

Asymmetric Dynamics of Partisan Warriors in YouTube Comments

Keyeun Lee¹, Sang Jung Kim²

¹Seoul National University

²University of Iowa

kieunp@snu.ac.kr, sangjung-kim@uiowa.edu

Abstract

Cross-cutting commenting on social media is often imagined as a path to deliberation, yet exposure to opposing views frequently fuels hostility. To explain this dynamic, we introduce the concept of partisan warriors—commenters who cross ideological lines primarily to launch uncivil attacks against out-partisans. We analyze a large corpus of YouTube comments (N= 1,854,320) surrounding the 2024 U.S. second presidential debate. After filtering for toxicity and active participation, we use large language models to identify attack targets and operationalize partisan warrior behavior. Our analysis highlights four dynamics. First, cross-cutting commenters do not exhibit greater civility than those who remain within their ideological camps (RQ1). Second, audience reactions diverge by ideology: conservative audiences tended to reward hostile attacks on out-group leaders, whereas liberal audiences offered no comparable incentives and at times penalized such attacks (RQ2). Third, partisan warriors are notably more prevalent in conservative-leaning channels than in liberal ones; commenters restricted to conservative spaces were substantially more likely to engage in partisan warrior behavior compared to their liberal-only counterparts (RQ3). Finally, regarding environmental triggers, robustness checks suggest that this participation is an ecological phenomenon driven largely by channel-level heterogeneity rather than transient responses to individual video titles (RQ4). By shifting attention from the prevalence of incivility to its targets, rewards, and structural drivers, this study advances understanding of how partisan hostility is enacted and sustained in online spaces.

Introduction

Social media platforms are widely criticized for reinforcing echo chambers and filter bubbles, limiting exposure to opposing perspectives along partisan lines (Bakshy, Messing, and Adamic 2015; Flaxman, Goel, and Rao 2016; Sunstein 2018). Against this backdrop, cross-cutting commenters—those who engage across partisan divides—are often seen as embodying a more democratic ideal: the promise of deliberation, tolerance, and dialogue (Garrett 2009; Mutz 2002; Wojcieszak and Mutz 2009). Yet in practice, cross-commenting often manifests as antagonism and incivility toward out-partisans, raising questions about whether such encounters bridge divides or deepen partisan hostility. Rather

than bridging divides, cross-cutting commenters may function as performers of partisan hostility (Çetinkaya et al. 2025; Wu and Resnick 2021; Xia et al. 2025).

While the literature on incivility has primarily focused on the degree of hostility in individual posts or comments (Coe, Kenski, and Rains 2014; Rossini 2022; Theocharis et al. 2016; Yudhoatmojo, Balci, and Blackburn 2024), a fuller understanding requires examining its targets (Matsui, Murayama, and Yoshida 2024). Incivility toward political out-groups reflects proactive denigration, making it important to distinguish between commenters who remain within ideologically aligned spaces and those who deliberately cross partisan lines to criticize opponents (Kim 2018).

Despite growing interest in political incivility and cross-cutting contact, a key gap remains: we lack a systematic way to identify commenters who cross ideological boundaries primarily to attack out-partisans. Prior work has examined the prevalence of incivility and the structure of cross-cutting interactions (Wu and Resnick 2021; Çetinkaya et al. 2025), but it has paid little attention to the targets of incivility or to the audience responses that may incentivize and sustain such behavior.

We address this gap by introducing the concept of *partisan warriors*—commenters who leave uncivil remarks directed at out-partisans in opposing channels. This conceptualization matters because crossing into hostile territory is atypical in selective-exposure environments and signals a confrontational form of partisan expression (Garrett 2009; Rossini 2022; Stroud 2010). Studying partisan warriors allows us to move beyond measuring the amount of incivility to examining who is targeted and how partisan identity is performed in cross-cutting contexts. In this paper, we compare the prevalence of partisan warriors among commenters active in conservative channels, cross-commenters, and commenters active in liberal channels on YouTube.

Beyond commenters themselves, we also consider the broader platform dynamics that shape partisan conflict. The toxicity of videos may encourage or deter partisan warriors, while the ways in which audiences reward or sanction comments reveal whether partisan hostility is amplified or constrained (Theocharis et al. 2016; Gao et al. 2024). Together, these dynamics clarify not only the behavior of commenters but also the broader processes through which platforms sustain or exacerbate partisan conflict.

Building on this motivation, our study seeks to unpack the layered relationships that structure partisan discourse: the channels, the commenters who engage with them, the videos channels host, the comments those videos generate, and the audiences who react. This multi-level perspective allows us to move beyond isolated analyses of incivility and instead trace how partisan hostility is produced, expressed, and rewarded across the platform. Accordingly, we focus on four research questions:

RQ1 Is there a difference in the level of incivility between cross-cutting commenters and commenters who remain within ideologically aligned spaces?

RQ2 How do audiences in conservative and liberal channels reward or sanction uncivil comments directed at out-partisans?

RQ3 How does the prevalence of partisan warriors differ across conservative-channel-only commenters, cross-commenters, and liberal-channel-only commenters?

RQ4 How does the toxicity of videos influence the participation of partisan warriors in comment sections?

Related Work

Cross-cutting Exposure and Incivility

Cross-cutting exposure—the encounter with opposing political views—has long been theorized as essential for democratic deliberation (Mutz 2002). Online platforms make such contact observable at scale, as users encounter dissonant perspectives in comment threads and social media feeds (Bakshy, Messing, and Adamic 2015). Yet its democratic value remains contested. Some studies suggest that exposure fosters tolerance and reduces extremity (Mutz 2002), while others find it provokes defensiveness, selective avoidance, or withdrawal (Stroud 2010). More recent work shows that contact with opposing views can heighten political polarization (Bail et al. 2018).

These tensions are especially visible in comment sections, where cross-cutting engagement often devolves into antagonism and incivility rather than reasoned dialogue (Masullo, Tenenboim, and Lu 2023). In some cases, both deliberation and incivility rise together, suggesting that cross-cutting encounters are double-edged (Rossini and Maia 2021). Recent work on Twitter similarly shows that cross-partisan interactions are relatively rare and, when they occur, tend to be negative rather than constructive (Çetinkaya et al. 2025).

Prior work has documented strong connections between toxic or hostile language in political discourse and affective polarization on social media (Lerman et al. 2024). Interactions between users with opposing ideological affiliations are more likely to exhibit negativity and toxicity, consistent with affective polarization in online exchanges (Lerman et al. 2024; Simchon, Brady, and Van Bavel 2022). Recent computational research further emphasizes the role of toxic and uncivil speech in shaping polarized dynamics in cross-partisan interactions. Large-scale audits of social media ranking systems show that engagement-based algorithms can amplify hostile and divisive content, motivating interest in alternative ranking approaches that prioritize prosocial or

constructive interactions (Milli et al. 2024). Related efforts, including the Twitter Prosocial Ranking Challenges, examine whether re-ranking content based on prosocial attributes can reduce exposure to toxic speech and improve the quality of cross-partisan engagement (Human-Compatible AI Institute 2024; Saltz, Jalan, and Acosta 2024). These studies highlight how platform design and content ranking shape whether exposure to opposing viewpoints leads to deliberation or incivility.

Beyond explicitly political contexts, prior research shows that the discourse of highly engaged partisans is often characterized by heightened toxicity even when political topics are not salient (Mamakos and Finkel 2023). Such findings further motivate a behavioral focus on incivility, as antagonistic rhetoric may emerge independently of issue content or informational disagreement. Uncivil cross-partisan engagement can be understood as an expression of partisan identity rather than a response to specific political claims.

Building on this scholarship, our work directly compares the incivility of cross-cutting commenters with that of commenters operating within partisan-aligned channels. By focusing on observable rhetorical behavior, we provide systematic evidence on whether cross-cutting engagement is uniquely associated with heightened incivility.

Audience Reactions and the Rewarding of Incivility

Comments from both cross-cutting commenters and those in partisan-aligned channels are not consumed in isolation; they are evaluated by audiences whose responses—likes, up-votes, and replies—signal which voices are legitimized within partisan communities. Prior work shows that incivility can be rewarded as authentic or entertaining (Rossini 2022), but it can also be sanctioned as norm-violating (Theocharis et al. 2020). Platform affordances further shape these dynamics: algorithmic amplification often privileges “engaging” content regardless of its tone, thereby elevating hostile participation (Ribeiro et al. 2020).

Recent research adds to this picture by modeling how different categories of incivility affect conversational outcomes and user engagement (Gao et al. 2024), suggesting that audiences play an active role in sustaining or constraining hostile discourse. Our contribution lies in examining whether partisan hostility is rewarded differently when comments are directed at political ingroups or outgroups, using LLMs to attribute incivility to its partisan targets and to connect these patterns with audience responses.

Targets of Incivility and Partisan Warriors

Research on political incivility has largely examined its forms and prevalence. Early typologies distinguished insults, name-calling, and pejorative speech (Papacharissi 2004; Coe, Kenski, and Rains 2014), while later work emphasized its multidimensional character (Stryker, Conway, and Danielson 2016; Pendzel et al. 2024). Computational social science has extended these approaches by operationalizing toxicity for large-scale detection, often through machine learning models such as the Perspective API, which assigns toxicity scores based on linguistic cues (Almerekhi et al. 2019; Gao et al. 2024).

Far fewer studies investigate who is targeted by incivility. Beyond identification, incivility also serves functional roles when directed at specific groups: it can polarize by delegitimizing opponents (Gervais 2015; Theocharis et al. 2020), but also mobilize by energizing supporters and signaling authenticity (Rossini 2022). To advance this line of work, we leverage large language models (LLMs), which enable scalable annotation of incivility across millions of comments while preserving contextual nuance. Unlike dictionary approaches or toxicity classifiers, LLMs can incorporate surrounding conversational context and be prompted to recognize who is targeted, allowing us to distinguish hostility toward out-partisans from generic uncivil speech. Recent work has begun to demonstrate the feasibility of using LLMs for nuanced annotation tasks in social media research (Farjam, Meyer, and Lohkamp 2025), underscoring their potential for advancing incivility studies.

We introduce the concept of **Partisan Warriors (PWs)**: commenters who leave uncivil remarks directed at out-partisans in opposing channels. Conceptually, PWs are closely related to political trolls, as both engage in antisocial behavior toward political opponents and contribute to polarized political discussions on social media (Fichman and Akter 2023). Prior work on political trolling documents practices such as provocation and the disruption of political discourse through uncivil interactions with opposing partisans (Fichman and Akter 2023; Simchon, Brady, and Van Bavel 2022). At the same time, political trolling is often operationalized to include the dissemination of misinformation or disinformation, as well as strategic or deceptive behavior intended to manipulate political conversations (Fichman and Akter 2023; Blinder and Ware 2024). In contrast, our operationalization of partisan warriors does not incorporate the spread of mis- or disinformation. Instead, we focus on observable rhetorical behavior on social media, irrespective of informational accuracy or strategic intent.

Additionally, while studies of political trolls frequently distinguish between domestic and foreign actors, particularly in the context of coordinated influence campaigns (Blinder and Ware 2024), we do not classify commenters by origin. Our analysis centers on behavioral manifestations of incivility toward partisan out-groups, rather than assumptions about identity, motivation, or geographic provenance.

Crossing into hostile territory is rare in selective-exposure environments, yet those who do so often embody a confrontational form of partisan expression. Similar patterns are noted in prior work on cross-partisan engagement on YouTube, where conservatives were more likely to comment on liberal channels (Wu and Resnick 2021). Our contribution is to formalize this behavior with the concept of partisan warriors and measure its prevalence across different types of commenters, moving from isolated cases to a systematic framework for identifying targeted incivility.

Video Content, Toxicity, and Commenters

Channels are not neutral backdrops as the videos they host shape the discursive climate and the commenters they attract. Prior research has demonstrated that content frames shape political attitudes and conversational tone. Studies of

YouTube in particular find that algorithmic curation and radicalizing content can steer users toward more extreme forms of engagement (Ribeiro et al. 2020). Toxic or inflammatory videos create an environment that may embolden hostile participation while deterring deliberative voices.

Jiang, Robertson, and Wilson (2019) demonstrate that partisanship and misinformation on YouTube interact with commenting behaviors, reinforcing divides. Nevertheless, there is a dearth of work systematically connecting video toxicity to the composition of commenters. Understanding this link is essential for explaining why some channels attract partisan warriors more than others. We extend this line of research by linking video-level toxicity to the presence of partisan warriors in comment sections, leveraging LLMs to scale and nuance toxicity detection at the content level.

Data and Methods

Data Collection

Channel Selection and Classification Our study leveraged the US Partisan Media Dataset developed by Wu and Resnick (2021), which provides a comprehensive classification of YouTube channels based on their political orientation. From this dataset, we selected channels that had been previously classified by Media Bias/Fact Check (MBFC)¹ as either left-leaning or right-leaning sources. The final sample comprised 551 channels, with 293 classified as left-leaning and 258 as right-leaning.

Video Data Collection Video data were collected using the YouTube Data API V3 during a strategically selected timeframe surrounding the second U.S. presidential debate of 2024 (September 10, 2024). We focused on the week immediately following the debate because such high-profile political events act as a common external shock, triggering heightened engagement and rapid shifts in discourse. This window provides a unique opportunity to examine how commenters’ behaviors—such as patterns of toxicity and attack targets—evolve in response to the same political stimulus. To ensure content relevance, videos were included only if they contained “Trump” or “Harris” in their titles or descriptions, directly linking the content to the two main candidates. This approach yielded a total of 2,920 videos across 160 channels (86 left-leaning and 74 right-leaning channels), with left-leaning channels contributing 1,710 videos and right-leaning channels contributing 1,210 videos.

	Total	Liberal channels	Conservative channels
Channels	160	86	74
Videos	2,920	1,710	1,210
Comments	1,854,320	1,294,519	559,801

Table 1: Dataset overview by channel leaning.

¹<https://mediabiasfactcheck.com/>

Comment Data Collection We collected comment data using the YouTube API during a strategically defined temporal window to capture immediate reactions and discussions surrounding the presidential debate. The collection period spanned two weeks (September 10-24, 2024), beginning with the debate itself and extending through the subsequent week when related videos were published and comments actively accumulated. This timeframe ensured comprehensive coverage: we captured both debate-related videos uploaded in the week following the event and allowed an additional week for viewer engagement. This approach is particularly important as early commenters often play a disproportionate role in shaping video discourse and represent the most politically engaged participants (Xiao and Khazaei 2019; Tan et al. 2016; Miyazaki et al. 2024). By restricting our window to this two-week period, we focused on the most salient and immediate forms of electoral engagement.

We conducted data retrieval retrospectively between October 26 and October 28, 2024, approximately one month after the defined analysis window. Prior research on moderation bias suggests that although moderation decisions are generally unbiased when controlling for hate speech and extremeness, certain political spaces exhibit a higher prevalence of toxic attributes, which can lead to disproportionate removal rates for conservative content (Jiang, Robertson, and Wilson 2020). Although the one-month gap between the analysis window and data retrieval could have been influenced by platform moderation and channel-level curation, our dataset reflects the “surviving” discourse—namely, comments that remained visible at the time of collection.

For each video, we collected only top-level (root) comments that are not replies to other comments; these reflect viewers’ initial reactions to the video content and serve as the starting points of discussion threads. The final dataset comprised 1,854,320 comments from 677,238 unique commenters. As shown in Table 1, 1,294,519 comments were written on videos from liberal channels, while 559,801 comments were written on videos from conservative channels. While this disparity partly reflects differences in the number of videos per channel type, it is also influenced by the differential moderation dynamics described above. Consequently, our analysis likely provides a conservative, lower-bound estimate of the true extent of toxicity and partisan hostility, as the most extreme instances were presumably removed prior to our data collection.

Cross-Cutting Commenter Identification

To categorize commenters, we followed prior work that classifies users based on their activity across ideologically distinct channels (Han et al. 2023). We distinguished between liberal-only commenters (who posted comments exclusively on liberal-leaning channels), conservative-only commenters (who posted comments exclusively on conservative-leaning channels), and cross-cutting commenters (who posted comments on both channels). Table 2 shows the distribution of commenters and comments by type. Notably, although cross-cutting commenters represented just 11.5% of all individuals, they produced fully one-third of all comments, underscoring their outsized role in shaping discourse across

Commenter Type	Commenters	Comments
Only-Liberal	411,893 (60.8%)	919,376 (49.6%)
Only-Conservative	187,365 (27.7%)	323,212 (17.4%)
Cross-Cutting	77,980 (11.5%)	611,732 (33.0%)
Total	677,238	1,854,320

Table 2: Distribution of commenters and comments by commenter type.

partisan boundaries. To further assess the robustness of this categorical classification, we also examined a continuous measure of cross-ideological engagement based on the fraction of comments posted on liberal-leaning channels; the results, reported in Appendix, support the validity of our grouping strategy.

Detecting Attack Targets in Comments using LLMs

Data Filtering Pipeline To focus our analysis on persistent patterns of political incivility, we applied a two-step filtering pipeline to the raw dataset.

First, we isolated comments most likely to contain hostile or derogatory language by measuring toxicity using Perspective API’s toxicity model (Lees et al. 2022). Consistent with prior work on political incivility (Gervais, Dye, and Chin 2025), we applied a 90th percentile threshold (≥ 0.578) to retain only the most uncivil subset of discourse.

Second, to ensure our analysis captured behavioral patterns rather than isolated incidents, we restricted the sample to *active users*, defined as those who posted three or more comments that met the toxicity threshold. This step is crucial for establishing a reliable user-level stance in subsequent sections; a minimum of three data points prevents the classification of users based on a single, potentially anomalous outlier (Wu and Resnick 2021).

The intersection of these two filters—high toxicity and repeated engagement—yielded a final dataset of 62,598 comments for attack target analysis. While this excludes non-toxic comments and one-off posters, this focused sampling strategy is necessary to reliably model the specific phenomenon of sustained partisan hostility.

Model Selection We employed GPT-4.1, a state-of-the-art large language model, to identify attack targets in each comment. LLMs offer advantages for this task given their ability to process contextual nuance at scale, particularly important for analyzing contemporary political discourse (Heseltine and Clemm Von Hohenberg 2024).

Taxonomy Development Prompt Design To capture the nuances of political incivility during the election cycle, we developed an 18-category taxonomy that extends Rossini (2021)’s framework. The taxonomy was designed to reflect the personality-driven nature of contemporary U.S. politics

Comment	Channel Name	Channel Leaning	Video Title	Predicted Attack Target(s)
“Hey Trump, don’t you dare grab that hag’s hand”	Fox News	Conservative	Doocy: Kamala Harris has some explaining to do	Kamala Harris
“Notice they always used ‘apparent assassination attempt’ because he was a radical left Biden/Harris supporter brainwashed by the radical far left media garbage propaganda”	MSNBC	Liberal	Journalist who interviewed suspect of apparent Trump assassination attempt last year speaks out	Supporters of Kamala Harris, Mainstream media
“Keep your mouth shut, puppet!”	The Washington Examiner	Conservative	“I’m Talking Now”: Trump Quotes Harris After She Interrupts #shorts	Unclear
“Claiming Haitians eat people’s pets is basically calling them less than human. DonOld is utterly deranged. Maybe if he stopped binging right-wing trash TV, he wouldn’t come off sounding like such an idiot.”	CNN	Liberal	“Knock It Off”: Ohio AG Reacts to Uptick in Threats After JD Vance’s False Claims About Immigrants	Donald Trump, Mainstream Media

Table 3: Examples of attack target classification of comment.

and to distinguish between attacks on political leaders, their supporters, and broader groups. Categories include individual political actors (e.g., Donald Trump, Kamala Harris, Joe Biden), political groups (e.g., Republicans, Democrats, conservatives, liberals), protected or marginalized groups (e.g., immigrants, racial or gender identity groups, religious communities), institutional actors (e.g., debate moderators, mainstream media, government bodies), and an *Unclear/Other* category for residual cases. To implement this taxonomy, we prompted GPT-4.1 to classify only the *Attack target(s)*—entities explicitly or implicitly addressed with uncivil language—while excluding civil criticism or neutral mentions. Also, we adopted a conservative contextual resolution strategy for vague references (e.g., “he,” “they”) (Gan, Poesio, and Yu 2024). The LLM was permitted to consult video metadata only when strong evidence supported a clear resolution; otherwise, the reference was coded as *Unclear*. This procedure ensured consistent operationalization of the taxonomy while avoiding over-interpretation. Examples of comment attack target classifications can be found in Table 3, and the GPT-4.1 prompt example used for attack target detection is provided in the Appendix.

Validation To establish reliability, two authors trained in political communication independently coded 100 comments, yielding a Jaccard similarity of 0.865, indicating strong human agreement. The Jaccard Index was chosen because it is well-suited for multi-label classification tasks, where comments may include multiple targets (Gouk, Pfahringer, and Cree 2016; Levandowsky and Winter 1971). We then compared human and LLM outputs on a stratified sample of 626 comments (1% of the dataset), testing temperature values from 0.1 to 0.5 with five iterations each. Agree-

ment peaked at temperature 0.3 ($Jaccard = 0.724$), demonstrating substantial human–LLM alignment.

Final Setup. Based on these results, we fixed the decoding temperature at 0.3 for all subsequent classifications, ensuring an optimal trade-off between consistency and contextual flexibility. This validation confirms that our LLM-based approach provides a scalable and reliable method for mapping incivility targets in large-scale political discussions. As shown in Appendix Figure 8, the distribution of attack targets also varies by commenter type depending on the channel in which they posted.

Operationalizing Partisan Warriors

Commenter Type	Stance Label(%)		
	Anti-Lib.	Anti-Con.	Unclear
Only-Liberal	6.9	92.5	0.6
Only-Conservative	69.2	30.0	0.8
Cross-cutting	29.3	69.6	1.1

Table 4: Cross-tabulation of commenter type and inferred stance.

Stance Classification We first classified commenters according to the partisan alignment of their uncivil attacks. Each attack target identified by the LLM was mapped into an ideological polarity. Categories referencing Supporters of Donald Trump, Donald Trump, Republicans, or conservatives in general (IDs 1, 4, 5, and 8) were coded as **anti-conservative attacks**, while those referencing Supporters

of Kamala Harris, Kamala Harris, Democrats, or liberals in general (IDs 2, 3, 6, and 9) were coded as **anti-liberal attacks**. Each comment could increment both dimensions if multiple categories were targeted simultaneously. These values were then aggregated at the commenter level, producing for each user the total count of anti-liberal ($A_u^{(l)}$) and anti-conservative ($A_u^{(c)}$) attacks.

To quantify the partisan orientation of a user’s hostility, we calculated the **Stance Score**. This metric represents the normalized difference between a user’s attacks on liberals versus conservatives:

$$\text{stance_score}_u = \frac{A_u^{(l)} - A_u^{(c)}}{A_u^{(l)} + A_u^{(c)}}$$

where $A_u^{(l)}$ and $A_u^{(c)}$ represent the total number of anti-liberal and anti-conservative attacks by user u , respectively. The score ranges from -1 (exclusively attacking conservatives) to $+1$ (exclusively attacking liberals), with values near zero reflecting mixed or balanced attack behavior. Commenters who produced no attacks were assigned missing values.

For categorical analysis, we discretized the `stance_score` into three labels: *Anti-Liberal* ($\text{stance_score} \geq 0.2$), *Anti-Conservative* ($\text{stance_score} \leq -0.2$), and *Unclear* ($-0.2 < \text{stance_score} < 0.2$). The ± 0.2 threshold was selected based on sensitivity analysis spanning 0.2 to 0.8 (see Appendix Table 6), ensuring that labeled users exhibited a distinct direction of hostility rather than a balanced mix of attacks.

Partisan Warrior Classification We operationalized **Partisan Warriors (PWs)** as users who enter ideologically opposing spaces to attack that group (e.g., an anti-liberal user attacking liberals within a liberal channel).

Unlike general partisan commenters who may vent frustration within their own echo chambers, PWs are defined by the intersection of their **Inferred Stance** and their **Commenting Arena**. A user was classified as a PW if they met the following criteria:

- **Directional Hostility:** The user must have a clear stance (either Anti-Liberal or Anti-Conservative).
- **Oppositional Context:** The user must deploy these attacks in an environment ideologically aligned with their target.

Based on their channel activity, users were classified as follows:

- **Single-Side Commenters:** Users active on only one side of the media ecosystem were labeled PWs only if they engaged in counter-attitudinal attacks (i.e., invading the opposing camp).
 - Users on *liberal channels* were labeled PWs if they had an **Anti-Liberal** stance.
 - Users on *conservative channels* were labeled PWs if they had an **Anti-Conservative** stance.
- **Cross-Cutting Commenters:** Users active on both liberal and conservative channels were labeled PWs if they possessed *any* distinct stance.

Analytical Overview by Research Question

Because each research question targets a distinct analytical goal, the unit of analysis and the subset of data used vary across RQ1–RQ4. We provide a visual overview of this methodological pipeline in Appendix Figure 7.

RQ1 and RQ2 investigate broad patterns of incivility and audience response without relying on the specific Partisan Warrior (PW) classification. RQ1 provides a baseline comparison of incivility across commenter types using toxicity scores for *all collected comments*. RQ2 examines audience endorsement mechanisms, maintaining the *individual comment* as the unit of analysis. Here, we model the number of likes received by targeted toxic comments as a function of their attack targets, retaining comment-level granularity to capture variation in audience rewards.

In contrast, RQ3 and RQ4 specifically analyze the PW phenotype, requiring aggregated user-level or interaction-level data. RQ3 focuses on the prevalence of PWs, necessitating a *commenter-level analysis* ($N = 12,884$ unique commenters). For this inquiry, we aggregate comments posted by active commenters to classify each unique commenter based on their overall cross-posting behavior and stance, comparing the proportion of identified PWs across different commenter groups.

Finally, RQ4 investigates environmental triggers for PW participation using a dataset of *unique commenter-within-video* observations ($N = 57,644$). To ensure statistical independence and avoid pseudoreplication from hyper-active users, if a user posted multiple comments on the same video, their activity was treated as a single participation event. Using this deduplicated dataset, we employ multilevel modeling to estimate the likelihood that a participating commenter was classified as a PW, with video-level properties (e.g., title toxicity) as predictors and random intercepts to account for clustering within videos and channels.

Results

RQ1: Incivility of Cross-cutting vs. Single-side Commenters

To address RQ1, we examined whether cross-cutting commenters differ in toxicity from single-side commenters who comment only on liberal-leaning or only on conservative-leaning channels. The average toxicity scores were similar for cross-cutting ($M = 0.247$, $SD = 0.215$) and liberal-only commenters ($M = 0.252$, $SD = 0.225$), whereas conservative-only commenters showed a lower mean toxicity ($M = 0.215$, $SD = 0.210$).

Given the skewed distribution of toxicity, we employed a Kruskal–Wallis test, which indicated significant differences among the three groups, $\chi^2(2) = 7,203.4$, $p < .001$. Post-hoc Dunn’s tests with Bonferroni correction showed no significant difference between cross-cutting and liberal-only commenters ($p = .32$), but both groups were significantly more toxic than conservative-only commenters ($p < .001$ for both comparisons). Taken together, these results indicate that cross-cutting commenters are neither more civil nor more uncivil than liberal-only commenters, while both




Figure 1: Distribution of toxicity scores by commenter group. Violin plots show the full distribution, with boxplots overlaid for medians and interquartile ranges. Cross-cutting commenters are active on both liberal- and conservative-leaning channels, while only-liberal and only-conservative commenters denote activity restricted to liberal- or conservative-leaning channels. While the overall test indicates significant differences across groups, post-hoc comparisons show no significant difference in toxicity between cross-cutting and only-liberal commenters.

cross-cutting and liberal-only commenters exhibit higher toxicity than conservative-only commenters.

To interpret these scores, it is helpful to examine representative examples. Comments near the conservative-only mean (0.21) often featured slogan-based dismissals or rhetorical sarcasm, such as “*What spice goes best with liberal tears?*” or “*Klamedia obviously doesn’t care! Never vote Democrat!*” Similarly, comments near the cross-cutting mean (0.25) frequently employed descriptive criticism or mockery, such as “*Trump is a fugitive from the law*” or “*Send him to Mars with Elon Musk ... They can keep themselves busy stroking their egos.*”

These examples suggest that, despite statistically significant differences in average toxicity, the substantive differences in linguistic intensity across groups are relatively modest. Rather than reflecting sharp contrasts in severity, toxicity scores in this range appear to capture broadly similar forms of partisan hostility expressed through different rhetorical styles. This finding highlights a key limitation of relying solely on aggregate toxicity scores: such measures do not distinguish between different behavioral orientations underlying hostile speech. In other words, it is not necessarily the *level* of toxicity that differentiates engagement patterns, but potentially *who* is being targeted.

RQ2: Asymmetric Audience Rewards for Incivility

To examine whether the targets of uncivil comments influenced audience endorsement, we modeled the number of likes received by each uncivil comment using negative binomial regression, which is appropriate for over-dispersed

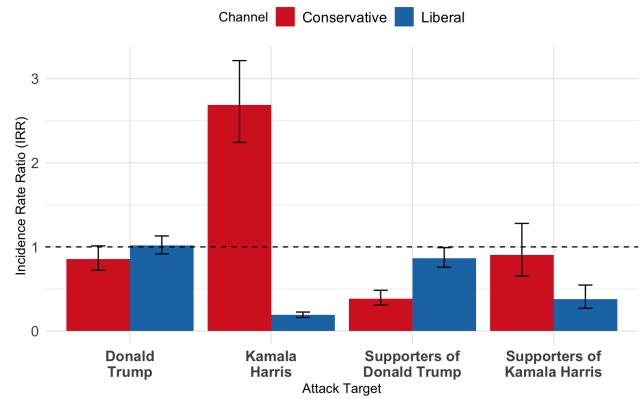


Figure 2: Incidence Rate Ratios of audience likes for uncivil comments across targets and channel leanings, estimated via negative binomial regression. Error bars indicate 95% confidence intervals.

count outcomes. Separate models were fit for liberal-leaning and conservative-leaning channels. In each model, the dependent variable was the count of likes, and the key predictor was the attack target category. The reference category was set to *Unclear* targets. We report incidence rate ratios (IRRs), obtained by exponentiating the model coefficients, with 95% confidence intervals (Figure 2).

On liberal channels, results indicate that uncivil comments directed at *Supporters of Donald Trump* received significantly fewer likes than unclear attacks ($IRR = 0.87$, $p = .036$). Comments targeting *Supporters of Kamala Harris* were even less endorsed ($IRR = 0.38$, $p < .001$), and those directly attacking *Kamala Harris* received the least support ($IRR = 0.19$, $p < .001$), amounting to an 81% reduction relative to the unclear baseline. In contrast, uncivil comments targeting *Donald Trump* obtained a similar level of likes as unclear attacks ($IRR = 1.02$, $p = .72$). These results suggest that left-leaning audiences systematically withdraw support from uncivil comments aimed at ingroup figures, while showing no significant preference between attacking Trump versus issuing vague, non-specific attacks.

A distinctively different pattern emerged on conservative channels. Uncivil comments targeting *Supporters of Donald Trump* were strongly disfavored ($IRR = 0.39$, $p < .001$), attracting only about 39% of the likes received by unclear attacks. By contrast, uncivil comments directed at *Kamala Harris* were highly rewarded, receiving nearly 2.7 times as many likes as the unclear baseline ($IRR = 2.69$, $p < .001$). Attacks on *Supporters of Kamala Harris* ($IRR = 0.90$, $p = .55$) and *Donald Trump* himself ($IRR = 0.86$, $p = .072$) did not significantly differ from unclear attacks.

Thus, right-leaning audiences strongly endorse uncivil comments against Harris while penalizing attacks on their own ingroup supporters. Together, these findings reveal that the relationship between uncivil discourse and audience endorsement is contingent on both the ideological leaning of the channel and the specific attack target. This asymmetric incentive structure—where out-group attacks are rewarded in

some spaces but not others—provides the environmental context for analyzing the prevalence of *Partisan Warriors* in the subsequent section (RQ3).

RQ3: Partisan Warrior Prevalence by Commenter Type

To examine whether PW tendencies vary across commenter types, we conducted a chi-squared test of independence between commenter type (Cross-cutting Commenters (Cross-Commenters), Only-Liberal Commenters, Only-Conservative Commenters) and PW status at the unique commenter level ($N = 12,884$). Cross-Commenters are active on both liberal- and conservative-leaning channels, while Only-Liberal Commenters and Only-Conservative Commenters denote activity restricted to liberal- or conservative-leaning channels.

The distribution differed substantially across groups, $\chi^2(2) = 9,222.1, p < .001$, with a large effect size (Cramér’s $V = .85$). As expected given our operational definition, nearly all Cross-Commenters were classified as PWs (93.6%). In contrast, PW proportions diverged sharply between single-leaning commenters: only 6.6% of Only-Liberal Commenters were PWs, whereas 27.3% of Only-Conservative Commenters fell into this category.

To further quantify these differences, we ran a logistic regression predicting PW status from commenter type (reference: Only-Conservative Commenters). Results showed that Cross-Commenters had markedly higher odds of being PWs (OR = 38.80, 95% CI [32.6, 46.4], $p < .001$), while Only-Liberal Commenters had dramatically lower odds (OR = 0.19, 95% CI [0.16, 0.22], $p < .001$). Model performance was strong ($R^2_{\text{Tjur}} = .72$), indicating that commenter type alone accounted for a large portion of variance in PW behavior.

Taken together, these results confirm that Cross-Commenters overwhelmingly meet the partisan warrior definition, which is consistent with the operationalization itself. More importantly, we observe a striking asymmetry between the two single-channel groups. While the majority of Only-Conservative Commenters were non-warriors, they were more than four times as likely to engage in partisan warrior behavior compared to their Only-Liberal counterparts (27.3% vs. 6.6%). This imbalance highlights that PW activity is not evenly distributed across ideological camps but is disproportionately concentrated among commenters who primarily engage on conservative channels.

Commenter Type	Not PW	PW
Cross-cutting	380 (6.4%)	5,534 (93.6%)
Only-Liberal	5,627 (93.4%)	397 (6.6%)
Only-Conservative	688 (72.7%)	258 (27.3%)
Total	6,695	6,189

Table 5: Distribution of PW classification by commenter type (Unique commenter-level analysis).

RQ4: Effects of Video Title Toxicity on Partisan Warrior Participation

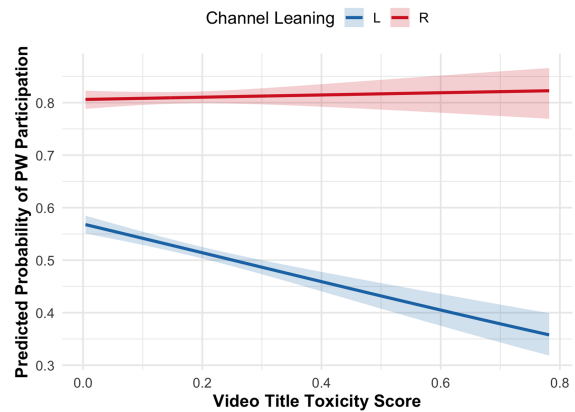


Figure 3: Predicted probabilities of PW participation by video title toxicity and channel leaning, based on the Baseline Model. Shaded regions represent 95% confidence intervals. The plot illustrates the apparent ideological asymmetry: toxic titles appear to deter PWs on liberal channels (blue) but have a negligible or slightly mobilizing effect on conservative channels (red).

Given the disproportionate concentration of Partisan Warriors in conservative spaces identified in RQ3, we next investigate the environmental triggers that might drive this engagement. Specifically, to examine whether toxic video titles are associated with the likelihood of PW participation, we fitted a series of generalized linear mixed models (GLMMs) with a binomial link function using the dataset of commenter-video pairs ($N = 57,644$). Descriptive statistics within this dataset highlight a baseline asymmetry: while 79.3% of participation events on conservative channels involved PWs, this rate was markedly lower on liberal channels (45.1%). In contrast, the average toxicity of video titles was nearly identical across channel ideologies (liberal: $M = 0.164, SD = 0.143$; conservative: $M = 0.167, SD = 0.141$), suggesting that aggregate title framing does not differ systematically by political leaning.

Baseline Association We first estimated a baseline model including a random intercept for each video to account for unobserved video-level heterogeneity ($AIC = 72,531$). In this model, video title toxicity was negatively associated with PW participation on liberal-leaning channels ($Estimate = -1.10, SE = 0.15, p < .001$), suggesting that toxic titles deterred PW participation. Conversely, the interaction term between title toxicity and channel leaning was positive and significant ($Estimate = 1.24, SE = 0.30, p < .001$). This indicated a strong ideological asymmetry: while toxic titles appeared to suppress PW activity on liberal channels, this effect was effectively offset on conservative channels, suggesting a potential mobilization effect in the baseline specification.

Robustness Check: Channel Identity However, considering that videos are embedded within channels that possess distinct community norms and audience bases, we conducted a robustness check to determine whether these effects were driven by immediate video cues or stable channel characteristics. We extended the model by adding a random intercept for *channel identity* to control for channel-level heterogeneity.

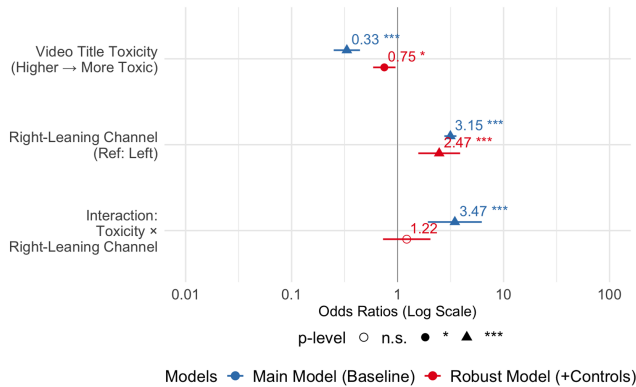


Figure 4: Comparison of fixed effects between the Main Model (Baseline, Blue) and the Robust Model (Red). Error bars represent 95% confidence intervals. Notably, the interaction term is significant in the Baseline Model but becomes non-significant in the Robust Model, confirming that the apparent asymmetry is driven by channel-level heterogeneity.

The extended model demonstrated a significantly better fit to the data ($\Delta AIC = 547$, $\chi^2(1) = 548.5$, $p < .001$), indicating that channel-level variation is a critical predictor of PW behavior.

Crucially, after controlling for channel identity, the ideological asymmetry observed in the baseline model vanished. The interaction between title toxicity and channel leaning was no longer statistically significant ($Estimate = 0.20$, $SE = 0.26$, $p = .450$). This implies that the differential response to toxicity observed in the baseline model was an artifact of channel composition rather than a true interaction with video framing.

Furthermore, the main effect of title toxicity remained negative and statistically significant, though substantially weaker in magnitude ($Estimate = -0.29$, $SE = 0.12$, $p = .020$). Within channels, higher video title toxicity is associated with a modest decrease in PW participation, indicating that toxic framing does not mobilize PWs once stable channel norms are taken into account.

Summary These findings suggest that PW participation is primarily an ecological phenomenon structured at the channel level. The apparent association between toxic titles and PW activity on conservative channels is explained by the fact that PWs tend to inhabit specific channels that have higher baseline levels of engagement. Once the stable channel context is accounted for, the specific framing of an individual video title has minimal impact on mobilizing Partisan Warriors. Thus, PW behavior appears to be governed more by

established community norms and channel ecosystems than by transient content triggers.

Conclusion and Discussion

This study, grounded in the immediate aftermath of the second 2024 U.S. presidential debate, offers a granular examination of the conditions under which partisan hostility unfolds on YouTube. By shifting the analytical lens from aggregate toxicity to the directionality and context of attacks, we challenge existing assumptions about cross-cutting interaction and reveal the structural dynamics that sustain “Partisan Warrior” behavior.

First, our investigation into cross-cutting commenters challenges the idealized view that traversing ideological divides fosters deliberation (RQ1). We found that these users were not distinctively civil; rather, their toxicity levels were indistinguishable from liberal-only commenters and significantly higher than conservative-only commenters. This pattern should be interpreted with caution, as the data reflect post-moderation content collected approximately one month after the analysis window, and prior work suggests that moderation may differentially affect the visibility of toxic comments across political spaces (Jiang, Robertson, and Wilson 2020).

At the same time, the qualitative similarity in rhetorical intensity across groups underscores the critical limitations of relying solely on aggregate toxicity scores. In a polarized landscape where ideological mobility often signals hostility rather than dialogue, generic metrics fail to distinguish between internal policing and out-group aggression. This motivated our shift toward operationalizing the Partisan Warrior construct, which defines users not merely by how toxic they are, but by whom their attacks target.

Examining the surrounding environment helps explain why such behavior persists (RQ2). Our analysis of audience endorsement revealed an asymmetry in incentive structures. While audiences in both liberal and conservative spaces penalized in-group attacks—reflecting a shared preference for internal cohesion—the rewards for out-group hostility diverged sharply. On conservative channels, uncivil attacks targeting Kamala Harris received nearly three times as many likes as baseline uncivil comments, forming a potent engagement loop that validates hostile behavior. No comparable reward structure emerged on liberal channels, where attacks on Donald Trump received no significant engagement boost. This pattern suggests that the online “marketplace of ideas” is not neutral: hostility toward the out-group is socially reinforced in conservative spaces, whereas comparable behavior in liberal spaces lacks similar validation.

Taken together, these incentive asymmetries align with the structural differences in Partisan Warrior prevalence across ideological spaces (RQ3). Conservative channels appear to function as contested battlegrounds, characterized by frequent hostile incursions and strong social rewards for retaliatory behavior. In these environments, resident users are disproportionately likely to adopt Partisan Warrior tactics; indeed, only-conservative users exhibited a Partisan Warrior rate (27.3%) more than four times higher than their only-liberal counterparts (6.6%). We interpret this pattern as con-

sistent with an ecological dynamic in which environments marked by higher perceived antagonism and asymmetric rewards are more likely to sustain resident hostility. By contrast, liberal spaces more closely resemble homogenous safe havens, where lower levels of perceived threat coincide with weaker incentives for out-group aggression and a minimal presence of resident warriors.

Finally, our analysis of environmental triggers clarifies the locus of this hostility (RQ4). Although baseline models suggested a link between video title toxicity and PW participation, robustness checks revealed that these patterns were largely explained by channel-level heterogeneity. This indicates that Partisan Warrior behavior is not a transient reaction to immediate framing cues (e.g., a toxic video title) but an ecological phenomenon driven by channel-specific norms. This interpretation aligns with prior work showing that ideological concentration on social media is driven more by user self-selection and network homophily than by algorithmic ranking or content-level triggers (Bakshy, Messing, and Adamic 2015). Moreover, research on online moderation demonstrates that toxic behavior is not solely an individual attribute but is strongly conditioned by the norms of the communities in which users participate (Chandrasekharan et al. 2017). Given evidence that users entering political communities rapidly adapt their behavior to local norms (Rajadesingan, Resnick, and Budak 2020), the observed asymmetry in PW participation is best understood as emerging from enduring, channel-specific ecosystems—shaped by the incentive structures identified in RQ2—rather than from the framing of individual videos.

These findings extend the understanding of online political discourse by demonstrating that partisan hostility is not a uniform feature of the internet but a phenomenon structured by specific environmental incentives. For platform governance, this implies that interventions must move beyond content-level policing to address ecosystem-level dynamics.

First, on the algorithmic side, ranking systems that may inadvertently gamify hostility—by rewarding or amplifying comments that attack out-groups—require recalibration. Downranking uncivil comments that disproportionately target partisan opponents could dismantle the reward loops that fuel partisan warrior dynamics. Second, given that hostility is rooted in channel ecosystems (RQ3, RQ4), moderation policies should focus on identifying and prioritizing channels that foster hostile norms for ecosystem-level interventions, rather than merely flagging individual assets. Ultimately, mitigating partisan warfare requires reshaping the structural “battlegrounds” and incentive systems that allow such behavior to flourish.

Limitations and Future Work

While this study provides insights into the asymmetric nature of online partisan hostility, several limitations should be noted, opening avenues for future research. First, our analysis is restricted to YouTube and a curated set of liberal and conservative channels in the U.S., focusing on comments after the second presidential debate. This period represents a highly salient and politically charged moment, and the observed PW patterns may reflect dynamics amplified during

election cycles rather than stable features of everyday discussion. Platform architecture and user demographics also vary, so comparative analyses across platforms, cultural contexts, and routine versus high-salience periods are needed.

Second, our reliance on automated tools for toxicity detection and target classification introduces bias. Such models often miss nuanced forms of incivility, leading to underestimation or mischaracterization. Combining large-scale computational methods with human annotation or qualitative analysis would improve accuracy and reveal more subtle rhetorical strategies.

Third, our definition of partisan warriors is behavioral—crossing ideological lines to post uncivil comments directed at identifiable partisan targets—and infers but does not explain intent. Some users may see themselves as debating, others as trolling or seeking attention. Surveys or interviews with commenters could clarify these motivations and provide a richer theoretical account.

Fourth, our findings are correlational and capture a single political moment. While toxic video titles are associated with more partisan warriors in conservative spaces, causal direction remains unclear. Longitudinal research tracking the same channels and commenters across election and non-election periods would help assess the temporal stability of PW behavior and disentangle short-term mobilization effects from enduring interaction patterns.

Fifth, future work could extend our analysis by modeling the impact of platform recommendation and visibility dynamics on commenting behavior. For example, integrating data on recommendation pathways, ranking positions, or personalized exposure signals may clarify how algorithmic curation interacts with content toxicity and partisan orientation. Such approaches would complement our content-focused analysis by accounting for exposure in shaping engagement patterns.

Finally, our dataset reflects only comments available through YouTube’s API at the time of collection. Because YouTube routinely removes content that violates its guidelines, the most extreme comments were likely already deleted. As a result, our analysis is necessarily conservative: we capture patterns among comments that remained visible rather than the full spectrum of hostility. Future work could address this by combining API-based datasets with real-time or pre-moderation collection strategies, enabling a more complete assessment of how toxicity unfolds before moderation.

Ethical Statement

This research adheres to established ethical guidelines for social media research. We analyzed only publicly available comments collected through YouTube’s official API, in full compliance with its terms of service. To protect user privacy, all identifiers were anonymized, and the analysis focused on aggregate patterns rather than individuals. Following the Association of Internet Researchers (AoIR) guidelines, all example quotes in Table 3 were paraphrased to preserve meaning while ensuring anonymity and preventing traceability.

We recognize that, as with any study of partisan dynamics, findings could be misrepresented or selectively cited to

intensify political conflict. For example, the asymmetry in partisan warrior behavior might be oversimplified to stigmatize one side, or insights about environmental triggers misused to design manipulative content. However, the intention of this work is explicitly diagnostic and constructive. Our analysis aims to illuminate mechanisms of online hostility so that platforms, policymakers, and scholars can better understand and mitigate polarization. By foregrounding methodological transparency and careful interpretation, we seek to minimize misuse while maximizing the potential to inform healthier and more inclusive online political dialogue.

Acknowledgements

We are grateful to Jason Greenfield (Computer Science PhD student, Princeton University) for his significant insights and helpful comments, which substantially improved the quality of this manuscript.

References

- Almerekhi, H.; Kwak, H.; Jansen, B. J.; and Salminen, J. 2019. Detecting Toxicity Triggers in Online Discussions. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media*, 291–292. Hof Germany: ACM. ISBN 9781450368858.
- Bail, C. A.; Argyle, L. P.; Brown, T. W.; Bumpus, J. P.; Chen, H.; Hunzaker, M. B. F.; Lee, J.; Mann, M.; Merhout, F.; and Volfovsky, A. 2018. Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37): 9216–9221.
- Bakshy, E.; Messing, S.; and Adamic, L. A. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239): 1130–1132.
- Blinder, S.; and Ware, J. 2024. In Sweden, a Domestic Troll Farm Tests Democracy. GNET Research Insight. Accessed 2025-01-05.
- Chandrasekharan, E.; Pavalanathan, U.; Srinivasan, A.; Glynn, A.; Eisenstein, J.; and Gilbert, E. 2017. You Can't Stay Here: The Efficacy of Reddit's 2015 Ban Examined Through Hate Speech. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW): 1–22.
- Coe, K.; Kenski, K.; and Rains, S. A. 2014. Online and Uncivil? Patterns and Determinants of Incivility in Newspaper Website Comments. *Journal of Communication*, 64(4): 658–679.
- Farjam, M.; Meyer, H.; and Lohkamp, M. 2025. A Practical Guide and Case Study on How to Instruct LLMs for Automated Coding During Content Analysis. *Social Science Computer Review*, 08944393251349541.
- Fichman, P.; and Akter, S. 2023. Trolling asymmetry toward Republicans and Democrats and the shift from foreign to domestic trolling. *Telematics and Informatics*, 82: 101998.
- Flaxman, S.; Goel, S.; and Rao, J. M. 2016. Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opinion Quarterly*, 80(S1): 298–320.
- Gan, Y.; Poesio, M.; and Yu, J. 2024. Assessing the Capabilities of Large Language Models in Coreference: An Evaluation. In Calzolari, N.; Kan, M.-Y.; Hoste, V.; Lenci, A.; Sakti, S.; and Xue, N., eds., *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 1645–1665. Torino, Italia: ELRA and ICCL.
- Gao, Y.; Qin, W.; Murali, A.; Eckart, C.; Zhou, X.; Beel, J. D.; Wang, Y.-C.; and Yang, D. 2024. A Crisis of Civility? Modeling Incivility and Its Effects in Political Discourse Online. *Proceedings of the International AAAI Conference on Web and Social Media*, 18: 408–421.
- Garrett, R. K. 2009. Echo chambers online?: Politically motivated selective exposure among Internet news users. *Journal of Computer-Mediated Communication*, 14(2): 265–285.
- Gervais, B. T. 2015. Incivility Online: Affective and Behavioral Reactions to Uncivil Political Posts in a Web-based Experiment. *Journal of Information Technology & Politics*, 12(2): 167–185.
- Gervais, B. T.; Dye, C.; and Chin, A. 2025. Incivility or Invalidity? Evaluating Perspective API Scores as a Measure of Political Incivility. *American Politics Research*, 53(3): 266–274.
- Gouk, H.; Pfahringer, B.; and Cree, M. 2016. Learning Distance Metrics for Multi-Label Classification. In Durrant, R. J.; and Kim, K.-E., eds., *Proceedings of The 8th Asian Conference on Machine Learning*, volume 63 of *Proceedings of Machine Learning Research*, 318–333. The University of Waikato, Hamilton, New Zealand: PMLR.
- Han, J.; Lee, Y.; Lee, J.; and Cha, M. 2023. News comment sections and online echo chambers: The ideological alignment between partisan news stories and their user comments. *Journalism*, 24(8): 1836–1856.
- Heseltine, M.; and Clemm Von Hohenberg, B. 2024. Large language models as a substitute for human experts in annotating political text. *Research & Politics*, 11(1): 20531680241236239.
- Human-Compatible AI Institute. 2024. The Prosocial Ranking Challenge. Online research initiative. Accessed 2025-01-05.
- Jiang, S.; Robertson, R. E.; and Wilson, C. 2019. Bias Misperceived: The Role of Partisanship and Misinformation in YouTube Comment Moderation. *Proceedings of the International AAAI Conference on Web and Social Media*, 13: 278–289.
- Jiang, S.; Robertson, R. E.; and Wilson, C. 2020. Reasoning about Political Bias in Content Moderation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(09): 13669–13672.
- Kim, J. W. 2018. Online incivility in comment boards: Partisanship matters - But what I think matters more. *Computers in Human Behavior*, 85: 405–412.
- Lees, A.; Tran, V. Q.; Tay, Y.; Sorensen, J.; Gupta, J.; Metzler, D.; and Vasserman, L. 2022. A New Generation of Perspective API: Efficient Multilingual Character-level Transformers. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '22, 3197–3207. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393850.

- Lerman, K.; Feldman, D.; He, Z.; and Rao, A. 2024. Affective Polarization and Dynamics of Information Spread in Online Networks. *npj Complex Systems*, 1(8).
- Levandowsky, M.; and Winter, D. 1971. Distance between Sets. *Nature*, 234(5323): 34–35.
- Mamakos, M.; and Finkel, E. J. 2023. The social media discourse of engaged partisans is toxic even when politics are irrelevant. *PNAS Nexus*, 2(10): pgad325.
- Masullo, G. M.; Tenenboim, O.; and Lu, S. 2023. “Toxic atmosphere effect”: Uncivil online comments cue negative audience perceptions of news outlet credibility. *Journalism*, 24(1): 101–119.
- Matsui, A.; Murayama, T.; and Yoshida, M. 2024. Hater Is Not the Only Source of Toxic Online Communication, But Also Fan. *Proceedings of the Eighteenth International AAAI Conference on Web and Social Media*.
- Milli, S.; Carroll, M.; Wang, Y.; Pandey, S.; Zhao, S.; and Dragan, A. D. 2024. Engagement, User Satisfaction, and the Amplification of Divisive Content on Social Media. *arXiv preprint*.
- Miyazaki, K.; Uchiba, T.; Kwak, H.; An, J.; and Sasahara, K. 2024. The impact of toxic trolling comments on anti-vaccine YouTube videos. *Scientific Reports*, 14(1): 5088.
- Mutz, D. C. 2002. Cross-cutting Social Networks: Testing Democratic Theory in Practice. *American Political Science Review*, 96(1): 111–126.
- Papacharissi, Z. 2004. Democracy online: civility, politeness, and the democratic potential of online political discussion groups. *New Media & Society*, 6(2): 259–283.
- Pendzel, S.; Lotan, N.; Zoizner, A.; and Minkov, E. 2024. A Closer Look at Multidimensional Online Political Incivility. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 14881–14896. Miami, Florida, USA: Association for Computational Linguistics.
- Rajadesingan, A.; Resnick, P.; and Budak, C. 2020. Quick, Community-Specific Learning: How Distinctive Toxicity Norms Are Maintained in Political Subreddits. *Proceedings of the International AAAI Conference on Web and Social Media*, 14: 557–568.
- Ribeiro, M. H.; Ottoni, R.; West, R.; Almeida, V. A. F.; and Meira, W. 2020. Auditing radicalization pathways on YouTube. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 131–141. Barcelona Spain: ACM. ISBN 9781450369367.
- Rossini, P. 2021. More Than Just Shouting? Distinguishing Interpersonal-Directed and Elite-Directed Incivility in Online Political Talk. *Social Media + Society*, 7(2): 20563051211008827.
- Rossini, P. 2022. Beyond Incivility: Understanding Patterns of Uncivil and Intolerant Discourse in Online Political Talk. *Communication Research*, 49(3): 399–425.
- Rossini, P.; and Maia, R. 2021. Characterizing Disagreement in Online Political Talk: Examining Incivility and Opinion Expression on News Websites and Facebook in Brazil. *Journal of Deliberative Democracy*, 17(1).
- Saltz, E.; Jalan, Z.; and Acosta, T. 2024. Re-Ranking News Comments by Constructiveness and Curiosity to Promote Healthier Dialogue. *arXiv preprint*.
- Simchon, A.; Brady, W. J.; and Van Bavel, J. J. 2022. Troll and divide: the language of online polarization. *PNAS Nexus*, 1(1): pgac019.
- Stroud, N. J. 2010. Polarization and Partisan Selective Exposure. *Journal of Communication*, 60(3): 556–576.
- Stryker, R.; Conway, B. A.; and Danielson, J. T. 2016. What is political incivility? *Communication Monographs*, 83(4): 535–556.
- Sunstein, C. R. 2018. *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press. ISBN 9781400890521 9780691180908.
- Tan, C.; Niculae, V.; Danescu-Niculescu-Mizil, C.; and Lee, L. 2016. Winning Arguments: Interaction Dynamics and Persuasion Strategies in Good-faith Online Discussions. In *Proceedings of the 25th International Conference on World Wide Web*, 613–624. Montréal Québec Canada: International World Wide Web Conferences Steering Committee. ISBN 978-1-4503-4143-1.
- Theocharis, Y.; Barberá, P.; Fazekas, Z.; and Popa, S. A. 2020. The Dynamics of Political Incivility on Twitter. *Sage Open*, 10(2): 2158244020919447.
- Theocharis, Y.; Barberá, P.; Fazekas, Z.; Popa, S. A.; and Parnet, O. 2016. A Bad Workman Blames His Tweets: The Consequences of Citizens’ Uncivil Twitter Use When Interacting With Party Candidates: Incivility in Interactions With Candidates on Twitter. *Journal of Communication*, 66(6): 1007–1031.
- Wojcieszak, M. E.; and Mutz, D. C. 2009. Online Groups and Political Discourse: Do Online Discussion Spaces Facilitate Exposure to Political Disagreement? *Journal of Communication*, 59(1): 40–56.
- Wu, S.; and Resnick, P. 2021. Cross-Partisan Discussions on YouTube: Conservatives Talk to Liberals but Liberals Don’t Talk to Conservatives. *Proceedings of the International AAAI Conference on Web and Social Media*, 15: 808–819.
- Xia, Y.; Monti, C.; Keller, B.; and Kivelä, M. 2025. Integrated or Segregated? User Behavior Change After Cross-Party Interactions on Reddit. *Proceedings of the International AAAI Conference on Web and Social Media*, 19: 2044–2061.
- Xiao, L.; and Khazaei, T. 2019. Changing Others’ Beliefs Online: Online Comments’ Persuasiveness. In *Proceedings of the 10th International Conference on Social Media and Society*, 92–101. Toronto ON Canada: ACM. ISBN 978-1-4503-6651-9.
- Yudhoatmojo, S.; Balci, U.; and Blackburn, J. 2024. *A Data-Driven Analysis of Sovereign Citizens Movement on Telegram*. US: ICWSM.
- Çetinkaya, Y. M.; Ghafouri, V.; Suarez-Tangil, G.; Such, J.; and Elmas, T. 2025. Cross-Partisan Interactions on Twitter. *Proceedings of the International AAAI Conference on Web and Social Media*, 19: 324–340.

Paper Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures?
Yes, and we took several steps to address privacy concerns: (1) we analyzed only publicly posted comments accessible via YouTube's official API, (2) we focused on aggregate patterns rather than individual users, (3) we removed direct user identifiers in our analysis, and (4) our research aims to understand democratic discourse patterns rather than profile individuals.
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? Yes
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? Yes, see [Introduction](#)
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? Yes
 - (e) Did you describe the limitations of your work? Yes, see [Limitations and Future Work](#)
 - (f) Did you discuss any potential negative societal impacts of your work? Yes, see [Ethical Statements](#)
 - (g) Did you discuss any potential misuse of your work? Yes, see [Ethical Statements](#)
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? Yes, see [Ethical Statements](#)
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? NA
 - (b) Have you provided justifications for all theoretical results? NA
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? NA
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? Yes
 - (e) Did you address potential biases or limitations in your theoretical framework? Yes, see [Limitations and Future Work](#)
 - (f) Have you related your theoretical results to the existing literature in social science? Yes, see [Related Work](#)
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? Yes, see [Conclusion and Discussion](#)
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? NA
 - (b) Did you include complete proofs of all theoretical results? NA
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? NA
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? NA
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? NA
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? NA
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? NA
 - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? NA
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
 - (a) If your work uses existing assets, did you cite the creators? Yes, see [Methods](#)
 - (b) Did you mention the license of the assets? NA
 - (c) Did you include any new assets in the supplemental material or as a URL? No
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? Yes. Data used was curated from publicly available sources.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes, and we conform to the [YouTube policies](#). See [Ethical Statements](#).
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? NA
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? NA
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
 - (a) Did you include the full text of instructions given to participants and screenshots? NA
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA
 - (d) Did you discuss how data is stored, shared, and de-identified? NA

Appendices

A. Robustness Check: User Categorization using Continuous Metric τ

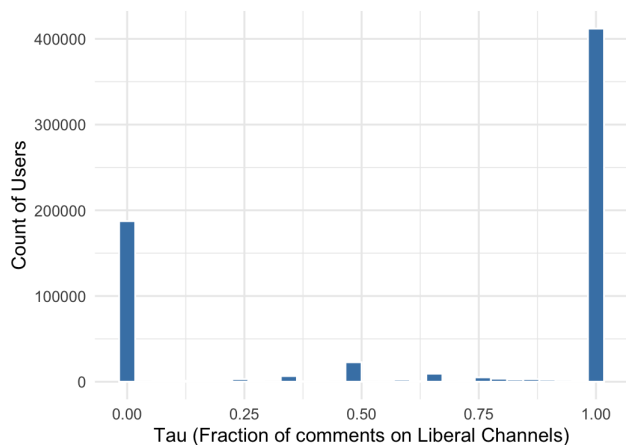


Figure 5: Distribution of τ for All Users. The overall population exhibits a bimodal distribution, justifying categorical separation.

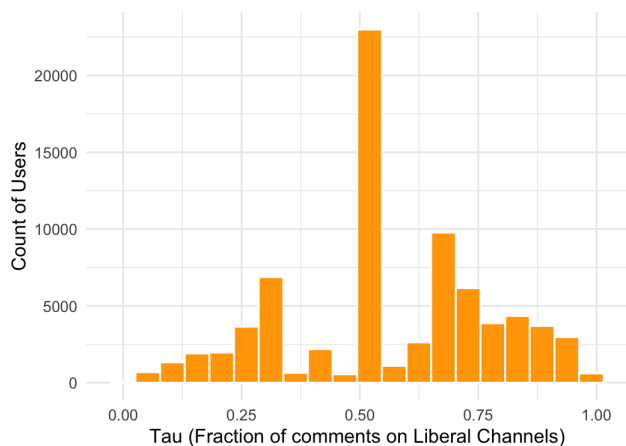


Figure 6: Distribution of τ among Cross-Commenters. The distribution peaks around $\tau = 0.5$, indicating balanced engagement across ideological contexts.

To assess the robustness of our categorical user definitions (L-Only, C-Only, and Cross-Commenters), we additionally computed a continuous engagement metric, τ , defined as the fraction of a user’s comments posted on liberal-leaning channels. Under this definition, $\tau = 1$ corresponds to exclusive participation on liberal channels, whereas $\tau = 0$ indicates exclusive participation on conservative channels.

Figure 5 presents the distribution of τ for the full commenter population ($N = 12,884$). The distribution is strongly bimodal, with the majority of commenters concentrated at the extremes ($\tau = 0$ or $\tau = 1$). This pattern

indicates that most commenters engage almost exclusively within a single ideological context, supporting the use of categorical classifications to capture dominant participation patterns.

Figure 6 shows the distribution of τ among commenters classified as Cross-Commenters. Rather than clustering near either extreme, this group exhibits a broad distribution of τ values with a clear mode around $\tau = 0.5$. This indicates that Cross-Commenters tend to engage substantially with both liberal- and conservative-leaning channels, rather than predominantly favoring one side with minimal cross-ideological activity.

Taken together, these results suggest that ideological engagement on YouTube is best characterized by discrete participation regimes rather than a smooth continuum. While τ provides a useful descriptive measure, the underlying user behavior is highly polarized, and the categorical distinctions employed in our analysis reflect meaningful and empirically grounded differences in commenting patterns.

B. Characterizing ‘Unclear’ Commenters

To address the nature of users classified as *Unclear* ($-0.2 < \text{stance_score} < 0.2$), we examined their prevalence and behavioral patterns. As shown in Table 4, these users constitute a negligible fraction of the active user population across all groups.

Qualitative inspection reveals that these users are typically not “trolls” attacking both sides, but rather commenters whose hostility falls outside the specific Trump-Harris binary defined in our taxonomy. Their comments generally fall into two categories:

1. **Ambiguous References:** Comments relying on pronouns without sufficient context for the model to resolve a specific target (e.g., “*This lady sucks*,” “*He is an idiot*”). While likely partisan in context, they lack the explicit naming required for our conservative coding strategy.
2. **Peripheral Targets:** Attacks directed at entities other than the main presidential candidates, such as media organizations or other political figures (e.g., “*CNN is pathetic... Sick of the lies*,” “*ABC is so biased and disgusting*,” “*Nikki [Haley] you are a two-faced snake!*”).

Thus, the classification of *Unclear* largely reflects the strict operationalization of our target taxonomy (focusing on the Trump vs. Harris rivalry) rather than a lack of partisan sentiment among these users.

C. Prompt example for detecting attack target from comments

You are an assistant who identifies the target of uncivil political YouTube comments and classifies the target.

Your job is to:

- Identify only the ATTACK target(s) | entities mentioned or implied in the comment with uncivil language. Uncivil

language is a feature of discussion that conveys an unnecessarily disrespectful tone toward the discussion forum, its participants, or its topics.

- For each target, classify it into one or more of the following categories:
 0. Unclear - The referent is absent, or cannot be confidently resolved, even with metadata.

1. Supporters of Donald Trump
2. Supporters of Kamala Harris
3. Kamala Harris (as an individual)
4. Donald Trump (as an individual)
5. Republicans
6. Democrats
7. Other political actors
8. Conservatives in general
9. Liberals in general
10. Individual or group(s) with a particular religion
11. Individual or group(s) with particular gender identities
12. Individual or group(s) with particular race/ethnicity
13. Immigrants
14. Debate moderators
15. Mainstream media or journalist(s) associated with mainstream media (Specify the media)
16. Government Bodies (White House)
17. Joe Biden
18. Others - Other specified targets that cannot be classified from 1 - 17.

- For vague references (e.g., "he", "she", "they"), ONLY resolve them using the video metadata when there is strong evidence for a specific resolution. If there is any ambiguity, classify the target as "Unclear" (category 0).

- ONLY identify targets that are being attacked with uncivil language, not just mentioned or criticized in civil terms.

...

D. Methodological Pipeline

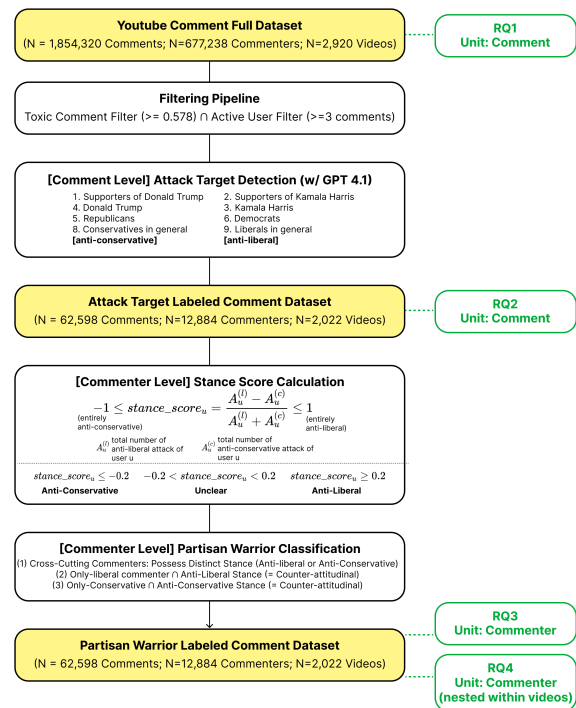


Figure 7: Methodological pipeline showing the complete process from YouTube comment collection (N=1,854,320) to partisan warrior classification.

E. Supplementary Analyses

Threshold T	N	Stance Distribution (%)			PW Rate (%)			χ^2 (dof=2)	
		Anti-Lib.	Anti-Con.	Unclear	Overall	OnlyL	OnlyR		Both
0.2	12,884	20.4	73.6	6.0	48.0	6.6	27.3	93.6	9222.1***
0.3	12,884	20.4	73.5	6.2	47.9	6.5	27.2	93.4	9198.9***
0.4	12,884	20.1	72.9	7.1	47.4	6.4	27.0	92.4	9016.8***
0.5	12,884	20.1	72.8	7.1	47.4	6.4	27.0	92.3	9009.9***
0.6	12,884	19.9	72.1	8.0	46.8	6.3	26.4	91.2	8820.1***
0.7	12,884	19.7	71.3	9.0	46.1	6.2	26.2	90.0	8602.4***
0.8	12,884	19.6	70.9	9.6	45.8	6.1	25.7	89.4	8496.3***

Notes: Anti-Lib. = anti-liberal, Anti-Con. = anti-conservative. PW Rate indicates the proportion of commenters classified as partisan warriors within each commenter type. χ^2 values are from chi-squared tests of independence between commenter type and PW status. *** $p < .001$.

Table 6: Robustness check of stance thresholding. Distribution of stance labels and partisan warrior (PW) rates across thresholds. All proportions are computed from the same set of 12,884 commenters.

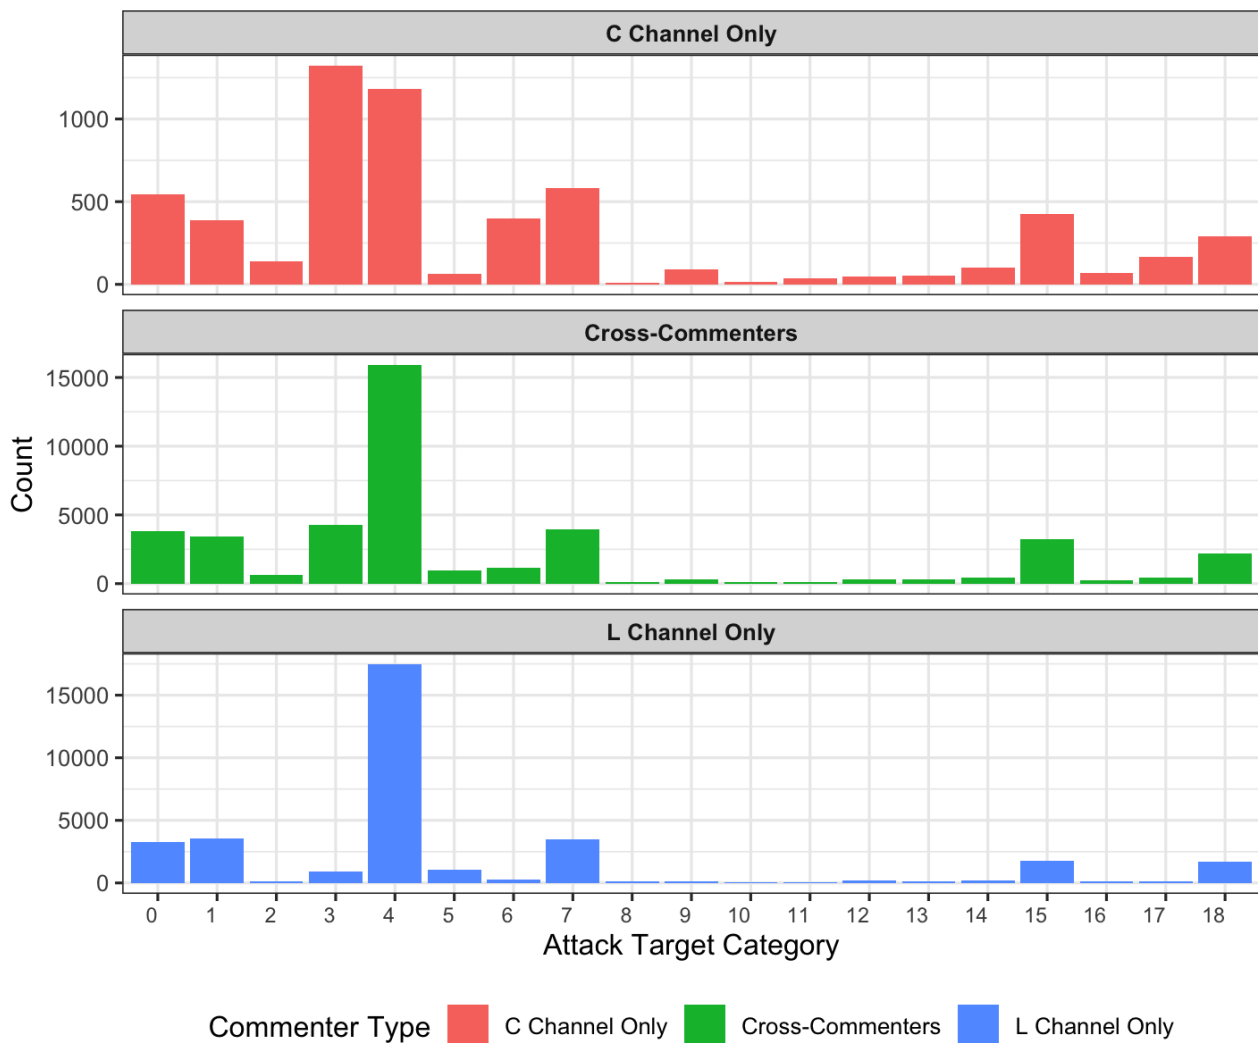


Figure 8: Distribution of total attack targets across commenter types, defined by the channel on which comments were posted (C Channel Only, L Channel Only, Cross-Commenters). Bars represent the total number of uncivil attacks directed at each target category.