Comparing the Language of QAnon-Related Content on Parler, Gab, and Twitter

Andrea Šipka
andrea.sipka@dfki.de
University of Zurich
Zurich, Switzerland
Data Science and its Applications,
German Research Center for AI
Kaiserslautern, Germany

Anikó Hannák
hannak@ifi.uzh.ch
University of Zurich
Zurich, Switzerland

Aleksandra Urman
urman@ifi.uzh.ch
University of Zurich
Zurich, Switzerland

ABSTRACT

Parler, a "free speech" platform popular among conservatives, was taken offline in January 2021 due to the lack of moderation of harmful content. While other popular social media platforms were also used to spread conspiratorial, hateful and threatening content, Parler suffered the most consequences in the aftermath of the 2020 US presidential elections, having been singled out in the news coverage. Through a comparative study, we identify differences in content using #QAnon across three social media platforms, Parler, Twitter, and Gab, focusing on the volume, the amount of anti-social language, and the context of QAnon-related content over a month-long period. While the number of posts is the highest on Parler, this could be attributed to the differences in the use of hashtags on the platforms, which has consequences for other analyses. In our analysis, Parler exhibits the highest levels of anti-social language, while Gab has the highest proportion of #QAnon posts with hate terms. To get at qualitative differences in the posts, we perform analysis of named entities and narratives, focusing on differences in the focus of conversations and the levels of anti-social language of posts mentioning different groups of political figures.

CCS CONCEPTS

• **Human-centered computing → Social media**; • **Computing methodologies → Discourse, dialogue and pragmatics**.

KEYWORDS

cross-platform analysis, language, social media platforms, political discourse

ACM Reference Format:
Andrea Šipka, Anikó Hannák, and Aleksandra Urman. 2022. Comparing the Language of QAnon-Related Content on Parler, Gab, and Twitter. In *14th ACM Web Science Conference 2022 (WebSci '22), June 26–29, 2022, Barcelona, Spain.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3501247.3531550

WebSci '22, June 26–29, 2022, Barcelona, Spain

ACM ISBN 978-1-4503-9191-7/22/06…$15.00
https://doi.org/10.1145/3501247.3531550

1 INTRODUCTION

Months following the 2020 US presidential elections were characterised by the claims of the President of the United States Donald Trump and his supporters that the election was "stolen" from him. These claims were exacerbated by conspiratorial content that has found a home on many social media platforms. The months-long period of legal, political, and cultural turbulence that followed resulted in the January 6th insurrection on the US Capitol. It is often posited that many social media platforms are home to rampant hateful and damaging content, in particular in the political and ideological context.

The scrutiny is brought upon more and more social media companies, although the focus of scrutiny shifts based on public perception. Media, public and policymakers often look only at individual platforms and only when journalists, researchers or whistleblowers point to something that manages to garner enough public attention. Similarly, following the events of January 6th, various media outlets attempted to uncover details around its organisation. Various reputable media outlets identified Parler, a social networking site launched in 2018 as a "free speech alternative" to more established platforms such as Twitter, as the place where the organisation took place [11, 28]. In connection to these allegations, Parler was removed from the most popular app stores and many service providers pulled their support, ultimately resulting in Parler being taken offline on January 10, 2021. Parler returned online on February 15, 2021, having found alternative service providers, although with all the previous data removed.

Many studies have been published on political discourse, conspiracy theories, and hateful language, but similarly to public scrutiny, they have primarily focused on single platforms (e.g. [1, 7, 33]). To make a fair judgment about the difference in discourse, tone, and anti-social language (e.g. hateful, toxic, threatening language), the scientific community needs to conduct studies comparing multiple platforms, yet comparative studies are rare, in particular those encompassing both alternative and mainstream platforms [15]. While Parler was undoubtedly one of the homes for the type of content that has created a fertile ground for the events that have unfolded, there has been legal and anecdotal evidence of other platforms being used to spread election-related conspiratorial content, fake news, and to organise protests and riots following the election [5].

We aim to partially address the existing knowledge gap with the present study encompassing three platforms - Parler, Twitter, and Gab. We choose Twitter and Gab as comparison points to Parler for two reasons. Firstly, the three platforms are very similar in basic functionality, primarily micro-blogging platforms with the same mechanism of posting and engaging with the content. Secondly, as insufficient content moderation was frequently cited as a reason for removing Parler's access to services, a comparison with platforms that have opposing types of moderation policies is particularly relevant - Twitter has been engaging in moderation of types of content that are subject of our investigation [30], while Gab is marketing itself in similar ways to Parler ("championing free speech") and has been reported as one of the platforms that Parler users migrated to following the shutdown [32].

We study a sample of the content on the aforementioned three websites in the month-long period running up to the shutdown of Parler. This time frame encompasses multiple events of interest: disputes over election results, protests including the January 6 riot, Senate run-off election in Georgia, the confirmation of Biden as the winner, Twitter banning Trump, and the announcement of the shutdown of Parler. We choose to analyse posts from all three platforms that have been posted in this period with #QAnon, denoting a conspiracy theory suggesting that "Trump has been battling against a satan worshipping global child sex-trafficking ring and an anonymous source called 'Q' is providing secret information about the ring" [37]. Beyond just being one of the most popular hashtags on Parler prior to the shutdown [1], QAnon-related hashtags were amongst the top hashtags used in profile descriptions of users spreading misinformation related to the 2020 presidential election on Twitter [8], and among the top 10 hashtags in a sample of 600 million US election-related tweets [10]. The followers of QAnon have been spreading conspiracy theories related to the elections and participated in the storming of the Capitol [31].

1.1 Research Questions

Our aim is to compare Parler, Twitter, and Gab posts that have used the same hashtag in the same time period, to gain an understanding of differences and similarities between the three platforms in the context of the discourse surrounding #QAnon in the run-up to, and for a few days following, the Capitol riot. We analyse the volume of posting and the number of users who post to gauge the activity and interest levels in relation to QAnon across the three platforms. In addition, as lack of moderation of threatening and harmful content was a frequent criticism of Parler, we analyse multiple aspects of anti-social language (such as hateful, threatening, and toxic language). This allows us to draw conclusions on how Parler compares to functionally similar moderated and unmoderated platforms. To increase understanding of which figures, events and similar entities are mentioned in posts, we look at named entities used across the platforms. In particular, as the chosen period includes major political events in the US, we test if the language differs in posts mentioning groups of political figures. We look at the posts about politicians from opposing political parties (as Gab and Parler are considered right-wing), and posts mentioning politicians of different genders[1] (as previous research reported higher levels of incivility on social media towards women in politics [21]). Finally, we analyse the prevalent narratives on the three platforms to gain a deeper understanding of differences and similarities in the conversation. We formulate these aims as research questions:

- **RQ1**: How does the volume of #QAnon posts, and the number of users posting, vary across the three platforms?
- **RQ2**: How do the three platforms compare with respect to the prevalence of anti-social language?
- **RQ3**: Are there differences in the relative prevalence of themes and political figures mentioned between the three platforms, and in the use of anti-social language related to those mentions?
- **RQ4**: What are the prevalent narratives on the three platforms, and how do they compare?

[1]male and female only, due to a lack of posts mentioning genderqueer politicians in the sample

2 RELATED WORK

Twitter. Twitter is one of the most, if not the most, studied social network. Regarding QAnon-related content in particular, a study of banned Twitter users found that "QAnon is well-positioned at the centre of the political hashtag community" of banned users [7]. Numerous studies have been conducted on hate speech on Twitter. For the present research, of particular relevance are studies of hate speech related to violent or political events. A study of hate speech and white nationalist language on Twitter around and after the 2016 US presidential election reveals "no evidence of an increase in hate speech before or after the election", while noting that there are short time periods where the level of hate rises, finding "evidence of tens of thousands of tweets containing hate speech and white nationalist rhetoric on Twitter" [23].

Gab. To the best of our knowledge, no studies examining QAnon on Gab have been conducted, although a study of topic evolution on the platform found that by the start of 2018, the discourse on Gab has switched to alt-right political topics, and in particular posts related to QAnon [16]. Related to language on the platform, previous research on Gab found that 5.4% of Gab posts contain at least one hate word, and that the most prevalent points of discussion on Gab are news, events, and conspiracy theories [33]. 90% of posts on Gab were found to have toxicity scores less than 0.7 (on a scale from 0 to 1), although the authors note that there are users who "abuse the lack of moderation to spread hate" [14].

Parler. Most of the research on Parler is still undergoing peer review, and to the best of our knowledge, nobody examined QAnon or anti-social language on the platform. However, studies conducted include reports of the number of users on Parler having more than doubled within weeks around the 2020 US election, and that Parler "has emerged as a space in which accounts that have been suspended by Twitter Safety continue to communicate with their audiences" [27]. A study supplementary to a release of a large Parler dataset has found that Parler has experienced growth in user base in close proximity to "online censorship on mainstream platforms like Twitter, as well as events related to US politics" [1]. In addition, the study reports that QAnon is the 8th most popular hashtag on Parler.

Cross-platform studies. allow for a fair comparison and deepen our understanding of the roles different platforms play in the ecosystem of information. The latter is particularly important when it relates to phenomena damaging to society (e.g. conspiracy theories) or to individuals (e.g. anti-social language). However, few cross-platform studies have focused on conspiracy-related content, anti-social language or hate speech. Notable exceptions include an analysis and comparison of language related to QAnon, across sites collecting posts claimed to be written by "Q", Twitter, 4chan, 8chan, Reddit, and Voat [2]. The study suggested that the QAnon community can find a home for their content on mainstream platforms, and that bans on one platform "do little to slow growth on others". The same study reports that posts written by "Q" are less toxic than posts by QAnon communities on sites such as Voat and 4chan, although the study does not capture the same measure for Twitter.

Several comparative studies have been published on social media engagement with the Covid-19 pandemic, some focusing on conspiratorial content. A study of Covid-19-related conspiracy theories on 8kun and Gab found that 24% of the Covid-related posts on Gab contain conspiracy theories, and that 57% of randomly selected user profiles contain conspiracy theory content, such as QAnon, although that the prevalence of such content is higher on 8kun [36]. Another study has focused on the effect of moderation, contrasting Facebook, Twitter, Reddit, and 4chan [18], finding that "content moderation on Twitter was less effective than on the other platforms", potentially attributed to the fact that content on Twitter spreads within the first hours of being posted. Finally, a study on the emergence of sinophobic behaviour on web communities in the face of the pandemic on Twitter and 4chan has focused on content analysis, including racist slurs, as a type of hate speech [26].

Other existing scholarship involving multiple platforms and hate speech includes a comparative study of hate speech on Twitter and Reddit around attacks involving Arabs and Muslims as perpetrators or victims [17]. The authors observed that "extremist violence tends to lead to an increase in online hate speech", with the biggest increase seen in messages advocating for the violence. A comparison of Gab and 4chan's Politically Incorrect board (/pol/) in the context of antisemitism, found evidence of increasing antisemitism around political events such as the US presidential election of 2016 [35].

3 DATA AND PLATFORMS

Data Collection. We collect all posts with the #QAnon hashtag (not case sensitive) from all above-mentioned platforms in the period between December 7, 2020 and January 10, 2021. While #QAnon is not the only hashtag used by QAnon supporters, we believe that it is the most intuitive one to search for by someone not yet well-versed in this conspiracy theory. We obtained Parler data via Parlance API [6], which allowed collection of the same data that would have been accessible if we had used Parler's search feature to search for the same hashtag. We collected data for the same hashtag, for the same time period from Gab using garc library [25], and from Twitter using Twitter's Academic full archive search API, via the academictwitteR library [4, 29]. The data for all three platforms included the text of the post, as well as metadata (such as information on author, time, date).

For all three platforms, we do not analyse re-posts that are duplicate but only re-posts made with additional comments. Additionally, since Parler and Gab allow comments on posts, while Twitter does not, we exclude comments from the analysis to ensure comparability. We note that language in comments may differ from the language in posts, which warrants a separate study.

Platform	Nr. of posts	Nr. of users	Posts per user (μ)
Twitter	12 325	5861	2.18
Parler	78 892	4648	17.52
Gab	6 708	501	13.97

Table 1: Number of collected vs analysed posts, number of users, and average number of posts per user.

Hashtags have the same functionality on all three platforms, allowing users to add keyword-like metadata to their posts, signifying a topic. All three platforms enable their users to add any number of hashtags to their posts, but the hashtags count within the character limit. However, these character limits vary across platforms, the maximum length of posts on Twitter, Parler and Gab being 280, 1000, and 3000 characters, respectively. As a consequence, hashtagging behaviour is different across the platforms which is important to keep in mind throughout our analysis and interpretation of the results (for example, as we will show later, the ratio of hashtags to text across the three data sets varies substantially).

Data-related limitations. We highlight three limitations in particular. Firstly, we limit our analysis to posts written in English to ensure the comparability across the three platforms. In our data, 96.9% of Parler, and 95.9% of Gab are posts in English, and while Twitter contains a larger proportion of non-English posts (19%). Secondly, as Twitter has been moderating QAnon-related content since at least late July 2020 [30], it is likely that there is a body of Tweets that have been removed and are not available for analysis. It is reasonable to assume that this moderated content would contain a higher amount of anti-social language, so we consider the results for Twitter a lower bound for the platform. Given the largely unmoderated nature of Parler and Gab, we assume that only a small number of posts would have been removed, if any, and that our Parler and Gab data truthfully represent the conversation on those platforms, as it happened. Finally, while the one month period of analysis helps us achieve higher internal validity and covers events that could increase interest in QAnon, the limited volume of data could effect analysis of language undertaken in RQ3 and RQ4.

Ethical considerations. As neither Parler nor Gab have documented APIs, we have used third-party APIs for both. This data was, at the time of collection, publicly available to anyone with an account on the platform (Parler), or anyone on the internet (Gab). Our data collection was non-intrusive as it did not affect users or platforms. The collection of Twitter data is in line with the platform's Terms of use. While the original data for all platforms included usernames, bios, and similar information which could lead to the identification of individuals, we have excluded these fields from the analysis, hence anonymising the data. The posts were only collected for accounts that were not private.

	Characters			Words		Hashtags	
	max	μ	σ	μ	σ	μ	σ
Twitter	280	154.8	82.3	23.3	13.3	4.6	4.7
Parler	1000	561.6	310.9	56.9	31.7	39.3	25.3
Gab	3000	512.6	456.3	61	59.2	26.3	31.7

Table 2: Means (μ) and standard deviations (σ) for number of characters, words and hashtags in posts

4 RQ1: ACTIVITY, USERS, AND POSTS

To understand how the volume of posts, number of users participating in the discussion, and posting styles vary across the three platforms, we utilise metadata extracted from posts, which allow us to summarise temporal posting activity, number of users posting, and similar descriptive statistics. In addition to analysing the metadata, we extracted various features from the text, such as hashtags, and text length descriptive statistics, to capture differences in posting styles. We highlight that our study aims to establish how platforms compare to each other, rather than how QAnon-related posts compare to users' overall activity on each platform. We do not draw conclusions about how #QAnon compares to non-#QAnon content, or other hashtags.

Table 1 shows the volume of posts with #QAnon for the observation period on each platform. On Parler, the number of posts was much higher than on Gab and Twitter. Though at first glance that indicates a much higher prevalence of #QAnon-related discussions on Parler as compared to Gab and Twitter, the observation might be partially explained by the differences in the way hashtags are used across platforms. The discrepancies in the number of users who posted about #QAnon between platforms were smaller than the differences in the sheer volume of posts with this hashtag. Though Parler contained substantially more posts with #QAnon than Twitter, the number of unique users posting with the hashtag was higher on the latter (Table 1). This difference could imply that Parler and Gab users become more invested in QAnon-related discussion once they become involved, posting more frequently about the topic. An alternative explanation is that Parler and Gab users include this hashtag in posts that are not explicitly about QAnon, using the hashtag to boost the visibility of their content, and upon qualitative observation we suspect this to be the case.

While Twitter has been a well-known platform with a large and established user base for many years, Parler and Gab have reported a substantial growth of their user base in the run up and the aftermath of the 2021 election. For this reason, we look at the ages of the accounts that have posted with #QAnon in our sample, finding a large difference. The average account age at the time of posting on Twitter was 2781 days, on Parler 131 days, and on Gab 675 days. To examine whether this discrepancy is attributed to the difference in the age of the platform alone (Twitter was founded in 2006, Gab in 2016, and Parler in 2018) or to external events, we further scrutinised the dates when users joined each platform. We found that 19.8% of #QAnon posts on Parler and 18.8% on Gab were created by users who have been on the platform for less than 1 month. On Twitter, the share of such users is only 1.7%. This implies

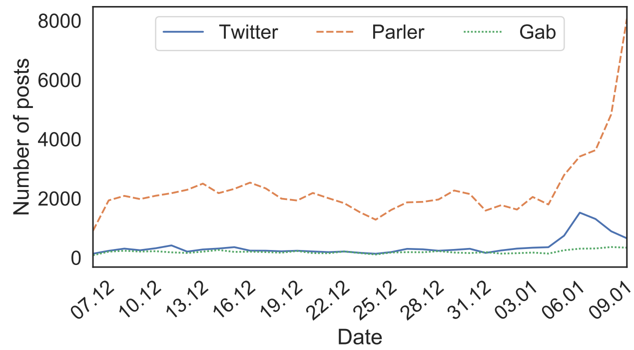

Figure 1: Daily frequency of #QAnon posts

that the differences in the average account age is attributed not only to the age of the platform but also external events such as the election and events that followed: of the users posting with the #QAnon, 1.9% have joined Twitter after the Election (November 3, 2020), compared to 56.1% on Parler, and 40.1% on Gab.

However, the apparent growth in the number of users who joined Gab and Parler during the one-month observation period was not accompanied by the bursts of #QAnon posting activity. The number of posts with the hashtag was relatively stable across the three platforms up until early January, as depicted in Figure 1. Then both Twitter and Parler but not Gab saw a rise in popularity of the hashtag around January 6, with the conversation on Parler further intensifying in volume in the run-up to the platform going offline.

In Table 2 we report the means and standard deviations for character counts, word counts, and the number of hashtags. Unsurprisingly, posts on Gab and Parler are longer on average as the two platforms have higher character limits than Twitter. The higher character limit leaves the users of the two "alt-tech" platforms more space for hashtags, with the average number of hashtags being much higher on Gab and Parler than on Twitter, and 75.5% of Parler posts having more hashtags than other words, compared to 51.7% for Gab, and 17.5% for Twitter. The extra character limit is employed for particularly active hashtagging by Parler users. The mean words to hashtag ratio on Parler is 1.44, compared to 6.94 on Gab and 10.15 on Twitter. This points to a very different way in which hashtags are used across the platforms. Even though hashtags have the same intended use across all platforms, the differences suggest that each platform's user base developed their own hashtagging culture. In fact, upon a qualitative inspection of a sample from all three, we noticed that many Parler posts used "hashtag walls" - blocks of many continuous hashtags, not necessarily related to the post itself. The observed propensity of Parler users to include a high number of hashtags, not always related to the immediate content of the post, might partially explain why the volume of posts on Parler with #QAnon is much higher than on Twitter and Gab.

5 RQ2: ANTI-SOCIAL LANGUAGE

To compare the three platforms with respect to the prevalence of anti-social language, we capture measures of anti-social language and hate speech, using similar methodology as many other studies

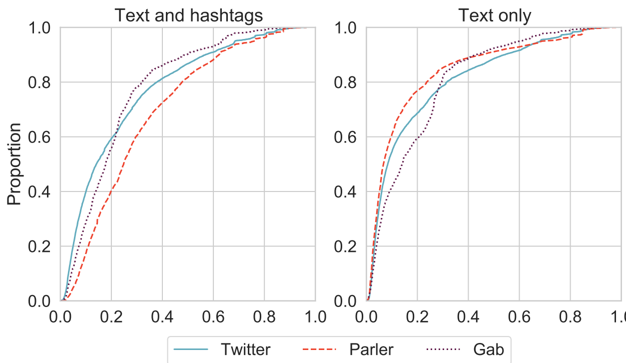

Figure 2: CDFs of severe toxicity for text only, and for text and hashtags combined.

- Perspective API [2, 9, 19, 34] and Hatebase lexicon of hate words [16, 24, 33].

We used Perspective API to extract features that act as proxies of anti-social language [13]. While Perspective exhibits bias in certain contexts (e.g. classifying language predominantly used by African Americans as more toxic than the language used by white people [22]), studies found that it outperforms other available tools on similar texts [34]. For each feature of interest, when queried against text, Perspective provides a score between 0 and 1, denoting the probability that the post is as the feature describes. In our comparison we look at the distribution of scores to capture the differences. We perform statistical testing using a two-sample Kolmogorov-Smirnov test to examine the difference in score distributions. All results we report below are statistically significant at p<0.001. We select only features from Perspective which are tested across multiple domains and trained on significant amounts of human-annotated comments: Severe Toxicity (A very hateful, aggressive, disrespectful comment or otherwise very likely to make a user leave a discussion), Threat (intention to inflict pain, injury, or violence against an individual or group), Identity attack (negative or hateful comments targeting someone because of their identity), and Insult (insulting, inflammatory, or negative comment towards a person or a group of people).

To examine the effect of hashtags on anti-social language measures, due to the difference in hashtagging behaviour across the platforms discussed in the previous section, we query Perspective three times for each post. Firstly, we query it on the post cleaned of URLs and mentions, secondly having additionally removed all hashtags, and finally on hashtags alone. The cross-platform differences in hashtagging have important implications for the results in relation to the prevalence of anti-social language. In Figure 2, we show cumulative distributions of severe toxicity scores across the three platforms for full posts and for posts with hashtags removed. The figure shows that without the hashtags, Parler shows lower or similar toxicity to Twitter and Gab. However, once hashtags are added, Parler consistently scores higher on severe toxicity.

We further explored the relationship between the presence of hashtags and the posts' classifications. In Figure 3, we show the distribution of differences in Severe Toxicity scores between full

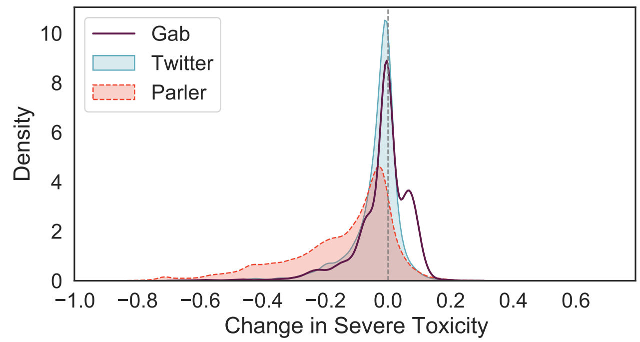

Figure 3: Impact of adding hashtags on severe toxicity. Negative values indicate that a post is more toxic with hashtags.

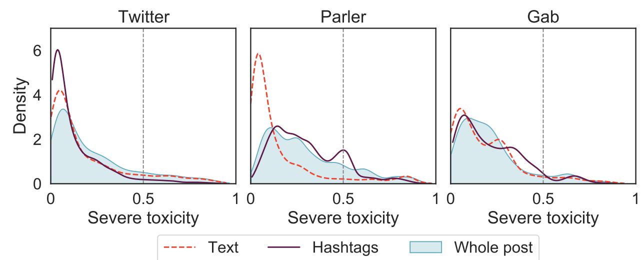

Figure 4: A kernel density estimate of severe toxicity of different post components - hashtags, text, and hashtags and text combined.

posts and posts with hashtags removed, by platform. Values smaller than 0 indicate that a post is assigned a higher toxicity score once hashtags are added to the analysis. As the longer left tail for Parler indicates, Parler posts with hashtags included are assigned higher severe toxicity scores compared to the same posts without hashtags. Severe toxicity scores for different post components across all three platforms are in Figure 4. Parler's hashtagless texts, compared with hashtags only or combined (texts and hashtags), are overwhelmingly not severely toxic, and compare to scores of tweets. This means that posts themselves are not more toxic on Parler, yet there seems to be a custom on the platform to add hashtag blocks which are more severely toxic. We don't know whether and how people are affected by these hashtag blocks, but we find it noteworthy that a culture on this platform has evolved in which users add considerably more negative labels to their posts.

Mean scores for severe toxicity across time on all three platforms are presented in Figure 5. The addition of hashtags changes the overall picture in this case as well. For full posts (with hashtags), Parler consistently has a higher mean severe toxicity score assigned than Twitter and Gab, which have similar scores to each other. On the other hand, with the hashtags removed, Parler consistently has lower mean severe toxicity scores than both Twitter and Gab, except for a large uptick in toxicity seen after January 6.

Cumulative distribution functions (CDFs) for Insult, Identity attack, and Threat features are presented in Figure 6. The model identifying threatening language shows that posts on all three

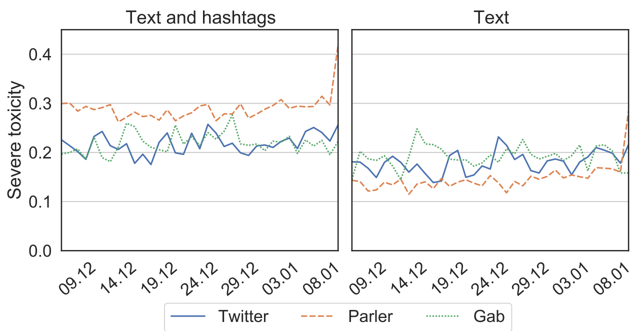

Figure 5: Daily mean severe toxicity of posts, with hashtags, and without

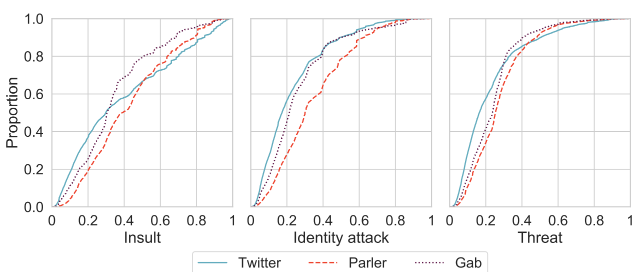

Figure 6: CDFs for Insult, Identity attack, and Threat models

platforms do not have a high probability of being threatening, with platforms having a similar proportion of posts more likely than not to be threatening (with a probability over 0.5). The identity attack model shows Parler posts scoring the highest, with Twitter and Gab being very similar. Finally, the platforms only show minor differences when it comes to insults, although a lower proportion of Gab posts are likely to be insulting than posts on Twitter and Parler.

While Perspective allows us to capture measures such as toxicity, and it takes into account words that are considered hate speech, we separately measure hate speech specifically, as the most extreme type of anti-social language. We perform keyword-based classification by using Hatebase lexicon of hateful terms [12]. Hatebase collects words and phrases considered hateful across multiple categories, such as ethnicity, nationality, religion, gender, class, etc. We matched cleaned text (without links, mentions, or hashtags) from all three platforms with words and phrases in the lexicon.

Upon inspection, a few ambiguous words have resulted in a high number of false positives (e.g. "Apple" is in Hatebase lexicon as it can be used to signify "An American Indian who is 'red on the outside, white on the inside."', but all occurrences of the word in our corpus referred to the company Apple Inc., or the fruit apple). For this reason, we have omitted a number of words[2] from the analysis. We chose words to remove by manually going through each occurrence of a hate word, and removing it from a dictionary if all occurrences of it were used in a non-hateful context. Gab takes the leading position in terms of the share of posts with hate words despite scoring similarly, or even lower, than Twitter and Parler on previously discussed anti-social language measures. Percentage of posts that contain at least one hate word for Parler, Twitter and Gab are 2.66%, 2.8% and 4.77% respectively.

To conclude, most indicators show that the language of most posts across all three platforms should not be considered anti-social. An exception is a higher probability of posts being insulting on Twitter and Parler. Language on Parler appears marginally worse in terms of severe toxicity and identity attacks than the other two platforms, although we have demonstrated that this result is greatly affected by the prolific hashtagging of Parler users. While most Gab posts have a lower probability of being anti-social than Parler, and even Twitter for some indicators, they also have the highest percentage of hate speech occurrence.

6 RQ3: THEMES AND POLITICAL FIGURES

To compare similarities and differences in mentioned themes (e.g. individuals, locations, phrases), we use named entity recognition [20] on clean text (after removing links, mentions, and hashtags). We discarded some categories of entities, such as numbers and percentages, as they were of no interest for the analysis. The resulting entity list required cleaning, as terms signifying the same entity are not always automatically recognised as the same (e.g., "January 6", "Jan 6", "6. January").

Following the removal of irrelevant entities, we have obtained 12759 named entities from Twitter, 90769 from Parler, and 15937 from Gab. 50.8% of Twitter posts, 42.7% of Parler posts, and 54.7% of Gab posts contain at least one entity. The number of unique entities extracted from Twitter is 3937, from Parler is 17439, and from Gab is 5066, suggesting that the conversation on Twitter is more focused around a few subjects, while on Parler and Gab there is a higher diversity of entities discussed in connection to #QAnon. Many entities are popular across all platforms (Donald Trump, United States, Georgia, Twitter, Joe Biden, American(s), Democrat). Even the terms which appear in the top 20 on only two of the three platforms (e.g. "Republican" and "GOP" on Twitter and Parler) are still just outside of the top 20 on the third (25th, and 33rd most popular on Gab, respectively). This suggests that the discourse in relation to #QAnon on all three platforms was centred largely around similar subjects.

To infer the differences between popular entities, we first removed entities used by less than 10 users (on either platform). This is to filter out unusual and erroneous entities which are used by one user repeatedly (e.g. "q and the plan to save the world"). We present the top 30 entities that are relatively more popular on one platform (in the top 50 of the most popular entities), as compared to the other two, in Table 3. If $r(x)_t, r(x)_p, r(x)_g$ represent the ranking of popularity of term (x) on Twitter, Parler and Gab, the maximum difference is calculated as

$$\max(|r(x)_t - r(x)_p|, |r(x)_t - r(x)_g|, |r(x)_p - r(x)_g|)$$

The results are indicative of potential differences in the focus of the discussions on the three platforms. For instance, mentions of the Capitol, Jake Angeli ("QAnon Shaman"), Antifa, Nazis as well as Kelly Loeffler and Marjorie Taylor Greene were more prevalent on

[2] ABC, Afro-Saxon, Anglo, Ann, apple, banana, Becky, bird, boos, bubble, bucks, Charlie, chief, coconuts, egg, egg-plant, frog, girl, guinea, lefties, mock, pancakes, pepper, Pepsi, property, queen, skinny, snowflake, sole, spikes, Tommy, Yankee, yellow

Entity	Rank			Entity	Rank		
	Twi	Par	Gab		Twi	Par	Gab
K. Loeffler	**49**	84	634	Texas	134	**23**	**21**
M.T. Greene	**44**	563	590	Michigan	172	**36**	59
J. Angeli	**17**	412	362	Pennsylvania	89	**31**	**38**
A. Babbitt	**35**	174	109	California	101	**44**	79
A. Jones	**37**	208	132	Russian	**31**	81	88
J. Epstein	**32**	74	140	CCP	152	**28**	**18**
JFK	**46**	150	155	Nazi	**41**	130	118
Dave	271	506	**45**	2. Amendment	-	**46**	349
Brian	498	376	**47**	Elec. college	**43**	129	65
Gab	551	54	**12**	DOJ	164	90	**40**
Amazon	146	79	**50**	Constitution	80	**45**	**25**
The republic	404	**32**	240	defense	-	159	**30**
Deep State	170	195	**23**	Antifa	**15**	**30**	72
Dominion	129	**34**	**32**	Capitol	**12**	66	60
MSM	82	58	**29**	Yesterday	**19**	**47**	76

Table 3: 30 terms with the biggest difference in popularity

Twitter than on Parler or Gab, while the latter two platforms saw higher popularity of Pennsylvania (probably connected to the vote count in the state) and Dominion voting machines (according to conspiracy theories, a company which aided the "stealing" of the election), as well as Deep state (Gab), MSM ("mainstream media", Gab), 2nd Amendment (Parler) and defense (Gab).

To analyse if there is a difference in the use of anti-social language when discussing selected groups of political figures, we examined entities occurring at least 20 times in the whole corpus manually, and classified them into groups of interest. We form four groups, with political figures divided by gender and party[3]. We consider either the party membership, or service in an administration, when dividing by party lines. As in RQ2, we ensure that we only report results stemming from different distributions (tested using a two-sample K-S test, at p<0.001, excluding posts that include both groups being compared to ensure the independence of samples).

Table 4 shows the differences with regard to the posts mentioning political figures (female, male, Republican and Democrat) and two presidential candidates. While Donald Trump and male Republican politicians were mentioned consistently more (in terms of the share of posts mentioning them) than their Democratic counterparts across all platforms, there was a divergence in posts about female politicians. There was a higher share of posts about female Republicans than Democrats on Twitter, while on Gab and Parler, the situation was reversed. We also observe significant differences in the mean Perspective scores of posts mentioning different groups of politicians (female vs male; Republican vs Democrat). On all three platforms, posts mentioning female (vs male) politicians, Democrats (vs Republicans) and Trump (vs Biden) scored higher on average for anti-social language features that exhibited significant cross-group differences (Table 5).

In conclusion, the entity-based analysis shows that while most prevalent named entities across the platforms are similar, there are differences between the mainstream platform and the alt-tech platforms. Twitter saw more mentions of high profile individuals considered QAnon supporters (Jake Angeli, Marjorie Taylor Greene), while more Parler and Gab posts used conspiratorial terms (Dominion, Deep state, Mainstream media). On all three platforms, posts mentioning female (compared to male) politicians, Democrats (compared to Republicans) and Donald Trump (compared to Joe Biden) score higher on all anti-social language features.

[3]excluding Trump and Biden, who are analysed separately

Group	Percent of posts			Number of posts		
	Twi	Par	Gab	Twi	Par	Gab
Female Republicans	1.18	0.62	0.66	145	487	44
Female Democrats	0.5	1.21	1.95	62	956	131
Male Republicans	2.08	2.69	3.91	256	2126	262
Male Democrats	0.54	1.3	1.43	66	1022	96
Donald Trump	9.71	9.6	8.59	1197	7572	576
Joe Biden	0.95	2.54	2.76	117	2000	185

Table 4: Percentage, and the number, of posts containing mentions of politician groups (excluding Trump and Biden)

	Feature	Party		Gender		Candidate	
		μ(D)	μ(R)	μ(F)	μ(M)	μ(B)	μ(T)
Twi	Ide.Att.	0.23	0.21	-	-	-	-
	Insult	-	-	-	-	0.44	0.50
Par	Ide.Att.	0.35	0.29	0.35	0.30	0.33	0.34
	Insult	0.47	0.41	0.48	0.41	0.42	0.51
	S. Toxic	0.30	0.25	0.31	0.26	0.26	0.32
	Threat	0.26	0.24	-	-	0.24	0.26
Gab	Ide.Att.	0.29	0.20	0.26	0.23	-	-
	Insult	0.43	0.33	0.45	0.34	-	-
	S. Toxic	0.21	0.16	0.20	0.17	-	-
	Threat	-	-	-	-	0.22	0.26

Table 5: Mean Perspective scores mentioning Democrats (D) or Republicans (R), female (F) or male (M) politicians, and Biden (B) or Trump (T).

7 RQ4: NARRATIVE ANALYSIS

While NER-based analysis allows us to measure the prevalence of terms on each platform, as well as undertake anti-social language analysis related to groups of interest, it does not offer insight into the *context* in which those terms are mentioned. Simply knowing that something, for example QAnon, is mentioned a lot does not allow us to capture if one platform is overwhelmingly critical of QAnon, while the other is supportive. To deepen our understanding of content across the three platforms, we analyse narratives and inter-connected terms, using an adjusted Relatio library, and the method presented alongside it [3]. We use Ash et al. operationalisation of narratives as triples of words or phrases that take the form *Agent - Verb - Patient* (Figure 7).

The narrative analysis pipeline is presented in Figure 8. In preprocessing, we split posts into sentences, as the method can not reliably connect narratives spanning multiple sentences. Sentences

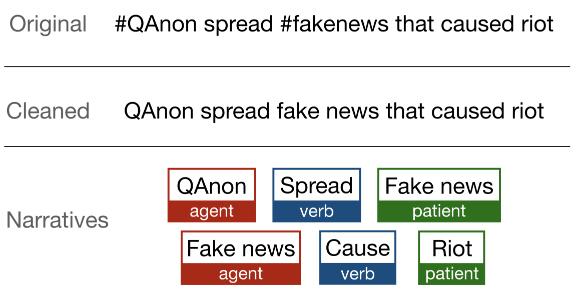

Figure 7: An example sentence being split into narratives

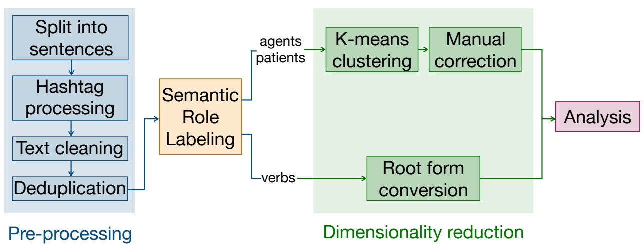

Figure 8: Narrative analysis pipeline

can contain incomplete narratives (which we omit from analysis), or one or more complete narratives, such as the example we give in Figure 7. Special features of social media posts, such as hashtags, required adjustment of Ash et al. pre-processing. While some users use hashtags out of context, to boost visibility of their posts (in our case, especially on Parler as discussed previously), hashtags can be used in context, and simple removal would result in loss of information. To overcome this, we make best effort to remove hashtag blocks while preserving hashtags that are likely used in context. As some hashtags contain multiple words (e.g. #fakenews in Figure 7), we manually establish how to split hashtags that occur more than five times on either platform. This maximises the information extracted from data, as we do not simply discard posts such as the one given in Figure 7. Finally, we clean text by removing mentions, links, emoji, and stopwords from posts, and remove duplicate posts made by the same person to ensure that spamming does not affect our analysis.

After pre-processing, we use semantic role labelling to identify roles in a sentence, extracting building blocks for narratives: verbs, agents and patients. This results in a vast vocabulary consisting of 100103 unique agents and patients, and 8213 unique verbs. We reduce dimensionality by grouping agents and patients into similar themes with k-means clustering on GloVe embeddings. As our sample was not large enough to obtain a small number of high-quality clusters, we have manually corrected some of the categories (e.g. to ensure that mentions of Trump, Biden, and President are in separate clusters), resulting in 440 clusters representing agents and patients. We convert verbs to their root form, reducing the dimensionality of verbs to 3510. While dimensionality reduction results in some loss of information, it is necessary due to language diversity. A summary of the number of sentences and narratives is available in Table 6.

	Twitter	Parler	Gab
Analysed posts	12759	81456	6997
Deduplicated posts	11412	55532	5076
Sentences	21306	111457	16122
Complete narratives	6629	32220	5546
Posts with complete narratives	4631	19656	2898

Table 6: Summary on the number of narratives extracted

Of the 44395 complete narratives, 31180 are unique. The reason for this is a high number of unique verbs that connect Agents and Patients. While we reduced the dimensionality of verbs from 8213 to 3510, we could not reliably group them according to the similarity of action they represent. The overlap between unique narratives between the platforms is presented in Figure 10. 96% of narratives appear on only one of the three platforms, with only 56 narratives appearing on all three. Parler shares more narratives with Gab than with Twitter, despite Gab having fewer unique narratives than Twitter. This indicates a higher similarity of the conversations happening on the Alt-Tech platforms compared to the mainstream platform.

The low overlap in narratives suggests that themes frequently appearing together differ substantially on the three platforms. To understand what are frequently co-occurring agents and patients, and how platforms differ in this regard, we model Agents and Patients as nodes in an undirected network. This network is agnostic to what verb connects the two but allows us to look at the most popular Agents and Patients on platforms, and which terms tend to appear together in a narrative. We present giant connected components for the 30 most frequent Agent-Patient pairs in Figure 9. The three platforms differ substantially when it comes to what is central to the conversation - conversations on Twitter (Fig. 9a) are highly centred on QAnon, which is expected given how we collected the data. Terms connected to QAnon represent political entities - such as media, government, Donald Trump, and public personalities (mainly overlapping with those observed in RQ3). In addition, we note the mention of Capitol, and that Twitter users linked it to QAnon. In stark contrast, QAnon does not even appear in the top 30 connections on Parler (Fig. 9b). This is yet another indication of different use of hashtags on Parler, where #QAnon was frequently used as part of hashtag blocks, rather than in the conversation context. Donald Trump and other public personalities take central roles of conversation on Parler, and the core is well-connected, except for "Location", "Book", and "One" (representing a person) making an appearance due to the high popularity of quote "the man who reads nothing at all is better educated than the man who reads nothing but newspapers" amongst Parler users. Gab network (Fig. 9c) is more similar to Parler's, although public personalities are more connected than Trump himself, with QAnon making an appearance in relation to lockdowns, public personalities, and popular phrase used by Trump supporters "Drain the swamp".

Finally, we present the ten most prevalent narratives on all three platforms in Table 7. These results provide some context in which Agent-Patient pairs presented in Figure 9 appear together, although we note that even the most popular narratives are not used many

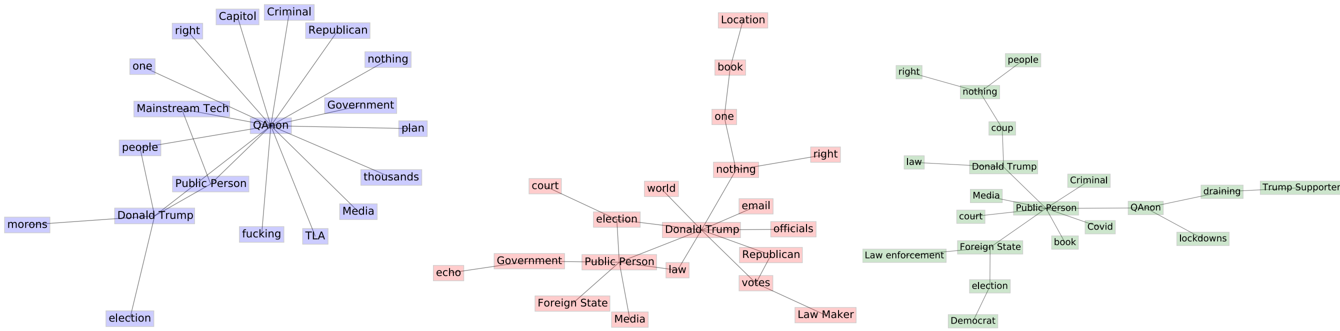

(a) Twitter (b) Parler (c) Gab

Figure 9: Giant connected components of the 30 most frequently co-occurring Agents and Patients

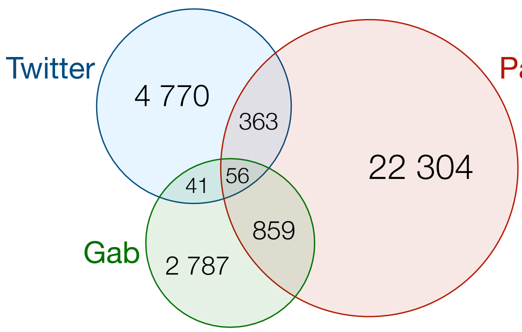

Figure 10: Overlap of unique narrative between platforms

Twitter	Parler	Gab
supporter show balls	email send trump	trump supporter keep draining
republican push qanon	person read book	qanon keep draining
qanon storm capitol	person read nothing	nothing wake people
platform ban public person	evil destroy world	science use technology
trump delude morons	art support bot	trump invoke insurrection
qanon include legit	book honor location	nothing stop coming
trump meet qanon	government follow echo	activities affect military
thousands believe qanon	nothing stop right	cia infiltrate us state
qanon arrest crime	deep state engineer covid	morons prove qanon
trump supporter enter gov.	evil destroy world	lockdowns prove qanon

Table 7: 10 most prevalent narratives

times (some of the narratives in the top 10 appear as few as 5 times), a result of the high diversity of verbs used. QAnon is central to the discussion on Twitter, appearing in 6 top narratives, which are indicative of republicans' and Trump's relationship with QAnon, and the role that Twitter users believe QAnon had in the storming of the Capitol. On the other hand, Parler mostly sees the popularity of quotes (such as the example above) and calls to action by other Trump supporters - asking them to email him or follow and re-post. Gab has the highest level of conspiratorial ideas breaking into top 10 narratives - CIA infiltrating local state governments, suggesting that nobody can prove QAnon is fake, and calling on Trump to invoke the Insurrection Act in response to the "election being stolen".

While the three platforms proved to be vastly similar when tested for anti-social language, prevalent topics of conversation, and which groups of political figures attract more anti-social language, narrative analysis reveals considerable differences in the context of mentions of these figures and how central QAnon itself is to the discussion. Twitter centres the conversation on QAnon, often mentioning it in a negative context; Parler focuses on expressing support for Trump and making calls to action towards like-minded users, while Gab narratives hint at a higher prominence and more central role of conspiratorial ideas.

8 CONCLUSIONS

We have compared discourse around #QAnon on Twitter, Gab, and Parler in a month preceding the storming of the US Capitol on January 6, 2021., using a variety of computational methods to capture different aspects of posts - from measures of anti-social language, to themes at the core of the discourse. Our findings show that the volumes of posting with this hashtag differ drastically across platforms, with Parler having the highest volume of posts. While more unique users have posted using #QAnon on Twitter, users on Parler and Gab made considerably more posts per user on average. We found that the prevalence of anti-social language on the three platforms varies depending on the measure. While Twitter and Parler emerge as leaders in terms of the distribution of posts with anti-social language based on the analysed Perspective API features, Gab has the highest proportion of posts with hate words. While our exploratory analysis suggests that including re-posts does not alter the conclusions, the data available for re-posts is not robust enough for a definitive conclusion.

The analysis of the most frequent named entities across the three platforms revealed important similarities and differences between them. There are overlaps in the most popular named entities across the platforms, suggesting that #QAnon-related discourse mentioned largely similar themes during the observation period (e.g. Trump, the US, Washington). Entities enjoying higher popularity on Twitter compared to the other two platforms include individuals salient in the US political context in connection to QAnon (e.g., Marjorie

Taylor Greene or Jake Angeli) as well as "Antifa" and "Nazi", suggesting divisive stance in relation to #QAnon. On Gab and Parler, popular entities suggest a focus on discussions around the election results and related conspiracies (e.g., Dominion voting machines or Deep State) and the right to bear arms.

We observe differences in anti-social language measures in posts mentioning different political groups or individuals. Though the prevalence of such posts differs across platforms, as well as the prevalence of anti-social language in posts about different groups of politicians, we find that on all platforms posts mentioning female politicians, Democrats, and Donald Trump score higher on anti-social features than those with mentions of their male or Republican counterparts, or Joe Biden. Since our analysis focused on a very specific period and topic, it is unclear whether this observation can be generalised. We suggest that it is worthwhile to examine the prevalence of anti-social language in posts about different political groups in further cross-platform research.

Finally, our analysis of terms that appear together in posts, and the narratives they appear as part of, indicates that the core of discussions related to #QAnon differs substantially across the three platforms. While Twitter focuses on QAnon and the most prevalent narratives related to it are critical of it, Parler focuses on supporting Donald Trump, while Gab sees a bigger focus on other political figures, as well as Trump, and has a higher prevalence of conspiratorial content amongst its most popular narratives.

Limitations. Our study has limitations in addition to the ones listed in the Data section. Firstly, we focused on a single hashtag that, despite being one of the most popular in relation to QAnon conspiracy theory, still reflects overall discourse surrounding QAnon only partially. Nonetheless, due to the prominence of this hashtag, we argue that the collected data reflects the most dominant aspects of QAnon-related discourse. Secondly, while the month-long time frame our study encompasses is appropriate for analysis of events that see a heightened volume of activity, a further study examining the development of QAnon-related discourse over time could be relevant, comparing the observations during the periods of high activity, as well as in more routine periods.

Implications. The implications of our findings are three-fold. First, while the differences between the three platforms exist in our sample, they do not exactly align with widespread assumptions. Given the press Parler has received, and the consequences it has suffered in part related to the anti-social language, one might expect that Parler would exhibit a much higher prevalence of threatening, or toxic language than, for example, Twitter. Yet, our results show that platforms are largely similar in this regard, and that platforms can go further in limiting particularly harmful aspects of anti-social language. This also highlights the need for cross-platform comparative studies. The research community is uniquely equipped with methods, independence, and thoroughness required to make fair judgements to inform public perceptions, which might be currently largely based on preliminary analyses. Second, we show that more nuanced methods are necessary for comparative cross-platform studies, as single measures or simple metrics often do not reveal the differences, especially when observing multifaceted phenomena such as language and discourse. Finally, as we have seen with hashtags, seemingly identical functionalities can be used differently across different platforms, and pre-processing choices we make can substantially affect the conclusions. This highlights the importance of the decisions we make at each step of our analyses, including pre-processing, as well as how transparently we communicate them in our studies, both in cross-platform research, and beyond.

REFERENCES

[1] Max Aliapoulios, Emmi Bevensee, Jeremy Blackburn, Barry Bradlyn, Emiliano De Cristofaro, Gianluca Stringhini, and Savvas Zannettou. 2021. An Early Look at the Parler Online Social Network. *arXiv preprint arXiv:2101.03820* (2021).

[2] Max Aliapoulios, Antonis Papasavva, Cameron Ballard, Emiliano De Cristofaro, Gianluca Stringhini, Savvas Zannettou, and Jeremy Blackburn. 2021. The gospel according to Q: Understanding the QAnon conspiracy from the perspective of canonical information. *arXiv preprint arXiv:2101.08750* (2021).

[3] Elliott Ash, Germain Gauthier, and Philine Widmer. 2021. Text Semantics Capture Political and Economic Narratives. *arXiv preprint arXiv:2108.01720* (2021).

[4] Christopher Barrie and Justin Chun-ting Ho. 2021. academictwitteR. https://github.com/cjbarrie/academictwitteR

[5] Thomas Brewster. 2021. Sheryl Sandberg Downplayed Facebook's Role In The Capitol Hill Siege—Justice Department Files Tell A Very Different Story. *Forbes* (2021).

[6] Castle Lemongrab. 2020. Parlance API. https://github.com/castlelemongrab/parlance.

[7] Farhan Asif Chowdhury, Lawrence Allen, Mohammad Yousuf, and Abdullah Mueen. 2020. On Twitter Purge: A Retrospective Analysis of Suspended Users. In *Companion Proceedings of the Web Conference 2020*.

[8] EIP. 2020. *Repeat Offenders: Voting Misinformation on Twitter in the 2020 United States Election.* Technical Report. Election Integrity Partnership.

[9] Mai ElSherief, Vivek Kulkarni, Dana Nguyen, William Yang Wang, and Elizabeth Belding. 2018. Hate lingo: A target-based linguistic analysis of hate speech in social media. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 12.

[10] Emilio Ferrara, Herbert Chang, Emily Chen, Goran Muric, and Jaimin Patel. 2020. Characterizing social media manipulation in the 2020 U.S. presidential election. *First Monday* 25 (2020).

[11] Sheera Frenkel. 2021. The storming of Capitol Hill was organized on social media. *The New York Times* (2021).

[12] Hatebase Inc. 2021. Hatebase API. https://hatebase.org.

[13] Jigsaw. 2021. Perspective API Attributes and languages. https://developers.perspectiveapi.com/s/about-the-api-attributes-and-languages.

[14] Lucas Lima, Julio CS Reis, Philipe Melo, Fabricio Murai, Leandro Araujo, Pantelis Vikatos, and Fabricio Benevenuto. 2018. Inside the right-leaning echo chambers: Characterizing gab, an unmoderated social system. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*.

[15] Ariadna Matamoros-Fernández and Johan Farkas. 2021. Racism, hate speech, and social media: A systematic review and critique. *Television & New Media* 22, 2 (2021), 205–224.

[16] Reid McIlroy-Young and Ashton Anderson. 2019. From "welcome new gabbers" to the pittsburgh synagogue shooting: The evolution of gab. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 13.

[17] Alexandra Olteanu, Carlos Castillo, Jeremy Boy, and Kush Varshney. 2018. The effect of extremist violence on hateful speech online. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 12.

[18] Orestis Papakyriakopoulos, Juan Carlos Medina Serrano, and Simon Hegelich. 2020. The spread of COVID-19 conspiracy theories on social media and the effect of content moderation. *Harvard Kennedy School Misinformation Review* 10 (2020).

[19] John Pavlopoulos, Jeffrey Sorensen, Lucas Dixon, Nithum Thain, and Ion Androutsopoulos. 2020. Toxicity Detection: Does Context Really Matter?. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.

[20] Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. Stanza: A Python Natural Language Processing Toolkit for Many Human Languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.

[21] Ludovic Rheault, Erica Rayment, and Andreea Musulan. 2019. Politicians in the line of fire: Incivility and the treatment of women on social media. *Research & Politics* 6, 1 (2019).

[22] Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A Smith. 2019. The risk of racial bias in hate speech detection. In *Proceedings of the 57th annual meeting of the association for computational linguistics*.

[23] Alexandra A Siegel, Evgenii Nikitin, Joanna Sterling, Bethany Pullen, Richard Bonneau, Jonathan Nagler, Joshua A Tucker, et al. 2019. Trumping Hate on Twitter? Online Hate Speech in the 2016 US Election Campaign and its Aftermath. *Quarterly Journal of Political Science* 16, 1 (2019).

[24] Leandro Silva, Mainack Mondal, Denzil Correa, Fabrício Benevenuto, and Ingmar Weber. 2016. Analyzing the targets of hate in online social media. In *Proceedings*

of the International AAAI Conference on Web and Social Media, Vol. 10.

[25] Chris Stevens. 2020. garc: Python and Command-Line Interface for Gab.com API. https://github.com/ChrisStevens/garc.

[26] Fatemeh Tahmasbi, Leonard Schild, Chen Ling, Jeremy Blackburn, Gianluca Stringhini, Yang Zhang, and Savvas Zannettou. 2021. "Go eat a bat, Chang!": On the Emergence of Sinophobic Behavior on Web Communities in the Face of COVID-19. In *Proceedings of the Web Conference 2021*. 1122–1133.

[27] David Thiel, Renée DiResta, Shelby Grossman, and Elena Cryst. 2021. *Contours and Controversies of Parler*. Technical Report. Stanford University.

[28] Craig Timberg and Drew Harwell. 2021. Pro-Trump forums erupt with violent threats ahead of Wednesday's rally against the 2020 election. *The Washington Post* (2021).

[29] Twitter. 2021. Twitter Academic Research API. https://developer.twitter.com/en/solutions/academic-research/

[30] TwitterSafety. 2020. Twitter's announcement regarding moderation of QAnon. https://twitter.com/TwitterSafety/status/1285726277719199746.

[31] US Department of Justice. 2021. Three Men Charged in Connection with Events at U.S. Capitol. https://www.justice.gov/usao-dc/pr/three-men-charged-connection-events-us-capitol

[32] Jason Wilson. 2021. Rightwingers flock to 'alt tech' networks as mainstream sites ban Trump. *The Guardian* (2021).

[33] Savvas Zannettou, Barry Bradlyn, Emiliano De Cristofaro, Haewoon Kwak, Michael Sirivianos, Gianluca Stringini, and Jeremy Blackburn. 2018. What is gab: A bastion of free speech or an alt-right echo chamber. In *Companion Proceedings of the The Web Conference 2018*.

[34] Savvas Zannettou, Mai ElSherief, Elizabeth Belding, Shirin Nilizadeh, and Gianluca Stringhini. 2020. Measuring and characterizing hate speech on news websites. In *12th ACM Conference on Web Science*.

[35] Savvas Zannettou, Joel Finkelstein, Barry Bradlyn, and Jeremy Blackburn. 2020. A quantitative approach to understanding online antisemitism. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14.

[36] Jing Zeng and Mike S Schäfer. 2021. Conceptualizing "Dark Platforms". Covid-19-Related Conspiracy Theories on 8kun and Gab. *Digital Journalism* (2021), 1–23.

[37] Ethan Zuckerman. 2019. QAnon and the Emergence of the Unreal. *Journal of Design and Science* 1, 6 (2019).