Vaccine 40 (2022) 2209–2214

Contents lists available at ScienceDirect

Vaccine

journal homepage: www.elsevier.com/locate/vaccine

The impact of Facebook's vaccine misinformation policy on user endorsements of vaccine content: An interrupted time series analysis

Jiayan Gu [a], Avi Dor [b,c], Kun Li [b], David A. Broniatowski [d,e], Megan Hatheway [a], Lailah Fritz [f], Lorien C. Abroms [a,e,*]

[a] Department of Prevention and Community Health, Milken Institute School of Public Health, The George Washington University, Washington, D.C., United States
[b] Department of Health Policy and Management, Milken Institute School of Public Health, The George Washington University, Washington, D.C., United States
[c] National Bureau of Economic Research, Cambridge, MA., United States
[d] Department of Engineering Management and Systems Engineering, The George Washington University, Washington, D.C., United States
[e] Institute for Data, Democracy & Politics, The George Washington University, D.C., United States
[f] Department of Psychology, University of Michigan, Ann Arbor, MI, United States

A R T I C L E I N F O

Article history:
Received 19 July 2021
Received in revised form 24 January 2022
Accepted 15 February 2022
Available online 1 March 2022

Keywords:
Vaccinations
Anti-vaccine
Social media
Misinformation

A B S T R A C T

Objectives: To evaluate the impact of Facebook's vaccine misinformation policy in March 2019 on user endorsements of vaccine content on its platform.
Methods: We identified 172 anti- and pro-vaccine Facebook Pages and collected posts from these Pages six months before and after the policy. Using interrupted time series regression models, we evaluated the policy impact on user endorsements (i.e., likes) of anti- and pro-vaccine posts on Facebook.
Results: The number of likes for posts on anti-vaccine Pages had decreased after the policy implementation (policy = 153.2, p < 0.05; policy*day = –0.838, p < 0.05; marginal effect at the mean = –22.74, p < 0.01; marginal effect at the median = –24.56, p < 0.01). When the number of subscribers was considered, the policy effect on the number of likes for anti-vaccine posts was much smaller, but still statistically significant (policy = 4.849, p < 0.05; policy*day = –0.027, p < 0.05; marginal effect at the mean = –0.742, p < 0.01; marginal effect at the median = –0.800, p < 0.01). There was no policy effect observed for posts on pro-vaccine Pages.
Conclusions: Our analysis suggested that Facebook's March 2019 vaccine misinformation policy moderately impacted the number of endorsements of anti-vaccine content on its platform. Social media companies can take measures to limit the popularity of anti-vaccine content by reducing their reach and visibility. Future research efforts should focus on evaluating additional policies and examining policies across platforms.

1. Introduction

Vaccines are one of the most safe, consequential, and efficient public health interventions to reduce mortality and morbidity associated with vaccine-preventable diseases in the modern medicine [1]. Yet, the anti-vaccination movement, fueled by a mass of unsubstantiated, misleading, and manipulated messages, has reemerged in recent years, undermining decades of public health efforts in the prevention of measles and other infectious diseases [2]. It is of concern that in the face of the COVID-19 pandemic, misinformation and disinformation regarding SARS-CoV-2 vaccines are having a resurgence on the Internet including social media [3], inducing a decline in intention to vaccinate [4]. The spread of misleading vaccine claims is detrimental to addressing the COVID-19 pandemic, as well as future public health outbreaks [5]. The World Health Organization (WHO) has listed vaccine hesitancy, defined as a delay in acceptance or refusal of vaccines despite availability of vaccination services, as one of the top ten major threats to global health since 2019 [6,7].

Current studies using computational and analytic techniques have documented anti-vaccine sentiment in posts across different social media platforms including Twitter [8,9,10], YouTube [11], Pinterest [12,13], Instagram [14,15], and Facebook [16–19]. Common anti-vaccine themes include safety and effectiveness concerns

* Corresponding author at: Milken Institute School of Public Health, 950 New Hampshire Avenue NW, Washington, D.C., 20052, United States.
E-mail address: lorien@gwu.edu (L.C. Abroms).

https://doi.org/10.1016/j.vaccine.2022.02.062
0264-410X/

of vaccination, promotion of alternative immunization, civil liberties, conspiracy theories, and morality or ideological arguments [9,18,20]. Such content often exists in information silos or echo chambers with few challenges from evidence-based sources [21]. A system-level analysis of nearly 100 million Facebook users has suggested that anti-vaccine clusters, although numerically smaller than pro-vaccine clusters in overall size, are more highly entangled with undecided clusters compared to pro-vaccine clusters and have the potential to expand in the next decade [17].

Historically, social media companies have resisted addressing misinformation on their platforms. However, as the spread of misleading anti-vaccine information has been linked to systematic, organized, and manipulative efforts [8], and in response to a groundswell of public pressure following several measles outbreaks in 2019, new policies have been introduced by platforms to combat anti-vaccine information. On March 7, 2019, Facebook announced its first attempt to address vaccine misinformation by implementing the following new policies: 1) reducing the ranking of anti-vaccine groups and Pages in News Feed and Search and not including them in recommendations; 2) rejecting ads that include vaccine misinformation and removing targeting advertising options related to vaccine misinformation; and 3) connecting people through banners to more authoritative vaccine information (e.g. CDC) when they come across vaccine misinformation [22]. These policies aimed to reduce the reach and visibility of vaccine misinformation without directly removing anti-vaccine content. Since this initial announcement, and with the outbreak of COVID-19, Facebook announced more aggressive content policies including its intention to remove false claims related to vaccines in December 2020 [23].

While these public announcements are consistent with social media companies' declared responsibility of protecting public health, the effectiveness of these policies on improving the vaccine information ecosystem on social media remains unclear. The goal of the present study was to model the impact of Facebook's initial vaccine misinformation policy implemented in March 2019 on the endorsements of posts on anti-vaccine Pages in comparison with posts on pro-vaccine Pages. Facebook Pages are public profiles specifically created for businesses and non-profits, brands, celebrities, and causes such as vaccination [24]. Posts on Pages come from account owners and can include status updates, links, events, photos, and videos [24]. When someone likes or follows a Page, they will start seeing updates from that Page in their News Feed, with posts with a greater number of likes, shares and comments appearing higher up in the News Feed [25]. We chose to focus on Pages because, unlike groups, every Facebook user can access and interact with content on Pages without joining a public group or requesting to join a private group.

The March 2019 policy would presumably affect the endorsements of anti-vaccine content in the following ways: following the introduction of the policy, users who subscribed to anti-vaccine Pages would be less likely to see anti-vaccine posts in their News Feed; and users who searched anti-vaccine keywords would be less likely to find anti-vaccine Pages in their search results. In addition, if users visited an anti-vaccine Page, banners on the top would direct them to authoritative information on vaccines (e.g., CDC), making Pages less likely to be subscribed to or liked by new users (for an example see Appendix A, available as a supplement). Because these changes should lower the reach and visibility of anti-vaccine content on Facebook and, as such, make users less likely to endorse (e.g., "like") anti-vaccine posts, we sought to test the hypotheses that Facebook's 2019 policy had reduced the number of likes of posts on anti-vaccine Pages, but had little or no effect on posts on pro-vaccine Pages.

2. Methods

The data collection for the analysis was conducted between September 6 and November 23, 2020 using CrowdTangle, a public insights tool owned and operated by Facebook [26]. Our data collection and coding consisted of 3 steps. First, we extracted all vaccine posts between 2017 and 2019 to identify Facebook Pages frequently posting vaccine content. Second, we manually coded Pages as either pro- or anti-vaccine based on their profile information (i.e., account name, description, and banner information) or 10 randomly sampled posts if the account information was not clear. Third, we retrieved all posts published by eligible Pages six months before and after the policy and applied an interrupted time series analysis to model the policy effect on the number of likes of posts on pro- and anti-vaccine Pages separately.

2.1. Data collection

We first used a keyword search approach to collect public posts that contained vaccine-related keywords (see Appendix B, available as a supplement) between January 1, 2017 and December 31, 2019 through CrowdTangle on September 6 and 7, 2020. The data collection was conducted in the United States without any restrictions on the geography. The query resulted in a total of 947,778 posts published by 113,838 Facebook Pages after removing duplicates. We then calculated the number of vaccine-related posts each Page had published in each month. A Page was included for further analysis if it met either of the following criterion: (1) the account name had any of the substrings of "vaccin" or "vax" or "immune"; or (2) the account had published at least 3 vaccine-related posts in >80% of months starting from its first non-zero month. A total of 402 Pages were included for the next stage of coding.

2.2. Coding of Pages

To determine the overall view of each Page towards vaccines, a codebook (see Appendix C, available as a supplement) was developed on the basis of previous work [8,18,20]. Subcategories under anti-vaccine views included unsafety and ineffectiveness, alternative immunization, civil liberties, conspiracy, censorship and media cover-up, marketing and advertising, and morality, religion and ideology [8,18,20]. Subcategories under pro-vaccine views included advocacy and education, misinformation refutation, anti-vaccine exposure, sarcasm, and official accounts of vaccine organizations [8,18,20]. Two trained annotators (J.G and L.F) independently coded each English language Page based on the account name, banner information, and description, with non-English Pages removed (n = 10). A Page would be coded as pro- or anti-vaccine if its profile information fell into any subcategory as indicated in the codebook. Otherwise, a Page would be coded as undecided if the profile's position was unclear based on the available information or excluded if it did not focus on vaccines. The two annotators achieved 93.77% agreement (Cohen's κ = 0.905; SE = 0.020, 95% confidence interval [CI] = 0.866, 0.944), suggesting almost perfect agreement [27]. Discrepancies were reconciled by the third reviewer (L.A). For Pages coded as undecided (n = 13), a random sample of 10 vaccine-related posts were selected and manually coded by two annotators (J.G and M.H) using the same codebook to decide the Page's position. The two annotators achieved 97.63% agreement (Cohen's κ = 0.956; 95% confidence interval [CI] = 0.913, 0.998) in this round and discrepancies were reconciled by a third reviewer (L.A). There were 220 Pages removed after this step because they did not focus on vaccines (e.g., chemtrail conspiracy, personal page, news media) (n = 156), were non-relevant (e.g., job hiring, buy & sell) (n = 49), or were no longer available (n = 15).

The final list consisted of 172 Pages, including 66 pro-vaccine and 106 anti-vaccine Pages (for the full list see Appendix D, available as a supplement). Subsequently, we retrieved all posts on anti-vaccine Pages (n = 63,145) and pro-vaccine Pages (n = 27,405) between September 1, 2018 and September 28, 2019 from CrowdTangle on November 23, 2020, right before Facebook announced that they would start removing anti-vaccine content. The metadata included each post's creation time, URL, interaction metrics, account information, and subscriber count at the time of data retrieval.

2.3. Statistical analysis

To determine the effect of Facebook's vaccine misinformation policy from March 7, 2019 on the endorsements of posts on anti- and pro-vaccine Pages, we choose to select the number of likes of a post as a proxy for endorsements in our models because likes are the lowest effort and the most certain method of indicating support for a post and affect the likelihood that others will see the post in their News Feed. We implemented an interrupted time series design with a series of segmented regression models to model the impact of policy on the changes in the number of likes of anti- and pro-vaccine posts respectively, controlling for Page fixed effects and month fixed effects (for statistical details see Appendix E). The first set of linear regression models estimated the change in the number of like counts for posts between periods before and after the policy was implemented. However, this set of models did not adequately take into consideration the heterogenous dispersion of like counts across different Pages. Thus, we conducted a second set of segmented regression models where the outcome was the number of like counts per 1000 subscribers. By normalizing the outcome, we ruled out the possibility that posts by Pages with larger number of subscribers had extremely large like counts away from the sample mean. We applied the number of subscribers as a scalar in the dependent variable instead of controlling for it in the model because the Page fixed effects and the number of subscribers were highly correlated. Standard errors were clustered at page-level, which allowed residuals of posts within the same Page to have non-zero correlations in estimation, to account for both heteroskedasticity and autocorrelation [28,29].

While our linear models performed well, we were also concerned with the heavy-tailed and positively skewed distribution of like counts. As a sensitivity analysis, we conducted interrupted time series analyses for pro- and anti-vaccine Pages that allowed trends to vary between the pre- and post-policy periods. To model the outcome of like counts, we ran negative binomial regressions, as appropriate for discretely measured counts. For the outcome of like counts adjusted for the number of subscribers (i.e., number of like counts per 1000 subscribers), we estimated Poisson regressions which allowed for more granular distributions [30,31,32].

Finally, in all sets of models, we also estimated the marginal effect of policies at both mean and median time points. All analyses were conducted using STATA version 14. Given that the findings were generally consistent across all models, we report the results from the nonlinear count data models separately in an Appendix.

3. Results

Tables 1 presents the characteristics for pro- and anti-vaccine Pages and posts. There were almost twice as many anti-vaccine Pages compared with pro-vaccine Pages on Facebook, and the number of posts published by anti-vaccine Pages was 2.3 times more than pro-vaccine Posts over the observation time. At the time of data retrieval, anti-vaccine Pages had an average of 26,745 subscribers with the smallest Page of 78 subscribers and the largest Page of 233,892 subscribers. In contrast, pro-vaccine Pages had an average of 39,694 subscribers with the smallest Page of 74 subscribers and the largest Page of 467,485. The mean number of likes for an anti-vaccine post was lower than for a pro-vaccine post, but after scaling by the number of subscribers, anti- and pro-vaccine posts had comparable mean likes.

The results of segmented linear regression models are shown in Table 2. In the first set of linear regression models, which evaluated changes in the number of likes for posts before and after the policy, there was a statistically significant impact for anti-vaccine posts (policy = 153.2, $p < 0.05$; policy*day = −0.838, $p < 0.05$) after the policy was announced. The marginal effects indicate that the number of likes of anti-vaccine posts was reduced after the policy (−22.74 likes at the mean, $p < 0.01$; −24.56 likes at the median, $p < 0.01$). However, no significant policy effect was observed for pro-vaccine posts (policy = 134.3, $p > 0.1$; policy*day = −0.889, $p > 0.1$). The scatter plots of like counts for anti- and pro-vaccine with corresponding trend lines are shown in Fig. 1.

In the second set of linear regression models when the number of subscribers was considered, anti-vaccine posts still had a statistically significant, but much smaller, change in the number of likes per 1000 subscribers (policy = 4.849, $p < 0.05$; policy*day = −0.027, $p < 0.05$). The marginal effects of the policy on the number of likes per 1000 subscribers for anti-vaccine posts were −0.742 likes ($p < 0.01$) at the mean and −0.800 likes ($p < 0.01$) at the median. Likewise, there was no policy effect observed for pro-vaccine posts (policy = −11.31, $p > 0.1$; policy*day = 0.051, $p > 0.01$).

The parallel results for the negative binomial regression and Poisson regression models were included in the Supplementary Material (see Appendix F). Given the strong similarity between the linear results and the count-based estimation, we opted to represent the linear regression models for expositional convenience.

Table 1
Summary of Characteristics of Pro- and Anti-vaccine Pages and Posts.

	Anti-vaccine			**Pro-vaccine**		
	Total	***Pre-policy***	***Post-policy***	***Total***	***Pre-policy***	***Post-policy***
No. of Pages	106			66		
Mean subscribers (SD)	26,745 (46,117)			39,694 (89,401)		
No. of posts	63,145	25,514	37,631	27,405	11,654	15,751
Mean Likes (SD)	65.01 (163.33)	64.15 (187.33)	65.6 (144.82)	83.72 (550.36)	87.07 (478.79)	81.25 (597.83)
Mean Likes per 1000 subscribers (SD)	2.91 (6.21)	2.22 (4.87)	3.38 (6.93)	3.09 (40.8)	2.72 (7.08)	3.37 (53.47)

Table 2
Segmented Regression Results for Change in the Number of Likes of Posts on Anti- and Pro-Vaccine Pages following Facebook's 2019 Vaccine Misinformation Policy.

Models	Linear			
	# of Likes		# of Likes/1000 subscribers	
	Anti	Pro	Anti	Pro
Coefficients				
policy	153.2**	134.3	4.849***	−11.31
	(62.86)	(184.2)	(1.166)	(10.68)
day	0.504**	0.634	0.0174***	−0.0178
	(0.172)	(0.473)	(0.00412)	(0.0257)
day*policy	−0.838**	−0.889	−0.0266***	0.0505
	(0.328)	(0.920)	(0.00646)	(0.0539)
Marginal effects of the policy at				
Mean day (=210)	−22.74***	−52.35**	−0.742**	−0.703
	(7.204)	(23.92)	(0.305)	(0.757)
Median day (=212)	−24.56***	−54.28**	−0.800**	−0.593
	(7.422)	(25.68)	(0.315)	(0.856)
Observations	63,142	27,404	63,142	27,404
R-squared	0.147	0.103	0.230	0.008
Page FE	YES	YES	YES	YES
Month FE	YES	YES	YES	YES

a. Linear regressions adjusted for page fixed effects and month fixed effects.
b. Cluster-robust standard errors in parentheses (*** p < 0.01, ** p < 0.05, * p < 0.1). Standard errors were clustered at page-level.

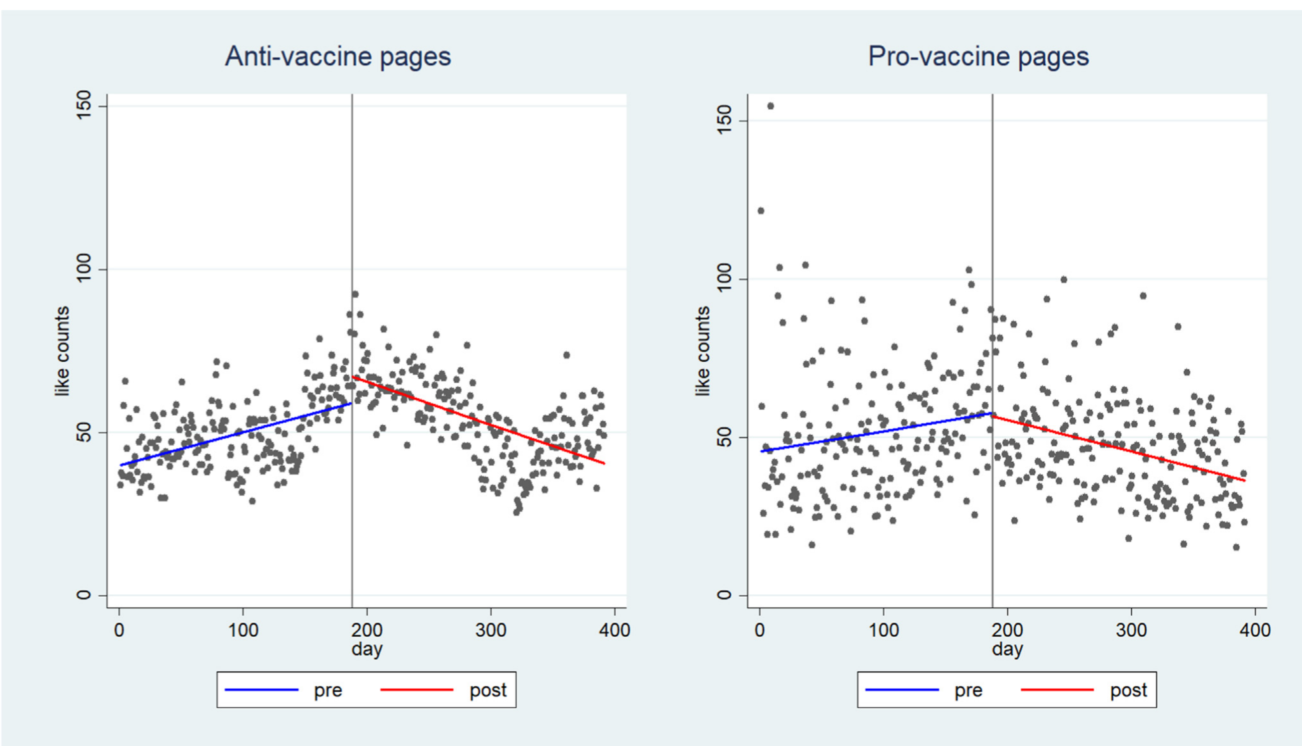

Fig. 1. Mean daily like counts before and after Facebook's vaccine misinformation policy (09/01/2018–09/28/2019). Observations with the top 1 percentile of like counts were trimmed and mean daily like counts were computed for visualization purposes only; Unadjusted linear trends were fitted for pre- and post-policy periods respectively.

4. Discussion

This study empirically evaluated the impact of a social media platform's vaccine misinformation policy based on a large number of vaccine-related posts. We found that the number of anti-vaccine Pages and posts were far greater than pro-vaccine Pages and posts on Facebook. Our findings suggest that Facebook's vaccine misinformation policy, announced on March 7, 2019, slightly, but statistically significantly, reduced the number of likes of posts per 1000 subscribers on anti-vaccine Pages during the following 6-month period, but had little or no effect on the number of likes per 1000 subscribers of pro-vaccine posts.

Anti-vaccine content has previously flourished on Facebook partly because of the platform's search results and recommendation algorithms that encouraged homogeneous and personalized information [33]. The reduction in the number of likes for anti-vaccine posts was therefore expected; Facebook's vaccine misinformation policy aimed to reduce the ranking of anti-vaccine Pages and their content in News Feed and Search so that subscribers were presumably less likely to come across the posts published by anti-vaccine Pages [34]. When such features are discouraged or disabled by the site, anti-vaccine posts will become less influential and harder to reach the audience. The banners directing users to external authoritative vaccine sources on anti-vaccine Pages may also deter new users from endorsing the Page when they come across it.

Although the effect of Facebook's vaccine misinformation policy was statistically significant, the effect size was relatively small after scaling for the number of subscribers and the volume of anti-vaccine posts remained steady after the policy. There was still a large amount of anti-vaccine content (i.e., 37,631 anti-vaccine posts) generated on Facebook after the policy. Thus, simply reducing the reach and visibility of anti-vaccine posts may have helped alleviate the rampant spread of anti-vaccine content somewhat, but may not be effective in qualitatively addressing the problem, especially among a loyal anti-vaccine audience. Facebook may be aware of the limited effect of their policy, as in December 2020, they introduced a more stringent policy that sought to not only downgrade misinformation, but remove it. As misinformation continues to be a problem on social media, future research should be conducted to evaluate the effectiveness of such additional policy measures.

A few limitations need to be noted. First, the exact date of policy implementation was not confirmed by Facebook. The policy date used in our analysis (March 7, 2019) was inferred from the publication date of Facebook's official announcement on these policies. However, our results were robust to changes in this inferred date upon sensitivity analysis (see Appendix G). Second, 15 Pages became unavailable during the process of data collection between September and November in 2020. For instance, Stop Mandatory Vaccination, which was one of the largest anti-vaccine Pages on Facebook was suspended by Facebook because of QAnon during the data collection and we were not able to include the data after its removal. Moreover, there were likely some additional Pages removed either by Facebook or by the Page's admins due to increased scrutiny on vaccine misinformation that were not detected by our methods. Although the data collection was completed before Facebook started systematically removing COVID-19 vaccine misinformation in December 2020, it was unclear exactly how many Pages were removed or suspended by the admins or due to other policies by Facebook that might have inadvertently affected vaccine-related content (such as the October 2020 QAnon ban) [35]. Once a Page was removed or suspended, we were no longer able to retrieve data on its posts from CrowdTangle, which may have led to an underestimation of the engagement of anti-vaccine content on the platform. Thus, our results apply primarily to Facebook Pages that are still extant. Third, we used the action of like as the only metric for measuring the endorsement of vaccine posts on the platform without considering other available metrics such as share and comment. We opted to be conservative about the outcome as it is less clear whether these other actions (e.g. share and comment) reflect endorsements or disagreements with a post on social media. However, we acknowledge it is one of our limitations that should be taken into consideration when interpreting the result.

It is undeniably positive to see that leading social media companies have started taking action to combat health misinformation. While platforms may have different features, currently social media companies have mainly adopted two types of approaches. The first type of approach is to reduce or block the availability of anti-vaccine content. For example, Instagram has banned hashtags associated with false claims such as #vaccinecauseaids [7]. The other type of approach is to provide alternative authoritative information to users when they come across potential misinformation. Facebook's initial vaccine misinformation policy in March 2019 was a combination of both approaches and is shown to have effectively reduced the number of likes of anti-vaccine posts on its platform by changing their recommendation algorithms, rejecting ads, and connecting users to authoritative sources; however, the effect size was quite small. Our results emphasize that social media companies are able to take an active role in addressing the spread of misleading claims by adjusting certain platform features [13]. However, the small effect size also highlights the difficulty for social media companies in improving the established health misinformation ecosystem and points to the need for consideration of other policies such as removing misinformation.

In summary, this study demonstrates the feasibility of empirically evaluating the effectiveness of vaccine misinformation policies on social media platforms and provides a framework for the analysis of public health related information policies. Future research efforts may expand to include more platforms and future policies in order to apply similar approaches to monitoring health-related misinformation on social media.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Guanghao Qi, PhD for related discussions.

Funding

This work was supported by the John S. and James L. Knight Foundation to the GWU Institute for Data, Democracy, and Politics. The funder played no role in the design and conduct of the study; collection, management, analysis, and interpretation of the data; preparation, review, or approval of the manuscript; and decision to submit the manuscript for publication.

Author Contributions

J. Gu, A. Dor, D.A. Broniatowski and L.C. Abroms conceived of the presented idea. J. Gu designed the study, collected, and analyzed the qualitative data and wrote the first draft of the article. K. Li and A. Dor designed and conducted the statistical analyses. M. Hatheway and L. Fritz assisted with the qualitative coding. All authors critically revised the manuscript.

Appendix A. Supplementary material

Supplementary data to this article can be found online at https://doi.org/10.1016/j.vaccine.2022.02.062.

References

[1] Ten Great Public Health Achievements — United States, 2001–2010. Accessed January 10, 202 https://www.cdc.gov/mmwr/preview/mmwrhtml/mm6019a5.htm.

[2] Dubé E, Vivion M, MacDonald NE. Vaccine hesitancy, vaccine refusal and the anti-vaccine movement: influence, impact and implications. Expert Rev Vacc 2015;14(1):99–117. https://doi.org/10.1586/14760584.2015.964212.

[3] Mian A, Khan S. Coronavirus: the spread of misinformation. BMC Med 2020;18 (1):89. https://doi.org/10.1186/s12916-020-01556-.

[4] Loomba S, Figueiredo A de, Piatek SJ, Graaf K de, Larson HJ. Measuring the Impact of Exposure to COVID-19 Vaccine Misinformation on Vaccine Intent in the UK and US. medRxiv. Published online October 26, 2020:2020.10.22.20217513. https://doi.org/10.1101/2020.10.22.20217513.

[5] Burki T. The online anti-vaccine movement in the age of COVID-19. The Lancet Digital Health 2020;2(10):e504–5. https://doi.org/10.1016/S2589-7500(20)30227-2.

[6] Stahl J-P, Cohen R, Denis F, Gaudelus J, Martinot A, Lery T, et al. The impact of the web and social networks on vaccination. New challenges and opportunities offered to fight against vaccine hesitancy. Med Mal Infect 2016;46(3):117–22. https://doi.org/10.1016/j.medmal.2016.02.002.
[7] Burki T. Vaccine misinformation and social media. The Lancet Digital Health 2019;1(6):e258–9. https://doi.org/10.1016/S2589-7500(19)30136-0.
[8] Broniatowski DA, Jamison AM, Qi S, AlKulaib L, Chen T, Benton A, et al. Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. Am J Public Health 2018;108(10):1378–84. https://doi.org/10.2105/AJPH.2018.304567.
[9] Jamison A, Broniatowski DA, Smith MC, Parikh KS, Malik A, Dredze M, et al. Adapting and extending a typology to identify vaccine misinformation on Twitter. Am J Public Health 2020;110(S3):S331–9. https://doi.org/10.2105/AJPH.2020.305940.
[10] Chen T, Dredze M. Vaccine images on Twitter: analysis of what images are shared. J Med Internet Res 2018;20(4):e130. https://doi.org/10.2196/jmir.8221.
[11] Tang L, Fujimoto K, Amith M, (Tuan), et al. "Down the Rabbit Hole" of vaccine misinformation on YouTube: network exposure study. J Med Internet Res 2021;23(1):e23262. https://doi.org/10.2196/23262.
[12] Guidry JPD, Carlyle K, Messner M, Jin Y. On pins and needles: how vaccines are portrayed on Pinterest. Vaccine 2015;33(39):5051–6. https://doi.org/10.1016/j.vaccine.2015.08.064.
[13] Guidry JPD, Vraga EK, Laestadius LI, Miller CA, Occa A, Nan X, et al. HPV vaccine searches on Pinterest: before and after Pinterest's actions to moderate content. Am J Public Health 2020;110(S3):S305–11. https://doi.org/10.2105/AJPH.2020.305827.
[14] Kearney MD, Selvan P, Hauer MK, Leader AE, Massey PM. Characterizing HPV vaccine sentiments and content on Instagram. Health Educ Behav 2019;46(2_suppl):37S–48S. https://doi.org/10.1177/1090198119859412.
[15] Basch CH, MacLean SA. A content analysis of HPV related posts on instagram. Human Vacc Immunotherapeut 2019;15(7-8):1476–8.
[16] Schmidt AL, Zollo F, Scala A, Betsch C, Quattrociocchi W. Polarization of the vaccination debate on Facebook. Vaccine 2018;36(25):3606–12. https://doi.org/10.1016/j.vaccine.2018.05.040.
[17] Johnson NF, Velásquez N, Restrepo NJ, Leahy R, Gabriel N, El Oud S, et al. The online competition between pro- and anti-vaccination views. Nature 2020;582(7811):230–3. https://doi.org/10.1038/s41586-020-2281-1.
[18] Hoffman BL, Felter EM, Chu K-H, Shensa A, Hermann C, Wolynn T, et al. It's not all about autism: The emerging landscape of anti-vaccination sentiment on Facebook. Vaccine 2019;37(16):2216–23. https://doi.org/10.1016/j.vaccine.2019.03.003.
[19] Broniatowski DA, Jamison AM, Johnson NF, Velasquez N, Leahy R, Restrepo NJ, et al. Facebook Pages, the "Disneyland" measles outbreak, and promotion of vaccine refusal as a civil right, 2009–2019. Am J Public Health 2020;110(S3):S312–8. https://doi.org/10.2105/AJPH.2020.305869.
[20] Kata A. A postmodern Pandora's box: anti-vaccination misinformation on the Internet. Vaccine 2010;28(7):1709–16. https://doi.org/10.1016/j.vaccine.2009.12.022.
[21] Chou W-Y-S, Oh A, Klein WMP. Addressing health-related misinformation on social media. JAMA 2018;320(23):2417–8. https://doi.org/10.1001/jama.2018.16865.
[22] Combatting Vaccine Misinformation. About Facebook. Published March 7, 2019. Accessed June 2, 2021. https://about.fb.com/news/2019/03/combatting-vaccine-misinformation/.
[23] Keeping People Safe and Informed About the Coronavirus. About Facebook. Published December 18, 2020. Accessed June 2, 2021. https://about.fb.com/news/2020/12/coronavirus/.
[24] What's the difference between a profile, Page and group on Facebook? | Facebook Help Center. Accessed June 2, 2021. https://www.facebook.com/help/337881706729661.
[25] News Feed ranking, powered by machine learning. Facebook Engineering. Published January 26, 2021. Accessed June 2, 2021. https://engineering.fb.com/2021/01/26/ml-applications/news-feed-ranking/.
[26] CrowdTangle Team. CrowdTangle. Facebook. Menlo Park, California, United States. List ID: [1442378, 1442380]. Accessed June 2, 2021. https://www.crowdtangle.com//.
[27] Landis JR, Koch GG. The measurement of observer agreement for categorical data. Biometrics 1977;33(1):159–74. https://doi.org/10.2307/2529310.
[28] Abadie A, Athey S, Imbens GW, Wooldridge J. When should you adjust standard errors for clustering? (No. w24003). 2017; National Bureau of Economic Research.
[29] Angrist JD, Pischke JS. Mostly harmless econometrics. Princeton University Press; 2008.
[30] Deb P, Trivedi PK. The structure of demand for health care: latent class versus two-part models. J Health Econ 2002;21(4):601–25. https://doi.org/10.1016/S0167-6296(02)00008-5.
[31] Cameron AC, Trivedi PK. Microeconometrics: methods and applications. Cambridge University Press; 2005.
[32] Silva JS, Tenreyro S. The log of gravity. Rev Econ Stat 2006;88(4):641–58. https://doi.org/10.1162/rest.88.4.641.
[33] Chiou L, Tucker C. Fake news and advertising on social media: a study of the anti-vaccination movement. National Bureau Econ Res 2018. https://doi.org/10.3386/w25223.
[34] Wardle C, Singerman E. Too little, too late: social media companies' failure to tackle vaccine misinformation poses a real threat. BMJ 2021;372:. https://doi.org/10.1136/bmj.n26n26.
[35] Facebook bans QAnon across its platforms. NBC News. Accessed June 2, 2021. https://www.nbcnews.com/tech/tech-news/facebook-bans-qanon-across-its-platforms-n1242339.