

# Winds of Change: Impact of COVID-19 on Vaccine-Related Opinions of Twitter Users

Soham Poddar<sup>\*1</sup>, Mainack Mondal<sup>1</sup>, Janardan Misra<sup>2</sup>, Niloy Ganguly<sup>1,3</sup>, Saptarshi Ghosh<sup>1</sup>

<sup>1</sup> Indian Institute of Technology, Kharagpur, India

<sup>2</sup> Accenture Labs, Bangalore, India

<sup>3</sup> Leibniz University of Hannover, Germany

\* sohampoddar@kgpian.iitkgp.ac.in

## Abstract

Administering COVID-19 vaccines at a societal scale has been deemed as the most appropriate way to defend against the COVID-19 pandemic. This global vaccination drive naturally fueled a possibility of Pro-Vaxxers and Anti-Vaxxers strongly expressing their supports and concerns regarding the vaccines on social media platforms. Understanding this on-line discourse is crucial for policy makers. This understanding is likely to impact the success of vaccination drives and might even impact the final outcome of our fight against the pandemic. The goal of this work is to improve this understanding using the lens of Twitter-discourse data. We first develop a classifier that categorizes users according to their vaccine-related stance with high precision (97%). Using this method we detect and investigate specific user-groups who posted about vaccines in pre-COVID and COVID times. Specifically, we identify distinct topics that these users talk about, and investigate how vaccine-related discourse has changed between pre-COVID times and COVID times. Finally, for the first time, we investigate the change of vaccine-related stances in Twitter users and shed light on potential reasons for such changes in stance.

## Introduction

The ongoing COVID-19 pandemic has affected close to 500M people till now (April 2022) and caused more than 6M deaths. Since December 2020 / January 2021, multiple pharmaceutical companies (e.g., AstraZeneca, Pfizer, Moderna) have put forward vaccines that are claimed to reduce the chance of COVID infection and fatality.<sup>1</sup> Naturally, governments across the world have been procuring and administering these vaccines to their citizens since then. However, administering COVID-19 vaccines also has a key societal angle. To make these vaccines effective against COVID-19 at-scale and to eradicate the disease, almost everyone in the society needs to consent to be vaccinated. In other words, only with near societal-scale vaccination, can we achieve the “herd immunity” and eradicate the disease<sup>2</sup>. As of April 2022, a large portion of the population has already been vaccinated, however a fraction of people have still been hesitant and resisted taking the vaccines, speculating various

reasons including that the development of these vaccines might be rushed and propelled by monetary and political reasons (Chou et al. 2020). Their hesitancy persists in spite of Governments giving various incentives towards vaccination.

There exists a long-drawn debate about vaccines since long before the onset of the COVID-19 pandemic. The debate is between *Anti-Vaxxers* who believe that vaccines do more harm than good, and *Pro-Vaxxers* who support and promote the benefits of vaccines. Both the groups are known to actively propagate their views over social media (Mitra et al. 2016; Gunaratne et al. 2019). It is important to know what the present attitude of the Anti-Vaxxer and Pro-Vaxxer communities is towards COVID-19 vaccines, and whether there has been any change in these attitudes in the pandemic times. While there have been a few surveys and manual analyses that have attempted to answer these questions (Funk and Tyson 2020; Bonnevie et al. 2020), to our knowledge, there has not been any prior attempt to devise automatic methods of answering these questions at scale using data from social media. It is important for policy makers to get answers to these questions at multiple stages of the pandemic, e.g., to figure out what the main concerns about vaccines are at different stages, so that they can take intervening steps accordingly.

To our knowledge, this study is the first that (i) devises automated methods over data collected from social media (Twitter) to answer such questions, and (ii) conducts a long-term (over 2 years) analysis of people’s attitudes towards vaccines, both individually and collectively. Specifically, we investigate two key questions – (1) What attitudes towards COVID-vaccines are expressed by Anti-Vaxxers and Pro-Vaxxers on Twitter? Especially, are the anti-vax opinions about COVID-19 vaccines similar to the traditional anti-vax arguments that were prevalent in pre-COVID times?, and (2) Has the vaccine-related stance of some users changed over time due to the pandemic? If so, why?

We make three important contributions in this work. First, we develop a dataset of COVID-vaccine related tweets labelled according to their vaccine stances, as detailed in later sections. Training a classifier on these tweets, we created a method to *categorize Twitter users according to their vaccine-related stance* as expressed through their tweets both in pre-COVID and COVID times, with 97% precision. We use this method to detect and investigate specific user-

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup><https://tinyurl.com/WHO-vaccine-advice>

<sup>2</sup><https://tinyurl.com/WHO-COVID-immunity>

groups, e.g., Anti-Vaxxers and Pro-Vaxxers.

Second, we use topic modeling to identify 12 COVID-vaccine related topics/themes from the tweets of Pro-Vaxxers and Anti-Vaxxers. We further analyze the relative presence of these topics in vaccine-related online discourse in different time periods – pre-COVID, COVID and COVID-vax time periods. Through this analysis, we find that some anti-vax themes such as *the vaccine development being rushed* and *concerns about vaccines being ineffective* have become more frequently discussed in COVID times (compared to the pre-COVID times).

Third, we show that there exists a set of users who have changed their prior vaccine-related stance in the COVID times (from Anti- to Pro-Vaxxers, or vice-versa). Understanding this set of users, whom we call **stance-changed users**, can be important for the authorities attempting to maximize the coverage of vaccines. Though some prior works have shown that people have changed their stances towards COVID-19 vaccines (Funk and Tyson 2020; Bonnevie et al. 2020), no prior work has systematically explored such users and *why* they changed their stance. In this work, we characterize different sub-groups of stance-changed users and explore *possible reasons of such stance-change*. To this end, we uncover correlations between the stance-change of a user and life-events (e.g., loss of a close relative) as well as change in stance of the users’ social contacts.

We believe that the insights obtained from this study will help authorities better understand how vaccine-related stances of users are changing in pandemic times, and will help them to formulate effective policies to increase vaccine adoption by the masses. We make our dataset and classifier public to facilitate further research on this topic.<sup>3</sup>

## Related Works

The dichotomy of opinions on vaccines between *Anti-vaxxers* and *Pro-vaxxers* on social media has been identified in many prior works (Mitra et al. 2016; Gunaratne et al. 2019). Both these groups have been seen to be quite strong in their stance, each believing the other to indulge in conspiracies (Yuan et al. 2019; Cossard et al. 2020). We briefly survey a few relevant works in this section.

**Vaccine-stance detection methods:** Several studies have attempted to classify social media posts (e.g., tweets) based on their stance towards vaccines, using hashtag counts (Gunaratne et al. 2019), machine learning models (Mitra et al. 2016; Yuan et al. 2019) and neural models (Müller et al. 2020; Coffas et al. 2021). Some works have also provided labeled datasets for this classification. Some studies have also tried to classify the stances of *users* based on percentages of posted tweets of different stances (Mitra et al. 2016) as well as by community detection algorithms (Yuan et al. 2019; Gunaratne et al. 2019).

**Social media discourse on COVID-19 vaccines:** With the onset of the COVID-19 pandemic, people have naturally resorted to sharing their emotions on social media (Lwin et al. 2020). Tough times give rise to several conspiracy theo-

ries and the spread of misinformation about the disease has risen (Cinelli et al. 2020). The pandemic has also led to decline of public trust on the healthcare system and a corresponding huge increase in hesitancy towards COVID-19 vaccines (Johnson et al. 2020; Bonnevie et al. 2020).

A few studies have tried to understand which demographics of people are more likely to be hesitant towards the vaccines (Malik et al. 2020; Thunstrom et al. 2020; Funk and Tyson 2020). Some studies have also tried to understand the reasons behind the rise of this hesitancy (Jamison et al. 2020; Bonnevie et al. 2020; Praveen et al. 2021), some major reasons being distrust towards pharmaceutical industries and the potential side effects (which we also observe in this study, as detailed in later sections).

**This work:** Though some prior works showed that people have changed their stances towards COVID-19 vaccines, these works relied on surveys (Funk and Tyson 2020) or manual coding (Bonnevie et al. 2020). To our knowledge, no automated methods have been developed for such analyses at scale. Also, it is unclear *why* different users have changed their stance, and what are the characteristics of such users. This work is the first attempt to bridge this gap where we study how and why vaccine-related opinions of users have changed between pre-COVID times and COVID times.

## Data Collection

In this work, we planned to study tweets posted since Pre-COVID times. Since the Twitter API only allows collecting tweets posted in the last ~ 7 days, we used the Twint library<sup>4</sup> to collect old tweets. The Twint library has been used by several prior works for collecting older tweets (Dutta et al. 2021; Punuru et al. 2020; Nuzhath et al. 2020). We used this library to collect a total of 15.7 million English tweets from 4 million distinct users over the duration January 2018 – March 2021. We consider this duration divided into three stages, as described below.

**(1) Pre-COVID period (Jan 2018 to Dec 2019):** This period is before the COVID pandemic, when people talked about vaccines in general. To get a good coverage of vaccine-related tweets, we used some generic keywords such as ‘vaccine’, ‘vaccination’, ‘vaxxer’, etc. along with an extensive set of Pro-vax and Anti-vax hashtags provided by Gunaratne et al. (2019). Using these keywords, we collected 2.9M English tweets posted in this period.

**(2) COVID period (2020), and (3) COVID-vax period (Jan-Mar 2021):** The *COVID period* signifies the year 2020, when COVID-19 was declared as a pandemic and various organizations started working on developing vaccines. The *COVID-vax period* signifies the first quarter of 2021 (January–March 2021) when COVID-19 vaccines started to be administered globally. We consider these two time-periods separately since people’s perceptions about COVID-19 vaccines may have varied between these time periods (before and after a vaccine was available).

For these periods, in addition to the keywords used for Pre-COVID period, we also fetched tweets using some

<sup>3</sup><https://github.com/sohampoddar26/covid-vax-stance>

<sup>4</sup><https://github.com/twintproject/twint>

generic keywords such as ‘corona vaccine’ and ‘covid vaccine’ and the names of popular vaccines and their manufacturers, such as ‘astra zeneca’, ‘pfizer’, ‘moderna’, etc. We collected 6.6M English tweets from the COVID period (Jan–Dec 2020) and 6.2M English tweets from the COVID-vax period (Jan–Mar 2021). Evidently, the social media discourse on vaccines has increased manifold in 2021 after the administration of COVID vaccines has started.

**Checking for bias in tweets collected by Twint:** Since we did not use the official Twitter API to collect tweets, we check if our collected data contains any bias. To this end, we collected the tweets posted during June 2021 using both the Twitter API (fetched in real-time) and the Twint library (fetched in August 2021) using the same set of search-query terms. While the Twitter API collected 4.8M distinct tweets, the Twint library collected 745K distinct tweets during the same period. Thus, Twint collects 15.5% of all tweets collected by the Twitter API. Note that, Twint does not collect retweets; so we have not considered the retweets collected by the Twitter API. Next, we compared the set of tweets collected by the two methods as follow:

- **Representativeness w.r.t. time:** On each of our observed days, the fraction of tweets collected by the Twint library is relatively uniform (close to 15% of what is collected via Twitter API). This gives us confidence that the set of tweets collected by the Twint library is temporally representative.

- **Representativeness w.r.t. tweet popularity:** We measure the popularity of a tweet by its *retweet-count* that is stated in the tweet object. Among the tweets collected by Twitter API, 93.0% had less than 10 retweets and 99.1% had less than 100 retweets. In comparison, among the tweets collected by Twint, 90.4% had less than 10 retweets and 98.8% had less than 100 retweets. The numbers are quite similar for both and thus, the set of tweets collected by Twint does not have much bias towards more popular tweets.

Thus, though we could fetch only 15% of all the posted tweets, the collected tweets do not have much apparent biases. Also we believe that the stance-detection methods used in this work (described in the next section) are robust towards small variations in data w.r.t time or tweet popularity, due to the conservative thresholds we use in the classifiers.

Since the Twitter API does not allow access to tweets older than  $\sim 7$  days, we needed to use some library like Twint to collect older tweets, as has been done by several prior longitudinal studies (Dutta et al. 2021; Punuru et al. 2020; Nuzhath et al. 2020). Also, many prior studies have been based on the 10% and even the 1% Twitter random sample (provided by Twitter). Hence, we believe that the mostly unbiased 15% tweet sample (that Twint gave us) is sufficient and appropriate for the present study.

## Identifying Users’ Stance towards Vaccines

This section discusses our approach to classify users based on their stance towards vaccines during the three time periods (pre-COVID, COVID, and COVID-vax). Since a user’s stance may have changed over time, we do *not* attempt to do this classification based on the profile attributes of the user. Instead, we first develop a tweet-classifier to predict

Dataset	#tweets	Anti-Vax	Pro-Vax	Neutral
<i>Muller</i>	18.5k	10.5%	40.3%	49.3%
<i>Cotfas</i>	2,792	28.3%	35.5%	36.2%
<i>Cotfas + our data (majority)</i>	4,398	24.6%	38.2%	37.2%
<i>Cotfas + our data (unanimous)</i>	3,577	26.0%	38.7%	35.3%

Table 1: Summary of different datasets used for training and testing of the classifiers, each of them having 3 class labels.

the stance of individual tweets, and then use this classifier to identify a user’s stance during a certain time period from the tweets posted during that period.

## Developing a Tweet Classifier

**Annotating tweets and other datasets:** There exist a few publicly available datasets that have labelled tweets as *Anti-Vax*, *Pro-Vax* or *Neutral* according to their stance towards vaccines. Müller et al. (2019) provided a large set of tweets about vaccines before the onset of the COVID-19 pandemic, while Cotfas et al. (2021) provided a dataset of tweets about COVID-19 vaccines collected during a month between November and December of 2020. We call these two datasets the **Muller** and **Cotfas** datasets respectively.<sup>5</sup>

Since the Cotfast dataset only consists of tweets from one particular month, it might not cover different types of tweets posted over the entire discourse about COVID-19 vaccines. Hence we wanted to extend this dataset with tweets from a broader timeline. To this end, we randomly selected 1,700 tweets posted between January 2020 and March 2021, to be annotated into the three classes. For the annotation, we employed three annotators on the crowdsourcing platform Prolific (<https://www.prolific.co/>). To ensure reliable annotation, we employed only annotators who are proficient in English and have participated in more than 1000 tasks with 100% acceptance rate on the platform (top 1% annotators).

Among the 1700 annotated tweets, 94 tweets were labelled different by all three annotators, 785 tweets had unanimous agreement, and 821 tweets were labelled the same by two annotators and different by the other. We discarded the tweets where all three annotators disagreed. We create two datasets by merging our annotated tweets with the tweets in the Cotfas dataset – in one dataset, we added the 1,606 tweets having majority agreement to the Cotfas dataset; in the other dataset, we add the 785 tweets having unanimous agreement to the Cotfas dataset. Thus we have four datasets as summarized in Table 1.

**Classifier models:** To classify tweets into the three classes, we used a variety of classification models, including Support Vector Machines (SVM) with TF-IDF feature vectors of tweets, the supervised FastText classifier (Joulin et al. 2016) along with some state-of-the-art BERT-based models. Specifically, we tried out the standard ‘bert-large’ (Devlin et al. 2018), and the ‘covid-twitter-bert-v2’ (Müller

<sup>5</sup>Both these datasets only provide the tweet-ids, so we could only work with those tweets that have not been deleted earlier.

	Muller	Cofas	Cofas + our data	
			(majority)	(unanimous)
<i>SVM</i>	0.628	0.751	0.642	0.695
<i>FastText</i>	0.650	0.735	0.624	0.683
<i>BERT-large</i>	0.661	0.763	0.670	0.720
<i>CT-BERT</i>	<b>0.687</b>	0.768	0.744	0.797
<i>CT-BERT++</i>	-	<b>0.858</b>	<b>0.775</b>	<b>0.825</b>

Table 2: Average macro F1-score over stratified 5-fold cross-validation, for the classifiers on the four datasets. “CT-BERT++” is CT-BERT finetuned on the dataset by Muller et al, before cross-validation on the other datasets. The best scores are marked in bold.

Testing dataset	Training dataset	Training dataset		
		<i>Muller</i>	<i>Cofas</i>	<i>Our dataset</i>
	<i>Muller</i>	-	0.430	0.471
	<i>Cofas</i>	0.723	-	0.347
	<i>Our dataset</i>	0.583	0.475	-

Table 3: Average macro F1-score of the CT-BERT classifier trained and tested on different datasets.

et al. 2020) model which has been pre-trained specifically on COVID-related tweets (which we refer to as CT-BERT).<sup>6</sup> We used the implementations of these models provided by the HuggingFace library (Wolf et al. 2020). For each of the BERT-based models, we paired them with a classification head to perform three-class classification.

**Evaluation of tweet classification:** We performed stratified 5-fold cross validation on the four datasets, and evaluated the classifiers using the average macro-F1 scores across the 5-folds. Table 2 (first four rows) presents the cross validation performance of the classifiers. On each of the four datasets, ‘covid-twitter-bert-v2’ (CT-BERT) performs the best.

We also performed cross-dataset testing, where we trained the CT-BERT model on one dataset and tested on other datasets. Table 3 shows the macro-F1 scores of CT-BERT in cross-dataset setting. The model does not perform too well when trained and tested on different datasets, possibly because the time-scale of the datasets are very different – the Muller dataset is of pre-Covid times, Cofas dataset contains only tweets posted in November 2020, while our dataset spans January 2020 through March 2021 – and hence the nature of the tweets is somewhat different.

**Final classifiers for the whole collected data:** We intend to develop an accurate classifier for vaccine-related tweets posted over a long time span (January 2018 through March 2021). Since we saw that different datasets have different characteristics (as discussed above), we decided to use data from all the different datasets to design a more accurate classification scheme.

Since the Muller dataset is an order of magnitude larger than the other two datasets, we try a variant of the CT-BERT

<sup>6</sup>We also tried some other models including Random Forest and ‘Bertweet’ (Nguyen et al. 2020). These models performed worse than the reported ones in comparison; hence omitted for brevity.

Time Period	#tweets	Anti-Vax	Pro-Vax
Pre-COVID	2.9M	19.9%	47.0%
COVID	6.5M	20.0%	27.6%
COVID-vax	6.3M	17.1%	41.7%
Overall	15.7M	18.8%	36.9%

Table 4: Distribution of predicted stances of all our collected tweets during the different time periods.

model that is initially finetuned on the Muller dataset. We refer to this model as ‘CT-BERT++’ (with weights updated via finetuning over the Muller dataset). We evaluate the CT-BERT++ model over the other datasets (apart from Muller) via stratified 5-fold cross validation, whose results are shown in the last row of Table 2. The CT-BERT++ model performs better than CT-BERT (and all other classifiers) over all the COVID-specific datasets, achieving a macro-F1 score of 0.825 on the Cofas dataset extended with the unanimously labelled tweets in our dataset. On the Cofas dataset extended with our tweets having majority agreement, our model had a average macro-F1 score of 0.775, which is in line with the original work (Cofas et al. 2021).

Based on the results in Table 2, we decided to use two different classifiers to classify tweets from the pre-Covid times and the Covid-times. To predict stance of tweets from the *pre-COVID period*, we use the CT-BERT classifier fine-tuned over the Muller dataset (which achieved the best Macro F1-score over the Muller dataset in Table 2). To predict the stance of tweets from the *COVID period* and the *COVID-vax period*, we use CT-BERT++ finetuned again over the Cofas dataset extended with the 785 unanimously labelled tweets in our dataset.

Using these classifiers we predicted the stance of all the tweets we had collected; the distribution of stances is given in Table 4.

**Checking for classifier bias:** We check for any potential biases that the final classifier may have with respect to certain specific words in the data, using the *SPCPD bias word detection strategy* (Badjatiya et al. 2019). We ran the final classifier (CT-BERT++) on each individual *word* from the datasets, and checked which words are classified as Anti-Vax or Pro-Vax. Out of 37,127 distinct words in all the datasets combined, as many as 37,011 words are *not* labeled as Anti-vax/Pro-vax, indicating *no bias with respect to 99.7% of the words*. Only 116 words were predicted as Anti-Vax or Pro-Vax. From manual observation, we find all these words to be hashtags that are used almost exclusively with the corresponding stances. For example, hashtags such as ‘#refusePoisonVaccines’, ‘#dontTrustVaccines’, and ‘#vaccinesKill’ were labeled as Anti-Vax, whereas hashtags such as ‘#getVaccinated’ and ‘#vaccinesSaveLives’ were labeled as Pro-Vax. We believe these keywords are indeed strongly associated with the corresponding stances. Thus we did not find any unjustified bias in the classifier.

## Identifying Vaccine-Related Leaning of Users

**Identifying Pro-Vaxxers and Anti-Vaxxers:** We use the tweet-classifiers described above to identify the stance of

Time Period	Anti-Vaxxers	Pro-Vaxxers	Unidentified
Pre-COVID	14,189	57,905	91,800
COVID	46,872	54,771	362,765
COVID-vax	36,768	124,053	304,475

Table 5: Number of users identified as Pro-Vaxxers and Anti-Vaxxers during the different time periods. Last column states the number of users who posted 3 or more tweets during the corresponding time-period, but their stance could not be classified due to the classifier thresholds.

Time Period	Anti-Vaxxers	Pro-Vaxxers
Pre-COVID (2018-19)	50 / 50	50 / 50
COVID (2020)	50 / 50	48 / 50
COVID-vax (2021)	44 / 50	49 / 50

Table 6: Evaluation of user classification: number of correctly classified users (according to human annotators) out of 50 randomly sampled users from each group as identified by the classifier. Overall precision = 97% over the 300 users.

every tweet posted by a user. We then label a user as *Anti-Vaxxer* or *Pro-Vaxxer* during a specific time period if they posted at least 3 vaccine-related tweets during that period, out of which more than  $\tau$  fraction is labeled as *Anti-Vax* or *Pro-vax* respectively. We filter out users who posted less than 3 tweets as it is difficult to accurately estimate their stance from so few tweets. We use  $\tau = 0.7$  similar to (Mitra et al. 2016) to get good precision in detecting user-stances. Using this method, we identified several users as *Anti-Vaxxers* and *Pro-Vaxxers* during the three time periods, as summarised in Table 5.

**Evaluation of user classification:** To evaluate the user classification, we randomly sampled 50 users from each of *Anti-Vaxxer* and *Pro-Vaxxer* groups, as predicted by the classifier for each of the three time periods. Thus, overall, we selected 300 users for this evaluation. Two annotators proficient in English and use of Twitter checked all the tweets posted by these users during the corresponding time periods, along with the users’ Twitter profiles to determine their actual stance towards vaccines.

Table 6 shows the results of the user-classification evaluation. Both annotators agreed that all the 100 users from the pre-COVID period were classified correctly. For the COVID period, 98 were correct, the other two users being ambiguous (i.e., at least one of the annotators was unsure of the user’s actual stance). For the COVID-vax period, 93 out of the 100 users checked were correctly classified, while 6 Pro-Vaxxers were wrongly classified as Anti-Vaxxers (both annotators agreed in all cases). This relatively higher misclassification rate for the COVID-vax period is because during this period, many users criticize the government for mismanagement of vaccination programmes, and the classifier wrongly classified some of these users as Anti-Vaxxers.

Overall our method of user classification correctly identified the vaccine-related stance of 291 out of 300 users, giving a precision of 97%. We believe this performance is sufficiently good for the present study.

### Change in numbers of Anti-Vaxxers and Pro-Vaxxers:

Table 5 shows the variations in the number of Anti-Vaxxers and Pro-Vaxxers over the three time periods, while Table 4 shows the variations in distribution of Anti-Vax and Pro-Vax tweets. From these tables, it is evident that a lot more users have started discussing about vaccines since 2020, than in the pre-COVID times. Table 5 shows that there was a substantial increase in the number of Anti-Vaxxers in the COVID period (2020), owing to the uncertainty and delay in vaccination during that period. However, once vaccination started in the COVID-vax period, a much larger number of users are seen to post Pro-vax opinions (even after the misclassification of some Pro-Vaxxers as Anti-Vaxxers in the COVID-Vax period, as stated in Table 6).

It can be noted that, in this paper, the set of ‘Anti-Vaxxers’ may include a broader group of users than the traditional group studied in previous works prior to 2020. Anti-Vaxxers traditionally mean people who are against the intake of all vaccines in general. Whereas, in COVID times, people have shown different levels of hesitancy/unwillingness to take the COVID-19 vaccines (or even a particular COVID-19 vaccine) due to various reasons (discussed in the next section), even though they might not be against other vaccines.

**Summary:** We designed a highly precise user-classification method, using which we classified the vaccine-stance of tens of thousands of Twitter users, over the three time periods.

### Discourse Regarding Vaccines on Twitter

In this section, we identify major themes/topics of discussion on vaccines across the different time-periods, by Anti-Vaxxers and Pro-Vaxxers.

### Method for Identifying Vaccine-Related Topics

Given a particular set of tweets (e.g., the tweets posted by Anti-Vaxxers), we adopt the topic modeling-based method suggested in (Bozarth et al. 2020). The method involves four steps: (1) identifying candidate topics and topic-words through vanilla LDA, where we estimate the number of topics using topic coherence or perplexity scores (Newman et al. 2010), (2) manually mapping the identified topics and topic-words to a smaller subset of coherent topics and associated seed-words, (3) using Labelled-LDA (Ramage et al. 2009) with these seed-words as input, and finally (4) manually examining the topics generated by Labelled LDA and combining topics that share the same theme and removing topics without coherence. This approach finally produced 8 distinct *Anti-vax topics* from the tweets posted by Anti-Vaxxers and 4 *Pro-vax topics* from the tweets of Pro-Vaxxers. Overall, 80% of the tweets posted by Pro-vaxxers and 83% of tweets posted by Anti-vaxxers could be assigned a topic by the above method. We also computed the relative frequency of the topics in the three different time periods.

### Opinions Expressed by Different User-Groups

**Topics discussed by Anti-Vaxxers:** The topics discussed by the Anti-Vaxxers are shown in Table 7, along with their distribution over the three time periods (considering the tweets that could be assigned topics), and example tweets. We find

Pre-COVID period (2018-19)	COVID period (2020)	COVID-vax period (2021)
<b>Topic 1 [Health Concerns]:</b> Concerns about the ingredients and side effects of vaccines, including deaths caused. <i>Top Words:</i> die, people, reaction, kill, receive, safe, effect, health, reaction, danger		
[19.0%] Perhaps we now have the link between vaccination and autism spectrum disorder, the link being the inclusion of an aluminium adjuvant in the vaccine.	[15.1%] I won't take one derived from any aborted fetal tissues. Astra Zeneca does come from a cell line from aborted fetal tissues. No way I'm taking that.	[18.3%] Herd immunity would serve us better and protecting the elderly and vulnerable. No need for a vaccine or experimental gene therapy
<b>Topic 2 [Against Mandatory]:</b> The user is against mandatory vaccination – talks about freedom and choice. <i>Top Words:</i> want, force, mandatory, time, body, right, mandate, passport, refuse, court		
[17.6%] Anti-vaxxers are AGAINST gov't mandates and FOR informed consent and choice.	[12.1%] Biden is going to do his best to get a mandatory Covid vaccine . I feel very bad for what's going to happen	[13.0%] The first time I'm turned away from anywhere for not having a Vaccine Passport. My body my choice.
<b>Topic 3 [Big Pharma]:</b> Dissatisfied with the role of pharmaceutical companies – they are tricking people, trying to make money. <i>Top Words:</i> pfizer, big, company, ask, pharma, immune, moderna, money, drug, manufacture		
[17.4%] Big Pharma lines their pockets, doctors get bonuses for having 63% of their patients vaccinated	[13.3%] Sure, let's trust Bill Gates, a non-doctor billionaire who pays the media instead	[9.9%] If you suspect you have a slight breathing issue.... Run Away from astrazeneca vaccine...
<b>Topic 4 [Political]:</b> The vaccine is political in nature – Governments are promoting their own agenda. <i>Top Words:</i> trump, push, govern, lie, world, russian, gates, bill, fda, country		
[9.8%] US politicians unwilling to accept that vaccines cause patient harm WHY?	[12.2%] the CDC is spewing govt scripted propaganda and lies daily	[10.8%] no mention of opposing the terrible vaccine passports govt have plans for.
<b>Topic 5 [Ineffective]:</b> Vaccines don't work – Skeptical about the effectiveness of vaccines <i>Top Words:</i> work, will, get, stop, mask, lockdown, protect, effect, immune, mutate		
[8.6%] Merck scientists are in court over lying about the efficacy of Mumps portion.	[13.8%] haven't cured HIV in 40 years, so how could a vax for COVID be effective?	[16.2%] read an article explaining that the first Covid-19 dose is only 52% effective.
<b>Topic 6 [Rushed]:</b> Vaccine development is being rushed – Not enough testing has been done <i>Top Words:</i> test, experiment, long, year, trial, medic, rush, risk, study, approve		
[8.3%] This vaccine has NOT been tested for CARCINOGENIC potential, or for IMPAIRMENT OF FERTILITY.	[12.3%] What's the rush? Better to be safe than sorry with the politically timed vaccine	[11.9%] have you seen the side effects of the untested vaccine Why would anyone be jumping Covid 19 has a 99.8% survival rate
<b>Topic 7 [Shedding]:</b> Virus Shedding – People who take vaccines will transmit the disease to others <i>Top Words:</i> shed, people, flu, vax, live, mrna, fact, spread, cause, sick		
[13.2%] the recently vaccinated with the mmr can shed for upto 21 days through recently vaccinated children.	[9.6%] the vaccine will contain live virus, and people who get the vaccine, will shed the live virus and compound the problem.	[8.8%] these new 'variant' strains are the vaccinated shedding the virus, or the virus mutating because of the vaccine.
<b>Topic 8 [Deeper Conspiracy]:</b> Disease is a Hoax, Tracking Chips and talks of other conspiracies (other than just money making) <i>Top Words:</i> need, people, risk, control, believe, old, vulnerable, death, population, mask		
[6.0%] #Vaccines are for POPULATION CONTROL aka #eugenics.	[11.6%] COVID is about "control" & bringing in "Trojan horse" vaccines	[11.1%] implant us with tracking chips and drug us to become the brainwashed robots

Table 7: Broad topics posted by Anti-Vaxxers, along with the percentage of tweets in that topic posted during a time period (within square brackets) and excerpts from sample tweets from the corresponding time periods.

that most of the broad topics have remained the same in the COVID times as they were in the pre-COVID times. However, the different topics are being discussed in different proportions in different time periods.

Across all three time-periods, many Anti-Vaxxers express **[Health Concerns]** about the vaccines, which is the most discussed topics across all the time periods. In the pre-COVID period, Anti-Vaxxers talked about how measles and flu vaccines cause various health concerns such as Autism and adverse effects on children. In the COVID times, the Anti-Vaxxers talk more about the ingredients of the COVID-19 vaccines which they assume to contain fetal cells, and how they will alter the DNA of those who take the vaccines. Stemming from various concerns, they strongly oppose mandatory vaccination (the **[Against Mandatory]** topic in Table 7). The Anti-Vaxxers have been demanding freedom

of choice of taking vaccines since the pre-COVID period. An emerging subtopic in 2021 is the protest against 'vaccine passports' being required to visit various places.

Anti-vaxxers also frequently complain against the **[Big Pharma]** (large pharmaceutical companies) stating that they only care about profit (by pushing unnecessary vaccines), or do not care about harmful effects of their vaccines. This criticism of Big Pharma seems to have reduced in the COVID-vax period, with people focusing more on other topics (discussed below). A related topic is about the **[Political]** side of vaccines. This was talked about *slightly less in the Pre-COVID period*, with the discourse mostly limited to concerns that the Governments are not doing enough to prevent the big pharma from exploiting the people. However, in the COVID period, Anti-Vaxxers are a lot more vocal about the governments pushing vaccines to gain political support.

Pre-COVID period (2018-19)	COVID period (2020)	COVID-vax period (2021)
<b>Topic 1 [Want Vaccines]:</b> User wants to get vaccinated or has got the vaccine. <i>Top Words:</i> get, time, shot, feel, protect, thing, soon, receive, prevent, health		
[31.2%] I got my kid vaccinated because vaccines are a win-win for me.	[21.0%] Yes, I will get the COVID-19 vaccine. (No, you cannot change my mind.)	[21.6%] my husband and I have gotten both Covid vaccinations and we are fine.
<b>Topic 2 [Promote Vaccines]:</b> User is sharing positive emotions about vaccines. <i>Top Words:</i> dose, go, receive, second, get, old, schedule, care, global, free		
[26.8%] If you have certain health problems like heart or breathing conditions you can get a free flu vaccine.	[21.7%] Very good news on the Pfizer / BioNTech SE mRNA vaccine: early data shows 90% effectiveness!!	[25.2%] Starting April 5, Nevada is fast forwarding to opening up vaccinations to everyone 16+
<b>Topic 3 [Against Anti-Vaxxer]:</b> Showing disagreement with Anti-Vaxxers <i>Top Words:</i> people, need, mask, risk, wait, pandemic, right, care, protect, save		
[21.3%] fact - vaccine laws increase the vaccination rate. fact - anti-vaxers believe in junk science. fact - your twitter bio made me lolz	[28.4%] Anti-vaxx conspiracy bollocks preys on marginalized people living in news deserts. I heard a person wants to take ivermectin - not COVID vaccine.	[30.4%] anti-vaxxers block access to the largest #COVID vaccination site in LA shouting epithets at seniors waiting to get their #CovidVaccine
<b>Topic 4 [Supports Authorities]:</b> Thanking or supporting medical staff or frontline workers for vaccines. <i>Top Words:</i> vaccineswork, help, leader, thank, polio, prevent, today, support, distribute, save		
[20.7%] Thanks to funding and support from UNICEF, cold chain equipment, used to store vaccines	[28.9%] Thanks @user for being a champion in our mission to get COVID-19 tests & vaccines to everyone	[22.8%] Just had my Pfizer jab !! good NHS thank you all and the Armed forces at the Vaccination point.

Table 8: Broad topics talked about by Pro-Vaxxers, along with the percentage of tweets in that topic among a time period (in brackets) and excerpts from tweets from the corresponding time periods.

Especially in the US, the 2020 presidential elections gave rise to a huge debate where the government was blamed for pushing COVID-19 vaccines (reflected by the increased fraction of such tweets in 2020).

Another common Anti-vax topic is that vaccines are [**In-effective**] against the disease. In the pre-COVID period, users opined that vaccines are ineffective against flu and measles. In the context of COVID-19 vaccines, many Anti-Vaxxers users have expressed concern that these vaccines will be ineffective against the disease (e.g, since the virus can mutate). In the COVID-vax period even more users are concerned about the effectiveness, following reports of people affected by COVID even after getting vaccines. Another major concern is that the vaccines are being rushed by the government and the big pharma, and that not enough testing has been done before releasing the vaccines to the general public (the topic [**Rushed**] in Table 7). Discussion on this topic has become much more frequent in the COVID times.

Finally, a few Anti-Vaxxers also talk about some more controversial topics. One of them is about [**Shedding**], which says that vaccinated users transmit the disease to others. In pre-COVID period, this was a very frequent talking point against MMR and Flu vaccines. In the COVID and COVID-vax periods, users also talk about shedding the virus after taking COVID-19 vaccines, but the frequency has reduced as compared to the pre-COVID times. Additionally, some [**Deeper Conspiracy**] theories are especially being discussed in the COVID times, such as COVID-19 being made up to trick people into buying vaccines, presence of tracking chips in vaccines, population control, etc. These conspiracy theories are talked about a lot more after the onset of COVID (Cinelli et al. 2020).

**Topics discussed by Pro-Vaxxers:** The topics discussed by

Pro-Vaxxers are given in Table 8, along with their distribution across the three time-periods and some example tweets. Many users have actively stated that they wish to take (or have taken) vaccines ([**Want Vaccines**]) and are seen to post information that promotes the adoption of vaccines (the [**Promote Vaccines**] topic in Table 8). In the COVID-vax period, many Pro-Vaxxers have shared that they have taken the COVID-19 vaccine, stating that the side effects they experienced are mild. Pro-Vaxxers also [**Support Authorities**] by thanking them for development and distribution of vaccines. However, it is observed that the fraction of people actively talking about getting vaccines ([**Want Vaccines**]) has reduced since the onset of COVID, which shows some level of vaccine hesitancy among the masses. On the other hand, more Pro-vaxxers have actively expressed their disapproval against Anti-Vaxxers ([**Against Anti-Vaxxer**]) in the COVID and COVID-vax periods (as also observed in prior works (Yuan et al. 2019; Cossard et al. 2020)).

**Summary:** We identified broad topics discussed by Anti-Vaxxers and Pro-Vaxxers across different time periods. Our identified topics match with those identified by prior works (Mitra et al. 2016; Jamison et al. 2020). But to our knowledge, no prior work analyzed how Anti-vax / Pro-vax discussions have varied in pre-COVID times and COVID times. While the broad topics being discussed are similar in pre-COVID and COVID times, we observe differences in their relative frequencies, e.g., lot more complaints about vaccines being rushed/ineffective and deeper conspiracy theories in COVID times, compared to the pre-COVID times.

### Change in Vaccine-Stances of Users

Recall from earlier sections that we identified large numbers of Pro-Vaxxers and Anti-Vaxxers from the pre-COVID,

		pre-COVID period	
		Anti-Vaxxer	Pro-Vaxxer
COVID period	Anti-Vaxxer	2791	<b>218</b>
	Pro-Vaxxer	<b>21</b>	6853
		COVID period	
		Anti-Vaxxer	Pro-Vaxxer
COVID-vax period	Anti-Vaxxer	7855	<b>121</b>
	Pro-Vaxxer	<b>329</b>	9898

Table 9: Statistics of users who changed their vaccine-stance across the different time periods. The number of users who have changed their stance are marked in bold.

COVID and COVID-vax time periods (see Table 5). Many of these users actively posted tweets only during one of the time periods; hence we could detect the stance of a user during *multiple time-periods* only for a smaller set of users. We now consider those users for whom we could detect vaccine-related stance during at least two different time periods. Table 9 shows how their stances varied across the three periods.

Most users continue to follow the same stance towards vaccines. However, we find *675 users who have changed their stance towards vaccines*. We see 218 users who were pro-vaxxer in pre-COVID period (2018-19) have become Anti-Vaxxers in COVID period (2020), which corroborates with the rise of vaccine hesitancy in COVID times (Bonnie et al. 2020). Interestingly, a small set of 21 users did just the opposite – they turned from Anti-Vaxxers in pre-COVID times to Pro-Vaxxers in 2020. Again, people seem to be gaining back trust in vaccines in 2021, as we see 329 users shifting from Anti-Vaxxers (in 2020) to Pro-Vaxxers (in 2021). However, 121 users who were supporting vaccines in 2020 started showing hesitancy towards vaccines in 2021. In fact, there are 14 users who changed their stance multiple times – from Pro-Vaxxers in 2018-19 to Anti-Vaxxers in 2020 and again to Pro-Vaxxers in 2021.

**Checking for bots:** Automated bots abound on Twitter (Sayyadiharikandeh et al. 2020), and we wanted to ensure that we are focusing on human users. To this end, we used Botometer-v4 (Sayyadiharikandeh et al. 2020) which returns a score between 0 (very likely to be a human) and 1 (very likely to be a bot) for a Twitter user. We applied Botometer-v4 over all the 675 users who changed their stance. The tool assigned a score  $\geq 0.5$  to 68 users, implying that these users are more likely to be bots. Since the automated detection of bots often predicts a lot of false positives (Rauchfleisch and Kaiser 2020), we manually checked these 68 accounts. We found only *one* of these accounts to be automated, that automatically shares articles from an external personal blog. We ignore this account in subsequent analyses.

In the rest of this section, we focus on the 674 users who have changed their vaccine-related stance. We refer to this set of users as the *stance-changed users*.

**Is the stance change an effect of users posting different opinions about different vaccines?** It is possible that an individual user has different opinions about different COVID-19 vaccines. Hence there is a potential concern about the set

of stance-changed users – is the observed change in stance an artefact of users posting positive opinions about one vaccine at one point of time, and negative opinions about some other vaccine at some other point of time?

To address this concern, we check for the mentions of different vaccines (with different variations of their names and manufacturers, such as AstraZeneca/Oxford, Pfizer/BioNTech, Moderna, Sputnik, Sinopharm, etc.) in all the tweets posted by the stance-changed users. Out of the 674 distinct stance-changed users, 133 had posted at least one tweet about two different vaccines. Out of these, we found only 13 users (2% of 674) who posted a clear preference towards one vaccine, while rejecting/opposing another vaccine. We have given a few examples of tweets from such users in Table 11 (each row showing two tweets from the same user). We also noted that, all these users also expressed concerns about vaccines in general during one of the time periods, and posted in support of vaccines in general in some other time period. Thus, we conclude that all the 674 users whom we have identified, actually changed their *general* (not related to specific vaccines) stance towards vaccines at some point of time.

**Demonstrating the change for some individual users:** Table 10 shows samples of tweets posted by some of the stance-changed users. Each row of this table show two tweets posted by the same user during two different time periods. Some of the tweets seem to indicate potential reasons of why the users changed their stances. For instance, user U1 (who shifted from Anti-Vaxxers in pre-COVID period to Pro-Vaxxers in COVID period) experienced a life-changing event that may have changed his/her stance. Those who changed from Pro-Vaxxers to Anti-Vaxxers (user U3 & U4) during this time were getting skeptical of COVID-19 vaccines in specific due to various reasons, including political factors. Some users were hesitant about the COVID-19 vaccines in 2020; but with the actual rollout of vaccines in 2021, they are willing to get vaccinated (e.g., U5 & U6). Finally some users who were Pro-Vaxxers in 2020 also have turned Anti-Vaxxers in the COVID-vax period (U7 & U8), due to factors such as mistrust in how effective vaccines would be, and opposition to mandatory vaccination.

## Representativeness of Stance-Changed Users

In this section, we analyse the stance-changed users to understand *how representative they are of the general population of active Twitter users*. This question is important to ascertain the practical utility of the insights about these users (as reported later in the paper), for policy-makers.

As stated in earlier sections, we attempted to detect the stance of a user during a certain time-period only if that user posted at least 3 tweets during that period. Thus the stance-changed users have posted at least 3 tweets each in two different time periods. Hence, we compare these users with the population of active Twitter users who posted at least 6 tweets related to vaccines in the entire time period of analysis (Jan 2018 – March 2021). In total, there are 340K such active users.

Figure 1 compares social network properties (number of



	Pre-COVID period (2018-19): Anti-Vaxxers	COVID period (2020): Pro-Vaxxers
U1	Thank you, @USER. You know your son. Seizures followed vaccination. Vaccine safety has not been proven. Cases like yours demand further investigation.	My husband died after becoming ill with corona virus. I am glad a vaccine has been developed in record time. I intend to be vaccinated as soon as possible
U2	Having an autistic child destroyed my family. I took my healthy baby to the doctor that gave the vaccine that caused it.	@USER Thank you for getting us vaccines in Illinois. As a frontline healthcare worker I thank you.
	Pre-COVID period (2018-19): Pro-Vaxxers	COVID period (2020): Anti-Vaxxers
U3	A vaccine eradicated smallpox. If the government was trying to kill us by way of a vaccine, we wouldn't be living longer, and diseases would not be prevented	I will NEVER take a vaccine for COVID. This whole fiasco w/masks & destroying our economy over a virus w/ 99% survival rates, is ridiculous.
U4	Every child should be vaccinated I don't understand why parents wouldn't sad endangering others	Honestly this administration caused so much sickness and death I will not trust a vaccine. #LyingTrump
	COVID period (2020): Anti-Vaxxers	COVID-vax period (2021): Pro-Vaxxers
U5	I'm not paying Big Pharma for a vaccine against something that they invented in order to sell vaccines.	I've had the first shot of the Phizer vaccine. I'm a school bus driver. I'm in contact with kids daily and several times
U6	But even a vaccine won't stop the spread, so we could see these levels for years.	Today has been a great day. My grandad has told he's getting the Covid vaccine.
	COVID period (2020): Pro-Vaxxers	COVID-vax period (2021): Anti-Vaxxers
U7	Yes, it's bad, Covid is a nasty virus and I look forward to the vaccine rollouts.	Covid is not dangerous to children. Forcing them to take this vaccine is insane.
U8	@USER Once we have a vaccine it will become as irrelevant as the common cold. Until then....	They are of concern because some vaccines may be less than 50% efficient in protecting from some variants.

Table 10: Samples of tweets (excerpts) from some of the stance-changed users (U1-U8). Each row shows two tweets posted by the *same user* during two different time periods. We do not state any usernames/identifiers to protect the privacy of the users.

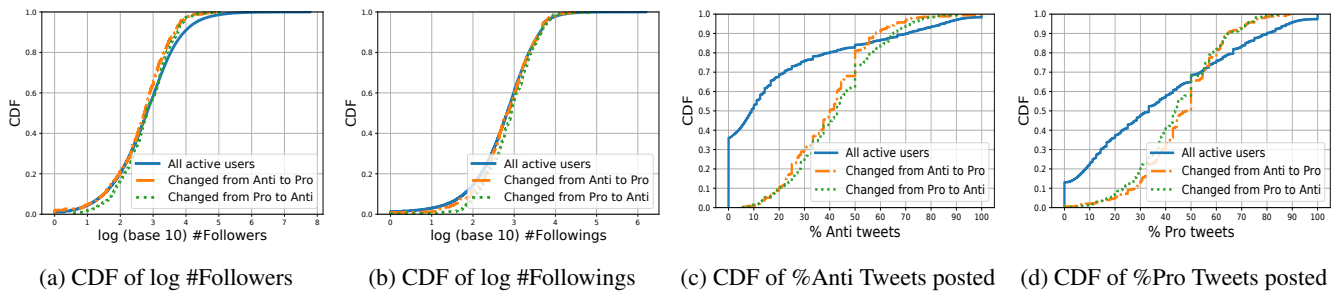


Figure 1: Comparison of stance-changed users with the general active Twitter population interested in vaccines, in terms of (a) number of followers, (b) number of followings, (c) % of Anti-vax tweets posted, and (d) % of Pro-vax tweets posted.

	Against a Vaccine	Supporting a vaccine
U1	declining <b>Pfizer</b> vaccine seeing peoples side effects	I'm more likely to trust the <b>Oxford</b> one
U2	I have concerns about the efficacy of the <b>As-traZeneca</b> vaccine	There's a shortage of <b>Pfizer</b> vaccines and the <b>Moderna</b> one won't be out till March
U3	I don't want the <b>Pfizer</b> vaccine. Let them have it.	PLEASE use the <b>Moderna</b> vaccine, not the <b>Pfizer</b>
U4	<b>Pfizer</b> has not been tested, at all, at these intervals.	<b>Oxford</b> vaccine is demonstrating lower transmission.

Table 11: Samples of tweets (excerpts) from some of the stance-changed users who posted against a particular vaccine, and supporting another vaccine.

followers and followings) and the tweet composition behaviour of the two user-groups. Figures 1a and 1b show that the followers and followings distributions of the stance-changed users and the active users are very similar.

Next, we compare the tweeting behavior of the two user-groups; Figure 1c and Figure 1d respectively compare the distributions of the fraction of Anti-Vax tweets and fraction of Pro-Vax tweets posted by an individual user (in the entire time period of analysis) in the two user-groups. We see that a large portion of the active user population (37%) did not post any Anti-Vax tweet at all (Figure 1c), while about 12% of the active user population did not post any Pro-Vax tweet at all (Figure 1d); also, only a small fraction of the active users post a large fraction of Anti-Vax tweets. In sharp contrast, the stance-changed users post higher fractions of Anti-vax as well as higher fractions of Pro-vax tweets, compared to the general active Twitter population interested in vaccines (who post a lot more neutral tweets on vaccines).

Thus, the stance-changed users are representative enough of the active Twitter population interested in vaccines in terms of social network properties, but they express their opinions (both Anti-vax and Pro-vax) much more than the general active Twitter population who post about vaccines.




Figure 2: Anti-vax topics posted by stance-changed users.

### Characterizing Stance-Changed Users

Now, we analyze the stance-changed users to understand what anti-vax topics they used to discuss before/after their stance change. To this end, we select four groups of stance-changed users – (G1) all the 21 users who changed from Anti-Vaxxer in pre-COVID period to Pro-Vaxxer in COVID period, (G2) 50 randomly selected users from those who changed from Anti-Vaxxer in COVID period to Pro-Vaxxer in the COVID-vax period, (G3) 50 random users from those who changed from Pro-Vaxxer to Anti-Vaxxer between the pre-COVID and COVID periods, and (G4) 50 random users who changed from Pro-Vaxxer in COVID period to Pro-Vaxxer in COVID-vax period. Two authors manually examined the vaccine-related tweets posted by these users from the time period when they were Anti-Vaxxers, to identify which anti-vax topics (from Table 7) they posted about.

Figure 2 compares these groups of users – the anti-vax topics are shown on the X-axis while the fraction of users from each group who posted tweets corresponding to the topics is shown on the Y-axis. Note that we did *not* see any tweet about deeper conspiracy theories (Topic 8 in Table 7) from any of these groups. Since all these users were/became Pro-Vaxxers at some point, it is likely that they were *not* staunch Anti-Vaxxers, and conspiracy theories may be an opinion prevalent among staunch Anti-Vaxxers only.

**Users who shifted from Anti- to Pro-Vaxxers:** We observe an interesting difference in the anti-vax opinions of the two groups of users who changed from Anti- to Pro-Vaxxers (G1 and G2). Most of the users in G1 – who were Anti-Vaxxers in pre-COVID period (2018-19) – opposed vaccines because of the health concerns, mostly in regards to their children being harmed. A fair number among them also were against mandatory vaccination, and against how the Big Pharma are trying to sell vaccines with side effects. However, the users in G1 did *not* post much about the vaccine development being rushed/untested or the role of politics. On the contrary, users in G2 – Anti-Vaxxers in COVID period (2020) but supporting COVID-19 vaccines (in 2021) – seem to care a lot more about the political nature of the vaccines, the potential ineffectiveness of the vaccines against COVID, and that the

vaccine development process was rushed.

**Users who shifted from Pro- to Anti-Vaxxers:** For the users in G3 – who became Anti-Vaxxers in the COVID period – the main concerns seem to be about the potential side-effects of vaccines (health concerns), and that the big pharma and politicians are exploiting the crisis to their advantage. Their other concerns are about vaccines being rushed into distribution, and that the vaccines would not be effective against the disease. On the other hand, among the users in G4 – who shifted from Pro-vaxxer in 2020 to Anti-Vaxxers in the COVID-vax period – most blamed the governments (more than the big pharma), stating that they will not take vaccines offered by specific governments and political leaders. When vaccinations actually started, these users started to tweet more against mandatory vaccination, and the ineffectiveness of the vaccines amidst rising mutations of the corona-virus. Health concerns about vaccines remains a common theme between both user-groups.

**Users who were Anti-Vaxxers in 2020 but Pro-Vaxxers in other time periods:** Finally, we compare the two groups G2 and G3 – both these groups were Anti-vaxxers in 2020; however, users in G2 became Pro-vaxxers in 2021, whereas users in G3 were Pro-vaxxers in 2018-19. Users in G3 tweeted more in 2020 about Health concerns and the Political side of vaccines, than the users in G2. In contrast, the users in G2 tweeted (in 2020) more about vaccines being Ineffective and Rushed. These observations signify that even though both groups had the same stance (Anti-Vax) in the COVID period (2020), their concerns were different.

### Exploring Correlation of Users’ Stance with Stance of their Neighbours

We explore if the social network has an effect of the changing stance of users. In particular we analyse the influence (if any) of the people whom a user is following at various points of time, on the changing stance of the said user.

**Method:** Currently, it is not possible on Twitter to know the creation date for social links, or to collect the list of followings of a user (also known as *friends* of the user) on a past date. Hence, we adopt the following method. For a particular user  $u$ , we collect all the current followings of  $u$ . Then we check the tweets posted by each of these followings during a particular time period (e.g., 2018-19, or 2020, or 2021). We compute what fraction of  $u$ 's followings were Pro-Vaxxers and Anti-Vaxxers during each of the three time periods, based on the vaccine-related tweets that they (the followings) posted in the respective time periods.

Thus, for user  $u$ , we first compute the fractions of Anti-Vaxxer and Pro-Vaxxer followings of  $u$  during a particular time period. Then we see how the fractions of Anti-Vaxxer and Pro-Vaxxer followings of  $u$  have *changed* from one time period to another (e.g., from 2018-19 to 2020, or from 2020 to 2021). Finally, we average this change in fractions of Anti-Vaxxer and Pro-Vaxxer followings of a user, over all users in a group (e.g., the groups G1, G2, G3, G4 introduced in the previous section).

**Observations:** Table 12 shows these changes. For instance, across all the users in G1 (who changed from Anti-Vaxxer in

User-groups who changed their stance	Change from 2018-19 to 2020 in		Change from 2020 to 2021 in	
	#Pro-followings	#Anti-followings	#Pro-followings	#Anti-followings
G1: <i>Anti (2018-19) to Pro (2020)</i>	<b>1.06</b> × (0.82, 1.29)	<b>1.94</b> × (1.03, 2.85)	<b>1.45</b> × (0.99, 1.91)	<b>0.70</b> × (0.47, 0.93)
G3: <i>Pro (2018-19) to Anti (2020)</i>	<b>1.04</b> × (0.88, 1.20)	<b>4.43</b> × (3.60, 5.27)	<b>1.67</b> × (1.47, 1.86)	<b>0.85</b> × (0.75, 0.94)
G2: <i>Anti (2020) to Pro (2021)</i>	<b>1.34</b> × (1.20, 1.48)	<b>4.38</b> × (3.43, 5.34)	<b>2.25</b> × (2.10, 2.40)	<b>0.70</b> × (0.60, 0.79)
G4: <i>Pro (2020) to Anti (2021)</i>	<b>1.38</b> × (1.19, 1.56)	<b>5.89</b> × (4.01, 7.78)	<b>1.49</b> × (1.30, 1.67)	<b>1.25</b> × (1.06, 1.45)
<i>Pro (2018-19) to Anti (2020) to Pro (2021)</i>	<b>0.72</b> × (0.56, 0.87)	<b>1.94</b> × (1.16, 2.73)	<b>2.17</b> × (1.61, 2.74)	<b>0.35</b> × (0.11, 0.58)

Table 12: Studying the change of stance among the followings of the stance-changed users – shown (in bold) are the changes in the number of anti-vax and pro-vax followings between different time periods, averaged over all users in a particular group. 95% confidence intervals are given beside each number in parantheses.

2018-19 to Pro-Vaxxer in 2020), the fraction of Pro-Vaxxers in their followings increased in 2020 to 1.06 times of the fraction in 2018-19, whereas the fraction of Anti-Vaxxer followings increased to 1.94 times over the same time-period. Below each number, we have also added the 95% confidence intervals (within parentheses) to demonstrate that the reported numbers are meaningful, and not skewed because of a few particular users. The first notable observation from Table 12 is that only a few of the changes are close to 1.0. This implies that, for the stance-changed users, the amount of Anti-Vaxxers and Pro-Vaxxers among their followings changed by quite a margin across different time periods.

In most cases, we also see correlation between the change of stance of the users with the change in the fraction of followings of the corresponding stance. Even though the overall number of Anti-Vaxxers increased in 2020, and Pro-Vaxxers in 2021 (as was seen in previous sections), we see a difference in the amount of increase in the fractions between the different user groups. For example, among the users in G3 (who shifted from Pro-Vax in 2018-19 to Anti-Vax in 2020), the fraction of Anti-Vaxxers in their followings increased by 4.43 times in 2020. In contrast, for the users in G1 (who shifted from Anti-Vax in 2018-19 to Pro-Vax in 2020), the increase of Anti-Vaxxers in their followings was only 1.94 times. Similarly, for the users in G4 (who changed from Pro-Vax in 2020 to Anti-Vax in 2021) and vice versa (G2), the increase in Pro-Vaxxers in the followings (from 2020 to 2021) was 1.49 times and 2.25 times respectively. The most interesting observation is for those users who changed their stance from Pro-Vaxxers to Anti-Vaxxers and again back to Pro-Vaxxers – there is an increase of 1.94 times in the fraction of Anti-Vaxxers among their followings between the pre-COVID and COVID periods, whereas the fraction of Pro-Vaxxers among their followings increased by 2.17 times between the COVID and the COVID-vax periods.

We also analysed the users who were Pro-Vaxxers or Anti-Vaxxers *throughout the three time periods* (details omitted for brevity). Even though the amount of Pro-Vaxxers and Anti-Vaxxers in their followings change over the time periods, the fraction of followings with the *same stance as the user* hugely surpasses the fraction of followings with the opposing stance throughout all time periods. For the users who are Pro-Vaxxers throughout, the fraction of Pro-Vaxxer followings was 80 times that of the fraction of Anti-Vaxxer followings at minimum (COVID period), and 230 times at maximum (pre-COVID period). Similarly, for the users who

are Anti-Vaxxers throughout, the fraction of Anti-Vaxxer followings was 4 times the fraction of Pro-Vaxxer followings at minimum (pre-COVID period) and 7 times at maximum (COVID-period).

Even though the overall number of users posting about vaccines may have increased during COVID times, the important point to note is that in most cases, there is a much greater increase in the fraction of followings with the same stance as the user, than the increase in followings with the opposite stance. Thus, the vaccine-stance of a user seems to be correlated with the vaccine-stance of his/her social neighborhood. Also, our findings hint that users who change their vaccine-stances have a non-homogeneous neighborhood across pre-COVID and COVID times.

## Conclusion

This work presents a first attempt to systematically analyse how a global event (the COVID-19 pandemic) has changed the vaccine-related stances/opinions of people from pre-COVID to COVID times. Our findings also hint at how the stance-change of a user can be correlated with factors such as major events in his/her life (e.g., death of a close relative) or changes in his/her social neighborhood.

**Implications of the work:** We believe that the present work is useful to the authorities in their quest to ensure large-scale adoption of vaccines, by educating and disseminating specific information to specific people in the society. Essentially, we provide a framework to use social media data to automatically detect the vaccine-stance of users at scale, and understand the primary vaccine-related concerns (if any) of users in different time periods. This framework can help the authorities in the following ways –

(i) **Devising targeted approaches to alleviate people’s vaccine-related concerns:** A policy maker cannot educate people unless they know the reasons/concerns behind their vaccine-related stance. We have seen that the anti-vax concerns are very different, from deep-rooted (e.g., vaccines kill people) to contextual (e.g., concerns regarding COVID-19 vaccines being ineffective and rushed). So, the policy makers need a targeted approach to educate people having different types of concerns.

(ii) **Identifying and reaching out to people who changed their stances:** We showed there is a population who have changed their vaccine-stances. The authorities can identify the people who can potentially move over from anti-vax to

pro-vax, and better nudge those people towards vaccination. Perhaps, the experiences of some of the users who changed from Anti-vaxxer to Pro-vaxxer, can be used to nudge others in the Anti-vax community towards vaccination. Also, the authorities can investigate why some originally pro-vax users have become hesitant to take COVID-19 vaccines, and which specific concerns need to be clarified for them.

(iii) **Utilizing social media for nudging people towards vaccination:** We found that the social connections can potentially influence the stance of users, which suggests the importance of using popular social media users in promoting more Pro-Vax information. However, the information needs to be carefully designed in accordance to the primary concerns of different target user groups, without which they might just be mistaken as propaganda. Also, policy makers can design targeted friend/content recommendation algorithms to break into today's Anti-vax echo chambers with pro-vax opinions.

**Future work:** We only explored the effects of the stances of a user's followings on the stance change of the user. Other forms of interactions such as retweeted or replied to/by can also have an effect on the stance change of users, and can be explored as a future work. Another future work would be to design predictive models that attempt to predict which users are likely to change their stance, e.g., from a user's posts and the stances of the followings of a user.

## Acknowledgements

The authors acknowledge the anonymous reviewers whose comments greatly helped to improve the paper. The project is partially supported by a research grant from Accenture Corporation. The first author (S. Poddar) is also supported by the Prime Minister's Research Fellowship (PMRF) from the Ministry of Education, Government of India.

## References

Badjatiya, P.; et al. 2019. Stereotypical bias removal for hate speech detection task using knowledge-based generalizations. In *The World Wide Web Conference*.

Bonnevie, E.; et al. 2020. Quantifying the rise of vaccine opposition on Twitter during the COVID-19 pandemic. *Journal of Communication in Healthcare*, 1–8.

Bozarth, L.; et al. 2020. Higher ground? How groundtruth labeling impacts our understanding of fake news about the 2016 US presidential nominees. In *Proc. ICWSM*.

Chou, W.-Y. S.; et al. 2020. Considering Emotion in COVID-19 vaccine communication: addressing vaccine hesitancy and fostering vaccine confidence. *Health communication*, 35(14): 1718–1722.

Cinelli, M.; et al. 2020. The covid-19 social media infodemic. *Scientific Reports*, 10(1).

Cossard, A.; et al. 2020. Falling into the echo chamber: the Italian vaccination debate on Twitter. In *Proc. ICWSM*.

Cotfas, L.-A.; et al. 2021. The Longest Month: Analyzing COVID-19 Vaccination Opinions Dynamics From Tweets in the Month Following the First Vaccine Announcement. *IEEE Access*, 9: 33203–33223.

Devlin, J.; et al. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Dutta, U.; et al. 2021. Analyzing Twitter Users' Behavior Before and After Contact by the Russia's Internet Research Agency. *Proc. CSCW*.

Funk, C.; and Tyson, A. 2020. Intent to get a COVID-19 vaccine rises to 60% as confidence in research and development process increases. *Pew Research Center*, 3.

Gunaratne, K.; et al. 2019. Temporal trends in anti-vaccine discourse on twitter. *Vaccine*, 37(35): 4867–4871.

Jamison, A. M.; et al. 2020. Not just conspiracy theories: Vaccine opponents and proponents add to the COVID-19 'infodemic' on Twitter. *Harvard Kennedy School Misinformation Review*, 1(3).

Johnson, N. F.; et al. 2020. The online competition between pro- and anti-vaccination views. *Nature*, 1–4.

Joulin, A.; et al. 2016. Bag of Tricks for Efficient Text Classification. *arXiv preprint arXiv:1607.01759*.

Lwin, M. O.; et al. 2020. Global sentiments surrounding the COVID-19 pandemic on Twitter: analysis of Twitter trends. *JMIR public health and surveillance*, 6(2): e19447.

Malik, A. A.; et al. 2020. Determinants of COVID-19 vaccine acceptance in the US. *EClinicalMedicine*.

Mitra, T.; et al. 2016. Understanding anti-vaccination attitudes in social media. In *Proc. ICWSM*.

Müller, M.; et al. 2020. COVID-Twitter-BERT: A Natural Language Processing Model to Analyse COVID-19 Content on Twitter. *arXiv preprint arXiv:2005.07503*.

Müller, M. M.; et al. 2019. Crowdbreaks: Tracking health trends using public social media data and crowdsourcing. *Frontiers in public health*, 7: 81.

Newman, D.; et al. 2010. Automatic Evaluation of Topic Coherence. In *Proc. NAACL*.

Nguyen, D. Q.; et al. 2020. BERTweet: A pre-trained language model for English Tweets. In *Proc. EMNLP: System Demo*.

Nuzhath, T.; et al. 2020. COVID-19 vaccination hesitancy, misinformation and conspiracy theories on social media: A content analysis of Twitter data. *SocArXiv*.

Praveen, S.; et al. 2021. Analyzing the attitude of Indian citizens towards COVID-19 vaccine-A text analytics study. *Diabetes & Metabolic Syndrome: Clinical Research & Reviews*.

Punuru, A.; et al. 2020. Cultural Norms and Interpersonal Relationships: Comparing Disclosure Behaviors on Twitter. In *Conference Companion Publication, CSCW*.

Ramage, D.; et al. 2009. Labeled LDA: A supervised topic model for credit attribution in multi-labeled corpora. In *Proc. EMNLP*.

Rauchfleisch, A.; and Kaiser, J. 2020. The False positive problem of automatic bot detection in social science research. *PloS one*.

Sayyadharikandeh, M.; et al. 2020. Detection of novel social bots by ensembles of specialized classifiers. In *Proc. CIKM*.

Thunstrom, L.; et al. 2020. Hesitancy towards a COVID-19 vaccine and prospects for herd immunity. *Available at SSRN 3593098*.

Wolf, T.; et al. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proc. EMNLP: System Demo*.

Yuan, X.; et al. 2019. Examining emergent communities and social bots within the polarized online vaccination debate in Twitter. *Social Media+ Society*, 5(3).