Multilingual Analysis of Narrative Properties in Conspiracist vs Mainstream Telegram Channels

Katarina Laken[1,2], Matteo Melis[3], Marcos Garcia[2], Sara Tonelli[1]

[1] Fondazione Bruno Kessler, Trento, Italy
[2] Centro Singular de Investigación en Tecnoloxías Intelixentes (CiTIUS), Universidade de Santiago de Compostela
[3] Department of Linguistics, Cognitive Science and Semiotics - Aarhus University
{alaken,satonelli}@fbk.eu, mmls@cc.au.dk, marcos.garcia.gonzalez@usc.gal

Abstract

Conspiracist narratives posit an omnipotent, evil group causing harm throughout domains. However, modern-day online conspiracism is often more erratic, consisting of loosely connected posts displaying a general anti-establishment attitude pervaded by negative emotions. We gather a dataset of 300 conspiracist and mainstream, Telegram channels in Italian and English and use the automatic extraction of entities and emotion detection to compare structural characteristics of both types of channels. We create a co-occurrence network of entities to analyze how the different types of channels introduce and use them across posts and topics. We find that conspiracist channels are characterized by anger. Moreover, co-occurrence networks of entities appearing in conspiracist channels are more dense. We theorize that this reflects a narrative structure where all actants are pushed into a single domain. Conspiracist channels disproportionately associate the most central group of entities with anger and fear. We do not find evidence that entities in conspiracist narratives occur across more topics. This could indicate an erratic type of online conspiracism where everything can be connected to everything and that is characterized by a high number of entities and high levels of anger.

1 Introduction

Although many studies focus on hate speech, fake news, and political and religious extremism in isolation, these are connected phenomena. Indeed, false information is often used to spread hate against political enemies (Berk and Filatova, 2019; Kim and Kesari, 2021; Hameleers et al., 2022; Vergani et al., 2022), and rumors are often based on underbelly feelings about minorities (Darmstadt et al., 2019; Vicari et al., 2024). This connection is exemplified by *conspiracism*: a belief in 'sinister, all-powerful conspiratorial forces' that are behind all bad things happening in the world (Bale, 2007). Conspiracism is typical for extremist groups that use conspiracies as a justification for their beliefs (Rousis et al., 2022).[1]

Social media and the Internet provide a platform for all kinds of extremism, including conspiracism. Algorithms favor content that is interesting and engaging, regardless of truth value or offensiveness. Unfortunately, undesirable content is often particularly interesting and engaging (Vosoughi et al., 2018; Van Prooijen et al., 2022). Conspiracy theories, for example, are compelling due to their narrative structure, that creates a greatly simplified and satisfactory projection of a world that is in reality complex and erratic (Bale, 2007; Bleakley, 2023). Conspiracist narratives tend to emphasize people and actants over events and actions (Introne et al., 2020), tying everything back to the one group of antagonists that is held responsible for everything bad in the world, combining actants and events from different domains in one overarching narrative (Tangherlini et al., 2020). On the other hand, Rosenblum and Muirhead (2019) and Pilati et al. (2024) identify a new kind of conspiracism, based more on flaming, attacks, and memes, and less on in-depth discussions of the underlying worldview.

Many studies addressing hate speech, misinformation, and conspiracy theories deal with these topics at the level of individual posts, considering them in isolation. This is a limitation, because social media users and channels that spread a political ideology or conspiracist narrative almost never repeat their whole worldview in a single post (Tangherlini, 2017; Allington et al., 2023).

To address this issue, this paper aims to analyze conspiracist content at the channel level, allowing us to compare the properties of the narra-

[1] We refer to extremism as "an ideological movement, contrary to the democratic and ethical values of a society, that uses different methods, including violence (physical or verbal) to achieve its objectives" (Torregrosa et al., 2023).

Proceedings of the The 9th Workshop on Online Abuse and Harms (WOAH), pages 426–441
August 1, 2025

tives pushed in conspiracist and mainstream channels. We focus our research on Telegram, a social medium noted for its wide range of content (La Morgia et al., 2021). We use well-established NLP approaches to characterize different aspects related to entities, topics, and emotions. While some of these features have been analyzed in the past, we are, to the best of our knowledge, the first to combine them to characterize different types of Telegram channels. Moreover, we perform a multilingual analysis (in Italian and English).

This work is structured around the following research question: *What are the narrative characteristics that are typical for conspiracist Telegram channels compared to mainstream ones?* We address this question by exploring *i)* the emotions associated with the different types of channels, *ii)* co-occurrence networks of named entities to infer structural properties of the way these are introduced and mentioned throughout each channel type and *iii)* the distribution of entities over topics in the text. We use the term 'narrative' to refer to 'the ways in which we construct disparate facts in our own worlds and weave them together cognitively in order to make sense of our reality' (Patterson and Monroe, 1998, p.315). Rather than doing narrative extraction, we develop tools to characterize general tendencies of the narrative developed in each type of channel, such as network analysis to characterize the way a channel introduces entities, and emotion analysis to understand the affect associated with each type of channel and with specific entities in the text.

We introduce language-independent tools for the automatic narrative analysis of social media discourse. The way we use network analysis for the comparison of different types of political content draws inspiration from the work done by for example Tangherlini et al. (2020); Bleakley (2023); Zhao et al. (2024), but this line of research is still very new (Amalvy et al., 2024). In general, this paper explores new ways to analyze large amounts of text data as a whole to infer structural properties of online social media discourse.

2 Background

2.1 Emotion analysis

Many have noted the role that emotions play in the spread and popularity of content (Doroshenko and Tu, 2023). Because messages that elicit extreme emotional responses are shared more often and get pushed in many social media feeds, political actors are incentivized to produce more extremist content (Marino et al., 2024). Extremist groups use emotionally appealing narratives to attract followers (Frischlich et al., 2018). Conspiracy theories also have a particularly high entertainment value (Van Prooijen et al., 2022) and tend to frame debates in terms of belief, rather than science (Reiter-Haas et al., 2024).

Generally, extremist discourse online is characterized by anger and more general negativity (Figea et al., 2016; Ajala et al., 2022). On the other hand, Dragos et al. (2022) found that extremist data seems more angry, and mainstream data more sad. Doroshenko and Tu (2023) found that far-right narratives evoked more enthusiasm and hope compared to centrist appeals. Conspiracist content is also often found to be relatively negative. For example, Zollo et al. (2015) studied the dynamics of affect in Italian Facebook comments under conspiratorial and scientific debate. They found that conspiratorial debates had more overall negative emotions, as did all types of polarized debates. Fong et al. (2021) compare tweets containing conspiracy theories to tweets regarding 'science debates'. Conspiracy theorists used significantly more words associated with negativity than science influencers. Cosgrove and Bahr (2024) compared scientific, conspiratorial, and 'general discussion' on Reddit and Twitter. Conspiracist discourse was related to high levels of anxiety, anger, power and death, which were positively correlated with user engagement. Liu et al. (2024) found that features regarding the emotions present in a text help an LLM to distinguish conspiracist data. Their analysis showed that conspiracy discourse triggered anger, fear and sadness.

2.2 Narrative structure

Ideologically varied extremist groups show great overlap in the types of narratives they employ, such as a paranoid style where the in-group is posited as the victim of a group of evil actants (Johnson, 2018). The posited opponents, such as the government, migrants, or the Jews, are perceived as enemies of the ingroup (Bonetto and Arciszewski, 2021). Extremist groups often create narratives that flatten the complexity of reality (Della Sala, 2010).

Conspiracist worldviews share a lot of these characteristics, such as black and white thinking and a thoroughly evil enemy. Compared to mainstream political narratives they focus disproportionately on actants, rather than actions or events (Introne et al., 2020). Moreover, these actants are seen as the force behind all evil in the world, reducing complex issues to comfortingly simple explanations (Bale, 2007).

Miani et al. (2022) create networks of co-occurrences of LDA-based topics and keywords in both conspiracist and non-conspiracist documents. Narrative networks of conspiracist social media posts were densely connected, and that conspiracy documents showed a large heterogeneity of topics while also being on average more similar to each other than non-conspiracy documents. In other words: they talk about different things, but keep repeating the same story. This aligns with Tangherlini et al. (2020), who compare the narrative network of an untrue conspiracy theory (Pizzagate) to that of a true conspiracy (Bridgegate). They find that true conspiracy has a very simple community structure, with most actants coming from a single domain. These results found a practical implementation in the work of Shahsavari et al. (2020), who study the emergence of conspiracist narratives in the early stages of the Covid-19 pandemic. They use graph representations of actants in a narrative to find actants that fill the role of threat.

Other research points towards a new type of conspiracism emerging with the advent of social media. Rosenblum and Muirhead (2019) and Pilati et al. (2024) point out the emergence of 'anti-science' rather than 'pseudo-science', a conspiracist movement based not around intricate narratives, but presenting itself as a mix of highly emotional, 'meme-like', loosely connected messages spreading hate and mistrust towards any kind of authority and establishment. Pilati et al. (2024) connect this to the characteristics of social media, that incentivizes a more erratic discourse. Interactions in traditional conspiracist message boards are slower, and there is no algorithm-induced incentive to grab people's attention. This incentivizes a more traditional conspiracist narrative, based on pseudo-scientific arguments.

Telegram has characteristics of both traditional Internet media and modern social media. Although the medium is based around followers, there is no 'feed' where users see popular content that could appeal to their interest. We expect this to favor a more traditional kind of conspiracist narrative. On the other hand, the app is known for trolling, extremism, and frequent reposting of messages, which could cultivate a more anti-scientific attitude. Based on this, we expect to see a mix of both traditional pseudo-scientific conspiracism and more emotion-based, erratic anti-scientific conspiracism.

3 Methodology

3.1 Data

We focus on Telegram, a platform often criticized for hosting large amounts of harmful content (Schulze et al., 2022; La Morgia et al., 2021; Hoseini et al., 2023). It can be used for private messaging, group chats, and public or private channels that broadcast messages to a large group of subscribers. We started by using TGDataset (La Morgia et al., 2025), a dataset of over 120,000 channels collected in 2022 that is publicly available for research purposes. We filtered out channels with less than 1,500 posts and channels not in English or Italian using the Langdetect package,[2] leaving us with 10,118 potentially relevant channels. We used seed words to find potentially conspiracist channels. For each channel whose username contained one of those seed words, we randomly sampled 10 posts that were annotated by the first author to classify the channel as either conspiracist, mainstream, or doubt/neither. We excluded all channels that did not clearly fall in either category, as well as channels that were not in English or Italian, exclusively discussed COVID-19/vaccines, or that mainly contained posts consisting of only 1 sentence (not counting URLs). We then collected all channels that were either reposted by or from the channels we previously identified, and annotated those with the same method We defined as conspiracist messages that presuppose or explicitly mention the presence of a powerful, evil group that is being covered up by the 'mainstream' or the 'establishment'. We labeled a channel as conspiracist if there was at least one message that explicitly espoused a conspiracy theory; borderline cases were not included. We found that finding mainstream channels posed a bigger challenge than finding conspiracist channels. We complemented the seedword strategy employed for the conspiracist channels with a manual search through both TGDataset and the Telegram web interface.

This resulted in a sample of 300 Telegram channels (see Table 1). Out of 435 channels identified as possibly conspiracist, we identified 104 English conspiracist channels and 52 Italian ones; we randomly downsampled to 100 and 50 respectively

[2]https://pypi.org/project/langdetect/

We used the public Telegram API in March 2025 to get more messages from all channels in this sample that were still available under the same username. In terms of content, the mainstream class contains a more diverse set of channels than the conspiracist class. We made a subdivision of this class into political/news, crypto/economy, and science/facts channels. For conspiracist channels, we did not select for specific conspiracy theories or ideologies, but our sample is heavily biased towards the QAnon, Sabmyk, and far-right conspiracist sphere.

We preprocessed the data by removing user mentions, URLs, and duplicate sentences. We split the posts in paragraphs of at least 40 characters for further processing; paragraphs that were shorter were concatenated with the next paragraph. This was done to limit influence of post length, that varied considerably among channels; moreover, we found that many channels (of both types) made 'summary posts', in which they mentioned several events or pieces or information without drawing any connection between them. Including these posts would have a disproportionate effect on especially the entity graphs. We capped each channel at 2,000 randomly sampled paragraphs, because especially the entity graph analysis would be heavily impacted by a large variation in the amount of posts considered.

Lang	Type	N channels	Av. tokens
En	Conspiracy	100	59.20
	Mainstream	100	66.22
	Political	52	67.81
	Science/facts	17	60.17
	Crypto/Econ	31	66.89
Ita	Conspiracy	50	105.72
	Mainstream	50	70.99
	Political	31	68.09
	Science/facts	15	78.70
	Crypto/Econ	4	64.58
Total		300	78.70

Table 1: Statistics per channel type and language. Average number of tokens is per paragraph.

3.2 Identification of topics

We manually identified lists of terms most related to five major topics, based on the output of this model: *economy and crypto*, *war*, *climate change*, and *migration*. We used a keyword match to map posts to one or more of these topics (see Appendix A).[3] This resulted in more relevant posts per key topic. Our keywords did not include named entities.

Unsurprisingly, the types of channels differ in the extent to which they talk about each topic. The most commonly discussed topics for all types of channels were *covid*, *economy*, and *war*, with *migration* and *climate change* being of lesser importance. Mainstream and conspiracist channels have on average a roughly equal ratio of posts discussing *war* (mainstream M = 0.13, SD = 0.18, conspiracy M = 0.12, 0.07), but mainstream channels included more discussions about *economy* (M=0.24, SD=0.29) than conspiracist channels (M=0.10, SD=0.04). Conspiracist channels included more discussions of *Covid* (M = 0.36, SD = 0.13) compared to mainstream channels (M = 0.19, SD = 0.16).

3.3 Emotion analysis

We used the multilingual MilaNLP emotion recognition model (Bianchi et al., 2022) to extract the emotions from the text data. This model recognizes sadness, anger, joy, and fear. We split the paragraphs in sentences using the NLTK sentence tokenizer (Bird et al., 2009), extracted three random sentences, both to save computational resources and to limit the influence of post length. We kept the top-scoring emotion for each sentence if it had a confidence score of >0.5. For each paragraph, we counted every emotion that appeared at least once.

3.4 Entity extraction

The goal of this analysis is to see which entities are central and typical for the different types of channels and how specific entities are positioned in the channel narrative. We designed a system for entity extraction and (semi-) linking with several steps. We first used the multilingual NER tool by Tedeschi et al. (2021) for entity extraction. We postprocessed the output of this tool and selected all entities (of any type) with a certainty score of at least 0.7, as through some experimenting we found this to be a reasonable threshold. We then employed the Flair POS tagger (Akbik et al., 2018)

[3]We experimented with BERTopic, (Grootendorst, 2022), using various settings and preprocessing steps, but found that the topics did not align across channel types; moreover, the topic model assigned the vast majority of posts to the -1, or 'trash' topic. Manual inspection of the posts assigned to this topic revealed that they did contain content and that they should have been assigned to some other topic. We therefore resigned to this more crude method

to identify adjectives that were not in a merged entity, or in an entity recognized as a person or a location, lemmatized them and (back-)translated them as a normalizing step.

We selected all entities corresponding to a person or a location that were recognized by the NER tool with a certainty score of at least .99 and that were at least 4 letters long. We chose this threshold because we found that including shorter tokens led to lots of false positives, especially with acronyms, but we manually added some very frequent string matches with less than 4 letters, like 'EU' and 'USA'. We then tokenized all paragraphs using NLTK (Bird et al., 2009) and checked for each token if it appeared in this list. The rationale was that not all channels were equally consistent with capitalization, and the NER module turned out to rely quite heavily on capitalization; this approach allowed us to catch some entities that were not originally covered by the NER tool. We also removed the channel name from the entities of each channel.

We used the Wikidata API (Vrandečić and Krötzsch, 2014) to link as many of the preprocessed entities as possible to a Wikidata page. This allowed us to merge mentions like 'Donald Trump', 'Trump', 'Donald J. Trump' all to the same entity of 'Donald Trump'. We only kept entities without a Wikidata entry if they had a certainty score of at least 0.8.

3.5 Entity graph creation

For each channel, we created co-occurrence networks of the entities in order to map its structural characteristics using the *Networkx* package (Hagberg et al., 2008). Each entity mentioned in a channel represents a node in the network; the edges represent co-occurrences (in the same paragraph). Edges were weighted for amount of co-occurrences.

4 Analysis

4.1 Channel emotions

We conducted a Mann-Whitney U rank test for non-normally distributed data and found a significant difference between conspiracy and mainstream channels (for $\alpha = 0.05$) in emotionality for all emotions apart from joy, with an especially large difference for anger and sadness. Results are reported in Table 2. The mainstream channels show a much larger standard deviation for all emotions, indicating a larger variability, whereas the conspiracy channels are more uniform.

Emotion	Channel	M	SD	M-W U
Fear	Cons	21.33	4.87	9046.00*
	Main	27.12	14.62	
Anger	Cons	63.98	9.01	20284.00**
	Main	34.70	18.79	
Sadness	Cons	17.24	4.95	13085.00*
	Main	15.88	8.28	
Joy	Cons	52.22	9.26	10265.50
	Main	56.13	22.17	

Table 2: Mean (M) and standard deviation (SD) of the % of posts expressing the emotions fear, anger, joy, and sadness, as well as the results of the Mann-Whitney U (M-W U) statistical tests. p values: * <=0.05, ** <= 0.001

Splitting the mainstream sample in the categories 'political', 'crypto/economy', and 'science/facts' showed that this diversity did not solely stem from the wider range of channel typologies in the category 'mainstream', as all individual subcategories still showed a much higher SD compared to the conspiracy class (Table 3).

Emotion	Channel type	M	SD	M-W U
Fear	Conspiracy	21.33	4.87	
	Political	31.85	14.13	3209.00**
	Crypto/economy	19.60	10.41	3353.0**
	Science/facts	23.08	15.47	2485.0
Anger	Conspiracy	63.98	9.01	
	Political	43.70	16.88	10476.0**
	Crypto/economy	21.85	13.35	5178.0*
	Science/facts	25.39	15.97	4630.0**
Sadness	Conspiracy	17.24	4.95	
	Political	18.09	8.50	5941.5**
	Crypto/economy	10.47	5.14	4510.0**
	Science/facts	16.07	7.96	2633.5
Joy	Conspiracy	52.22	9.26	
	Political	41.80	15.33	9110.0**
	Crypto/economy	78.15	14.28	350.0**
	Science/facts	69.24	15.33	805.5

Table 3: Emotions present in the conspiracist channels compared to the different types of mainstream channels present in our sample. p values: * <0.05, ** <= 0.001

The more fine-grained subdivision of the mainstream class also showed that the emotionality for conspiracy was not different for all types of channels. Conspiracist channels expressed more anger than all types of mainstream channels. Although conspiracist channels expressed more fear than mainstream channels related to cryptocurrencies and economics, they expressed less fear than mainstream political channels, and there was no significant difference with channels relating to science and facts. We did not see a difference between conspiracist channels and mainstream channels as a whole in the amount of joy they expressed, but when looking at the individual categories, we found that the conspiracist channels expressed significantly *more* joy than the mainstream political channels, but significantly *less* joy than the other two types of mainstream channels included in our sample.

We hypothesized that these differences might be due to the inherent emotionality related to certain topics (i.e. words like 'war', 'contagion' etc. carrying an inherently negative affect). To gauge the influence of topic on emotionality, we also compared the emotions between conspiracist and mainstream channels for posts associated with each of our five topics (section 3.2; results reported in B). We also calculated an adjusted tf-idf metric to find which words were most informative for each emotion (we calculated these separately for English and Italian). We lemmatized all words using the simplemma module for Python (Barbaresi, 2025). We represented each post with its set of unique lemmas. For each language, for each emotion, we calculated a metric of informativeness using formula 1.

$$I_t = \frac{|\{p_E : t \in p\}|}{|t_L|} * -log_{10}\left(\frac{|\{p_L : t \in p\}|}{|t_L|}\right) \quad (1)$$

1: Formula that gives the information value of a term t for an emotion E in a language L

Using this metric, we extracted the 10 most informative words for each emotion for each language (see appendix C). For joy, we found a rather big difference between English and Italian. Where the English channels associated with joy seem to be disproportionately associated with the crypto sphere, employing words like 'crypto', 'token', and 'trading', the Italian channels associated with joy use words like 'science' and 'world'. Sadness was associated with words relating to death for both languages, whereas fear showed words pertaining to world events like the pandemic and war. Anger associated with words relating to Trump, the government, elections and freedom.

4.2 Narrative network analysis

Metric	Type	M	SD	M-W U
N nodes	Cons Main	2090.25 2047.59	528.09 1145.68	8940.00*
Density	Cons Main	1.156e-2 1.110e-2	0.259e-2 0.625e-2	8469.0**
Weigh. density	Cons Main	0.017e-2 0.022e-2	0.006e-2 0.021e-2	10787.0
Transitivity	Cons Main	0.227 0.174	0.050 0.067	4933.0**
ACC	Cons Main	0.790 0.770	0.038 0.096	11513.0

Table 4: Mean (M), standard deviation (SD) and Mann-Whitney U statistic for metrics of the entity co-occurrence graphs. p values: * <=0.05, ** <= 0.001

We created co-occurrence networks of entities for all channels in our sample except for four channels which were excluded because they contained less than 500 nodes).We then computed different network metrics comparing conspiracy and mainstream channels. Results are reported in Table 4.

Similarly to the emotion occurrences, we tend to see a larger variability (in terms of standard deviation) for mainstream than for conspiracist data. This pattern persisted when splitting out the three types of mainstream channels (see table 9 in appendix E).

A main structural difference is that even though conspiracist networks in our sample have significantly more nodes, they are also slightly but significantly denser than the mainstream networks, when typically a higher number of nodes results in a less dense network. Our edge weights represent frequency of co-occurrence, and weighted density was calculated as the sum of the edge weights divided by the number of possible edges $n(n-1)$, with n = number of nodes.The fact that the weighted density does not significantly differ between the two types of channels thus indicates that the higher density of the conspiracy networks is not due to a greater tendency to mention entities together.

The conspiracist tendency to connect more different entities with each other is also reflected in the transitivity metric, which indicates the ratio of existing triangles to possible triangles; namely, it gives an indication of 'if A and B and A and C are connected, what is the likelihood of C and B being connected as well'. This ratio is significantly higher for conspiracist than for mainstream channels. On the other hand, the average clustering coefficient (ACC) does not differ significantly between the two types of channel networks. The clustering coefficient gives a ratio of the neighbours of a node that are also neighbours among themselves; the ACC is the average of this metric over all nodes in a graph. The transitivity metric puts more weight on high degree vertices, because it calculates the ratio of existing trials on possible triangles, and high degree nodes are part of more potential triangles. Average clustering coefficient, however, averages the clustering coefficient over nodes, giving the same weight to the clustering coefficient of each node. Thus, conspiracy channels having (much) higher transitivity than mainstream channels, but not significantly higher ACC could indicate the existence of a group of very high degree nodes, whose neighbouring nodes are also to largely connected.

Because conspiracist narratives tend to connect a small group of evil conspirors to actants and events from different domains, we hypothesized that conspiracist channels tend to refer to the same set of entities more often than mainstream channels, as well as have them more connected. We found indications for this in the density, transitivity, and ACC metrics of the conspiracist channels. We hypothesized that this might indicate a highly connected community of nodes central in the network that correspond to this group of conspirors.

We used the Louvain algorithm for community detection, a method for detecting groups in a network where nodes are more connected to each other than to the rest of the network. It starts with each node in its own group, then repeatedly merges groups if doing so increases the number of connections inside groups compared to between them. This process continues until no better grouping can be found. For each graph, we ran the Louvain algorithm 3 times (resolution = 2.5), and took the community (of at least 5 nodes) with the highest degree centrality (i.e. the largest fraction of non-community members connected to community members).

These top communities tended to be large in general, significantly (p = 0.031) more so for conspiracy (M = 51.93, SD = 36.59) than for mainstream (M = 49.85, SD = 58.75) channels. We compared the density of the top community over all runs with the density of the network overall. We found that the density of the most central community is not significantly different (p = 0.696) for conspiracist (M = 0.154, SD = 0.094) than for mainstream (M = 0.177, SD = 0.154) channels. However, transitivity was significantly higher ($p < 0.001$) for the conspiracist channels (M = 0.337, SD = 0.153) compared to the mainstream channels (M = 0.270, SD = 0.177).

4.3 Entities over topics

Miani et al. (2022) found that content promoting conspiracy theories is heterogeneous in topics, but that these texts are more similar to each other than mainstream content is. We hypothesize that this happens because in conspiracist content the same entities are mentioned across topics to a larger extent than mainstream content (Bleakley, 2023). We operationalize this intuition by calculating the Jaccard similarity $J(A,B) = \frac{|A \cap B|}{|A \cup B|}$ for all the sets of entities for each pair of topics (see section 3.2) with at least 100 posts for each topic. We average all of these similarities per channel in order to get a metric of entity overlap. The mainstream channels on average showed a slightly higher (M=0.664, SD=0.263) Jaccard similarity over all entities than the conspiracy channels (M=0.648, SD=0.263), but this difference was not statistically significant. Both conspiracist and mainstream channels unsurprisingly showed a higher average Jaccard similarity when only taking into account entities from the top communities. Mainstream channels had higher (M=0.740, SD=0.247) Jaccard similarity than conspiracist channels (M=0.706, SD=0.264) when only taking into account top communities, and this difference was statistically significant (U = 48203.5, p = 0.025).

4.4 Top community entities

In order to inspect which entities were associated with the top communities for each channel type, we implemented the same informativity metric used in section 4.1 (see formula 1), where l = entity, p = top community, and E = channel label (conspiracy or mainstream). Appendix D shows the most informative and frequent entities per category. Whereas mainstream channels in both English and Italian tend to mention countries and mainstream establishments, the conspiracy channels mention people and more miscellaneous entities. For English, Covid-19 seems to be mainly a topic of the conspiracist channels, whereas in Italian, this appears to be a topic for the mainstream channels as well. Moreover, the Italian channels displayed entities like "NATO" and "nazi", with the most relevant word

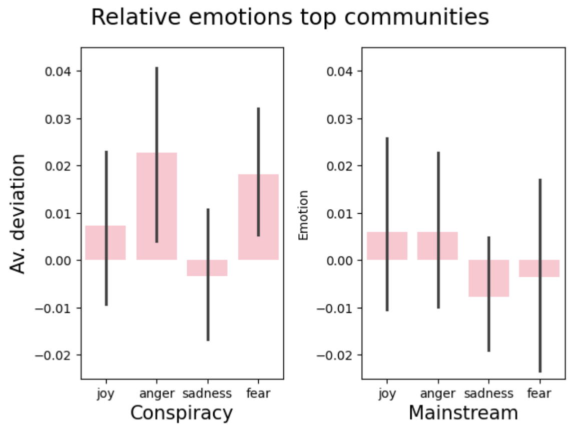

Figure 1: Relative disproportion of emotions for the top community compared to channel mean

being "anti", underlying the anti-establishment attitude of these conspiracy groups (Pilati et al., 2024). Finally, the English channel having more entities related to international political (e.g., Syria, Ukraine, etc.), while the Italians focus on internal politics (Rome, Milan, Naples, etc.).

Given the central role that conspiracy theories attribute to the evildoers who are behind all bad events, we investigated whether a post mentioning an entity corresponds to a stronger emotional value. Using the same strategy as in Section 4.1, for every emotion, we calculated the difference in ratios between posts expressing an emotion and posts expressing the same emotion, also containing an entity. The results of the conspiracist channels reported in Figure 1 showcase a disproportion in anger and fear when the post contains a top community entity. The emotions for the mainstream channels do not deviate from the channel's emotions as a whole.

5 Discussion

The classical view of conspiracism is that they are all-encompassing theories that tie back actants and events from different domains to an omnipotent group of evil conspirors (Bale, 2007). Several data-driven studies have found evidence for this type of narrative structure (Tangherlini, 2017; Miani et al., 2021; Bleakley, 2023). However, social media is also said to have brought about a more erratic, post-truth type of conspiracism, that is less about providing an alternative theory and more about emotion and anti-establishment sentiments (Rosenblum and Muirhead, 2019). Pilati et al. (2024) found this type of conspiracism to be more typical of a social medium like Reddit, which is driven by a more like-based system than for example traditional internet message boards.

Our analysis provided evidence that conspiracist data on Telegram shows characteristics of both types of narratives. Throughout all analyses we find higher variability for the mainstream category of channels. Although our mainstream channels were much more varied, these differences tended to persist when we split out the mainstream channels in the three identified subcategories. This points towards a high level of uniformity among conspiracist Telegram channels.

We expected the conspiracist channels to be characterized by negative emotions, as found by Figea et al. (2016); Cosgrove and Bahr (2024) and others. We did find that conspiracist channels had much higher levels of anger than mainstream channels and marginally more sadness, but mainstream channels expressed more fear. This difference was only significant for the political/news mainstream channels, but not for the economy/crypto or science-related channels. This might be due to the topics covered by this channels, such as war and the COVID-19 pandemic. We hypothesize that these topics inherently include words associated with fear and sadness, as also illustrated by our informative word analysis.

We did not find a significant difference in the amount of joy expressed between the two types of channels. Our findings thus do not clearly support the idea that conspiracist content is more negative than mainstream content, but it is in line with earlier research identifying anger as a characteristic of extremist data (Dragos et al., 2022) and the relationship between conspiracism and anger as a personality trait (Szymaniak et al., 2023).

Analysis of the most informative words per emotion showed that words relating to Trump, elections, protest and the police were associated with anger, whereas topics like war and pandemic were more associated with sadness and fear. More research is needed to disentangle cause and effect here, as these findings can either indicate that conspiracist channels disproportionately discuss topics associated with anger, or that the way conspiracist Telegram channels communicate triggers the emotion ’anger’ in our model, which would lead to an over-representation of these topics in our informative word analysis.

We quantified the way the different types of channels combine entities by analyzing networks of entity co-occurrences. We found that the entity co-occurrence networks for conspiracist channels are slightly but significantly more dense than mainstream channels. We theorize this indicates the tendency of conspiracist discourse to create narrative cohesion by combining more different entities.

We split our channel networks in communities using the Louvain algorithm and identified 'top communities', i.e. the communities with the highest group degree centrality. These top communities were more transitive for the conspiracist category, but (unexpectedly) not more dense. We hypothesize that these communities indicate the 'conspirors' in the conspiracy narrative. This is in line with our finding that the entities associated with the top community do not deviate from the emotionality of the channel as a whole for mainstream channels, but for conspiracist channels these entities are disproportionately associated with anger and fear. When looking at the most typical entities per channel type, we found that mainstream channels tend to mention more states and established entities, whereas the conspiracist channels were typified by more miscellaneous entities.

Finally, we expected conspiracist narratives to use entities to 'stitch together' different types of domains into one overarching narrative (Miani et al., 2021; Bleakley, 2023; Introne et al., 2020). However, we did not find evidence for that; we did not find a greater overlap in entities between topics for conspiracy channels compared to mainstream channels. This is more in line with the erratic anti-scientific style of online conspiracism (Rosenblum and Muirhead, 2019; Pilati et al., 2024). On the other hand, it should be considered that our keyword matching method of extracting topics might not have been fine-grained enough. Another option is that many of our conspiracist channels were rather mono-thematic, centering around the COVID-19 pandemic, resulting in one domain characterized by a wide variety of entities. We recommend that future studies further explore the relationship between entities and topics in online conspiracist discourse.

5.1 Future work

While this work provides new insights in the way conspiracist channels on Telegram frame their narratives, it also opens the door to further questions. Our dataset almost exclusively contains QAnon-related and far-right conspiracy channels; follow-up investigation is necessary to determine how well these results hold up with other types of ideologies and conspiracies. Future research is also needed to better understand the influence of language and culture on the metrics we used to describe the narrative and stylistic tendencies of Telegram channels. Moreover, we only considered four basic emotions, but we know that certain types of extremism are associated with more complex emotional features like nostalgia (Farokhi, 2022).

We were not able to find direct evidence for domain stitching of topics through the repeated use of entities, even though this was found in earlier research, and the dense entity graph of the conspiracist channels should give reason to expect this. However, we used a very crude method to detect topics, and many of the conspiracy channels were relatively skewed towards one of the topics (covid). More advanced analysis is necessary for more conclusive evidence in this regard.

6 Conclusion

The goal of this paper was to analyze the narrative tendencies of conspiracist channels on Telegram. We collected a dataset of 300 conspiracist and mainstream channels in English and Italian, a pseudonymized version of which will be released upon request. We conducted emotion classification and named entity extraction. We extracted topics based on keyword matching and created co-occurrence networks of the entities in each channel. Our results show a narrative network where everything is connected to anything, and a group of central actants associated with fear and anger, but no clear evidence of actants stitching together different domains. We theorize that these characteristics contribute to a narrative structure that is cohesive due to actants co-occurring in different constellations, and attractive due to the emotional nature of the discourse.

Limitations

There are several limitations to this study that should be taken into account when interpreting our results. Even though our dataset was annotated manually, this annotation was based on a rather small sample of posts per channel, and carried out by one annotator. Although we made sure not to include borderline/doubtful cases to keep the sample clean, this procedure might have introduced noise in our dataset. Moreover, we did not do any in-depth analysis to the differences between Italian and English channels. Next, although our sample included different types of mainstream channels, we did not dig deeper into the differences between the different types of mainstream channels; we will leave this for future research. The same goes for the variety of topics in our sample of conspiracy channels. Our analysis also relies heavily on the results of named entity recognition and linking. However, these tasks are difficult, and noise and errors can change the outcome of the results. Our entity graphs relied on off-the-shelf tools for entity extraction and linking, but manual inspection of the entities showed that these are not infallible. We expect that this could influence especially the conspiracist channels, as these could be expected to use more non-normative language. Finally, our pipeline is language-agnostic, but not all tools are available for all languages. This makes it hard to potentially extend the approach to low-resource languages.

Ethics Statement

This work has been carried out to better understand conspiracist narratives, raising awareness on this kind of online conversations with the final goal to contrast them. Our data sources include both the TGDataset, a freely available dataset released for research purposes, and data that we collected ourselves using the Telegram API. The channel names have been removed from our dataset and user mentions have been also deleted. This pseudoanonymisation process has been performed to avoid the retrieval of the original messages from Telegram and the identification of the users who posted them. Nevertheless, we will release the dataset only upon request and for research purposes.

Acknowledgments

This work has been supported by the European Union's Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreements no. 101073351 (HYBRIDS) and no. 101167978 (DEMINE), as well as by the Galician Government (ERDF 2024-2027: Call ED431G 2023/04), and by a Ramón y Cajal grant (RYC2019-028473-I). We also thank the anonymous reviewers for their thoughtful and insightful comments and suggestions.

References

Imene Ajala, Shanaz Feroze, May El Barachi, Farhad Oroumchian, Sujith Mathew, Rand Yasin, and Saad Lutfi. 2022. Combining artificial intelligence and expert content analysis to explore radical views on twitter: Case study on far-right discourse. *Journal of Cleaner Production*, 362:132263.

Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual string embeddings for sequence labeling. In *COLING 2018, 27th International Conference on Computational Linguistics*, pages 1638–1649.

Daniel Allington, David Hirsh, and Louise Katz. 2023. Correlation between coronavirus conspiracism and antisemitism: a cross-sectional study in the united kingdom. *Scientific Reports*, 13.

Arthur Amalvy, Vincent Labatut, and Richard Dufour. 2024. Renard: A modular pipeline for extracting character networks from narrative texts. *Journal of Open Source Software*, 9(98):6574.

Jeffrey M Bale. 2007. Political paranoia v. political realism: On distinguishing between bogus conspiracy theories and genuine conspiratorial politics. *Patterns of prejudice*, 41(1):45–60.

A Barbaresi. 2025. Simplemma: a simple multilingual lemmatizer for python. https://github.com/adbar/simplemma DOI: 10.5281/zenodo.4673264. Computer software.

Enis Alonso Berk and Elena Filatova. 2019. Incendiary news detection. In *The Thirty-Second International Flairs Conference*.

Federico Bianchi, Debora Nozza, and Dirk Hovy. 2022. XLM-EMO: Multilingual Emotion Prediction in Social Media Text. In *Proceedings of the 12th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*. Association for Computational Linguistics.

Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc.".

Paul Bleakley. 2023. Panic, pizza and mainstreaming the alt-right: A social media analysis of pizzagate and the rise of the qanon conspiracy. *Current Sociology*, 71(3):509–525.

Eric Bonetto and Thomas Arciszewski. 2021. The creativity of conspiracy theories. *The Journal of Creative Behavior*, 55(4):916–924.

Tylor Cosgrove and Mark Bahr. 2024. The language of conspiracy theories: Negative emotions and themes facilitate diffusion online. *Sage Open*, 14(4):21582440241290413.

Alina Darmstadt, Mick Prinz, and Oliver Saal. 2019. *The Murder of Keira. Misinformation and Hate Speech as Far-Right Online Strategies*, pages 155–168. transcript Verlag, Bielefeld.

Vincent Della Sala. 2010. Political myth, mythology and the european union. *JCMS: Journal of Common Market Studies*, 48(1):1–19.

Larissa Doroshenko and Fangjing Tu. 2023. Like, share, comment, and repeat: Far-right messages, emotions, and amplification in social media. *Journal of Information Technology & Politics*, 20(3):286–302.

Valentina Dragos, Delphine Battistelli, Aline Etienne, and Yolène Constable. 2022. Angry or sad? emotion annotation for extremist content characterisation. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 193–201.

Zeinab Farokhi. 2022. Making freedom great again: Conspiracy theories, affective nostalgia and alignment, and the right-wing base grammars of the# freeedomconvoy. *Global Media Journal*, 14(1):67–92.

Leo Figea, Lisa Kaati, and Ryan Scrivens. 2016. Measuring online affects in a white supremacy forum. In *2016 IEEE conference on intelligence and security informatics (ISI)*, pages 85–90. IEEE.

Amos Fong, Jon Roozenbeek, Danielle Goldwert, Steven Rathje, and Sander Van Der Linden. 2021. The language of conspiracy: A psychological analysis of speech used by conspiracy theorists and their followers on twitter. *Group Processes & Intergroup Relations*, 24(4):606–623.

Lena Frischlich, Diana Rieger, Anna Morten, and Gary Bente. 2018. The power of a good story: Narrative persuasion in extremist propaganda and videos against violent extremism. *International Journal of Conflict and Violence (IJCV)*, 12:a644–a644.

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

Aric Hagberg, Pieter J Swart, and Daniel A Schult. 2008. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Laboratory (LANL), Los Alamos, NM (United States).

Michael Hameleers, Toni Van der Meer, and Rens Vliegenthart. 2022. Civilized truths, hateful lies? incivility and hate speech in false information–evidence from fact-checked statements in the us. *Information, Communication & Society*, 25(11):1596–1613.

Mohamad Hoseini, Philipe Melo, Fabricio Benevenuto, Anja Feldmann, and Savvas Zannettou. 2023. On the globalization of the qanon conspiracy theory through telegram. In *Proceedings of the 15th ACM Web Science Conference 2023*, pages 75–85.

Joshua Introne, Ania Korsunska, Leni Krsova, and Zefeng Zhang. 2020. Mapping the narrative ecosystem of conspiracy theories in online anti-vaccination discussions. In *International Conference on Social Media and Society*, pages 184–192.

Jessica Johnson. 2018. The self-radicalization of white men:"fake news" and the affective networking of paranoia. *Communication Culture & Critique*, 11(1):100–115.

Jae Yeon Kim and Aniket Kesari. 2021. Misinformation and hate speech: The case of anti-asian hate speech during the covid-19 pandemic. *Journal of Online Trust and Safety*, 1(1).

Massimo La Morgia, Alessandro Mei, and Alberto Maria Mongardini. 2025. Tgdataset: Collecting and exploring the largest telegram channels dataset. *arXiv preprint*.

Massimo La Morgia, Alessandro Mei, Alberto Maria Mongardini, and Jie Wu. 2021. Uncovering the dark side of telegram: Fakes, clones, scams, and conspiracy movements. *arXiv preprint arXiv:2111.13530*.

Zhiwei Liu, Boyang Liu, Paul Thompson, Kailai Yang, and Sophia Ananiadou. 2024. Conspemollm: Conspiracy theory detection using an emotion-based large language model. In *ECAI 2024*, pages 4649–4656. IOS Press.

Erik Bran Marino, Jesus M Benitez-Baleato, and Ana Sofia Ribeiro. 2024. The polarization loop: How emotions drive propagation of disinformation in online media—the case of conspiracy theories and extreme right movements in southern europe. *Social Sciences*, 13(11):603.

Alessandro Miani, Thomas Hills, and Adrian Bangerter. 2021. Loco: The 88-million-word language of conspiracy corpus. *Behavior research methods*, pages 1–24.

Alessandro Miani, Thomas Hills, and Adrian Bangerter. 2022. Interconnectedness and (in) coherence as a signature of conspiracy worldviews. *Science Advances*, 8(43):eabq3668.

Molly Patterson and Kristen Renwick Monroe. 1998. Narrative in political science. *Annual review of political science*, 1(1):315–331.

Federico Pilati, Tommaso Venturini, Pier Luigi Sacco, and Floriana Gargiulo. 2024. Pseudo-scientific versus anti-scientific online conspiracism: A comparison of the flat earth society's internet forum and reddit. *new media & society*, page 14614448241252593.

Markus Reiter-Haas, Beate Klösch, Markus Hadler, and Elisabeth Lex. 2024. Framing analysis of health-related narratives: Conspiracy versus mainstream media. *arXiv preprint arXiv:2401.10030*.

Nancy L Rosenblum and Russell Muirhead. 2019. *A lot of people are saying: The new conspiracism and the assault on democracy*. Princeton University Press.

Gregory J Rousis, F Dan Richard, and Dong-Yuan Debbie Wang. 2022. The truth is out there: The prevalence of conspiracy theory use by radical violent extremist organizations. *Terrorism and Political Violence*, 34(8):1739–1757.

Heidi Schulze, Julian Hohner, Simon Greipl, Maximilian Girgnhuber, Isabell Desta, and Diana Rieger. 2022. Far-right conspiracy groups on fringe platforms: A longitudinal analysis of radicalization dynamics on telegram. *Convergence: The International Journal of Research into New Media Technologies*, 28(4):1103–1126.

Shadi Shahsavari, Pavan Holur, Tianyi Wang, Timothy R Tangherlini, and Vwani Roychowdhury. 2020. Conspiracy in the time of corona: automatic detection of emerging covid-19 conspiracy theories in social media and the news. *Journal of computational social science*, 3(2):279–317.

Kinga Szymaniak, Marcin Zajenkowski, Krzysztof Fronczyk, Sarah Leung, and Eddie Harmon-Jones. 2023. Trait anger and approach motivation are related to higher endorsement of specific and generic conspiracy beliefs. *Journal of Research in Personality*, 104:104374.

Timothy R Tangherlini. 2017. Toward a generative model of legend: Pizzas, bridges, vaccines, and witches. *Humanities*, 7(1):1–19.

Timothy R Tangherlini, Shadi Shahsavari, Behnam Shahbazi, Ehsan Ebrahimzadeh, and Vwani Roychowdhury. 2020. An automated pipeline for the discovery of conspiracy and conspiracy theory narrative frameworks: Bridgegate, pizzagate and storytelling on the web. *PloS one*, 15(6):e0233879.

Simone Tedeschi, Valentino Maiorca, Niccolò Campolungo, Francesco Cecconi, and Roberto Navigli. 2021. WikiNEuRal: Combined neural and knowledge-based silver data creation for multilingual NER. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2521–2533, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Javier Torregrosa, Gema Bello-Orgaz, Eugenio Martínez-Cámara, Javier Del Ser, and David Camacho. 2023. A survey on extremism analysis using natural language processing: definitions, literature review, trends and challenges. *Journal of Ambient Intelligence and Humanized Computing*, 14(8):9869–9905.

Jan-Willem Van Prooijen, Joline Ligthart, Sabine Rosema, and Yang Xu. 2022. The entertainment value of conspiracy theories. *British Journal of Psychology*, 113(1):25–48.

Matteo Vergani, Alfonso Martinez Arranz, Ryan Scrivens, and Liliana Orellana. 2022. Hate speech in a telegram conspiracy channel during the first year of the covid-19 pandemic. *Social Media+ Society*, 8(4):20563051221138758.

Rosa Vicari, Or Elroy, Nadejda Komendantova, and Abraham Yosipof. 2024. Persistence of misinformation and hate speech over the years: The manchester arena bombing. *International Journal of Disaster Risk Reduction*, 110:104635.

Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *science*, 359(6380):1146–1151.

Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.

Wanying Zhao, Siyi Guo, Kristina Lerman, and Yong-Yeol Ahn. 2024. Discovering collective narratives shifts in online discussions. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, pages 1804–1817.

Fabiana Zollo, Petra Kralj Novak, Michela Del Vicario, Alessandro Bessi, Igor Mozetič, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. 2015. Emotional dynamics in the age of misinformation. *PloS one*, 10(9):e0138740.

A Topic seed words

Topic	Words
Covid	’\Wvaccin’, ’\Wcovid[*A-Za-z]’, ’corona’, ’doctor’, ’medic’, ’salute’, ’health’, ’sanitari’, ’myocarditis’, ’miocardite’, ’coronavirus’, ’pandemic’, ’pandemia’, ’infection’, ’outbreak’, ’restrictions’, ’restrizioni’, ’mask’, ’mascherin’, ’decessi’, ’tamponi’, ’virus’, ’booster’, ’rna’, ’guariti’, ’deceduti’, ’green pass’, ’farmaceutic’ ’pharma’, ’hospital’, ’ospedal’, ’jab’, ’terapie’, ’contagiati’, ’ricoveri’, ’lockdown’, ’vaxx’, ’no ?vax’, ’protein’
War	’russia’, ’peace’, ’\Wwars?\W’, ’\Wpace\W’, ’guerr[ae]’, ’soldat[io]’, ’\Wtruppe’, ’militar[yi]’, ’troops’, ’soldiers?’, ’\Wsanctions?’, ’\Wsanzion[ie]’, ’gasdotto’, ’rubli’, ’rubles?’, ’rubble’, ’nuclear’, ’weapons?’, ’radioactive’, ’missil[ei]s?’, ’hypersonic’, ’ipersonic’, ’ballistic’, ’rifle’, ’drone’, ’warplane’, ’biolab’, ’peacekeeper’, ’warmonger’, ’airstrike’, ’injured’, ’feriti’, ’occupier’
Economy	’cr[yi]pto’, ’exchange’, ’currency’, ’market’, ’\Wtrading’, ’\Wtrade’, ’asset’, ’payment’, ’bitcoin’, ’mining’, ’price’, ’value’, ’nasdaq’, ’\Wstock[s\W]’, ’investor’, ’dividend’, ’blockchain’, ’profit’, ’binance’, ’finance’, ’\Wtassi’, ’\Wrates’, ’\Weconom’, ’memecoin’, ’valut[ae]’, ’\Wmercat[oi]’, ’investitor[ei]’, ’finanz’
Migration	’migrant’, ’migrat’, ’borders?\W’, ’immigration’, ’immigrazione’, ’asilo’, ’citizenship’, ’cittadinanza’, ’refugees?’, ’racism’, ’razzismo’, ’discrimination’, ’discriminazione’, ’racial’, ’\Wrazzial’, ’rifugiat’, ’profughi’
Climate	’clima’, ’temperatur’, ’riscaldamento’, ’geoingegneria’

Table 5: List of seed words used to map each post to one or more topics

B Emotions over topics

Emotion	Topic	Channel type	N	M	SD	Mann-Whitney U	p
Fear	Covid	Conspiracy	106442	25.03	4.45	8448.50	>0.001
		Mainstream	55140	31.38	14.97		
	Economy	Conspiracy	28966	21.18	4.96	10452.00	0.288
		Mainstream	70941	23.54	12.30		
	War	Conspiracy	35888	26.43	5.34	8079.00	>0.001
		Mainstream	39132	33.19	17.24		
	Climate	Conspiracy	3355	29.94	11.67	7774.00	>0.001
		Mainstream	2948	38.49	23.30		
	Migration	Conspiracy	5909	19.18	9.81	6222.50	>0.001
		Mainstream	6903	31.69	19.80		
Anger	Covid	Conspiracy	106442	70.69	7.57	21320.00	>0.001
		Mainstream	55140	35.70	18.27		
	Economy	Conspiracy	28966	63.20	13.22	18625.50	>0.001
		Mainstream	70941	38.45	20.79		
	War	Conspiracy	35888	73.64	7.58	21002.00	>0.001
		Mainstream	39132	41.50	20.07		
	Climate	Conspiracy	3355	70.14	14.12	18481.50	>0.001
		Mainstream	2948	30.99	24.61		
	Migration	Conspiracy	5909	81.84	9.65	19486.00	>0.001
		Mainstream	6903	47.12	25.04		
Sadness	Covid	Conspiracy	106442	18.19	5.93	11364.00	0.880
		Mainstream	55140	18.59	9.91		
	Economy	Conspiracy	28966	16.40	4.78	13041.50	0.017
		Mainstream	70941	15.14	8.03		
	War	Conspiracy	35888	16.37	5.11	11653.50	0.592
		Mainstream	39132	16.98	12.34		
	Climate	Conspiracy	3355	14.82	10.53	10514.50	0.658
		Mainstream	2948	15.14	14.82		
	Migration	Conspiracy	5909	14.27	7.23	11585.00	0.281
		Mainstream	6903	14.58	13.15		
Joy	Covid	Conspiracy	106442	45.59	8.66	8798.00	0.001
		Mainstream	55140	55.08	22.89		
	Economy	Conspiracy	28966	55.03	12.92	9398.00	0.014
		Mainstream	70941	60.37	21.35		
	War	Conspiracy	46431	45.15	8.35	9313.5	0.010
		Mainstream	51082	50.90	23.58		
	Climate	Conspiracy	3355	43.08	17.31	5837.00	>0.001
		Mainstream	2948	59.46	24.55)		
	Migration	Conspiracy	5909	38.34	10.86	8792.00	0.006
		Mainstream	6903	49.45	28.62)		

Table 6: Table showing the differences in emotionality split out by topic and channel type

C Most informative lemmas over emotions

Language	Fear	Anger	Sadness	Joy
English	risk fear warn attack death pandemic infection area coronavirus military	protest police law court election mandate try criminal why freedom	die lose death dead injure suffer hospital sorry family loss	crypto token share trading great buy price market join free
Italian	rischiare (to risk) morte (death) situazione (situation) caso (case) ucraina (ukraine) decesso (demise) secondo (second) tassare (tax; rates) russo (Russian) coronavirus (corona virus)	governo (government) chiedere (ask) dire (say) chi (who) trump (Trump) stesso (self) volere (want) senza (without) presidente (president) andare (go)	morire (die) morta (dead) era (was; era) vita (life) morte (death) dopo (after) poco (little) decesso (demise) giorno (day) anno (year)	nostro (ours) grazia (grace) grande (big, great) video (video) mondo (world) primo (first) gruppo (group) mio (my, mine) scienza (science) parlare (talk)

Table 7: Table showing most informative word per emotion per language

D Entities top communities

Language	Conspiracy		Mainstream	
	Most informative	Most frequent	Most informative	Most frequent
English	vaccine covid-19 pfizer Food and drug administration Biontech tate Shahin Biontec boss Cheung Oahu	vaccine covid-19 pfizer Food and drug administration Biontech tate Shahin biontec boss Cheung Oahu	bitcoin russia united states price ukraine state mercury nasa unesco myanmar	bitcoin russia united states ukraine state price venezuela mercury nasa unesco
Italian	anti sion nato nazi senator crisanti dio ella mestre toulouse	anti sion nato senator nazi crisanti ella dio mestre toulouse	nato telegram gazzetta ufficiale rome czech republic australia covid-19 centers for disease control and prevention la7 State of Palestine	nato telegram la7 gazzetta ufficiale rome State of Palestine Naples Israele Senator Czech Republic

Table 8: Table showing top entities (by informativity and frequency) for each channel type and language.

E Network metrics split out by channel subtype

Metric	**Channel type**	**M**	**SD**	**Mann-Whitney U**	p
N nodes	Conspiracy	2090.24	528.09		
	Political	2311.93	1237.45	6078.50	0.884
	Crypto/economy	1612.56	736.64	1392.50	>0.001
	Science/facts	1818.10	1092.93	1469.5	0.003
Density	Conspiracy	1.156e-2	0.259e-2		
	Political	0.982e-2	0.535e-2	3613.00	>0.001
	Crypto/economy	1.306-e2	0.655	2874.00	0.249
	Science/facts	1.241e-2	0.743e-2	1982.0	0.305
Weigh. density	Conspiracy	0.017	0.006		
	Political	0.018	0.013	5209.00	0.054
	Crypto/economy	0.027	0.032	3373.00	0.003
	Science/facts	0.026	0.023	2205.00	0.864
Isolate ratio	Conspiracy	0.025	0.016		
	Political	0.027	0.038	4318.00	0.000
	Crypto/economy	0.032	0.036	2415.00	0.631
	Science/facts	0.043	0.038	2702.0	0.083
Transitivity	Conspiracy	0.227	0.050		
	Political	0.163	0.054	2047.00	>0.001
	Crypto/economy	0.167	0.061	1154.00	>0.001
	Science/facts	0.210	0.091	1732.0	0.047
ACC	Conspiracy	0.790	0.038		
	Political	0.773	0.107	7063.00	0.062
	Crypto/economy	0.777	0.080	2689.00	0.621
	Science/facts	0.755	0.083	1761.00	0.061

Table 9: Table showing network metrics split out by channel type