Vaccine Narratives on Telegram: A Data-Driven Study of Information Diffusion in Brazil

Fernando H. M. de Oliveira[1], **Cleyton Mário de O. Rodrigues** [1]

[1]University of Pernambuco (UPE)
Recife, Pernambuco, Brazil

fhmo@ecomp.poli.br, cleyton.rodrigues@upe.br

Abstract. *This study analyzes over 330,000 vaccine-related Telegram messages in Brazil (2020–2025) to investigate diffusion dynamics. Results show that diffusion is strongly associated with message exposure, with visibility dominating observed engagement patterns. Forwarding is consistently associated with higher redistribution and reach, while user reactions show limited association with propagation. Contrary to common assumptions, content volume and user activity are not consistently linked to higher engagement outcomes. These findings highlight Telegram's broadcast-like dynamics providing insights into large-scale information spread in low-moderation environments, and offering empirical support for public health strategies.*

1. Introduction

Brazil presents a critical context for examining how social media shapes vaccine discourse, where despite broad public recognition of immunization's importance, the widespread dissemination of misinformation has fueled anti-science narratives and entrenched vaccine skepticism [Massarani et al. 2020]. Rooted in structural distrust toward political institutions and the pharmaceutical industry, this phenomenon has significantly impeded progress toward community immunity [Khadafi et al. 2022]. The crisis was further exacerbated by intense political polarization, which accelerated the large-scale spread of vaccine disinformation following the first domestic COVID-19 case and subsequently undermined adherence to routine and childhood immunization programs through politically driven scientific denial [Galhardi et al. 2022][Wilson and Wiysonge 2020][Araújo et al. 2024]. Although digital platforms hold the potential to disseminate evidence-based health information, they simultaneously function as highly efficient channels for the rapid, large-scale propagation of misleading narratives regarding vaccine safety and efficacy [Araújo et al. 2024].

Within this ecosystem, Telegram emerged as a key platform for the dissemination of vaccine-related misinformation in Brazil. Its relatively low level of content moderation and large group functionalities make it particularly effective for message amplification and redistribution [Tokojima Machado et al. 2020]. Anti-vaccine discourse on Telegram is deeply intertwined with broader conspiracy ecosystems. These narratives have grown substantially, with conspiracy-related content increasing by 290% during the pandemic [Cugler de Moraes Silva 2024]. The migration of anti-vaccine communities to Telegram was further accelerated by content moderation policies on mainstream platforms, which led to the deplatforming of misinformation actors and their relocation to alternative spaces [Díaz and Ricaurte 2024].

Health authorities identified the spread of false information as a contributing factor to declining immunization rates, driven by emotionally charged narratives that lack scientific evidence [Domingues et al. 2020]. Moreover, research shows that misinformation propagates efficiently through mass-messaging platforms, reaching broad audiences and reinforcing vaccine hesitancy [Barkhad et al. 2026]. Vaccine discourse became deeply embedded in political identity, with public debates frequently influenced by political leadership and ideological divisions [Amaral et al. 2022]. As a result, vaccine-related discussions often extend beyond health considerations, encompassing broader sociopolitical narratives. In this context, understanding the structural dynamics of information diffusion on Telegram is essential for analyzing how vaccine-related narratives spread, evolve, and influence public perception.

This study focuses on the structural and behavioral patterns of vaccine-related discourse rather than verifying the factual accuracy of the content. It examines diffusion mechanisms within Telegram-based vaccine discourse in Brazil (2020–2025). In this context, diffusion is operationalized through aggregate indicators (e.g., views and forwarding counts), rather than explicit modeling of network structures or cascade pathways. Consequently, diffusion should be interpreted as observed redistribution patterns rather than fully reconstructed transmission processes.

The remainder of this study is organized as follows: Section 2 presents the related work. Section 3 describes the materials and methods, including the data source, preprocessing pipeline, and analytical framework. Section 4 details the findings derived from the quantitative analysis and metric evaluation. Section 5 outlines the limitations of the current study, and Section 6 provides concluding remarks and directions for future work.

2. Related Work

Telegram emerged as a hotspot for misinformation during critical political events, including elections in countries like Italy, Brazil, and the United States, challenging election integrity and promoting divisiveness [Alvisi et al. 2024][Cavalini et al. 2023]. Studies from Brazil demonstrate substantial political mobilization on Telegram in recent years, suggesting a mass migration from other mainstream platforms like Twitter, YouTube, and Facebook [Júnior et al. 2021]. Research comparing different cultural contexts has unveiled distinct trends and dynamics within conspiracy theory discussions across languages and regions [Alvisi et al. 2024]. Analysis of almost 5 million messages across 137,000 users in public Telegram groups revealed distinct curiosity stimulation profiles, dynamic transitions in user behaviors, and the varying influence of groups in amplifying curiosity, particularly around politically relevant topics [Figueiredo Vasconcelos et al. 2025].

While recent studies have explored behavioral transitions and curiosity profiles within the context of political mobilization, our research shifts the focus toward the structural dynamics of public health discourse over a broader longitudinal period (2020–2025). Unlike previous works that emphasize psychological drivers or the identification of specific 'opinion leader' roles (influencers, conversation starters, etc.) [Tucci 2023], this study deconstructs the technical nature of engagement on Telegram.

3. Materials and Methods

The reference dataset for this study is the 'Brazilian Social Media Anti-vaccine Information Disorder Dataset (2020–2025)', presented by Phillipe Cardenuto et al. (2026)[Cardenuto et al. 2026]. This dataset comprises a longitudinal collection of over 4 million multimedia messages retrieved from Telegram channels and groups. We filtered a subset of this database to include only vaccine-related messages in Portuguese, resulting in 330,567 records across 24 variables, which were used for the subsequent analysis.

The subset encompasses metadata related to message timestamps, content attributes, user and channel identifiers, engagement metrics, and information diffusion pathways. Key variables include temporal fields (*collected_date, date, and edit_date*), engagement metrics (*views, reactions,* and *n_forwards*), diffusion metadata (*forward_from, forward_from_views, forward_from_n_forwards, and reply_to*), content descriptors (*text_content, media_type, media_url, is_vaccine_related, and language*), and entity identifiers (user_id, channel_id, and message_id).

Prior to analysis, the dataset was preprocessed to ensure data quality and analytical consistency:

- Timestamp standardization: The date column was converted to a datetime format using strict coercion to handle malformed entries.
- Content filtering: Records were retained only if *language* == 'Portuguese', *text_content* was non-null, and text length exceeded 15 characters to exclude trivial or automated system messages.
- Missing value handling: Numerical engagement fields were imputed with zeros during aggregation to preserve distributional integrity for downstream calculations. Categorical fields like *media_type* were assigned a default 'Text Only' label where missing.
- Derived features:
  - *is_forwarded*: Boolean indicator derived from forward_from.notna()
  - *engagement*: Composite metric calculated as *views* + *reactions* + *n_forwards*. Because this metric includes views as a component, subsequent analyses involving engagement, particularly correlations with views, may reflect structural dependencies rather than independent relationships. Accordingly, engagement is interpreted as a composite exposure-dominated indicator rather than a purely interaction-based measure.
  - *text_length*: Character count of *text_content*
  - *month*: Monthly period extraction
  - *time_diff*: Inter-message latency in seconds, computed per channel using grouped temporal differencing
  - *length_bin*: Quartile-based categorization of text length (*short, medium, long, very_long*)

All data manipulation, statistical analysis, and visualization procedures were implemented in Python within a Google Colab cloud computing environment. The analytical workflow leveraged a suite of specialized libraries, such as urllib.parse, regex, and tldextract for robust URL extraction and domain normalization from unstructured text content.

Temporal dynamics of platform activity were assessed through systematic resampling and frequency aggregation techniques. Message timestamps were converted to datetime objects with error coercion, and activity counts were resampled at weekly and monthly frequencies to trace longitudinal trends in content production. User-level profiling was conducted by aggregating activity metrics per user_id, computing total message counts and channel diversity through the number of unique channels each user contributed to. Channel-level profiling ranked entities by total message volume and by the count of unique participating users, the latter serving as a proxy for channel audience size.

Weekly temporal evolution of active users and channels was assessed by resampling the dataset on the date column (freq='W') and aggregating unique counts of *user_id* and *channel_id* per period, visualized as time-series line plots. The distribution of content production and audience reach was characterized using empirical Cumulative Distribution Functions (CDF)[Hammerla et al. 2013] where message volume per channel was computed by grouping by *channel_id* and counting *message_id* entries, while channel audience size was estimated by counting unique *user_id* values per *channel_id*; both distributions were plotted with logarithmic x-axis scaling to effectively visualize heavy-tailed patterns and heterogeneity across the network.

The forwarding rate per channel was computed as the mean proportion of forwarded messages multiplied by one hundred. Content attributes were systematically quantified through media type distribution analysis, and external domain extraction. The proportional representation of *media_type* categories was computed following imputation of missing values with a 'Text Only' label, and visualized to assess content format preferences. URLs embedded within *text_content* were extracted using regular expressions matching the pattern https, and domains were normalized and parsed using *tldextract* python library to isolate second-level domains and top-level domains while filtering common subdomain prefixes such as *'www.'*; domain-level impact was assessed by aggregating total URL occurrences, unique URL counts, and distinct user contributions per domain.

Engagement and virality were evaluated through composite metric construction, distributional analysis, and normalization techniques. A composite engagement score was derived by summing *views, reactions, and n_forwards* for each message, with missing numerical values imputed as zero to preserve distributional integrity. Channel-level virality was assessed by normalizing aggregate engagement by channel audience size, computed as *engagement_per_user = engagement / num_users*, where *num_users* represents the count of unique *user_id* values per channel, enabling comparison of diffusion efficiency independent of raw audience scale.

Multivariate relationships among key metrics were quantified using Spearman correlation coefficients and stratified group aggregations. A global correlation matrix was computed across the variables *views, reactions, n_forwards, engagement, and text_length* to identify monotonic associations at the dataset level. Channel-specific correlation analysis was performed for channels with sufficient activity ($n > 30$ messages), calculating the correlation between views and n_forwards to identify platforms where visibility strongly predicts sharing behavior. External domain impact was ranked by mean aggregate engagement (*views* + *reactions* + *n_forwards*) to identify which linked sources drove the highest audience interaction.

The analysis adopts a descriptive statistical framework based on aggregation and correlation analysis. No inferential statistical tests (e.g., hypothesis testing, confidence intervals) or causal models are applied. As such, reported relationships represent observed associations within the dataset and should not be interpreted as statistically validated effects or causal mechanisms.

4. Results

Figure 1 presents three key aspects of Telegram activity in the dataset: the temporal evolution of weekly message volume, showing how user and channel participation fluctuates over time and reflects platform dynamics. Although the dataset spans from 2020 to 2025, the vast majority of message volume and engagement is concentrated between 2021 and 2023, coinciding with the peak of the COVID-19 immunization campaign in Brazil. Consequently, the results primarily reflect the dynamics of this critical public health period; the cumulative distribution of messages per channel, which reveals a highly skewed pattern where a small number of channels produce most of the content; and the cumulative distribution of users per channel, demonstrating that a few channels attract the majority of users, while most channels have limited reach. Together, these plots highlight a platform marked by concentrated content production and audience engagement, with a small set of influential channels driving most activity and visibility.

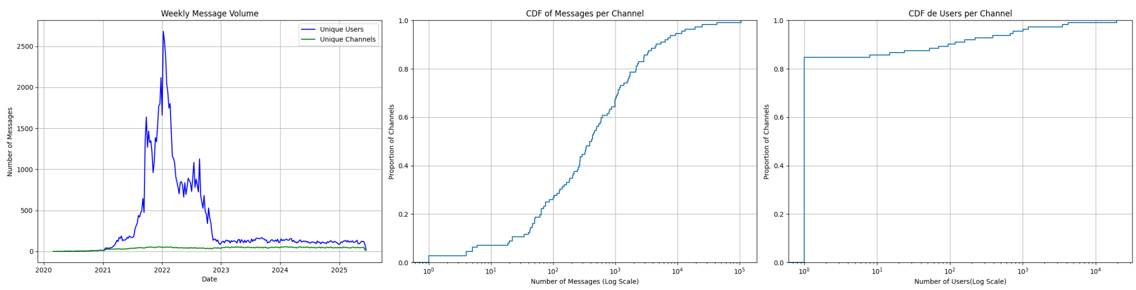

Figure 1. Temporal Evolution and Distribution of Telegram Messages and Users.

Table 1 presents correlation matrix between engagement metrics. Because engagement includes views, correlations involving this metric should be interpreted with caution, as they may reflect structural dependencies rather than an independent behavioral driver. Therefore, the relationship between views and engagement should be interpreted as evidence of an exposure-dominated metric rather than an independent causal effect of views on engagement.

A strong correlation between views and forwards indicates that higher exposure is associated with increased redistribution, suggesting a pattern where visibility and sharing are statistically coupled. In contrast, reactions exhibit weak correlations with forwarding and only moderate association with engagement, suggesting that observable interaction signals (e.g., reactions) are less strongly associated with redistribution metrics; however, this does not imply that user engagement is behaviorally irrelevant, only that it is not strongly captured by the available indicators. Text length exhibits negligible correlation with views and forwards, indicating it does not strongly influence visibility or sharing, although it shows a modest association with engagement, implying some contribution to interaction. Overall, diffusion dynamics suggesting patterns consistent with efficient redistribution dynamics, as reflected in aggregate forwarding and visibility metrics.

The relationship between content production and engagement (Table 2) is highly heterogeneous. Channels with a large number of messages do not consistently achieve

Table 1. Global correlation matrix of observed engagement metrics

	Views	Reactions	Forwards	Engagement	Text Length
Views	1.000	0.251	0.888	0.999	-0.017
Reactions	0.251	1.000	0.207	0.496	0.160
Forwards	0.888	0.207	1.000	0.890	0.004
Engagement	0.999	0.496	0.890	1.000	0.331
Text Length	-0.017	0.160	0.004	0.331	1.000

higher engagement. For example, Channel *71744f7659b8d983ad47* (388 messages) exhibits the highest average engagement (218,121). Channel *fd934b28f14024024031* (4 messages) achieves substantial engagement (32,559). This disparity indicates that engagement is not linearly proportional to message volume. Additionally, channels with high activity (e.g., *a69c7ef2333eb18639e6* with 430 messages) show comparatively lower engagement levels, further supporting the absence of a direct linear relationship. Larger channels do not consistently outperform smaller ones proportionally.

Table 2. Aggregate activity and mean metrics for high-volume channels

Channel ID	Messages	Views	Reactions	Forwards	Engagement
71744f...	388	216581.50	1224.76	314.95	218121.21
d2121a...	183	46560.65	359.73	855.80	47776.19
74f324...	71	35177.01	178.61	686.18	36041.80
27401c...	327	33367.66	228.66	301.05	33897.36
a69c7ef...	430	20062.15	41.72	253.02	20356.89

User-level analysis (Table 3) demonstrates that message production is not consistently associated with higher observed engagement outcomes, indicating a decoupling between activity levels and visibility-based metrics. The most active user (38,051 messages) shows relatively low average views (1,121), while less active users achieve substantially higher average views (4,268) and engagement. Additionally, users participating in multiple channels do not tend to exhibit higher engagement metrics compared to those restricted to a single channel.

The forwarding mechanism plays a central role in content propagation (Table 4). Messages with forwarding activity are associated with significantly higher mean engagement (3,753) compared to non-forwarded messages (1,407). Forwarded messages exhibit higher reach—with 3,682.19 average views compared to 1,382.86 for original posts while maintaining lower average reactions (5.35 vs. 9.08). This indicates that while forwarding effectively amplifies visibility, it does not necessarily drive higher levels of direct user interaction. The relationship between text length and engagement (Table 5) confirms a monotonic increase in performance, with average engagement rising from 693.79 for short texts to 3,062.24 for very long ones. While this may contrast with the modest global correlation reported in Table 1 (r=0.331), this discrepancy reflects the distinction

Table 3. Aggregate activity and mean metrics for the most active users

User ID	Messages	Channels	Avg Views	Avg Forwards	Engagement
4a74a4...	38051	2	1121.66	17.34	1144.42
2e9e49...	9242	1	105.39	1.88	107.52
59f557...	8318	2	2621.29	36.51	2662.62
319ff8...	6982	1	2727.63	31.53	2776.64
59b50c...	4353	2	4268.47	79.57	4348.04
ece83e...	3800	1	1605.36	21.93	1641.97

between high individual message variance and aggregate structural trends. At the level of individual posts, text length is a weak deterministic predictor of engagement because engagement is dominated by views, which show a negligible correlation with message size ($r = -0.017$). However, the binned analysis in Table 5 reveals a macroscopic pattern where greater narrative depth in longer messages tends to accumulate more intensive interaction suggesting that once visibility is achieved, longer content is more likely to be redistributed within the platform's broadcast architecture.

The domain analysis (Table 6) shows that high-engagement messages are associated with a diverse set of domains, including cos.tv video platform (high forwarding activity), brazilian military domains (*.mil.br*) with high views, and various media platforms (e.g., *tiktok.com)*. The dataset also reveals a concentration of shared links in a limited number of domains (Table 7), such as *Telegran (t.me), YouTube (youtube.com), and Rumble (rumble.com)*. Channel-specific correlation analysis (Table 8) reveals extremely high correlations between views and forwards, indicating a strong coupling between exposure and redistribution within individual channels. This pattern is consistent across multiple channels, suggesting a stable structural relationship between visibility and sharing behavior.

Table 4. Mean metrics for forwarded versus non-forwarded messages

Forwarded	Views	Reactions	Engagement
No	1382.86	9.08	1407.82
Yes	3682.19	5.35	3753.78

Table 5. Distribution of mean metrics across text length quartiles

Length	Views	Forwards	Engagement
Short	680.65	9.61	693.78
Medium	1480.51	19.66	1506.60
Long	2151.01	26.73	2186.89
Very Long	3001.34	46.55	3062.24

Table 6. External domains associated with high mean engagement

Domain	Views	Forwards	Engagement	Total URLs
fab.mil.br	193575.00	185.00	193774.00	1
eb.mil.br	171290.33	142.67	171445.00	3
cos.tv	80350.62	2159.89	70691.30	335
tiktok.com	66016.87	88.97	57241.20	35

Table 7. Aggregate frequency of external domains shared in the dataset

Domain	Total URLs	Unique URLs	Users
t.me	22878	2611	836
nih.gov	4759	823	161
youtube.com	3702	1191	683
twitter.com	2626	1057	278

5. Limitations

This study offers a large-scale quantitative view of vaccine discourse on Telegram but is subject to several important limitations. Its reliance on aggregate proxy metrics, particularly a composite engagement measure that includes views, introduces structural bias and inflates correlations, limiting interpretability. To address the structural redundancy observed between views and the composite engagement metric, future research should implement an active engagement indicator, calculated exclusively as the sum of reactions and forwards. The analysis is purely correlational, lacking causal inference, statistical testing, and control for confounding factors, which constrains the robustness of its conclusions. Additionally, diffusion is inferred from aggregate indicators rather than modeled through explicit network structures or cascades, and no semantic or contextual analysis of message content is performed. Methodological weaknesses such as absence of temporal causal modeling further restrict insights. Finally, dataset and platform-specific biases limit generalizability, meaning the findings should be interpreted as descriptive and exploratory rather than causal or predictive.

6. Conclusion

The findings of this study is consistent with prior evidence suggesting Telegram's structural capacity for rapid and scalable propagation. It supports that engagement in Brazil's vaccine-related Telegram discourse is strongly associated with exposure, reflecting both platform dynamics and metric construction. The minimal role of reactions, highlights a broadcast-oriented ecosystem characterized by passive consumption. Forwarding emerges as the central amplification mechanism, significantly boosting reach independent of content origin and confirming Telegram's role as a redistribution-driven platform, consistent with evidence that mass-messaging systems enable efficient misinformation propagation.

Furthermore, the absence of a linear relationship between message volume and

Table 8. Correlation between visibility and redistribution for selected channels

Channel ID	Correlation (Views vs Forwards)
a90cc9...	0.994
fcb493...	0.958
409bd1...	0.955
e2a748...	0.942

influence, combined with heterogeneous user-level metrics, indicates that impact depends more on circulation patterns than on posting frequency. Longer messages correlate with increased engagement and redistribution, suggesting that narrative richness enhances diffusion once exposure is achieved. Domain-level analysis reveals a hybrid media ecosystem with strong cross-platform integration, particularly involving YouTube and Twitter, aligning with post-moderation migration patterns. The consistently high correlation between views and forwards across channels underscores Telegram's structural capacity for rapid, scalable information propagation.

Collectively, these results reveal an exposure-driven system where platform architecture and forwarding dynamics are closely linked to observed patterns of information spread. Future research should incorporate longitudinal analyses to track how these structural patterns evolve across different phases of public discourse. Integrating semantic content analysis or narrative classification would further clarify which message frames drive redistribution. Additionally, network-based modeling of user interactions could uncover structural roles and key dissemination pathways, while expanded datasets and richer engagement metrics would enhance the robustness and generalizability of findings across platforms and contexts.

References

Alvisi, L., Tardelli, S., and Tesconi, M. (2024). Unraveling the italian and english telegram conspiracy spheres through message forwarding. In *International Conference on Advances in Social Networks Analysis and Mining*, pages 204–213. Springer.

Amaral, A. d. R., Jung, A.-K., Braun, L.-M., and Blanco, B. (2022). Narratives of anti-vaccination movements in the german and brazilian twittersphere: a grounded theory approach. *Media and Communication*, 10(2):144–156.

Araújo, D. A. B. S., Correia, L. L., Lima, P. L. G. d. S. B., Vasconcelos, S. C., Farías-Antúnez, S., Gomes, Y. V. C., Nogueira, D. L., Castro, M. C., and Machado, M. M. T. (2024). Coverage and determinants of childhood vaccination during the covid-19 pandemic in fortaleza, northeastern brazil: a longitudinal analysis. *Cadernos de Saúde Pública*, 40:e00074723.

Barkhad, A., de Almeida Santos, G., Roberto Costa Campos, S., Maria Almeida Braz, L., Picinini Freitas, L., Zinszer, K., de Souza, R., Waldron, I., Loeb, M., Luna, E., et al. (2026). Knowledge, attitudes, practices and perceptions of the eco-bio-social determinants of dengue transmission in são paulo, brazil: A mixed-methods study. *Tropical Medicine & International Health*, 31(1):58–79.

Cardenuto, J. P., Monari, A. C., Lopes, M., Filho, L. L., and Rocha, A. (2026). Brazilian social media anti-vaccine information disorder dataset - telegram (2020-2025). *arXiv.org*.

Cavalini, A., Malini, F., Gouveia, F., and Comarela, G. (2023). Politics and disinformation: Analyzing the use of telegram's information disorder network in brazil for political mobilization. *First Monday*.

Cugler de Moraes Silva, E. (2024). Antivax and off-label medication communities on brazilian Telegram: between esotericism as a gateway and the monetization of false miraculous cures. *arXiv e-prints*, page arXiv:2408.15308.

Díaz, E. P. and Ricaurte, P. (2024). Doctors for the truth: Latin american antivaccine oppositional cultures on telegram. *Big Data & Society*, 11(4):20539517241306359.

Domingues, C. M. A. S., Maranhão, A. G. K., Teixeira, A. M., Fantinato, F. F., and Domingues, R. A. (2020). The brazilian national immunization program: 46 years of achievements and challenges. *Cadernos de saude publica*, 36:e00222919.

Figueiredo Vasconcelos, F., Gomes Ferreira, C. H., Magno de Sousa, A., and Almeida, J. (2025). The role of curiosity in information dissemination on telegram: a case study from the 2022 brazilian elections. In *Proceedings of the 17th ACM Web Science Conference 2025*, pages 425–436.

Galhardi, C. P., Freire, N. P., Fagundes, M. C. M., Minayo, M. C. d. S., and Cunha, I. C. K. O. (2022). Fake news and vaccine hesitancy in the covid-19 pandemic in brazil. *Ciência & Saúde Coletiva*, 27:1849–1858.

Hammerla, N. Y., Kirkham, R., Andras, P., and Ploetz, T. (2013). On preserving statistical characteristics of accelerometry data using their empirical cumulative distribution. In *Proceedings of the 2013 international symposium on wearable computers*, pages 65–68.

Júnior, M., Melo, P., da Silva, A. P. C., Benevenuto, F., and Almeida, J. (2021). Towards understanding the use of telegram by political groups in brazil. In *Proceedings of the Brazilian Symposium on Multimedia and the Web*, pages 237–244.

Khadafi, R., Nurmandi, A., Qodir, Z., and Misran (2022). Hashtag as a new weapon to resist the covid-19 vaccination policy: a qualitative study of the anti-vaccine movement in brazil, usa, and indonesia. *Human Vaccines & Immunotherapeutics*, 18(1):2042135.

Massarani, L., Leal, T., and Waltz, I. (2020). The debate on vaccines in social networks: an exploratory analysis of links with the heaviest traffic. *Cadernos de Saúde Pública*, 36:e00148319.

Tokojima Machado, D. F., de Siqueira, A. F., and Gitahy, L. (2020). Natural stings: Selling distrust about vaccines on brazilian youtube. *Frontiers in Communication*, 5:577941.

Tucci, G. (2023). Identification of opinion leaders in a telegram network of forwarded messages. *arXiv preprint arXiv:2312.16528*.

Wilson, S. L. and Wiysonge, C. (2020). Social media and vaccine hesitancy. *BMJ global health*, 5(10).