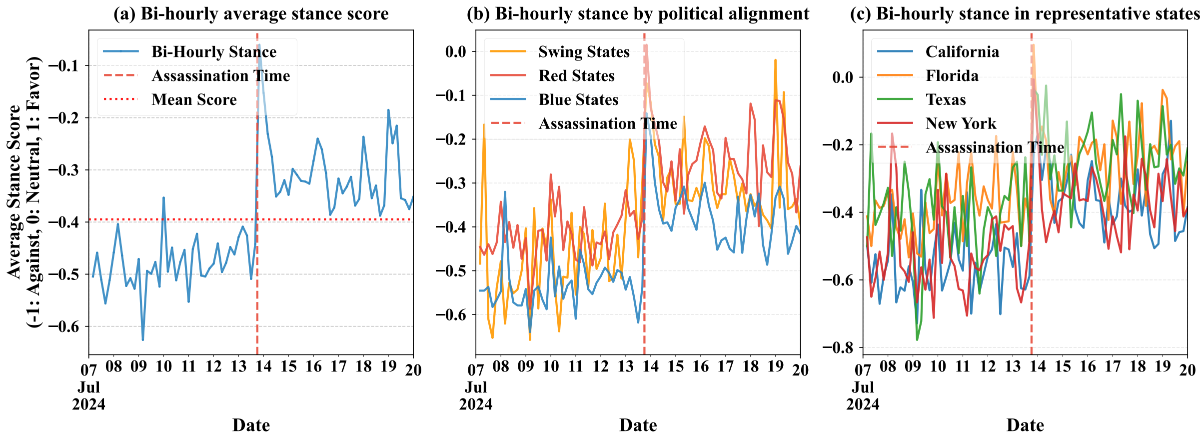

Figure 2: Temporal evolution of public stance following the attempted assassination.

	Dependent Variable: Averaged Public Stance Scores Toward Donald Trump		
	(1) Red vs. Blue States	**(2) Swing vs. Red States**	**(3) Swing vs. Blue States**
Intercept	-0.509 (0.035)***	-0.333 (0.029)***	-0.512 (0.027)***
Post	0.372 (0.048)***	0.419 (0.048)***	0.364 (0.044)***
Treatment	0.163 (0.043)***	-0.075 (0.086)	0.087 (0.040)*
Treatment_Post	0.040 (0.042)	-0.024 (0.057)	0.024 (0.150)
Controls	Yes	Yes	Yes
No. Observations	583	445	334
Adj. R-Squared	0.357	0.299	0.391

Table 5: DiD Major Estimation Results: Treatment vs. Control Groups

4.2.1 The Major Estimation Framework

DiD estimation results, shown in Table 5, assess the assassination attempt's effect on public stance across treatment-control group pairs. Model (1) designates red states as treatment and blue states as control; Model (2) uses swing states as treatment and red states as control; Model (3) compares swing states as treatment to blue states as control. The dependent variable in all models is the weighted public stance score toward Donald Trump.

The $Treatment\_Post$ coefficients across all three models are statistically insignificant, indicating that the assassination attempt did not disproportionately increase public stance scores toward Trump in treatment groups (red and swing states) relative to their control groups. Despite a general rise in stance across all groups, evidenced by significant positive $Post$ coefficients, the insignificant $Treatment\_Post$ terms suggest no differential stance shift.

These findings imply the assassination attempt did not intensify partisan affective polarization. Instead, the uniform increase in stance scores across models supports the sympathy hypothesis, suggesting a broad, non-polarized uptick in stance toward Trump post-event, unrelated to political affiliations.

4.2.2 Robustness Check

The placebo test, results documented in Table 6, assumes a fake intervention on July 11, 2024, two days before the actual event. This analysis aims to determine whether the observed effects on averaged public stance scores toward Donald Trump are attributable to the event itself or pre-existing trends and random noise. Across all models, the $Group\_Placebo$ coefficients are not statistically significant, indicating that the placebo did not produce a differential effect on public stance scores between treatment and control groups. This suggests that the observed effects in the main results are likely due to the actual assassination attempt rather than pre-existing trends or random noise. The significant $Placebo$ coefficients show a general upward trend in stance scores even before the event, but the lack of significance in the interaction terms confirms that this trend does not differ across groups in a way that mimics the main results. Thus, the main findings are robust and not driven by spurious factors.

4.3 RQ3: Topic Modeling

4.3.1 General Results

In this subsection, we described our general findings from topic modeling. In total, our topicGPT

Dependent Variable: Averaged Public Stance Scores Toward Donald Trump			
	(1) Red vs. Blue States	**(2) Swing vs. Red States**	**(3) Swing vs. Blue States**
Intercept	-0.483 (0.04)***	-0.301 (0.034)***	-0.487 (0.032)***
Treatment	0.161 (0.048)***	-0.070 (0.051)	0.100 (0.047)*
Placebo	0.271 (0.048)***	0.302 (0.048)***	0.260 (0.044)***
Treatment_Placebo	0.038 (0.045)	-0.032 (0.062)	0.005 (0.051)
Controls	Yes	Yes	Yes
No. Observations	583	445	334
Adj. R-Squared	0.315	0.244	0.325

Table 6: DiD Placebo Test Results: Treatment vs. Control Groups

Overall		**Before Assassination**		**After Assassination**	
Topic Name	**Number**	**Topic Name**	**Number**	**Topic Name**	**Number**
Elections, Voting, and Campaigns	24604	Elections, Voting, and Campaigns	11854	Political Violence and Extremism	20684
Political Violence and Extremism	21195	Allegations, Misconduct, and Scandals	8317	Elections, Voting, and Campaigns	12750
Political Leadership, Character, and Image	20283	Political Leadership, Character, and Image	7783	Political Leadership, Character, and Image	12500
Allegations, Misconduct, and Scandals	13072	Government Structure, Power, and Accountability	3662	Political Culture, Polarization, and Rhetoric	4966
Political Culture, Polarization, and Rhetoric	6789	Media, Journalism, and Public Perception	2849	Misinformation, Conspiracy Theories, and Extremism	4785

Table 7: Topic frequencies before and after the assassination

framework identified 26 high-level topics. Our top-5 topic list and its corresponding numbers for three time periods (overall, before and after assassination) are attached in Table 7, and the full topic list can be found in Table 12 in Appendix E.

What are people talking about on social media? As shown in Table 7, when people are talking about Trump, most posts center around *Elections, Voting, and Campaigns*, and other posts discuss Trump's political leadership or scandals. Other prominent themes include *Political Violence and Extremism*, with over 21,000 posts, especially after the assassination, and *Political Leadership, Character, and Image*, reflecting sustained attention on Trump's persona and conduct in office. Allegations, Misconduct, and Scandals also receive notable attention, with 13,072 posts. Although we could not verify the framework's accuracy quantitatively, the distribution of topics suggests that it produces coherent and interpretable results.

How do topics of interest change before and after the assassination attempt? In Table 7, we also display the number of topic changes after the assassination attempt. This result implies significant shifts in the thematic focus of posts before and after the assassination event. Prior to the assassination attempt, the most discussed topics were 'Elections, Voting, and Campaigns' and 'Allegations, Misconduct, and Scandals', indicating a strong public focus on the electoral landscape and ongoing controversies. 'Political Leadership, Character, and Image' also featured prominently, suggesting an interest in the personal and moral dimensions of political figures. Following the assassination attempt, the discourse became dominated by 'Political Violence and Extremism', reflecting heightened concern with the nature and implications of the attack. Elections and Political Leadership remained significant, but were accompanied by increased engagement with themes of 'Misinformation and Conspiracy Theories'.

Interestingly, while most themes remain relatively stable, some topics exhibit notable shifts in prominence. In particular, the top three rising topics, Political Violence, Political Leadership, and Misinformation, suggest a collective reorientation of public discourse in response to the assassination attempt. The surge in attention to Political Violence reflects heightened concern about the safety and stability of the political landscape, while the prominence of Misinformation further indicates that users were actively engaging in efforts to interpret, explain, or question the event, often through competing narratives and contested information. Conversely, the top three declining topics, Scandals, Government Structure, and Foreign Policy, suggest a temporary deprioritization of routine political controversies and institutional analysis. These results indicate that, in the immediate aftermath of the assassination attempt, public attention shifted away from systemic critiques and ongoing scandals toward more urgent concerns related to violence, leadership, and the reliability of information. This reallocation of focus reflects a broader pattern in crisis-driven discourse, where emotionally charged or high-stakes events tend to eclipse more procedural or policy-oriented discussions.

Overall, these results show how online discourse patterns shift following the assassination attempt. The patterns indicate a clear change in topic focus,

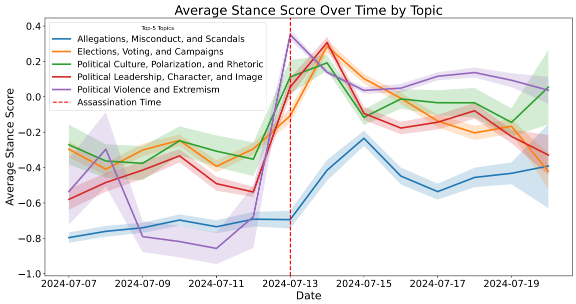

Figure 3: Performance evaluation of stance detection across different models.

with increased attention to the assassination event and supportive messages, while previously dominant topics like ‘scandals’ receive less attention. This suggests that significant political shock events, like an assassination attempt can temporarily redirect public discussion priorities on social media platforms.

4.3.2 Association with Stances

In this subsection, we examine the relationship between discussion topics and stance, asking which topics are most associated with the positive stance shifts identified in RQ1. To do so, we first analyzed the average stance scores for each topic before and after the assassination attempt. The changes of the top 5 topics are visualized in Figure 3, and the top three topics with the highest positive or negative stance score changes are presented in Table 8.

Topic	Score
Religion, Faith, and Politics	0.93±0.03
Political Violence and Extremism	0.83±0.05
Military, Defense, and National Security	0.75±0.05
Human Rights and Humanitarian Issues	-0.46±0.13
Crime, Justice, and Law Enforcement	-0.21±0.14
Energy Policy and Resources	-0.14±0.32

Table 8: Topics with Highest Positive/Negative Stance Score Changes, the ± represents 95% confidence intervals.

Among the five most frequent topics, we observe a general increase in stance scores, suggesting that the rise in positive stance is driven by posts across a range of themes. Notably, we find no evidence that shifts in negative stance are linked to polarization-related topics. The highest-scoring topics, Religion and Political Violence, reflect a broad outpouring of support, including prayers for Trump and collective condemnation of political violence. And in negatively changing topics, we witness non-polarizing topics and less stable variations. Taken together, these findings suggest that the increase in positive stance is grounded in affective, non-partisan discourse rooted in shared moral and emotional reactions, rather than ideological division.

5 Discussion

The findings of this study highlight the nuanced public response to the July 2024 assassination attempt on Donald Trump, with key insights emerging from stance detection, causal modeling, and topic modeling. We observe a significant but broadly uniform increase in positive stance across states and political affiliations, suggesting that the event elicited widespread empathy and a decline in negative sentiment. This pattern aligns with the sympathy hypothesis, which posits that acts of violence against prominent figures often generate public solidarity. The Difference-in-Differences (DiD) analysis further supports this interpretation, showing no significant interaction effects, indicating that the event did not exacerbate existing ideological divisions.

These findings complement those of Holliday et al. (2024), who reported a reduction in support for partisan violence and increased in-group attachment among Republicans, particularly MAGA Republicans. While their focus is on attitudinal shifts within partisan groups, our study emphasizes the broader stance landscape, revealing a national trend of reduced negativity that cuts across both political and geographic boundaries. Moreover, their observation that affective polarization remained unchanged is consistent with our finding that the stance shift, while positive, was non-polarizing.

Our topic modeling analysis reinforces these results by uncovering marked shifts in public discourse. Prior to the event, conversations centered on politically charged themes such as scandals. In the aftermath, discourse pivoted toward less contentious themes, including misinformation and expressions of concern captured in the political violence topic. These changes illustrate how high-profile crises can momentarily redirect public attention from divisive narratives to expressions of empathy and emotional engagement. Taken together, our findings and those of Holliday et al. (2024) underscore the complex interplay between stance, partisanship, and discourse during moments of political crisis. They point to both the unifying potential of such events and the persistent undercurrents of partisan identity that shape public reaction.

Limitations

This study encounters certain technical limitations with its LLM-based approaches. First, the Brandwatch API, used for sampling X data, may miss nuanced insights and overrepresent certain demographics (Blank, 2017). Second, the two-week analysis window, designed to isolate the assassination attempt's effect, restricts insight into longer-term stance trends. Finally, our NLP techniques have room for improvement: the aspect-based stance analysis and topic modeling, reliant on specific prompts and LLM testing, could be enhanced using advanced methods like Retrieval-Augmented Generation (RAG) (Lewis et al., 2020; Gao et al., 2023). Future research should address these technical gaps.

References

Joshua D. Angrist and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, Princeton, NJ.

Anthropic. 2024. Claude 3.5 sonnet. `https://www.anthropic.com/news/claude-3-5-sonnet`.

Lonna Atkeson and Cherie Maestas. 2012. *Catastrophic Politics: How Extraordinary Events Redefine Perceptions of Government*. Cambridge University Press.

Francesco Barbieri, Jose Camacho-Collados, Leonardo Neves, and Luis Espinosa-Anke. 2020. Tweeteval: Unified benchmark and comparative evaluation for tweet classification. *arXiv preprint arXiv:2010.12421*.

Jackson Bennett, Benjamin Rachunok, Roger Flage, and Roshanak Nateghi. 2021. Mapping climate discourse to climate opinion: An approach for augmenting surveys with social media to enhance understandings of climate opinion in the United States. *PLOS ONE*, 16(1):1–16.

Timothy Besley and Anne Case. 2002. Unnatural experiments? estimating the incidence of endogenous policies. *The Economic Journal*, 112(476):F231–F246.

Grant Blank. 2017. The digital divide among twitter users and its implications for social research. *Social Science Computer Review*, 35(6):679–697.

David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3(null):993–1022.

Johan Bollen, Huina Mao, and Xiao-Jun Zeng. 2011. Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1):1–8.

Richard A. Brody. 1991. *Assessing the President: The Media, Elite Opinion, and Public Support*. Stanford University Press, Stanford, CA. Hardcover.

David Card and Alan B. Krueger. 2000. Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania: Reply. *The American Economic Review*, 90(5):1397–1420.

Virginia Chanley. 2002. Trust in government in the aftermath of 9/11: Determinants and consequences. *Political Psychology*, 23:469–483.

Andrew J. Cherlin. 2021. White working-class support for trump. *Contexts*, 20(2):30–35. First published online May 26, 2021.

DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. Deepseek-v3 technical report. *Preprint*, arXiv:2412.19437.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.

R.E. Gilbert. 2013. Attempted assassination and presidential achievement: The case of ronald reagan. In *The world of biology and politics: Organization and research areas (Research in Biopolitics, Vol. 11)*, pages 161–183. Emerald Group Publishing.

Sandra González-Bailón, Ning Wang, Alejandro Rivero, Javier Borge-Holthoefer, and Yamir Moreno. 2014. Assessing the bias in samples of large online networks. *Social Networks*, 38:16–27.

Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.

Robert M. Groves, Floyd J. Fowler Jr., Mick P. Couper, James M. Lepkowski, Eleanor Singer, and Roger Tourangeau. 2009. *Survey Methodology*, second edition. Wiley Series in Survey Methodology. Wiley, Hoboken, NJ.

Marc Hetherington and Michael Nelson. 2003. Anatomy of a Rally effect: George W. Bush and the War on Terrorism. *PS: Political Science and Politics*, 36(1):37–42.

Derek E. Holliday, Yphtach Lelkes, and Sean J. Westwood. 2024. The july 2024 trump assassination attempt was followed by lower in-group support for partisan violence and increased group unity. *Proceedings of the National Academy of Sciences*, 121(49):e2414689121.

Leonie Huddy. 2003. Group identity and political cohesion. In David O. Sears, Leonie Huddy, and Robert Jervis, editors, *Oxford Handbook of Political Psychology*, pages 511–558. Oxford University Press.

Bengt Johansson, David Nicolas Hopmann, and Adam Shehata. 2021. When the rally-around-the-flag effect disappears, or: when the COVID-19 pandemic becomes "normalized". *Journal of Elections, Public Opinion and Parties*, 31(sup1):321–334.

Jonathan McDonald Ladd and Gabriel S. Lenz. 2009. Exploiting a rare communication shift to document the persuasive power of the news media. *American Journal of Political Science*, 53(2):394–410.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, Red Hook, NY, USA. Curran Associates Inc.

Lingyao Li, Zihui Ma, Lizhou Fan, Sanggyu Lee, Huizi Yu, and Libby Hemphill. 2024. Chatgpt in education: A discourse analysis of worries and concerns on social media. *Education and Information Technologies*, 29(9):10729–10762.

Stuart Mackintosh. 2014. Crises and paradigm shift. *The Political Quarterly*, 85.

John E. Mueller. 1970. Presidential popularity from Truman to Johnson. *American Political Science Review*, 64(1):18–34.

Gary O'Donoghue, Bernd Debusmann, and Matt Murphy. 2024. Secret service faces questions after trump assassination attempt.

OpenAI. 2024. Hello gpt-4o. *OpenAI Blog*.

Charles W. Ostrom Jr. and Dennis M. Simon. 1985. Promise and performance: A dynamic model of presidential popularity. *The American Political Science Review*, 79(2):334–358.

R. Pandarachalil, S. Sendhilkumar, and G. S. Mahalakshmi. 2015. Twitter sentiment analysis for large-scale data: An unsupervised approach. *Cognitive Computation*, 7:254–262.

Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1–2):1–135.

Chau Minh Pham, Alexander Hoyle, Simeng Sun, Philip Resnik, and Mohit Iyyer. 2024. TopicGPT: A prompt-based topic modeling framework. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2956–2984, Mexico City, Mexico. Association for Computational Linguistics.

David B. Pillemer. 2000. *Momentous Events, Vivid Memories*. Harvard University Press, Cambridge, MA.

POLITICO. 2024. Swing states 2024: Battleground states map, list & electoral votes. Accessed: 2024-12-12.

Robert D. Putnam. 2007. E Pluribus Unum: Diversity and Community in the Twenty-first Century. The 2006 Johan Skytte Prize Lecture. *Scandinavian Political Studies*, 30(2):137–174.

Miguel Rueda, Guy Grossman, and Shuning Ge. 2023. Do more disaggregated electoral results deter aggregation fraud? Research Paper 2023-6, MIT Political Science Department.

Jason Sokol. 2018. *The Heavens Might Crack: The Death and Legacy of Martin Luther King Jr.*, illustrated edition. Basic Books, New York, NY.

Henri Tajfel and John C. Turner. 2004. The social identity theory of intergroup behavior. In John T. Jost and Jim Sidanius, editors, *Political Psychology: Key Readings*, pages 276–293. Psychology Press.

U.S. Census Bureau. 2025. U.s. census bureau homepage. `http://www.census.gov/`. Accessed: May 17, 2025.

Claire Whitlinger and Joe Fretwell. 2019. Political assassination and social movement outcomes. *Sociological Perspectives*, 62(4):455–474.

Nicholas J. G. Winter. 2023. Hostile sexism, benevolent sexism, and american election. *Politics & Gender*, 19(2):427–456. Accessed: May 17, 2025.

Lu Xian, Lingyao Li, Yiwei Xu, Ben Zefeng Zhang, and Libby Hemphill. 2024. Landscape of large language models in global english news: Topics, sentiments, and spatiotemporal analysis. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, pages 1661–1673.

A Descriptive Statistics of Our X Dataset

In this section, we describe the general statistics of our analyzed X dataset. In total, our dataset consists of 122,526 posts and 68,600 unique users, with 1.79 posts per user. The detailed number of posts along time is shown in Table 9.

Additionally, we also attached three examples of our 'favor', 'against', and 'neutral' posts in Table 10.

B Prompt Template for Stance Detection

The following shows the prompt design for the stance detection given an X post.

Date	Post Count
2024-07-07	4989
2024-07-08	6560
2024-07-09	7248
2024-07-10	8615
2024-07-11	7181
2024-07-12	8392
2024-07-13	7128
2024-07-14	18343
2024-07-15	15259
2024-07-16	12056
2024-07-17	10170
2024-07-18	7780
2024-07-19	7693
2024-07-20	1112

Table 9: Number of Posts Per Day

System prompt:
You are a social science expert specializing in stance detection of social media posts. Your task is to analyze tweets about Donald Trump and classify their stance.

Prompt template: Background: On July 13, 2024, Donald Trump — former U.S. president and Republican nominee for the 2024 election — survived an assassination attempt during a rally in Butler, Pennsylvania. The attacker, Thomas Matthew Crooks, fired eight rounds from a rooftop using an AR-15–style rifle. Trump was wounded in the upper right ear.

Task: You will be shown a tweet. It might be related to this event. Only take the background into consideration when needed. You task is to classify the tweet's stance toward *Donald Trump* into one of the following categories:

- **Favor**: Shows support or approval of Trump.
  - Examples: praise, endorsement, expressions of sympathy or admiration.
- **Against**: Shows opposition or criticism of Trump.
  - Examples: blame, ridicule, hostile commentary, or celebration of the attack.
- **Neutral**: No clear stance or mixed opinions.
  - Examples: factual reporting, commentary on the event or others (e.g. the shooter, supporters), or conflicting views.

Important Guidelines:

1. Focus **only** on stances directed at Donald Trump.
2. Ignore stances directed at his family, supporters, or unrelated topics.
3. If a tweet includes multiple stances, choose the most dominant stance.
4. Your output must be only one word: favor, neutral, or against.
5. Do not include explanations or extra content — only the label.

Demonstration:
Tweet: I may not agree with Trump's politics, but no one deserves to be shot like that.
Output: favor

Tweet: Glad he got what he deserved. One less threat to democracy.
Output: against

Tweet: That shooter was a lunatic. What a tragedy all around.
Output: neutral

Now classify the following tweet: {}

C Geographical Distribution of Public Stance Across United States

In this subsection, we present our visualization of the geographical distribution of public stances across the United States before and after the assassination attempt, as displayed in Figure 4.

D Correlation Analysis between Demographic Factors and Average Public Stance Score

As a complement, we conducted correlation analyses (as shown in Table 11) to examine the relationship between socioeconomic and demographic factors and average public stances toward Donald Trump before and after the assassination attempt. States with higher Republican voter shares ($Republican_R$) displayed consistently more positive stance, with correlations of 0.448 before and 0.538 after the incident. Conversely, higher proportions of Democratic voters ($Democrat_R$) were associated with more negative stance, reflected in correlations of -0.367 and -0.462, respectively. Economic factors further highlight disparities in stance: areas with higher poverty rates ($Household_Below_Poverty_R$) and uninsured populations ($No_Insurance_R$) exhibited less negative stance, whereas states with higher median incomes (-0.451 before, -0.568 after) and bachelor's degree attainment rates (-0.495 before, -0.526 after) expressed increasingly negative stance toward Trump.

While these correlation results provide valuable insights, it is important to note that they are based on data from only 50 states, which limits the strength and generalizability of the findings. The relatively small sample size means that some relationships may be influenced by regional outliers or other confounding factors not captured in the analysis. However, the consistency of key trends, such as the positive association between Republican vote

Label	Example Post
Favor	Don Jr. calls out the media for their constant lies about Trump. Are they even capable of telling the truth?
Neutral	I have been critical of Trump and MAGA populism because I am a conservative, not because I want to give ground to the Left. I want the Right to do better. I believe we should hold ourselves to higher standards. Conservatism must remain committed to our constitutional order.
Against	Public service announcement: Project 2025 puts all of our heads on the guillotine. We can beat Trump together.

Table 10: Examples of Posts by Stance Label

share and stance, and the negative correlation with median income and bachelor's degree attainment, partially validates the reliability of our stance detection pipeline. These observed patterns align with broader sociopolitical divides documented in previous research, suggesting that our methodology is effective in capturing meaningful stance shifts within the dataset. This preliminary validation reinforces the pipeline's utility for analyzing larger-scale social media data and identifying nuanced patterns in public discourse.

Table 11: Correlation analysis of socioeconomic and demographic factors with stance before and after the assassination. Here, $\_R$ refers to rate.

Variable	Before	After
Total_Population	-0.193	-0.357
LUM_Race	0.173	0.018
Median_income	-0.451	-0.568
GINI	0.303	0.199
Democrat_R	-0.367	-0.462
Republican_R	0.448	0.538
No_Insurance_R	0.526	0.585
Household_Below_Poverty_R	0.560	0.595
HISPANIC_LATINO_R	-0.076	-0.125
White_R	-0.281	-0.177
Black_R	0.347	0.236
Indian_R	0.076	0.230
Asian_R	-0.337	-0.447
Under_18_R	0.252	0.381
Bt_18_44_R	0.033	-0.102
Bt_45_64_R	-0.246	-0.237
Over_65_R	-0.093	-0.040
Male_R	0.094	0.120
Bachelor_R	-0.495	-0.526
Population_Density	-0.066	-0.190
Unemployed_R	0.358	0.300

E Topic List

Our full topic list and each topic's corresponding number of posts are shown in Table 12.

F Artifact Use

Here we use DeepseekV3-0324 (DeepSeek-AI et al., 2025), which is released under the MIT license. We only use it for academic research purpose, following the intended use.

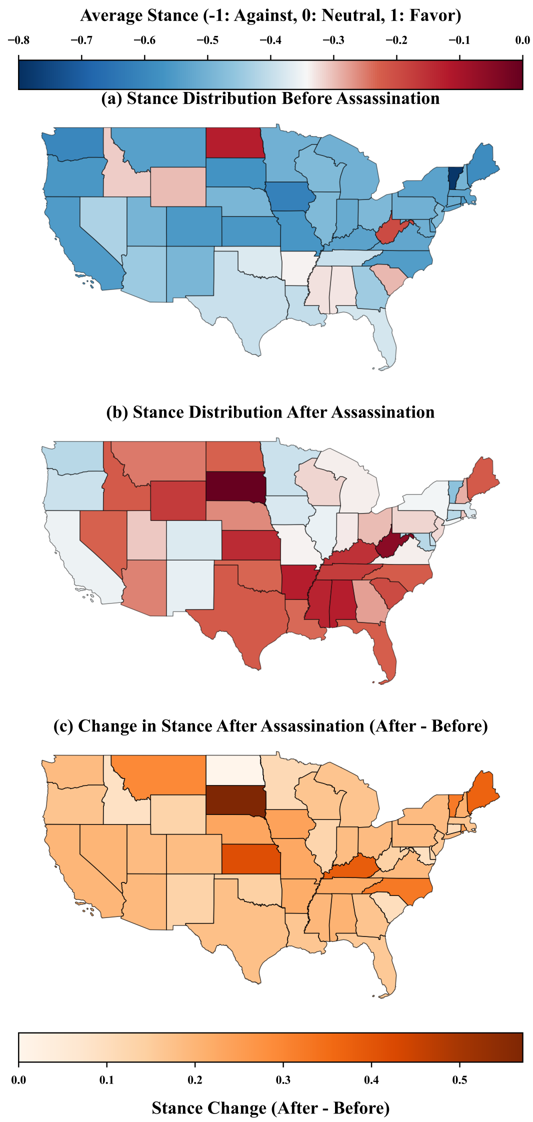

Figure 4: Geographic distribution of the public stance before and after the assassination.

Topic	Count
Elections, Voting, and Campaigns	24604
Political Violence and Extremism	21195
Political Leadership, Character, and Image	20283
Allegations, Misconduct, and Scandals	13072
Political Culture, Polarization, and Rhetoric	6789
Media, Journalism, and Public Perception	6220
Misinformation, Conspiracy Theories, and Extremism	6031
Government Structure, Power, and Accountability	5596
Abuse of Power and Government Corruption	3129
Foreign Policy and International Relations	3035
Domestic Economy, Fiscal Policy, and Employment	2314
Religion, Faith, and Politics	1747
Abortion and Reproductive Rights	1409
Civil Rights, Social Justice, and Equality	1167
Immigration and Border Security	1147
Crime, Justice, and Law Enforcement	1010
Social Media and Digital Platforms	973
Military, Defense, and National Security	796
Health Care and Public Health	427
Academic Freedom, Free Speech, and Censorship	334
Science, Technology, and Innovation Policy	332
None	244
Cybersecurity, Data Privacy, and Information Warfare	205
Education Policy and Reform	150
Human Rights and Humanitarian Issues	141
Energy Policy and Resources	91
Climate Change and Environmental Policy	85

Table 12: Topic Counts

G GenAI Statement

The authors used Cursor for coding support and ChatGPT for writing revisions as needed.