Social Dynamics and Mobilization Potential of Online Election Narratives

Kristen Engel, Tanushree Mitra, Emma S. Spiro

University of Washington

engelkri@uw.edu, tmitra@uw.edu, espiro@uw.edu

Abstract

The fragmentation of social media challenges how we might efficiently and effectively identify, understand, and counter harmful content. Prior work establishes frameworks for measuring problematic narratives, evaluating harms, and leveraging interdisciplinary theories and findings to design mitigating solutions. However, there is little understanding of these phenomena outside of mainstream platforms. This is particularly concerning given that alt-tech users have been observed to include insurrectionists, active shooters, and other extremists – often driven from mainstream platforms due to deplatforming and content moderation . Our work aims to characterize online narratives across alt-tech platforms. In particular, we highlight how rumoring and conspiracy theory narratives in the context of the 2022 U.S. elections impact social dynamics and inspire collective action. We gather a unique dataset of over 7,000 social media posts from Gab, Gettr, Parler, and Truth Social from which we derive prevalent narratives using natural language processing techniques. We then examine how the platform, affect, and engagement differ across context through the lens of narrative, social identity, and mobilization potential using mixed methods. Findings from our analyses show variation between how narratives support social identity conceptions of power and mobilization potential.

Introduction

Alt-tech social networks, which offer online users similar platform affordances to mainstream systems but with less restrictive (or no) content moderation, are growing in both prominence and participation. Yet there is limited understanding of how users engage in problematic discourse within and across alt-tech. Rumble, for example, is an alternative technology to YouTube; it provides an online platform to upload, stream, and comment on videos. With nearly 60 million monthly active users as of late 2023 and exclusive rights to stream the U.S. Republican primary debates, it is just one example of the growing role of alternative platforms (Rumble 2023; PBS 2023). Although alt-tech social media serve comparatively smaller and more homogeneous user bases than their mainstream counterparts, these platforms have been shown to host violent content and fervent users whose online activities have demonstrable effects on other online communities and offline events (Benner 2021; Horta Ribeiro et al. 2023). Some alt-tech users platform and promote dangerous speech and facilitate coordinated operations with more widespread impacts. The content and behavior on these platforms migrates from the fringes to the mainstream and manifests both online and offline, threatening democracy, safety, and information integrity (Davey and Ebner 2017; Russo et al. 2023). Indeed, regardless of the relative scale of harm targets and audiences, any environment where destructive beliefs and behaviors can spread and be operationalized is worthy of attention. Further the dynamics of how online users frame, interact with, and instigate through narratives on alt-tech is crucial to contextualizing sensemaking, trust, and security on social media in an increasingly fragmented landscape. It is therefore critical to understand how users spread and engage with problematic narratives on alt-tech, particularly in the case of events of national prominence such as elections.

Narratives provide a construct for people to engage with and sensemake around events, coming to individual and collective understanding in both offline and online environments (Wilson, Zhou, and Starbird 2018). While prior work considers many definitions and theories concerning narrative (Keith Norambuena, Mitra, and North 2023), here, we define a *narrative* as a representation of events that conveys content to serve a social function (Ryan 2007; Puckett 2016). This narrative, in turn, is just another perspective within a larger *discourse* comprising many narrative contexts (Puckett 2016). For example, a narrative might promote the perspective that ballots marked with a Sharpie would not count due to a larger effort to steal the election (Pentzold et al. 2022). It is one representation within a discourse of U.S. election integrity concerns. To enable cohesive narrative representation and interpretation, *framing* focuses attention, articulates coherence, establishes relationships, and inspires mobilization (Snow, Vliegenthart, and Ketelaars 2018).

Understanding the nuances of how people individually and collectively represent and engage with current events can provide insight into what narratives drive human interest, are misunderstood or manipulated, and are unanticipated. For example, analysis of online COVID-19 vaccines discourse would likely contain posts that are supportive, detracting, or uncertain. We can better characterize the nuances in vaccine discourse by understanding its narratives, such as a narrative of coordinated efforts to find doses or a narrative spreading false claims about mRNA vaccines altering DNA. Developing an understanding of subtleties within discourse has implications for better supporting platform management and evaluating audience perception and engagement. While prior work examines narratives on mainstream social media, less is known of how such narratives play out on and across alt-tech platforms. Our approach builds on prior work exploring narratives in election contexts with a case study of the U.S. midterm elections that contains complex discourse featuring misunderstandings, misinformation, and mistakes. This focus helps to understand the differing dimensions of an event with nuanced representations prevalent on alt-tech platforms during a defined period for study. Therefore, some findings may not generalize, however the methods and analysis dimensions were employed so that they might be applied to novel queries and cases.

In this work, we use quantitative and qualitative methods to characterize narratives within the discourse surrounding claims of 2022 U.S. midterm election process errors and interference in Arizona across four alt-tech platforms: Gab, Gettr, Parler, and Truth Social. This case study provides both insights into online narratives about a key event that gained national attention and a framework to demonstrate how these methods could be used to examine the interactions of social identity dynamics, mobilization indicators, and conspiracy theories with online narratives. Our dataset is derived from queries curated through real-time observation of trending false and misleading claims (Schafer et al. 2025). We look specifically at the case of ballot counting in Maricopa County, AZ where the Election Integrity Partnership (EIP) team tracked rumors of election worker incompetence and interference, as well as voter suppression. Our data comprise almost 7,000 alt-tech social media posts discussing these events for the week following the election, from November 8, 2022 to November 15, 2022. We standardize the associated content and engagement metadata across platforms. With these data, we aim to answer:

RQ1: What are the characteristics of prevalent election narratives on alt-tech platforms? By surfacing and assessing prevalent narratives, we contribute a refined understanding of how discourse resonates through the dimensions of cross-platform distribution, affective harms, and audience engagement. Using natural language processing and clustering techniques, we surface 9 narratives present on these platforms. We show distinction between narratives with respect to audience engagement and affective harms.

RQ2: How does social identity influence the framing of prevalent election narratives? Through examining the social dynamics present in these narratives, our work contributes insights into how power and identity is used to frame and motivate understanding. This helps to contextualize narratives with respect towards the role group identification plays in framing both understanding and collective action. We use a lexical-based approach to measure and compare conceptions of power between social identities. Although there is minimal variation between the articulated power differentials of in-group and out-group identities across platforms, we find variation between narratives. Of note, across platforms and narratives, post audiences were portrayed as having less power compared to post authors.

RQ3: What is the mobilization potential of prevalent election narratives? We consider the narratives through the framing of collective action mobilization and conspiracy theories, contributing an understanding of how these elements intersect within narratives. We use qualitative labeling of a sample of the data to apply and extend Snow et al.'s framing tasks to assess the theoretical potential for collective action mobilization from microblogging. We additionally label for the presence of conspiracy theory-related content. Our results indicate posts that support mobilization were the second most common. This subset of posts also have the highest instances of conspiracy theory content and comprise the majority of outlier high engagement posts. We also find that conspiracy theory content contains more motivational framing relative to other posts within narratives.

Our analytic approach to addressing these research questions centers a domain-agnostic and platform-agnostic framework. We integrate and iterate on existing methods to provide insight into how users across alt-tech platforms responded to an evolving election event. We analyze narrative differences across platforms, affective harms, audience engagement, social identity dynamics through group connotations of power, and mobilization potential through collective action frames. We close by highlighting how our results show a need for a more nuanced approach towards evaluating online discourse and opportunities for future work.

Background and Related Work

Understanding Online Narratives at Scale

Narratives are a central construct through which humanity forms beliefs, perceives the environment, and motivates action. As such, ongoing scholarly attention aims to understand how narratives can be extracted and characterized; this attention has only increased with the proliferation of digital data (Piper, So, and Bamman 2021). However, such work remains a challenge, as models of narrative representation can struggle to capture nuanced stories (Keith Norambuena, Mitra, and North 2023). Some work in this domain focuses on text extraction, often based in linguistic dependency parsing. One effective method includes the identification of Subject-Verb-Object (SVO) triplets, which interprets how agents interact with other entities (Samory and Mitra 2018). Jachim, Sharevski, and Pieroni (2021) leverage correspondence mapping between nouns and verbs to capture emerging narratives in an election context. Another approach includes mapping content to an embedding space and clustering semantically similar messaging (Dash et al. 2022). Our work connects these approaches, extending the SVO approach to consider the added elements of time and place followed by extracting embedding clusters to get more meaningful narrative cohesion.

Understanding Social Identity Dynamics

Social identity is central to understanding how individuals self-conceptualize and engage in interpersonal dynamics. Tajfel and Turner's (2004) work on social identity theory proposes that group membership motivates behavior, with a distinct comparison between one's in-group and out-groups. They hypothesize that unequal resources between groups prompt subordinates to work towards a more positive in-group identity. Conceptualizations of intergroup identities and dynamics center the relationships between groups, in particular subjective notions related to power such as status and legitimacy (McKeown, Haji, and Ferguson 2016). Individual action, interaction, and understanding is grounded in shared meanings and an intersubjective conception of the environment (Turner and Oakes 1986). Therefore, it is crucial to account for social identity and structure when attempting to characterize perceptions, notions of power and status, interactions, and actions. In particular, distinguishing between in-group (us) and out-group (them) provides a lens through which to understand the influences on and impacts of harmful behavior. Such framing is linked to biased and often prejudiced behavior (Brewer 1999). These identity constructions contribute to both individual and group consequences (Tajfel 1982). Social identity framing can both contribute to increasing engagement with problematic content and in promoting the targeting of identities and communities (Jolley, Meleady, and Douglas 2020; Engel, Phadke, and Mitra 2023). In addition to implications for perpetuating harms in online discourse, a negative social identity compared to out-groups can motivate individuals to take action to achieve a more positive social identity (Tajfel and Turner 2004). Indeed, narratives have the power to drive perceptions of both social comparison and group goals. This framing of social identity through narratives has implications for individual evaluations of their position and ensuing strategic decisions (Tajfel and Turner 2004). Our work is motivated by, and seeks to directly incorporate, this understanding of the nature and impact of social identity — we propose measuring the differences between narrative framing through this lens.

Characterizing Social Mobilization Online

Just as narratives mediate conception of social identity, they mediate social mobilization; group members seek to redefine perceptions through both re-framing of narrative, such as changing of values or assumed social status, and through action, such as generating conflict or attempting to change the conditions that are barriers to a more positive social identity (Tajfel and Turner 2004). To measure this, we draw on Snow, Vliegenthart, and Ketelaars's (2018) position that collective action frames, comprised of beliefs that legitimize and inspire social mobilization, have the power to move users from a passive to an active role, to take action, and to counter adversaries. Prior work applies this framework to study problematic content on X, known during this study period as Twitter. Phadke et al. (2018) contributed an annotation scheme of framing task dimensions in the context of hate speech, where *prognostic* framing might surface as solutions of violence or policy. Kavrakis (2023) found that Islamic extremist groups varied in use of framing tasks, despite sharing a common ideology. Other scholars have examined mobilization; Prochaska et al. (2023) provide a framework to examine how problematic content leads to informal and tactical mobilization in the context of online election discourse. Sternisko, Cichocka, and Van Bavel (2020) present a framework where the content and qualities of conspiracy theories present differing motivations for conspiracy theory beliefs and implications for social movements. We distinguish our work by presenting a narrative and platform-agnostic annotation scheme extending Snow et al.'s framework for evaluation online. We demonstrate the applicability of the proposed scheme in a case study of election content with varied narrative dimensions and across alt-tech platforms. Additionally, we apply this framework to measure outcomes across platforms, across narratives, and in the context of harmful outcomes and audience engagement.

Data

Our data is centered on a salient incident of election rumoring in Maricopa County, Arizona that captured national attention during the 2022 midterm elections in the U.S. It is comprised of microblogging posts authored on Gab, Gettr, Parler, and Truth Social between November 8 and November 15, 2022 that include key terms[1] related to a false claim that some ballots with printer issues were not included in the election results once collected in a secure box for central processing (Kelety 2022). (See Figure 1.) Election-related claim key terms were derived from EIP real-time monitoring of social media posts that may suppress or confuse voters, delegitimize elections, or interfere with election participation (Election Integrity Partnership 2022). The alt-tech posts, selected from platforms that center free speech over content moderation and offer functionality similar to Twitter, are collected using the Open Measures API (Open Measures 2023). We collected 7,732 social media posts, removing duplicates and posts with less than 2 words to yield 6,676 by 3,674 user accounts, with 28% of the posts on Gab, 48% on Gettr, 3% on Parler, and 21% on Truth Social. (See Table 1.)

The posts and associated metadata allow us to contrast narratives. To better understand the affective differences, we measured harms of post toxicity, identity attacks, insults, and threats using the Perspective API (Jigsaw 2024). The scores are probabilities of perceived harm, between 0 indicating *unlikely* and 1 indicating *likely*. A score of 0.5 represents uncertainty. This state-of-the-art model is trained on millions of online comments and manual labeled to evaluate harms.

RQ1: Prevalent Narratives

To better understand the nuanced discourse around this incident, we surface narratives by clustering extracted text patterns of posts mapped to an embedding space — several narratives involved similar events and entities but differed in focus or framing. All narratives are present across the four platforms studied, with no one narrative dominating. Nearly all narratives were driven by high levels of user participation, where power users did not show outsized effect. Engagement differences between narratives are statistically significant. Narratives experienced varying spikes in harmful content over time; these same narratives also had some of the highest engagement numbers. The narrative differences

[1]((“Maricopa” OR “Arizona” OR “ AZ ”) AND (“count” OR “tabulate” OR “ballot” OR “printer” OR “box” OR “tabulator”))

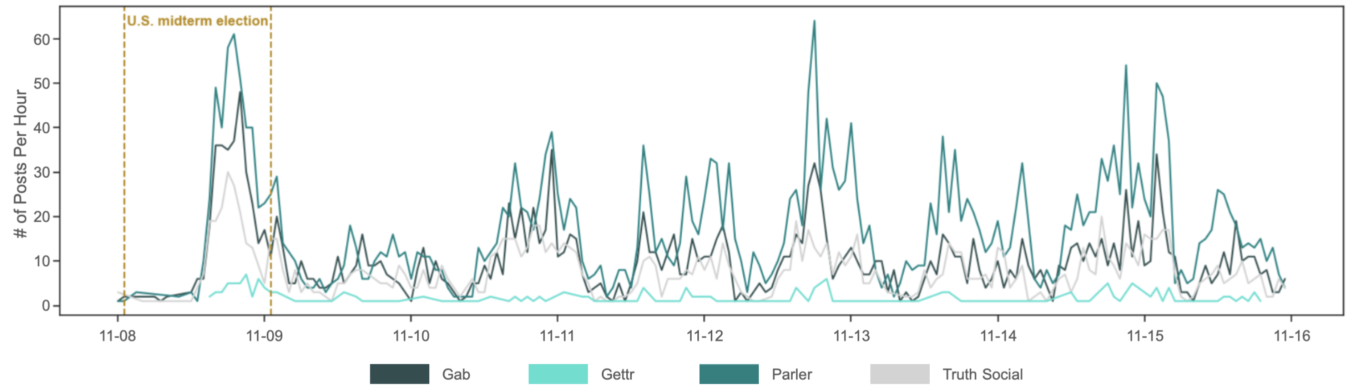

Figure 1: Claim-related posting activity by platform over the study period. The X-axis represents the post timestamp and the Y-axis indicates hourly post counts. The gold color highlights Election Day, and remaining colors correspond to platform.

Platform	**% of Data**	**# of Posts**	**% of Authors**	**# of Authors**	**Reaction Engagement Metric**	**Reply Engagement Metric**	**Reshare Engagement Metric**	**Engagement Scores**
Gab	28.26	1,912	12.51	454	Likes (*favouritescount*)	Replies (*repliescount*)	Reposts (*reblogscount*)	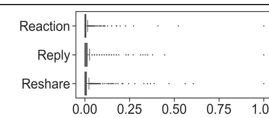
Gettr	48.00	3,248	22.79	827	Likes (*lkbpst*)	Comments (*cm*)	Reposts (*shbpst*)	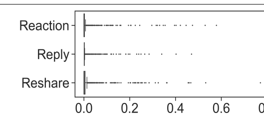
Parler	3.07	208	1.35	49	Upvotes (*upvotes*)	Comments (*comments*)	Echoes (*reposts*)	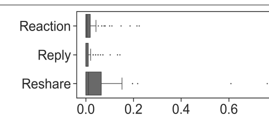
Truth Social	20.66	1,398	11.33	411	Likes (*favourites_count*)	Replies (*replies_count*)	ReTruths (*reblogs_count*)	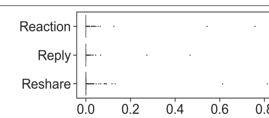

Table 1: Table of platform metrics. Counts are calculated from filtered data. Engagement scores are normalized within platform, with a value of 1 equivalent to the highest raw metric score and a value of 0 equivalent to the minimum raw metric score.

underline the importance of evaluating more refined representations within a discourse. Here, narratives provide insights into resonant or harmful components of the discourse that evaluation at the keyword-level would have missed.

Finding Prevalent Narratives

Post content extraction. To ensure clustering based on relevant information, we pre-process the post text, removing non-alphanumeric characters, URLs, and emojis as well as standardizing using casing and lemmatization. Samory and Mitra (2018) demonstrate the importance of agent-action-target triplets in uncovering conspiracy theory narratives, isolating common motifs. This approach, also used outside of the conspiracy domain, allows for differentiation between key elements of narrative structure to focus meaning (Gildea and Jurafsky 2002). For example, "We need to have vote at your precinct location again" is "we need vote". We expand this by incorporating spatial and temporal features to gain insights into *where* and *when* an agent did an action to a target. Using dependency parsing, we extract subjects, objects, roots, and prepositions to uncover the *what, when*, and *where* of surfaced narratives. Our method extracts "we need have vote at your precinct location". (See Figure 2a - 2c.)

Narrative extraction. We temporarily remove duplicate post extracts before clustering to improve performance; this content is reincorporated for accurate prevalence observations after narratives are uncovered. We focus on posts that met two criteria. First, to meet a meaningful threshold of content to embed, the post extract must have a subject, verb, and object. Second, to ensure content within scope, one of these three elements must be a keyword from our query. We train a Word2Vec model on our data — filtering text with an expanded query of the original query term subjects and a curated list from their top 25 semantically similar phrases. To further validate the expanded query scope, we also examine the top n-grams occurring in 50 or more posts for missing

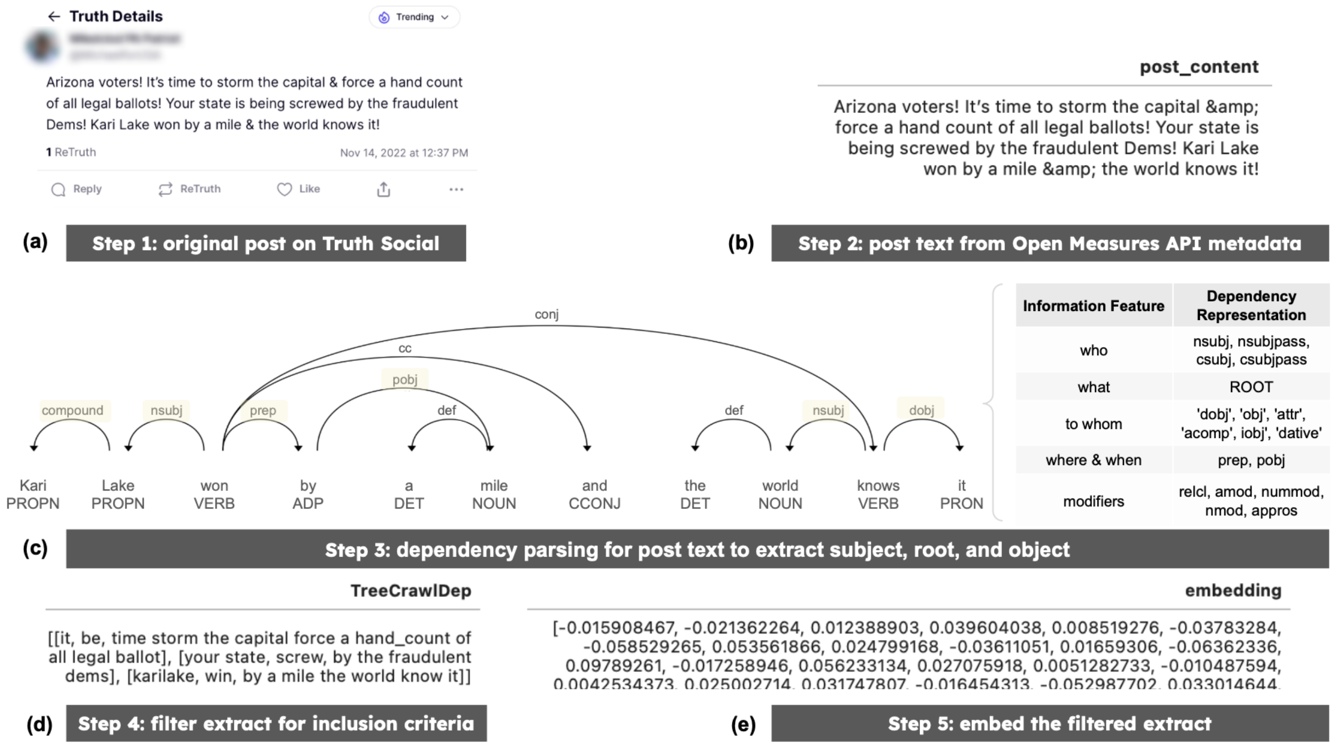

Figure 2: Post content preparation for finding prevalent narratives: (a) post in situ, (b) metadata for analysis, (c) dependency extraction from metadata, (d) filtering extracts, (e) embedding extracts

permutations. We then vectorize the post extracts using T5 Sentence Encoder (Ni et al. 2021). (See Figure 2d -2e.) We reduce the dimensionality of the embeddings by projecting into a two-dimensional space with UMAP, and use KMeans clustering to surface narratives from semantically similar posts. We determine model best fit by calculating distortion for a range of number of clusters, k = 2 to 200, minimizing for both using the elbow method. We adjusted hyperparameters, comparing results over multiple ablations, evaluating cluster topic cohesion by manual validation of random samples. Lastly, we reincorporate the duplicate posts to manually assess which clusters are prevalent.

RQ1 Results: Prevalent Narratives

We identify 10 clusters, each comprising between 4.95% and 13.96% of the total number of posts, with a ratio of posts to authors between 0.64% and 0.95%. Subsequent analysis focuses on the 9 containing 5% or more of the dataset. (See Appendix A, Table 3.) We refer to these as narratives, using short descriptions capturing the representation of election issues coherent in the cluster. We derive the descriptions from the narrative's top 50 occurring unigram, bigram, and trigram n-grams, manually validated through random samples.

Platform differences. We begin by examining the narrative distribution of posts and authors across platforms. (See Figure 3.) With the exception of *Election updates & statistics*, the total posts and total authors across narratives echoed the dataset-wide distributions. This anomaly is due to a few prolific Gab authors within the narrative who posted between 3 and 15 times the norm of the top authors on other platforms both within and across narratives. These author outliers were regular users deeply engaged with sharing and commenting on the evolving results throughout the week. Most narratives lacked dominant authors driving the discourse *within*, with a mean ratio of authors to posts of 0.89%, and *between*, with less than 20% of platform authors shared between any pair of narratives.

Affective differences. To better evaluate potential harms associated with narratives, we examine negative affect across multiple dimensions. To indicate harm in online content we measure posts for *toxicity*, or lack of civility, respect, and reason. To indicate potential harm to offline security we measure posts for *threats*, or intent to physically harm. Finally, to indicate potential personal harm we measure *insults*, or scorn and negativity, and *identity attacks*, or identity-based scorn and negativity. We identified differences between narratives as well as the distribution of scores over time. (See Appendix B, Figure 9.) We find that the majority of posts contain low likelihood of such harms, with only a few posts showing high likelihood of insults and identity attacks. The dimension with the highest likelihood in the largest quantity was toxicity across all narratives, with narratives discussing alleged voter suppression (*Voter suppression by ballot*, *Voter suppression by party*, and *Voter suppression by process*) and alleged rigged elections (*Rigged for Hobbs* and *Election meddling*) yielding greater quantities.

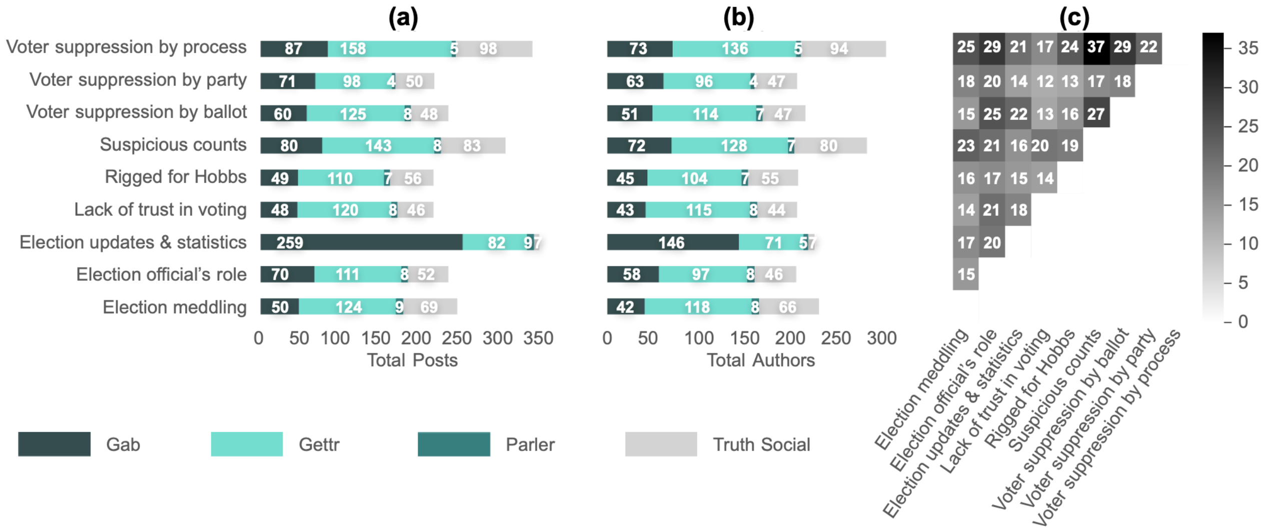

Figure 3: Narrative distribution of (a) posts by platform, (b) authors by platform, and (c) authors between platforms. For (a) and (b), color shows platform. For (c) color corresponds to a scale of author count. For all figures, the Y-axis indicates narrative. The X-axis indicates for (a) the total number of posts, (b) the total number of authors, and (c) narratives that share authors.

As more time passed following the election, toxicity also increased in posts discussing *Election updates and statistics*. During that same period it decreased in narratives expressing mistrust (*Suspicious counts* and *Lack of trust in voting*) and discussing *Election updates & statistics*. Insults were the second most prevalent harm across all narratives. We find that insults scores mostly followed the trends seen in toxicity scores, with slightly more uncertainty.

Engagement differences. We measure engagement across three dimensions: *reaction*, *reply*, and *reshare*. To enable cross-platform comparison, engagement scores are normalized within platform, where a 1 is equivalent to the highest raw metric score in the data and a value of 0 is equivalent to the minimum raw metric score. A *reaction* is each site's equivalent of a Twitter "like", a *reply* is each site's equivalent of a Twitter "reply", and a *reshare* is each site's equivalent of a Twitter "retweet". (See Table 1.) Across narratives, we find low engagement, with median engagement near 0. (See Appendix C, Table 4.) The differences in median engagement scores across narratives were statistically significant according to Kruskal-Wallis tests (*reaction*: $\chi^2 = 24.79, p < .05$, *reply*: $\chi^2 = 24.42, p < .05$, *reshare*: $\chi^2 = 34.13, p < .001$). Temporal analysis shows engagement outliers vary across clusters and over time. The highest recorded number of reactions, replies, and reshares were in *Election official's role*.

All narratives contained some posts with no engagement. We randomly sample low engagement data to assess corresponding authors reach and volume, and a majority of the posts indicate the "babble effect" observed in online rumoring, where low exposure and low volume give the impression of users "shouting into the void" (Arif et al. 2016). With respect to median engagement scores, we find that posts discussing the election (*Election updates & statistics* and *Election officials' role*) yielded the highest *reaction* and *reshare* engagement. All narratives had the same median *reply* engagement score of 0, with the exception of *Election updates & statistics*. This follows the narrative agnostic scores, dominated by low or no engagement outside of outlier posts usually authored by well-known entities with significant followings.

RQ2: Social Identity Dynamics of Narratives

To study how social identity is evoked in online narratives, we use a lexical-based method to evaluate how in-groups and out-groups — and in particular associated power — are represented in posts by platform and narrative. We find minimal variation in identity group power differentials across platforms, except Parler — which showed the minimum power measures for the out-group cohort and positive maximum power measures for the in-group cohort. Both platform and narrative-level analysis showed that post audiences were framed as comparatively less powerful than authors. These results affirm theories around how group identities and intergroup power dynamics can motivate groups to improve the conception of their collective social identities.

Characterizing Dynamics of Social Identities in Prevalent Narratives

Given increasing recognition of social identity as a key dimension of problematic narratives, we study in-group and out-group dynamics to compare posts' expressions of power (Robertson et al. 2022). While the composition of group identities can vary across narratives, the majority here are affiliated with a political conceptualization. This is concerning, as radicalization can occur when such groups are frustrated or their influence threatened (Klandermans 2014). This can lead to collective action, considered in RQ3.

Perdue et al. (1990) define in-group pronouns as the first person plural and out-group pronouns as the third person plural. We extend this definition to also capture singular pronouns to ensure we are capturing all instances of in-group and out-group entities speaking, being spoken to, or being spoken about. Further associated identifiers are surfaced through co-referencing. We examine in-group references through two dimensions: *In-group (Singular Excluded)* which represents the first person plural where the author is addressing the audience as members of their in-group and *In-group* which includes both first-person plural and singular where the author is either reflecting on themselves or their in-group. We examine out-group through one dimension: *Out-group* which represents the third person singular and plural where the author is addressing third parties.

We calculate power differential scores using Antoniak et al.'s (2023) RIVETER python package, which incorporates a lexicon of power connotation frames. This method enables us to evaluate over 1,700 authority-conferring verbs in the context of the verb's agent, theme, and sentiment. RIVETER also handles co-reference resolution and customized evaluation groups. The directional scores derived indicate the polarity of power associated with a verb, implying connotations of power dynamics between the verb's subject (agent) and direct object (event). The scores are scaled measures of power differentials where 1, 0 and -1 indicate positive, neutral and negative polarity, respectively. If John corrects Jane, then John has a positive power differential. If John describes Jane, then John has an equal power differential. If John echoes Jane, then John has a negative power differential. We use regular expressions to extract the scores that correlate with our custom evaluation groups[2]. We calculate three aggregate scores at the level of documents, narratives, platforms, and the entire corpora using RIVETER's *get_scores_for_doc* and *get_score_totals* methods, corresponding power measures for the aggregated data.

RQ2 Results: Social Identity Dynamics

Platform differences. Figure 4 shows the distribution of identity group power differentials across platforms. We find that with the exception of Parler, platform power differential closely align to the baseline identity group scores, all comprising negative power differentials (dashed line and corresponding annotation). The anomaly of Parler could be impacted by relatively smaller number of Parler posts, comprising only 3% of the data. This is likely due to the drop in user activity following its role in the January 6th insurrection and temporary deplatforming from app stores (Cameron 2023). It could also be impacted by platform-specific norms or the small number of posts. On Parler, in-group identities were portrayed with a positive power differential, compared to other platforms that show only negative power differentials. Additionally, Parler demonstrated both the maximum and minimum power differentials of all platforms. *In-group* had comparatively higher scores than *In-group (Singular Excluded)*, suggesting that the audience is likely addressed as being more negatively impacted by the post content relative to the author. This is supported by manual verification of post content, where the audience is often addressed as being harmed by the midterm election issues in Maricopa County.

Narrative differences. Compared to the baseline and platform-level scores, the narratives in the data show more variation in the expressed in-group and out-group power differentials. However, the trend of lower power differentials for *In-group (Singular Excluded)* messaging persisted. Another consistent finding was that in-group and out-group negative power differentials persisted, with the exception of the neutral score for posts referencing the narratives of *Voter suppression by ballot* addressing the *In-group (Singular Excluded)* and *Rigged for Hobbs* addressing the *In-group*. In Figure 4, the distribution of identity group power differentials across narratives show out-group power scores above the baseline in narratives related to the counting of votes (*Suspicious counts* and *Election meddling*). This slightly higher power differential might be attributed to the focus of many posts on the power exerted by third-parties over election outcomes. With respect to in-group cohorts, we find that exclusion versus inclusion of the first-person singular yields baseline-crossing deltas between narrative power scores. Narratives related to the counting of votes have lower-than-baseline power differentials when the authors also address themselves as within the *In-group* compared to as a collective with their audience (*In-group (Singular Excluded)*). This suggests that the authors are using posts to express feeling aggrieved by the process. The positive power differential for *In-group (Singular Excluded)* compared to a more neutral score might indicate a focus on messaging that the audience has the power to counter alleged voter suppression, a common sentiment found in manual validation of the narratives.

RQ3: Mobilization Potential of Narratives

The success of a narrative in theoretically prompting collective action depends on the extent to which Snow, Vliegenthart, and Ketelaars's (2018) three core *framing tasks* are present: *diagnostic* (highlighting the problem or the root cause of the problem), *prognostic* (presenting solutions to counter the problem or refuting the solutions of an adversary), and *motivational* (rationalizing or motivating a call to action). Yet the authors caution that such frames are subject to different interpretations. To better understand how narratives vary across platforms and to what extent they support mobilization functions, we qualitatively label posts with Snow et al.'s framework. Further, to better understand the dynamics between mobilization and conspiracy theories, we manually label posts for mentions of conspiracy theories. Our analysis informs on collective action frames in an evolving election event, highlighting that while posts describing the problem were most prevalent, posts with the theoretical potential to inspire mobilization were also common. We describe observed behavior of these posts that contain all the elements that drive collective action, showing a sharp drop after the initial catalyst (Election Day) that slowly increase as time passes with no resolution. Further, by examining

[2] Out-group: 3rd-person singular & plural, In-group: 1st person singular & plural, In-group (Singular Excluded): 1st person plural

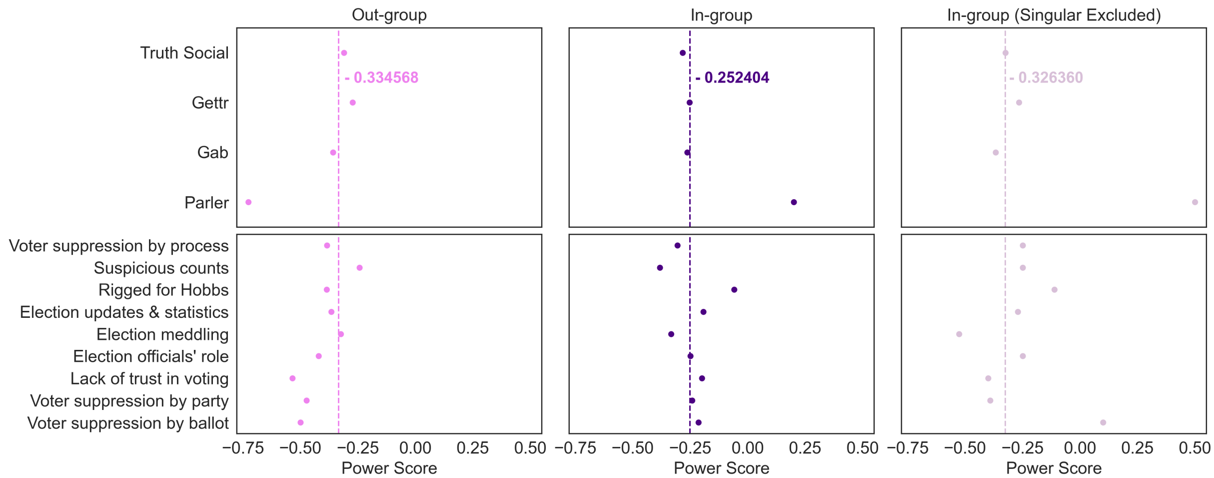

Figure 4: Distribution of identity group power differentials across platforms and narratives. Color represents identity group, the dashed line and corresponding annotation represents the identity group's power differential across all the data, and the X-axis represents the power score differentials. The Y-axis in the first row specifies platform and in the second row specifies narrative.

mobilization potential in the context of election conspiracy theories, we see that across platforms posts with conspiracy theories most often have the highest mobilization potential, highlighting how such theories might inspire action.

Characterizing Mobilization Through Collective Action Frames in Prevalent Narratives

Two graduate students independently reviewed and manually labeled 1% of the dataset as *diagnostic*, *prognostic*, *motivational*, or *not applicable*, with a Cohen's kappa of 0.72. This agreement outperforms prior work in this space with similar coding schemes (Phadke et al. 2018). Not all posts explicitly identified a problem or cause, however many referenced or alluded to them. We consider such implied problems as diagnostic if they were identified as Election Day problems by Election Integrity Partnership[3] reporting. (See Appendix D, Table 5.) All other implied problems are considered not applicable to ensure the codebook is generalizable to those without domain expertise. We expand on Snow et al.'s heuristics for determining framing by adding additional considerations given the social media context. Specifically, we consider the use of capitalization, punctuation, and emojis in expressing the motivational elements of *severity* and *urgency*. Posts were also annotated for references to *conspiracy theories*, with mention of (a) a person, a group, or an institution of conspiring entities take on some kind of action that targets someone or something for an evil, illegal, or harmful purpose or (b) a shorthand reference to a well-known election-related conspiracy theory[4], with a Cohen's kappa of 0.67. After establishing agreement, one author labeled 10% of the dataset where the distance to a narrative centroid is in the top 25% closest posts. This sampling technique allows us to focus the qualitative analysis on the posts most central to the most prevalent narratives.

Platform	Framing Task Annotation	% of Posts	% of Conspiracy Theory Posts
Gab	diagnostic	38.38	17.84
	diagnostic & prognostic	17.84	6.49
	diagnostic & motivational	17.84	11.35
	diagnostic & prognostic & motivational	25.95	18.38
Gettr	diagnostic	32.24	11.84
	diagnostic & prognostic	18.78	4.9
	diagnostic & motivational	23.27	16.73
	diagnostic & prognostic & motivational	25.71	16.33
Parler	diagnostic	42.86	7.14
	diagnostic & prognostic	21.43	7.14
	diagnostic & motivational	14.29	14.29
	diagnostic & prognostic & motivational	21.43	21.43
Truth Social	diagnostic	35.51	17.76
	diagnostic & prognostic	15.89	3.74
	diagnostic & motivational	22.43	13.08
	diagnostic & prognostic & motivational	26.17	16.82

Table 2: Table of platform framing task metrics.

RQ3 Results: Mobilization Potential of Narratives

Platform differences. We observe that nearly every post across platforms cited issues without offering solutions or motivation (*diagnostic*), showing authors found the issues in the case study to be problematic. (See Figure 2.) Posts that contain all three framing tasks (*diagnostic & prognostic & motivational*), which Snow et al. maintain are more likely to inspire collective action, were the second most common across platforms. With respect to conspiracy theory-related content, posts with all three framing tasks present were the most frequently occurring annotation. *Diagnostic & prognostic & motivational* comprised between 16% and 21%

[3]https://www.eipartnership.net/blog/about-eip-2022

[4]One World Order, deep state, voter fraud, Stop the Steal, "manufacturing of votes", "ballot drops", "slow rolling" of ballot counting to change results, and QAnon

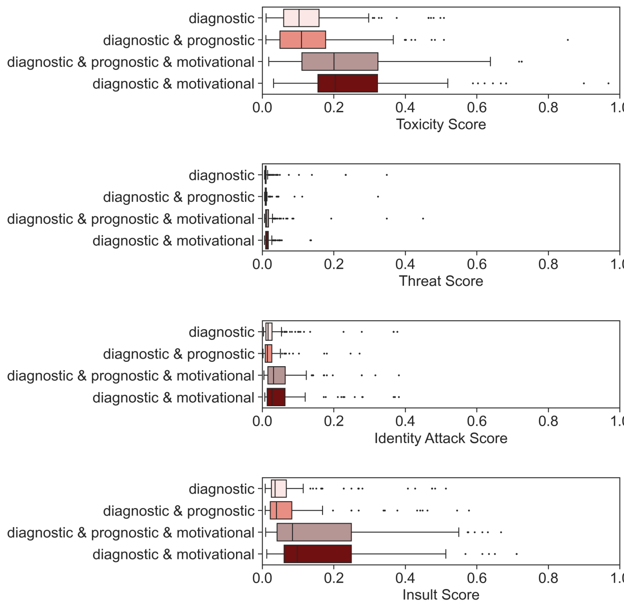

Figure 5: Distribution of toxicity, threat, identity attack, and insult harm scores by framing task. The Y-axis and color shows the framing task annotation. The X-axis shows the distribution of the harm score. Each facet is a harm.

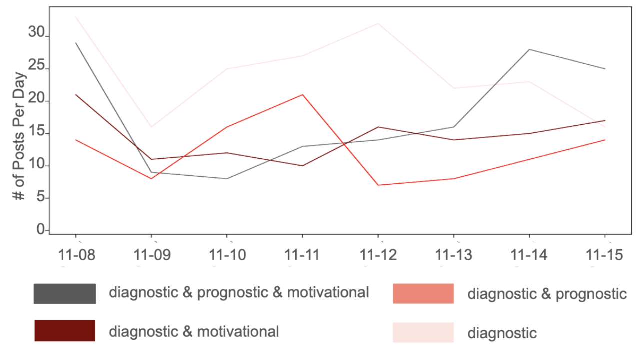

Figure 6: Distribution of posts by framing tasks over time. The Y-axis shows the daily post count, X-axis shows the timestamp, and color shows the framing task annotation.

across platforms, except for on Truth Social where such posts trailed *diagnostic* by less than 1%. This might indicate that posts with conspiracy theory-related content have a higher mobilization potential than those without.

Affective differences. Figure 5 depicts the distribution of post harm scores by framing tasks during the week after the election. The majority of posts limited to *diagnostic* framing were least likely to be considered harmful. We find that posts designated as more likely to be toxic or insults were those that included a motivational frame. Figure 6 shows that motivational posts were observed in smaller amounts in the data. Of note, diagnostic posts decreased as time elapsed following the midterms, while posts with theoretical mobilization potential increased. This might indicate that over time users were inspired to mention potential solutions and the need to address the issues. Yet despite the shift in framing tasks, harm score distributions remained consistent.

Engagement differences. As individual posts have the ability to go viral and be shared across platforms, it is important to evaluate audience engagement. While all framing tasks contained posts with outlier engagement, Figure 7 demonstrates that overall engagement remained low. Posts containing motivational framing failed to motivate audience engagement, showing the lowest levels of reactions and reshares but the highest numbers of replies. The differences in median engagement scores across framing tasks were not statistically significant according to Kruskal-Wallis tests (*reaction*: $\chi^2 = 1.63, p > .05$, *reply*: $\chi^2 = 0.22, p > .05$, *reshare*: $\chi^2 = 2.96, p > .05$). The engagement outlier posts contained the three framing tasks, with the exception of *diagnostic* and *diagnostic & prognostic*-framed *Election meddling* posts. Users engaged with conspiracy theory-driven narratives more often than those which were not. They were also more likely to use a *reaction* than to *reshare* or *reply*.

Narrative differences. Our evaluation of this manually validated narrative subset shows that while no one narrative is characterized by consistent inclusion of all three frames that together inspire collective action, individual posts often do. Posts are most often included *diagnostic* framing, with authors highlighting or commenting on perceived issues surrounding the incident. (See Figure 8(a).) Discussion of solutions to these issues are less frequently discussed, less than half as often as the problems across narratives. Similarly, *motivational* framing were also less frequently discussed, except for in posts concerning *Election officials' role*. Figure 8(b) shows that authors used *motivational* framing more frequently when discussing conspiracy theories in narratives. We also find that posts sometimes offered motivational framing that might inspire action in the absence of solutions, where a problem was identified and a need was justified.

Discussion

Small subsets, big concerns. Our study highlights the outsized role that individual posts can pose in perpetrating potential harm both online and offline. Despite smaller audiences, alt-tech content, and in particular high engagement content, still contain concerning attributes that are important to understand with respect to trust and safety. Content scored with a higher likelihood of harms (> 0.5) comprised 5%, 4%, <1%, and <1% respectively. This, combined with our finding that posts with some of the highest engagement were characterized by patterns of increased toxicity, insults, and conspiracy theory references, suggest the importance of giving serious consideration to all narratives irrespective of aggregate size or platform prominence. Any harmful content, in particular content which generates high engagement, can inspire behavior and consequences. Using the lens of narratives to evaluate information, scholars and practitioners can better understand what harmful and conspiracy theory-related frames are impacting and inspiring social media au-

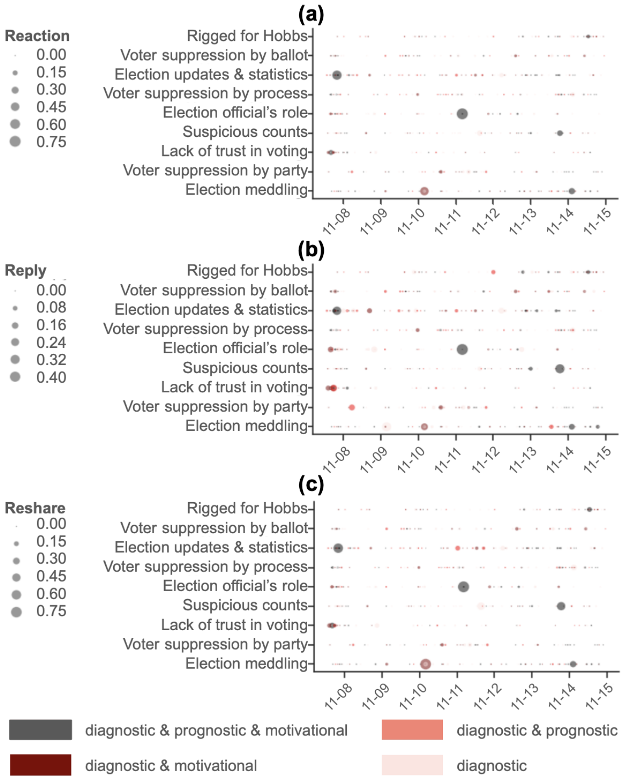

Figure 7: Post engagement by narrative over time. The X-axis shows timestamp, Y-axis shows narrative, color shows framing task annotation, and size shows engagement score for (a) reactions, (b) replies, and (c) reshares.

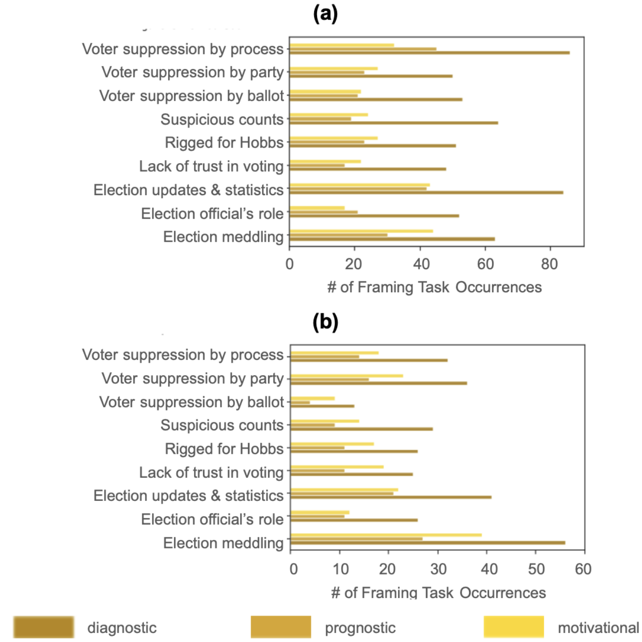

Figure 8: (a) Number of individual *framing task* occurrences observed across posts. (b) Individual *framing task* occurrences observed across posts that contain conspiracy theories. For both figures, the X-axis represents the number of *framing task* occurrences across each narrative specified on the Y-axis. Color represents the *framing task* annotation.

thors and audiences across platforms. This study indicates a need for continued work to better extract meaningful narratives from online data to better understand event discourse.

The negative framing of in-group power across the majority of platforms aligns with the tenets of the social identity theory of intergroup behavior (Tajfel and Turner 2004). By utilizing language positioning author and audience at a loss to out-groups, narratives with greater power differentials exploit framing that motivates allegedly subordinate groups towards actions to improve their social identity. Our findings show this behavior across platforms, with the exception of Parler (an outlier with less than 100 overall posts and the lowest percentage of motivational posts). We see this both within the prior work and our data, for example when a user expressed "another rigged election moves us closer to the demise of America. Everyone in every state should cure their ballot" (McKeown, Haji, and Ferguson 2016).

With respect to the mobilization potential of narratives, it is important to note that the majority of the higher engagement posts also contain the essential framing Snow et al. suggest inspires collective action. This is critically important given the harms found in such posts. Future research could establish if a causal relationship exists between online posts and measurable collective action outcomes with implications for understanding how audiences respond to different kinds of content. Another confounding and crucial consideration is that the majority of high engagement content was also conspiracy theory content — a now increasingly present element in the U.S. elections context. Future work might explore how collective action framing tasks are articulated in a conspiracy theory context compared to general content. Our work also broadens the consideration of mobilization potential changes over time, particularly in that certain nuances of online discourse can inspire collective action along different timelines. Indeed, understanding if these trends extend to longer study periods and different case studies will be an important avenue for further investigation.

Towards cross-platform and cross-disciplinary evaluation. Our results highlight the importance of distinguishing between nuanced narratives in online discourse across the dimensions of cross-platform distribution, affective harm, engagement, social dynamics, and framing. This has implications for surfacing and evaluating the spread of content that is both harmful online and has the potential to mobilize into offline harms. Our focus on tracking discourse and dynamics on alt-tech platforms concerning a complex and evolving incident narrative indicated distinct patterns with respect to harmful content, audience engagement, framing of intergroup expressions of power, and framing of conspiracy theories. By understanding beliefs and behaviors on alt-tech, we could alert to fringe or extremist beliefs, particularly those that might target *out-group* entities or mobilize into violence

— as content on Parler did during the Capitol Insurrection. By evaluating narratives from the perspective of both the differences in their characteristics as well as their social identity and collective action framing, we can find elements of high-impact discourse and help counter them. For example, "babble effect" posts in a narrative with mobilization potential might present less of individual concern, but in aggregate may indicate a cohesive movement. Likewise, threatening posts trigger concern despite low reach given prior indications of harm on these platforms materializing offline and the outsized impact of harm regardless of scale. Further, an understanding how affective harms and group social identity materialize can inform how it is mitigated. As the fragmentation of the social media landscape continues, an understanding of prevalent narratives across platforms, including fringe or nascent ones, and of their associated publics will become even more critical for understanding evolving norms, forecasting threats, and comprehending the public domain.

Ethical concerns and limitations. Our data collection is limited to posts containing pre-defined query terms, so we potentially miss related discourse with uncommon or subversive lexical variations. Prior work also uses keywords to study the dynamics and patterns of online information and discourse related to political narrative and election studies (Kennedy et al. 2022; Matatov, Naaman, and Amir 2022; Dash et al. 2021). To minimize both the loss of such content and noise, we derive our query from contextual experts on cross-platform content in the 2020 and 2022 U.S. elections.

To better surface prominent narratives in election-related discourse, we clustered extracted dependencies from posts based on assumptions of what content is central. We validated random samples to ensure extraction of core discourse when selecting dependencies, yet it is possible we excluded relevant content through the data cleaning or extraction processes. Although our work considers public domain data, metadata are not collected directly from platform-provided APIs. Without official documentation and access, our understanding is limited to retrieved data. To counter this, we examined metadata over time, manually assessing field names, content, and corresponding on-site metrics. We limited scope to fields consistent within and across platforms.

Users are potentially unaware their data contributes to research. We obscured usernames by hashing with salt and present aggregate results. A vector of misuse is the application of the methods without validation. We document steps to filter and validate. It is possible bad actors use our findings to manipulate engagement metrics. We suggest that potential gains for countering problematic content outweigh the risk.

Conclusions

We offer a platform and domain-agnostic approach for evaluating social media narratives, demonstrated through a contested election case study on alt-tech platforms and examining discourse through the dimensions of narrative, platform, affect, and engagement. We find that exploring discourse through a more nuanced lens of narratives highlights differences in conceptions of social identity dynamics and mobilization. This is critical for both evaluating public sensemaking and forecasting of intergroup conflict and collective action. Our findings inform future work examining narratives and associated affective and engagement behaviors across alt-tech. We contribute insights into how dynamics of power frame and motivate collective action in an election context.

Acknowledgments

This work was supported by the National Science Foundation (2120496), William and Flora Hewlett Foundation (2022-00952-GRA), and John S. and James L. Knight Foundation (G-2019-58788). We also thank Sydney DeMets.

References

Antoniak, M.; Field, A.; Mun, J.; Walsh, M.; Klein, L.; and Sap, M. 2023. Riveter: Measuring Power and Social Dynamics Between Entities. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*.

Arif, A.; Shanahan, K.; Chou, F.-J.; Dosouto, Y.; Starbird, K.; and Spiro, E. S. 2016. How information snowballs: Exploring the role of exposure in online rumor propagation. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*.

Benner, K. 2021. Parler says it sent the F.B.I. posts about threats to the Capitol ahead of Jan. 6. https://www.nytimes.com/2021/03/25/us/parler-fbi-capitol-attack.html. (Retrieved 4/30/23).

Brewer, M. B. 1999. The psychology of prejudice: Ingroup love and outgroup hate? *Journal of Social Issues*, 55(3).

Cameron, C. 2023. Parler, Social Media Site Sidelined After Jan. 6, Plans Return. https://www.nytimes.com/2023/12/18/us/politics/parler-relaunch-2024.html. (Retrieved 4/29/24).

Dash, S.; Arya, A.; Kaur, S.; and Pal, J. 2022. Narrative Building in Propaganda Networks on Indian Twitter. In *Proceedings of the 14th ACM Web Science Conference 2022*.

Dash, S.; Shekhawat, G.; Akbar, S. Z.; and Pal, J. 2021. Extremism & whataboutism: A case study on bangalore riots. *ACM CSCW Workshop on Addressing Challenges and Opportunities in Online Extremism Research*.

Davey, J.; and Ebner, J. 2017. The fringe insurgency. Connectivity, convergence and mainstreaming of the extreme right. *Institute for Strategic Dialogue*.

Election Integrity Partnership. 2022. The Election Integrity Partnership in 2022: Our Work Ahead. https://www.eipartnership.net/blog/about-eip-2022. (Retrieved 4/15/23).

Engel, K.; Phadke, S.; and Mitra, T. 2023. Learning from the Ex-Believers: Individuals' Journeys In and Out of Conspiracy Theories Online. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW2).

Gildea, D.; and Jurafsky, D. 2002. Automatic labeling of semantic roles. *Computational Linguistics*, 28(3).

Horta Ribeiro, M.; Hosseinmardi, H.; West, R.; and Watts, D. J. 2023. Deplatforming did not decrease Parler users' activity on fringe social media. *PNAS Nexus*, 2(3).

Jachim, P.; Sharevski, F.; and Pieroni, E. 2021. TrollHunter2020: Real-time detection of trolling narratives on

Twitter during the 2020 US elections. In *Proceedings of the 2021 ACM Workshop on Security and Privacy Analytics*.

Jigsaw. 2024. Perspective. https://www.perspectiveapi.com. (Retrieved 1/3/24).

Jolley, D.; Meleady, R.; and Douglas, K. M. 2020. Exposure to intergroup conspiracy theories promotes prejudice which spreads across groups. *British Journal of Psychology*, 111(1).

Kavrakis, K. 2023. Identity and ideology through the frames of Al Qaeda and Islamic State. *Terrorism and Political Violence*, 35(5).

Keith Norambuena, B. F.; Mitra, T.; and North, C. 2023. A survey on event-based news narrative extraction. *ACM Computing Surveys*, 55(14s).

Kelety, J. 2022. All votes counted in Maricopa County, despite online claims. https://apnews.com/article/fact-check-maricopa-county-ballots-uncounted-214752614581. (Retrieved 4/30/23).

Kennedy, I.; Wack, M.; Beers, A.; Schafer, J. S.; Garcia-Camargo, I.; Spiro, E. S.; and Starbird, K. 2022. Repeat spreaders and election delegitimization: a comprehensive dataset of misinformation tweets from the 2020 US election. *Journal of Quantitative Description: Digital Media*, 2.

Klandermans, P. G. 2014. Identity politics and politicized identities: Identity processes and the dynamics of protest. *Political Psychology*, 35(1).

Matatov, H.; Naaman, M.; and Amir, O. 2022. Stop the [Image] steal: The role and dynamics of visual content in the 2020 US Election Misinformation Campaign. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2).

McKeown, S.; Haji, R.; and Ferguson, N. 2016. Understanding peace and conflict through social identity theory. *Contemporary Global Perspectives*.

Ni, J.; Abrego, G. H.; Constant, N.; Ma, J.; Hall, K. B.; Cer, D.; and Yang, Y. 2021. Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. *arXiv preprint arXiv:2108.08877*.

Open Measures. 2023. Pro API. https://openmeasures.io. (Retrieved 2/8/23).

PBS. 2023. Here's how to watch the 4th Republican presidential debate. https://www.pbs.org/newshour/politics/heres-how-to-watch-the-4th-republican-presidential-debate. (Retrieved 1/12/24).

Pentzold, C.; Zuber, C.; Osterloh, F.; and Fechner, D. J. 2022. How to make sense of nonsense: Political absurdity and parodic memes in the# sharpiegate affair. *International Journal of Communication*, 16.

Perdue, C. W.; Dovidio, J. E.; Gurtman, M. B.; and Tyler, R. B. 1990. Us and Them: Social Categorization and the Process of Intergroup Bias. *Journal of Personality and Social Psychology*, 5(3).

Phadke, S.; Lloyd, J.; Hawdon, J.; Samory, M.; and Mitra, T. 2018. Framing hate with hate frames: Designing the codebook. In *Companion of the 2018 ACM Conference on Computer-Supported Cooperative Work & Social Computing*.

Piper, A.; So, R. J.; and Bamman, D. 2021. Narrative theory for computational narrative understanding. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.

Prochaska, S.; Duskin, K.; Kharazian, Z.; Minow, C.; Blucker, S.; Venuto, S.; West, J. D.; and Starbird, K. 2023. Mobilizing Manufactured Reality: How Participatory Disinformation Shaped Deep Stories to Catalyze Action during the 2020 US Presidential Election. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1).

Puckett, K. 2016. *Narrative theory*. Cambridge University Press.

Robertson, C. E.; Pretus, C.; Rathje, S.; Harris, E. A.; and Van Bavel, J. J. 2022. How social identity shapes conspiratorial belief. *Current Opinion in Psychology*, 47.

Rumble. 2023. Rumble Reports Third Quarter 2023 Results. https://corp.rumble.com/blog/rumble-reports-third-quarter-2023-results/. (Retrieved 1/12/24).

Russo, G.; Verginer, L.; Ribeiro, M. H.; and Casiraghi, G. 2023. Spillover of antisocial behavior from fringe platforms: The unintended consequences of community banning. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17.

Ryan, M.-L. 2007. *Toward a definition of narrative*. The Cambridge Companion to Narrative. Cambridge University Press.

Samory, M.; and Mitra, T. 2018. 'The Government Spies Using Our Webcams' The Language of Conspiracy Theories in Online Discussions. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW).

Schafer, J.; Duskin, K.; Prochaska, S.; Wack, M.; Beers, A.; Bozarth, L.; Agajanian, T.; Caulfield, M.; Spiro, E.; and Starbird, K. 2025. ElectionRumors2022: A Dataset of Election Rumors on Twitter During the 2022 US Midterms. *Journal of Quantitative Description: Digital Media*, 5.

Snow, D. A.; Vliegenthart, R.; and Ketelaars, P. 2018. The framing perspective on social movements: Its conceptual roots and architecture. *The Wiley Blackwell Companion to Social Movements*.

Sternisko, A.; Cichocka, A.; and Van Bavel, J. J. 2020. The dark side of social movements: Social identity, non-conformity, and the lure of conspiracy theories. *Current Opinion in Psychology*, 35.

Tajfel, H. 1982. Social psychology of intergroup relations. *Annual Review of Psychology*, 33(1).

Tajfel, H.; and Turner, J. C. 2004. The social identity theory of intergroup behavior. In *Political Psychology*. Psychology Press.

Turner, J. C.; and Oakes, P. J. 1986. The significance of the social identity concept for social psychology with reference to individualism, interactionism and social influence. *British Journal of Social Psychology*, 25(3).

Wilson, T.; Zhou, K.; and Starbird, K. 2018. Assembling strategic narratives: Information operations as collaborative work within an online community. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW).

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
   (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, please see Background and Related Work, Finding prevalent narratives, Characterizing dynamics of social identities in prevalent narratives, and Characterizing mobilization through collective action frames in prevalent narratives.
   (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, please see Ethical Concerns and limitations.
   (f) Did you discuss any potential negative societal impacts of your work? Yes, please see Ethical Concerns and limitations.
   (g) Did you discuss any potential misuse of your work? Yes, please see Ethical Concerns and limitations.
   (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, please see Ethical Concerns and limitations.
   (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? This study is not designed to test specific hypotheses and therefore does not articulate hypothesis for each research question. However, using a case study approach it motivates a general comparison of narrative elements and engagement across social media platforms. The motivation for these comparisons is discussed in the Related Work section of the manuscript.
   (b) Have you provided justifications for all theoretical results? Yes, cross-platform and narrative comparisons are motivated in the Introduction and Related Work sections of the paper.
   (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? The comparisons drawn here involve a null hypothesis of no difference across platforms or narratives.
   (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? For each result, we discuss a range of possible mechanisms that provide explanation for the finding.
   (e) Did you address potential biases or limitations in your theoretical framework? Fully addressing the limitations of the theoretical framework is beyond the limited scope of this paper, however, we ensure prior work on the topic is references throughout.
   (f) Have you related your theoretical results to the existing literature in social science? Yes, for each research questions we extend prior work and link results to core theories of social identity and collective action in the social sciences.
   (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? Yes, please see the Discussion and Conclusion sections of the manuscript.
3. Additionally, if you are including theoretical proofs...
   (a) Did you state the full set of assumptions of all theoretical results? N/A
   (b) Did you include complete proofs of all theoretical results? N/A
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)? N/A
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? N/A
   (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)? N/A
   (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? N/A
   (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? N/A
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 if possible, please see references for the Open Measures SMAT API, the Google Perspective API, and the RIVETER paper. We use exisitng dataset but to preserve anonymity during peer review we leave the source anonymous temporarily.
   (b) Did you mention the license of the assets? No
   (c) Did you include any new assets in the supplemental material or as a URL? We are not releasing new assets.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? Yes, please see Ethical concerns and limitations.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes, and we hashed usernames to provide a layer of anonymity. However, given post content

and other metadata some alt-tech users could be identified with some effort. Additionally, although the analysis outcomes do not include identifiable information or offensive content, the underlying data does include both. This is why we are not making the data readily publically available, but rather access controlled.

(f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? N/A

(g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? N/A

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? N/A

(b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? N/A

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? N/A

(d) Did you discuss how data is stored, shared, and deidentified? N/A

Appendix

A: Narrative Content and Statistics

Cluster	Description	Percent of Data	Authors to Posts	Example Post
Narrative 1	Rigged for Hobbs	8.708709	0.247748	First, we demand that the 200 thousand Republican ballots in Box 3 in Arizona be counted, not thrown in the garbage. Second, we demand that Hobbs explain why she ordered 30% of voting machines in Maricopa County to be out of service denying them ink for the printers. Third, we want Hobbs to explain her lie that the Republican ballots in Box 3 were mixed in with counted ballots, and therefore thrown out in the garbage.
Narrative 2	Voter suppression by ballot	9.309309	0.244813	God bless youse are in the entire first family. Karilake really need your help. Day of voting in Maricopa: 30% of the machines just won’t work. They ran out of ink. They ran out of toner. They gave these Daya voters the wrong sized ballots which got kicked back and had to be adjudicated. Well we’ll just put your vote in box number three. Woopsie Woopsie Dipsy we accidentally mixed bag three with ballasted of already been counted… These lying conniving dirty mother Humpers
Narrative 3	Election updates & statistics	13.963964	0.190476	Well isn’t this Shocking and Surprising. Arizona’s Maricopa County Elections Department has addressed an Election Day issue where dozens of polling stations ran out of paper ballots, impacting tens of thousands of votes.A joint statement from Chairman Bill Gates and Vice Chairman Clint Hickman said Arizona elections officials will investigate the incident completely and assured voter tabulations will not ultimately be impacted.”All ballots will be counted securely and accurately,” they claimed.Over the past 24 hours, we have learned more about the printer issue that caused some ballots to not be read at Vote Centers yesterday. While the issue impacted less than 7% of Election Day voters (about 17,000 ballots), we understand that for people who went through it, this was frustrating, inconvenient, and not how they pictured Election Day,” they said.Any Bet it is Not counting Republican Votes.Florida counted 7.5 mil ballots in 5 hours and these states are taking days to count a couple million? Embarrassing and backwards, plus inefficient. Sounds sketchy too.The Democrats are becoming more Desperate than Paul “Always Hammered “ Pelosi sitting in a Gay Bar at Closing Time.
Narrative 4	Voter suppression by process	14.114114	0.258621	Someone needs to go Arizona ballot counters, grab them by the scruff of their fucking neck and make them do their fucking job. NOW!!!
Narrative 5	Election officials’ role	9.459459	0.240664	This was the response by Maricopa County Elections Dept. to a voter Tuesday morning. Their ballot was among the 17k ballots that were dropped in Bin 3.The employee ’confirmed’ the 17k ballots would be THE FIRST to be tallied yesterday morning. It didn’t happen.Why are they lying to voters now?
Narrative 6	Suspicious counts	12.012012	0.248408	Interesting that the race in AZ in 2020 was called early.Now that the MAGA race is involved it takes DAYS to count the ballots! 🤔
Narrative 7	Lack of trust in voting	8.858859	0.265766	This is horse sh!t!! Sabotage & Massive Cheating Under Way In Arizona Anyone that is in Arizona that is at a polling place with malfunctioning tabulators that will not scan and count your ballot needs to go to another location with working tabulators to cast their vote. They should not leave their ballots in a box to be ’counted later’.
Narrative 8	Voter suppression by party	8.708709	0.260090	Proof of votes for Republicans by someone from ”Democrats” Woman Reports Finding Election Ballots in Ravine in the Santa Cruz Mountains..and Naveda and Arizona at the TEXAS machine scan vote for Democrats go to faster...OMG...So alot states ”cheating vote” makes for Republica’s Lose at midterm election.,while the Republicans are leading very high point
Narrative 9	Election meddling	9.909910	0.246032	Arizona voters! It’s time to storm the capital & force a hand count of all legal ballots! Your state is being screwed by the fraudulent Dems! Kari Lake won by a mile & the world knows it!

Table 3: Table displaying cluster content and statistics dominating 5% of more of the data.

B: Narrative Scores for *Toxicity*, *Identity Attack*, *Insult*, and *Threat* Over Time

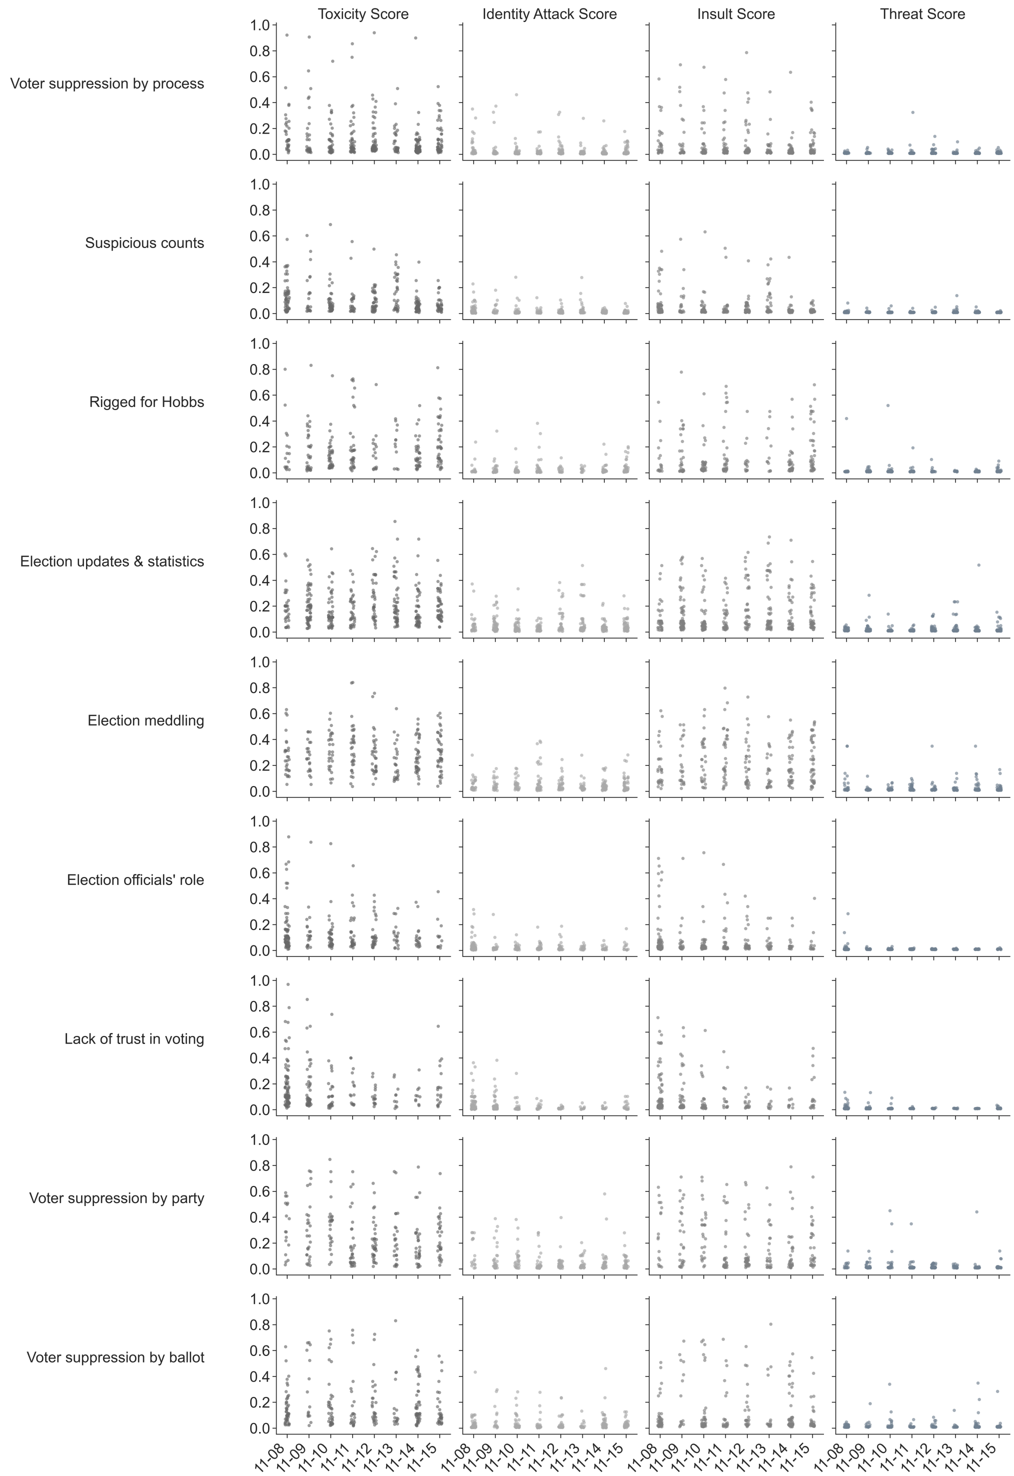

Figure 9: The X-axis shows the timestamp and both the Y-axis and color represent the harm score. The scores are scaled probabilities, between 0 being *unlikely* and 1 being *likely*.

C: Engagement Differences Across Narratives

Narrative	Reaction		Reply		Reshare	
	Median Score	***Standard Deviation***	***Median Score***	***Standard Deviation***	***Median Score***	***Standard Deviation***
Election updates & statistics	0.001148	0.046460	0.000000	0.036970	0.001927	0.057893
Voter suppression by process	0.000334	0.031942	0.001355	0.025721	0.000801	0.030333
Suspicious counts	0.000334	0.040199	0.001355	0.032307	0.000140	0.052303
Election meddling	0.000274	0.036551	0.001355	0.038082	0.000168	0.057636
Election officials' role	0.001148	0.051331	0.001355	0.039166	0.001927	0.059479
Voter suppression by ballot	0.000667	0.036705	0.001355	0.032373	0.001603	0.038504
Voter suppression by party	0.000334	0.015393	0.001355	0.021964	0.000056	0.014464
Lack of trust in voting	0.000334	0.026456	0.001355	0.037059	0.000056	0.041022
Rigged for Hobbs	0.000334	0.042644	0.001355	0.036510	0.000056	0.046269

Table 4: Table of narrative engagement score metrics. Engagement scores are normalized within platform, where a value of 1 equivalent to the highest raw metric score in the data and a value of 0 equivalent to the minimum raw metric score in the data.

D: Mobilization Codebook

Label	Description	Example
diagnostic	A diagnosis of some event or aspect as problematic and in need of repair or change and the attribution of blame or responsibility that addresses *"What is or went wrong?"* or *"Who or what is to blame"*. We consider a post diagnostic either when a problem is explicitly articulated or when a problem is implied that was covered by short suspense reporting by the EIP about Election Day.[a]	🚨 ATTN: ARIZONA 🚨 🚨 🚨 #ARIZONA - Some votes got kicked back.You need to CHECK YOUR BALLOT STATUS NOW:http://t.co/ZMvqDy0OdZ 🚨 And then CURE your ballotDeadline Is Wednesday 🚨 Make sure your vote gets counted-- Kari Lake Supporters Call For Military Intervention To Stop Election Fraud In Arizona Governor Race https://truthtent.com/kari-lake-supporters-call-for-military-intervention-to-stop-election-fraud-in-arizona-governor-race/
prognostic	A proposed solution to the problem, including a plan to carry it out and/or a refutation of the opponent's solution that addresses *"What should be done in response to the problem"*. We consider a post prognostic when a solution is explicitly articulated or when it asks specific entities for solutions to a perceived problem.	🚨 ATTN: ARIZONA 🚨 🚨 🚨 #ARIZONA - Some votes got kicked back.You need to CHECK YOUR BALLOT STATUS NOW:http://t.co/ZMvqDy0OdZ 🚨 And then CURE your ballotDeadline Is Wednesday 🚨 Make sure your vote gets counted-- Kari Lake Supporters Call For Military Intervention To Stop Election Fraud In Arizona Governor Race https://truthtent.com/kari-lake-supporters-call-for-military-intervention-to-stop-election-fraud-in-arizona-governor-race/
motivational	A call to arms or rationale for action that goes beyond diagnosis and prognosis, using a vocabulary of motive that provide prods to action such as 1) overcoming fear of risk, 2) attainment of a larger goal or public good by 1) accenting the severity of the problem, 2) the urgency of taking action now, 3) the probable efficacy of joining the cause, 4) the moral priority of doing so, or 5) the elevation of ones status that addresses *"Why we ought to do it"*. We consider a post motivational if it explicitly articulates these rationale or if it uses capitalization, punctuation, cursing, or emojis to express severity or urgency. Specifically, use of exclamation points, the use of all capital letters to yell (as opposed to headers or emphasis), expletives, and emojis that infer relevant meaning ‼️ 🚨 (urgent), 😡 😤 🤬 😖 😠 👿 😾 🗯️ (angry), ⏰ ⏳ ⏱️ (time - to include other timers and clock emojis), 🤯 💥 💣 🧨 (exploding), 🆘 ❌ 🛑 ⛔ 🚫 (red symbols for urgency and severity), ⁉️ ⚠️ ❕ (exclamation - to include other exclamation marks) .	🚨 ATTN: ARIZONA 🚨 🚨 🚨 #ARIZONA - Some votes got kicked back.You need to CHECK YOUR BALLOT STATUS NOW:http://t.co/ZMvqDy0OdZ 🚨 And then CURE your ballotDeadline Is Wednesday 🚨 Make sure your vote gets counted-- Kari Lake Supporters Call For Military Intervention To Stop Election Fraud In Arizona Governor Race https://truthtent.com/kari-lake-supporters-call-for-military-intervention-to-stop-election-fraud-in-arizona-governor-race/

Table 5: Table displaying our mobilization codebook. The labels of *diagnostic*, *prognostic*, *motivational*, and *not applicable* were sourced and extended from Snow et. al.'s collective action framing tasks (Snow, Vliegenthart, and Ketelaars 2018).

[a] EIP-highlighted election rumors: internet network issues indicate election fraud, Konnech and Dominion election infrastructure threaten results, ballot drops indicate election fraud, election officials compromise election through incompetence or intention, voting machine issued and counting offsite disenfranchise voters, obstruction of election facilities cover up tampering, slow counting of votes indicates fraud, voting infrastructure rigged to threaten election, poll worker collusion to influence election outcome, U.S. Postal Service collusion to influence election outcome, voter intimidation to influence election outcome, suspicious entities or lax security at polling locations threaten results