The Journal of the Southern Association for Information Systems

Volume 11 | Issue 1 Article 2

4-1-2024

The Celebrity Factor: Exploring the Impact of Influencers on COVID-19 Vaccine Sentiment through Bayesian Modeling of Time Series.

Abhishek Shah
Green Level High School, shahabhishek1729@gmail.com

Shweta Shah
Green Level High School, shwetashah857@gmail.com

Bill Rand
North Carolina State University, wmrand@ncsu.edu

Xiaoxia Champon
North Carolina State University, xzhao17@ncsu.edu

Recommended Citation

Shah, A., Shah, S., Rand, B., & Champon, X. (2024). The Celebrity Factor: Exploring the Impact of Influencers on COVID-19 Vaccine Sentiment through Bayesian Modeling of Time Series.. The Journal of the Southern Association for Information Systems, 11, 31-52. https://doi.org/doi:10.17705/3JSIS.00034

The Celebrity Factor: Exploring the Impact of Influencers on COVID-19 Vaccine Sentiment through Bayesian Modeling of Time Series.

Cover Page Footnote

The authors would like to thank Dr. Ana-Maria Staicu of NC State University for funding this project. Xiaoxia Champon is funded by the T32 NIH grant number HL079896. Dr. William Rand and Dr. Anthony Weishampel are funded by NSF Award 2020179

ABSTRACT

Online social networks allow for information to rapidly propagate throughout the world, and opinions expressed on such platforms can influence people's decisions. During the COVID-19 pandemic, many influential public figures used these social networks to share their opinions about the vaccines developed to combat the virus. Many influencers encouraged vaccination, and a considerable number also expressed doubt and skepticism over the efficacy of the vaccines. This study modeled the impact that eleven influencers' statements had on the overall sentiment towards COVID-19 vaccines, as expressed on Twitter. Sentiment is measured by collecting a series of publicly-available tweets made regarding the vaccine during the pandemic, and assigning each a sentiment score based on the VADER lexicon. Several models were used to analyze the impact of the influencers' statements, including linear, sequential and tree-based models. The results were obtained by constructing a Bayesian structural time series model based on each model's counterfactual estimate. The results found that influencers who share messages encouraging vaccination generally tend to increase the number of ”pro-vaccination” tweets over the next 20 days. Influencers sharing ”anti-vaccination” messages sometimes resulted in a decrease in anti-vaccine tweets, and other times in an increase over the next 20 days. The results from this study provide an introductory look into the complex issue of vaccine hesitancy and the effect of influencers on vaccine messaging, and inform public health strategy regarding this issue.

Keywords

COVID-19, Vaccine, Celebrity, Influencer, Vaccine hesitancy, time-series forecasting, vaccines, COVID, machine learning, celebrity influence, BSTS, sentiment analysis, VADER lexicon, Twitter, Endorsement, Influence

INTRODUCTION

The SARS–CoV–2 virus was first detected during December of 2019 in Wuhan, China (Zhu et al. 2020). The virus has since spread to most of the world and was declared a global pandemic by the World Health Organization on March 1, 2020 (Carvalho et al. 2021). To end the pandemic, governments around the world funded the development of a vaccine – the U.S. government alone spend roughly $13 billion on vaccines as of February 2021 (Bloom et al. 2021). mRNA–based COVID–19 vaccines, such as the COMIRNATY vaccine (developed by Pfizer inc., BioNTech SE, and Shanghai Fosun Pharmaceutical Co., Ltd.) and the Moderna vaccine, function by injecting a small amount of genetic material into the body, which instructs cells to produce a harmless replica of the target virus. The immune system then recognizes this foreign protein as a potential threat and develops a defense by producing antibodies, providing protection against future infections (Park et al. 2021). However, if governments are to achieve the goal of herd immunity (i.e., a state when a large part of a population has developed immunity to a disease), they need large proportions, an estimated 67%, of their populations to take the vaccines (Randolph & Barreiro 2020).

However, many individuals are hesitant to take the vaccine, citing beliefs (often based on misinformation) about the safety and efficacy of the vaccine to support their views (Afifi et al. 2021). Vaccine hesitancy leads to a delay in achieving herd immunity status despite enough vaccines having been produced – leading to further spread and mutation of the virus, and eventually even greater loss of life (MacDonald et al. 2015). Indeed, in 2019, the World Health Organization included vaccine hesitancy in the list of the top ten major threats to global health (Sweileh 2020). One approach governments around the world employed to promote vaccinations was to utilize the voice of prominent figures in popular culture to spread accurate information and encourage the public to get vaccinated (Rzymski et al. 2021). For instance, the Biden administration invited pop–star Olivia Rodrigo to the White House in mid–July, and conducted interviews with multiple internet personalities (e.g., Christina Najjar) to encourage young people to get vaccinated (Chansolme 2022); in Uganda, some prominent religious leaders were involved to achieve the same goals (Ssanyu 2022).

Celebrity Influence Celebrities have long held influence among the public, especially among younger age groups, i.e., adolescents and young adults, who are one of the most vaccine–hesitant groups (Adams et al. 2021). Celebrity advertising, for instance, has been a popular method of marketing used by companies for decades to reach these groups (Yannopoulos 2012). With the rise of social media platforms such as Twitter, Instagram, Facebook, Reddit, and others, celebrities have been able to spread their views further than ever before, due to the rapid diffusion rate of messages posted on their relatively massive networks (Selkie 2022). Recently, many popular celebrities and social media influencers have used their platform to push out statements supporting the vaccine, both independently and by partnering with public health authorities; however, a considerable number of celebrities have also pushed statements against the vaccine (Calac et al. 2022), potentially leading to increased vaccine hesitancy amongst their followers. This study set out to explore how celebrities and social media influencers have affected the public sentiment, as expressed on Twitter[1], regarding COVID–19 vaccines.

Theoretical Foundations The study was built on two foundational theories: the theory of planned behavior (Fishbein & Ajsen 1975) and the source credibility model (Hovland & Weiss 1952). The theory of planned behavior is a generalization over the theory of reasoned action (Fishbein 1979), which describes the factors that affect human decisions and behaviors. In its simplest state, the theory posits that *behavioral intention*, the likelihood of a specific behavior being performed, as being a linear combination of two factors:

1. *Attitudes*: The individual's *behavioral belief* of how probable possible outcomes are in response to performed behaviors and their *evaluation* of how favorable each outcome is evaluation.
2. *Subjective norms*: The individual's *normative beliefs* as to how each behavior would be perceived by society (which would include their peers as well as celebrities and people of importance), as well as that individual's *motivation to comply* with those perceptions.

This linear combination can be expressed by the following equation:

$$BI = w_A A + w_{SN} SN, \qquad (1)$$

where *BI* denotes the behavioral intention, *A* denotes the individual's attitudes, *SN* represents the subjective norms and w_A and w_{SN} represent the weights associated with each factor. However, the Theory of Reasoned Action assumes that the individual making the decision has complete volitional control over their behaviors, and that the individual perceives every possible behavior as being equally difficult to perform, which does not hold here. For instance, since many employers required their employees to be vaccinated in order to continue work, workers did not have complete volitional control over the vaccination decision. The theory of planned behavior arose as a solution to this limitation, which adds the following factor to the model:

3. *Perceived behavioral control*, the individual's perception of their ability to perform a behavior. The perceived behavioral control is determined by *control beliefs,* an individual's belief about factors that make some behaviors easier to perform than others, and the individual's *perceived power* for each control belief. The perceived behavioral control is related to the concept of self–efficacy (Bandura 1977), which describes an individual's belief of their own ability to perform a behavior.

With this addition, the model now expands to

$$BI = w_a A + w_{SN} SN + w_{PBC} PBC, \qquad (2)$$

[1] As of August 2023, Twitter has rebranded to X under X Corp., but will still be referred to as Twitter in this paper, staying consistent with the name of the platform when the study's data was collected.

where PBC denotes the perceived behavioral control, w_{PBC} denotes the weight associated with the PBC, and all other terms are unchanged from above. The three factors, A, SN and PBC are themselves proportional to the individual's behavioral, normative and control beliefs, respectively:

$$A \propto \sum_{i=1}^{n} b_i e_i, \qquad SN \propto \sum_{i=1}^{n} n_i\, m_i, \qquad PBC \propto \sum_{i=1}^{n} c_i\, p_i, \tag{3}$$

where, for behavior i, b_i denotes the strength of the behavioral belief associated with that behavior, e_i denotes the individual's evaluation of that behavioral belief, n_i denotes the strength of the normative belief associated with that behavior, m_i denotes the individual's motivation to comply with that normative belief, c_i denotes the strength of the control belief associated with that behavior and p_i denotes the perceived power of that control belief. Influencer messages on social media sites can affect all of these factors: prominent celebrities trusted by the public may change the public evaluation of these vaccines, a majority of popular figures sending messages one way or another may change people's normative beliefs, and influencers sharing their stories of the vaccination process may change individuals' perception of their self–efficacy (thereby affecting their control beliefs). This interaction is better explained by source credibility theory, which states the following:

1. The credibility of a message is dependent on the perceived credibility of the communicator.
2. Credibility is in turn dependent on the perceived honesty of the communicator and whether or not the audience considers them to be a valid source of information (Ohanian 1990).

Therefore, when social media users see posts by celebrities that they admire, or look up to, they trust the information provided and can be persuaded to believe in the celebrity's message. The more influence the account making the vaccine–related post has over the individual, the higher credibility the account has in that individual's lens; as a result, this study looks only at accounts with a large following (with "large" being more formally defined in "Data") on their social media platform since people tend to follow accounts/influencers whom they admire (Morton 2020).

The following section will discuss relevant literature to this topic and this study's contribution to the field. The process for collecting the tweet database and selecting influencers will then be described in "Data", and the models built to estimate each influencer's impact will be explained in "Models". The findings of this study will be examined in "Results" and "Discussion" and the conclusions and limitations of this study will be shared in "Conclusions and Future Work".

RELATED WORK

The emergence of social media platforms has revolutionized information dissemination and opinion formation on a global scale. Of particular importance is the role of platforms like Twitter in shaping public sentiment towards critical health issues, such as vaccine hesitancy. A study by Salathé and Khandelwal (2011) showed that analyzing sentiments on social media can help predict vaccination rates and even disease outbreaks in various communities. More recent studies, such as that by Bonnevie et al. (2020a) found that vaccine opposition posts increased significantly from late 2019 to mid–2020. However, since the first COVID–19 vaccine was not available until late 2020 (Webster 2021), the vaccine discussion on Twitter may have changed considerably since. Sutrave et al. (2021) also analyzed the sentiment on Twitter toward COVID–19 vaccines and found the overall sentiment on the platform was found to be positive, yet still with a considerable number of negative sentiment posts regarding the efficacy and adverse reactions caused by vaccines.

Many previous studies have shown that celebrities have the potential to play an important role in shaping the health landscape. An analysis by Hoffman and Tan (2013) concluded that celebrity messaging should be further examined as a tool to help public health authorities discredit misinformation and promote positive health practices. In a more specific context, the authors of Robinson (2003) examined AIDS in the African American community and found that in many cases, celebrities were indeed effective at promoting safe practices to avoid contracting HIV/AIDS. Alatas et al. (2019) examined a case study in Indonesia and found that certain elements of celebrity campaigns can make them very effective at promoting vaccinations and other positive health practices.

More recently, studies have also begun to look at how influencers fit into the vaccine sentiment landscape described earlier, although most focused on posts by anti-vaccine celebrities or posts sharing misinformation. Bonnevie et al. (2020b) showed that misinformation regarding vaccines has identifiable, upstream origins but has not been satisfactorily handled by health authorities. The authors of White et al. (2023) showed that most tweets referencing particular anti–vaccine celebrities during the COVID-19 pandemic had an overall negative tone and concluded that further research into the public response to celebrity messages is needed. One of the few studies to focus on both pro-vaccine and anti-vaccine influencers, Scannell et al. (2021) found that sharing references to celebrity statements was a common persuasive technique used by both pro–vaccine and anti–vaccine social media groups, although the impact of these references was not measured.

The above examples show that past research has looked at celebrity impacts on significant health decisions, examined how the public perception toward COVID–19 vaccines has changed over time, and even begun to study the origins and sentiments of anti–vaccine messages. However, there has been little research on studying the impacts of such messages from prominent figures, and especially on comparing the impacts of pro–vaccine and anti–vaccine influencers in the same context. This study aims to fill this gap, and in the process, provide valuable insight to public health officials as to how public health campaigns can be shaped by public figures.

DATA

In order to quantify the public opinion toward COVID–19 vaccines, messages regarding these vaccines posted on the Twitter platform (tweets) were collected using the Python package "snscrape". The following query was used: "*vaccine OR vaccines OR vaccinate OR vaccinated OR vaccination OR vaccinations OR vaccinated OR dose OR doses OR inoculate OR inoculation OR inoculations OR inoculated OR booster OR boosters OR vax OR anti–vax OR anti–vaxx OR antivax OR antivaxx OR antivaxxer OR anti–vaxxer OR antivaxer OR "breakthrough infection" OR "breakthrough infections" OR moderna OR pfizer OR biontech OR astrazeneca*" The following three constraints were applied on the query:

1. All collected tweets must have been written in English.
2. Tweets must have been posted in the United States.
3. The tweets included were allowed to include media, but in that case, the tweet must also have had some textual content that matched the query above.

By the end of the scraping process, 1,948,191 were scraped from as early as November 1, 2020, through to March 11, 2022 (with an upper limit of 4,133 tweets per day to ensure that data collected was distributed as evenly as possible throughout the relevant time period; otherwise, impacts would be tough to quantify since influencers' posts may have over or under–estimated impacts due to more or less data collected in the following days). Once this data was collected, each tweet was cleaned by removing all hashtags, mentions, URLs, media, special formatting characters embedded in our database's tweet representation (e.g., "\n" for a newline, or "U+0026" for an ampersand), non–ASCII characters, emojis and punctuation.

DOI: doi:10.17705/3JSIS.00034

To classify a tweet as pro–vaccination or anti–vaccination, a sentiment score was assigned to each tweet based on its processed text; tweets sharing messages promoting vaccines tend to express positive sentiments, whereas tweets sharing messages against vaccinations tend to express more negative sentiments (the actual boundary between the sentiments is formally defined below). To determine the sentiment scores, the rule–based model VADER (Valence Aware Dictionary and sEntiment Reasoner) was used (Hutto and Gilbert 2014). VADER's approach uses a dictionary, known as a lexicon, that stores a list of lexical features (words, emojis and other elements commonly found in text) with a sentiment score between 4 and –4 assigned to each feature (with higher scores representing increasingly positive sentiments). The VADER lexicon is optimized for abbreviated phrases, emojis, slang and other elements commonly found in tweets or other social media platform posts. To determine the sentiment score of a tweet, the score for each feature (known as the valence score for that feature) is determined from the lexicon and then summed according to five heuristic rules (that take into account words such as "not" that can change the entire meaning of the phrase, punctuation such as "!" that can indicate a more intense feeling, words written in uppercase for emphases, etc.). The sum is then normalized to compute the compound score, which is denoted here as *C* according to the following function:

$$C = \frac{x}{\sqrt{x^2 + \alpha}} \tag{4}$$

where x = sum of each valence score and α = normalization constant, which is usually 15 (Hutto 2013). Note here that although each word's score can vary from –4 to 4, the score of a tweet ranges only from –1 to 1 after the sum of the scores of its words is normalized as shown above. Once these compound scores were computed, tweets were then labeled as follows:

1. "Positive" tweets, where $C > 0.2$, which are likely to express pro–vaccination sentiments.
2. "Negative tweets, where $C < -0.2$, are likely to share messages against vaccinations or mandates.
3. "Neutral" tweets, where $-0.2 \leq C \leq 0.2$, tend to be factual statements expressing no clear sentiments or opinions, and are of no interest to this study (although they can potentially be useful and can still ignite responses which are quite "positive" or "negative").

These boundary/threshold values were chosen by randomly sampling nearly 1000 tweets from our dataset and hand labeling each as pro–vaccine, anti–vaccine, or neutral. After trying a series of boundary values (from ±0.01 to ±0.5 in increments of 0.01), –0.2 and 0.2 were found to result in the highest correlation between the hand–labelled classifications and the VADER generated ones.

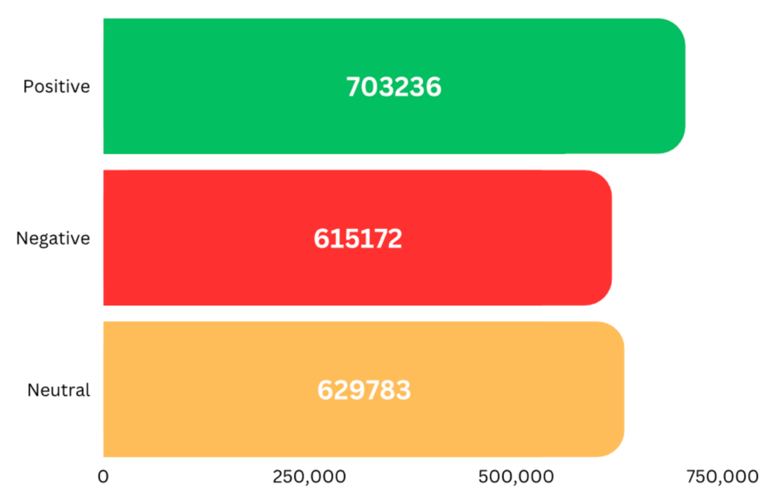

Figure 1. A chart showing the numbers of positive, negative and neutral tweets found in the database after applying the sentiment analysis. All three classifications had relatively similar numbers, with slightly more positive tweets than negative.

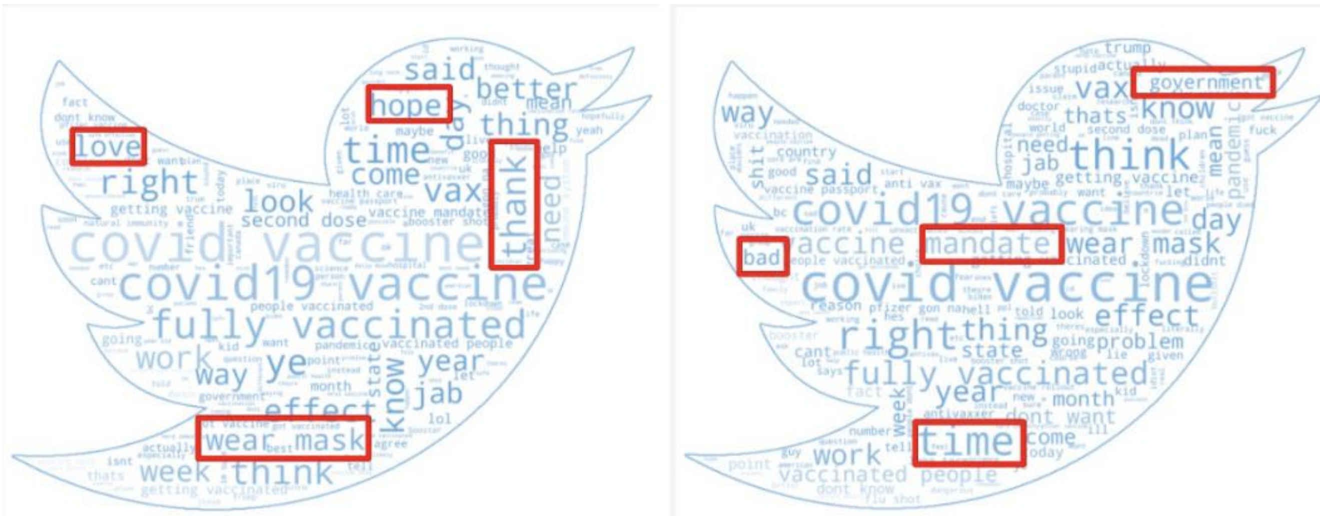

Figure 2. Word clouds displaying common words in both "positive" (left) and "negative" (right) tweets; the frequency of each word in the dataset is portrayed by its size and color intensity; some important differences are highlighted in red. These word clouds serve as a "proof of concept" that the sentiment analysis approach is fairly effective at identifying pro–vaccination and anti–vaccination tweets, as many of the differences highlighted are what would be expected between these two groups.

Two new datasets were built based on the collected tweets with one storing those with "positive" sentiment and one storing the "negative". Both datasets were formatted as time series of the number of tweets expressing each sentiment shared per day – this count was standardized and a log–transformation was applied. 36 calendar features were also added to each dataset, such as the month of year and the day of the week. These features encode information about seasonal patterns possibly present and aid the forecasting models (described in "Models") in better understanding the data – for instance, perhaps there tend to be more positive tweets on Saturdays. By the end of this process, both the positive and the negative datasets were 495 x 38 matrices, storing the date and tweet count with the 36 calendar features for 495 consecutive days.

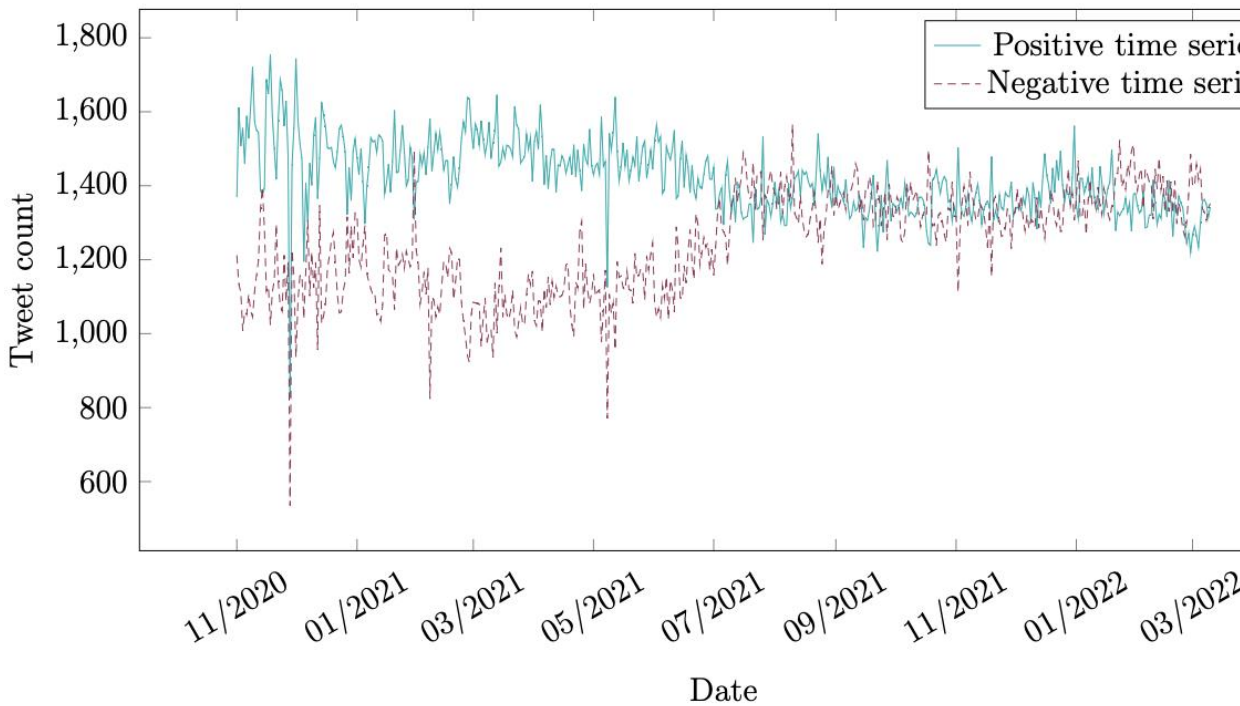

Figure 3. Plot showing the evolution of the number of positive (solid, in cyan) and negative tweets (dashed, in magenta) per day throughout the time window of the study.

The data collection phase also involved identifying influencers that have made statements in the past regarding COVID–19 vaccines that may have affected the public opinion. An *influencer* is, for the purposes of this study, defined as a social media account on either the Twitter or Instagram platforms (since many posts by prominent figures posted on Instagram heavily affected the conversation on Twitter, influencers were allowed to be on both platforms despite the responses being measured on Twitter) that meets the following criteria:

1. Wide reach: The account must have at least 100,000 followers on its respective platform and had been granted "verified" status on their platform as of the date of the study. According to source credibility theory (Hovland & Weiss 1952), the credibility of a message is directly dependent on the individual's perceived credibility of the individual sharing the message. The figures who tend to have a high perceived credibility among large segments of the population are celebrities and social media influencers (Hoffman & Tan 2015), and so we use the account's number of followers as a metric to determine their credibility and influence over the public.
2. Strong opinion: The account must have made a post relevant to COVID–19 vaccines, that, according to the same model described above, is categorized as either "pro–vaccine" or "anti–vaccine". If the post is categorized as "pro–vaccine", the influencer is said to be *positive* (regardless of whether or not other posts by the influencer express different opinions). Similarly, if the post is categorized as "anti–vaccine", the influencer is said to be *negative*.
3. Within window: The post described above must have initially shared between January 1, 2021, and February 20, 2022, inclusive, in order for the models (described in "Models") to have sufficient data before the post to build forecasts, as well as at least 20 days after the post to evaluate its impact.

4. No political involvement: The individual that the account represents must not be a politician or politically appointed, to minimize the effect of political polarization from affecting our results (of course, some political influence is unavoidable – some cases of this will be described in "Discussion"). This would exclude all U.S. senators, U.S. representatives and public health officials (such as Dr. Anthony Fauci).

This study tracked a total of eleven influential figures – six "positive" accounts: @SteveMartinToGo on Twitter, @BillGates on Twitter, @DollyParton on Twitter, @BigBird on Twitter, @johnlegend on Instagram and @americaferrera on Instagram, as well as five "negative" accounts: @evangelinelillyofficial on Instagram, @RobSchneider on Twitter, @chethanx on Instagram, @NICKIMINAJ on Twitter and @doutzen on Instagram. @SteveMartinToGo, @BillGates, @BigBird and @NICKIMINAJ were the only four influencers found within the tweet database that also met the three criteria above, and so the other seven influencers were chosen by examining well–publicized vaccine–related statements. The number of posts chosen and the manner of choosing were done similarly to White et al. (2023).

MODELS

To model the impact of a post by an influencer on the recorded sentiment scores, the following process was adopted:

1. Using only data prior to the selected influencer's post, a collection of 5 models (described below) were trained on the dataset matching the post's sentiment (i.e., models for a pro–vaccine post's impact are trained on the "positive" dataset, and vice versa).
2. Based on these models, 5 forecasts were generated for the next 20 days after the post (producing a separate forecast for each model). These forecasts are labelled *counterfactuals* i.e., estimates of what "would have happened" in the absence of the post. These counterfactual forecasts are, themselves, time–series with 20 time points (days) of data – they estimate the likely trajectory of public sentiment in the days following the post had it not been shared (for positive influencers, they represent a forecast of the number of positive tweets over the next 20 days, and vice versa).
3. The counterfactual along with the observed trajectory of the time series ("what actually happened" over the next 20 days) was then used as inputs to the "CausalImpact" *R* library to perform a Bayesian analysis on the two time–series and return distributions of the average and total difference between the counterfactual prediction (the predicted number of tweets matching the influencer's sentiment) and the observed values (the actual number of tweets matching the influencer's sentiment) over the next 20 days. It also computes the probability that the effect computed could have been observed due to chance.

With this methodology in mind, the impact analysis phase of this study boils down to a time–series forecasting problem to build the counterfactual estimates. Five different models were built, chosen by examining previous studies that benchmarked model performances on datasets similar to ours. Each model was evaluated based on its root mean squared error (RMSE). 80% of the data before the influencer's post was used to train the models, and the last 20% was used to evaluate their performance using the RMSE; after RMSE values were obtained for a model, the model was retrained on the full dataset (the 80% training set and the 20% testing set combined), and the forecasts for the period after the post were obtained. The RMSE was calculated in the following manner:

$$J = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}, \tag{5}$$

where $\hat{\boldsymbol{y}}$ denotes a 20–element vector of the model's counterfactual estimates, and $\mathbf{y}$ represents a 20–element vector of the true values of the tweet counts over the next 20 days. The following models were used to generate the counterfactuals (denoted y in the equations below):

1. A multiple linear regression model:

$$y = b_0 + b_1x_1 + b_2x_2 + b_3x_3 + \cdots + b_nx_n = b^{\intercal}\boldsymbol{x} \tag{6}$$

where $x_1, x_2, x_3, \ldots, x_n$ denote the regressors (in this instance, the 37 features described earlier). This approach was found to yield good results when used on a corpus of tweets similar to this study by Asul et al. (2010).

2. A Prophet model with regressors (Taylor and Letham 2018):

$$y = g(t) + s(t) + h(t) + \varepsilon(t) \tag{7}$$

where *g(t)* denotes non–period trends visible throughout the data, *s(t)* represents seasonal patterns, *h(t)* is used to account for the effects of holidays, and *ε(t)* denotes the error term, storing noise unaccounted for by the model. If the data seems to exhibit "non–linear, saturating growth" (which levels out at some carrying capacity *C*), *g(t)* can be most easily modeled by the equation:

$$g(t) = \frac{C}{1 + e^{-k(t-m)}} \tag{8}$$

where *C* denotes the carrying capacity, the growth rate is represented by *k*, and *m* denotes an offset. For non–logistic trends, below is a piece–wise linear trend model with *S* change–points (s_j, where $j = 1,\ldots,S$) at which the growth factor changes:

$$\boldsymbol{g}(\boldsymbol{t}) = (\boldsymbol{k} + \boldsymbol{a}(t)^{\intercal}\boldsymbol{\delta})t + (m + \boldsymbol{a}(t)^{\intercal}\boldsymbol{\gamma}), \tag{9}$$

where $\boldsymbol{a}(t) \in \{0,1\}^S$, such that $a_j(t) = 1$ for all $t \geq 1$ and $a_j(t) = 0$ otherwise. The change in the growth rate at s_j is denoted by δ_j and included in the vector $\boldsymbol{\delta}$; to ensure the function is continuous, $\boldsymbol{\gamma}$ is defined where $\gamma_j = -s_j\delta_j$. As before, the growth rate is represented by *k* and the offset parameter is denoted *m*. In this study, extra regressors were also added to the linear component of the model (the process for generating these regressors is described in more detail in the "Data" section). Due to the nature of the model, Prophet is exceptionally good at finding seasonal patterns in the data provided, as has been demonstrated with seasonal datasets in the past (Yasaman et al. 2022). As there is a high probability of seasonal patterns being present in this dataset, Prophet was chosen in this ensemble as well.

3. An Extreme Gradient Boosting (XGBoost) model:

Unlike the previous models, XGBoost builds an additive ensemble of *k* trees to compute the following prediction:

$$y = \phi(\boldsymbol{x}) = \sum_{k=1}^{K} f_k(\boldsymbol{x}), \qquad f_k \in \mathcal{F} \tag{10}$$

Here, CART (the space of trees), is represented by $\mathcal{F} = \{f(\boldsymbol{x}) = w_{q(x)}\}$, where $q : R^m \rightarrow T$ represents the tree structure, mapping a group of features to their leaf index, *T* represents the tree's leaf count and f_k represents a tree with structure *q* and weights $w \in R^T$. To update the model's functions based on previous examples, XGBoost works to minimize the following objective:

$$\mathcal{L}(\phi) = \sum_i l(\hat{y}_i, y_i) + \sum_k \left[\gamma T + \frac{1}{2}\lambda\left|\left|w\right|\right|^2\right], \tag{11}$$

and $\gamma T + \frac{1}{2}\lambda\left|\left|w\right|\right|^2$ is commonly abbreviated $\Omega(f)$. Here, $l(\widehat{y}_i, y_i) = \frac{1}{2}(\widehat{y}_i - y)^2$, but any differentiable l is acceptable. λ is the regularization parameter and γ controls the pruning of model, with both parameters working together to penalize extra complexity. This model is introduced and described extensively in Chen and Guestrin (2016), and various studies have shown that XGBoost performs at the level of, and often even better than, many state–of–the–art models for time–series forecasting problems (Alim et al. 2020; Gupta et al. 2022), including for social media time–series (Mubang and Hall 2022).

4. A Prophet model with gradient boosted errors:

   This model uses a two–pronged approach toward modeling the time series. First, a Prophet model as described above is used to approximately represent the time series (ignoring all extra regressors). Then, an XGBoost model is applied to the residual terms of the Prophet model. In this way, approaches (2) and (3) above are combined to generate a fourth model. This model is shown to perform better than many standalone, single models by Wang et al. (2022).

5. A ranked/weighted ensemble combining all four previous models based on their RMSE on the test set:

$$y = \theta_1 h_1 + \theta_2 h_2 + \theta_3 h_3 + \theta_4 h_4, \tag{12}$$

   where h_i represents the predictions of the ith model, and $\theta_i \in \{0.1; 0.2; 0.3; 0.4\}$ represents the weight of the ith model receives in the ensembled prediction. These weights were assigned in a manner where the model with the lowest RMSE (the best–performing model) receives a weight of 0.4, the model with the second lowest RMSE receives a weight of 0.3, the model with the second highest RMSE receives a weight of 0.2 and the model with the highest RMSE (the worst–performing model) receives a weight of 0.1 (so that $\sum \theta_i = 1$). This basic approach is shown and found to achieve higher accuracies than component models in Adhikari et al. (2015).

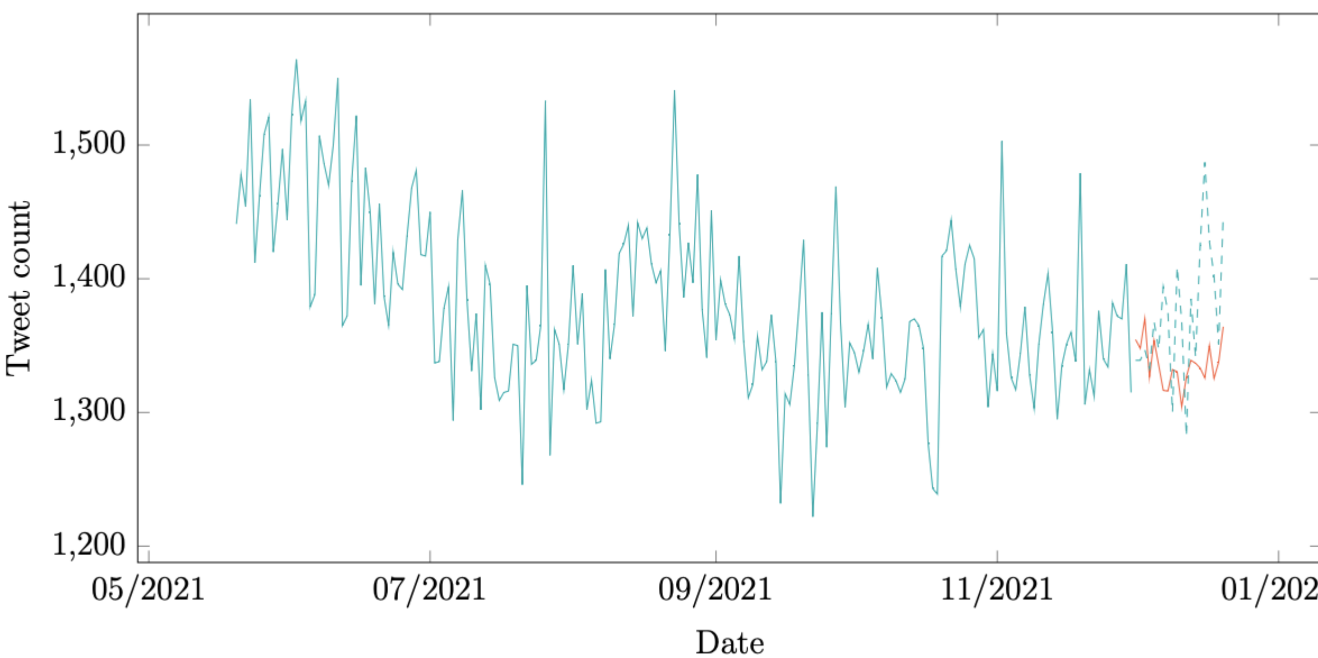

Figure 4. Plot showing the process to compute the impact for an imaginary influencer post on 12/01/2021 — all data points prior to this date (of which only the most recent 200 are shown in solid cyan for visibility purposes) would be used to train the models

DOI: doi:10.17705/3JSIS.00034

described above, which generate the counterfactual shown (in orange). It is then compared to the observed data (dashed cyan) to generate the impact estimates.

To perform the intervention analysis, two distinct time periods were defined: if d represents the date of an influence post regarding the vaccine, the 1) *pre–intervention period* is defined to run from November 1, 2020, to $d - 1$, and the 2) *post–intervention period* is defined to run from d to $d + 20$. Then, using the "CausalImpact" *R* package (Broderson et al. 2015), the information from these counterfactuals were combined with a Bayesian structural time series model. A structural, or state space, time series model takes the following form:

$$y_t = Z_t^T \alpha_t + \varepsilon_t \qquad \varepsilon_t \sim \mathcal{N}(0, H_t) \tag{13}$$

$$\alpha_{t+1} = T_t \alpha_t + R_t \eta_t \qquad \eta_t \sim \mathcal{N}(0, Q_t) \tag{14}$$

where y_t represents the observed data, and α_t represents the unobserved latent state. The parameters H_t and Q_t denote the variances of ε_t and η_t, respectively ($H_t = \sigma_\varepsilon^2$, $Q_t = \sigma_\eta^2$). The former of the two equations is referred to as the *observation equation* (as it depicts the relationship between the observed data and the latent state), while the latter is referred to as the *transition equation*, representing the relationship between lags of the unobserved state. A variety of models can be represented in state–space form, including ARIMA and Holt–Winters models. For more information on Bayesian structural time series and state space models, see Scott and Varian (2013). The priors for the Bayesian model were the distributions of model parameters, which were chosen to be Gaussian random walks with standard deviation 0.01 (This is the "CausalImpact" recommendation for relatively stable time series).

"CausalImpact" was used to estimate 3 values over the next 20 days: 1) the *average impact*, which was computed as the average daily difference between the observed number of tweets sharing the influencer's sentiment and models' estimates throughout post–intervention period; 2) the *total impact*, which was computed as the sum of the differences between the observed data and counterfactual estimate over the post–intervention period; and 3) the *p–value*, which represents the Bayesian one–sided tail area probability of obtaining the computed differences by chance.

RESULTS

Before interpreting the model results, it is important to make a few distinctions:

1. An influencer is said to have a *positive impact* when the estimated average and total impacts for an influencer are greater than 0. If the impact estimates are below 0, the influencer is said to have had a *negative impact*. This means that a "negative" influencer (based on the sentiment their vaccine–related post expresses) may have a positive impact (implying that the number of negative tweets increased overall in the post–intervention period), and vice versa.
2. An influencer is said to have had a *significant impact* when the p–value (Bayesian one–sided tail area probability) of their impact is below 0.050, i.e., significance value (α) = 0.05.

"Positive" influencers

The summary results (which are the results from the weighted ensemble described above) for the six "positive" influencers are shown below, with the model results for each individual influencer shown in Appendix A.

Influencer	RMSE	Average Impact	Total Impact	p–value
@SteveMartinToGo	0.785	620.02	12400.32	0.022

@BillGates	0.690	238.59	5671.80	0.001
@BigBird	0.665	12.10	241.97	0.013
@johnlegend	0.612	14.204	284.079	0.003
@DollyParton	0.709	51.317	1026.334	0.001
@americaferrera	0.614	45.685	913.692	0.001

Table 1. Weighted ensemble estimates for influencers sharing pro–vaccination messages.

For @SteveMartinToGo and @BillGates, all models agree that they have significant, positive impacts, albeit to varying degrees (for possible explanations for this variation, see Discussions).

The results for @BigBird and @johnlegend (see next page) were subject to slightly more variation than the preceding two influencers. For @BigBird, four out of five models seemed to agree on the tweet having a positive impact, and four out of five models also showed that the tweet had a significant impact. For @johnlegend, only three out of five models believed the post had a significant impact, although every model showed that the impact was positive (note however that the XGBoost counterfactual estimated an imperceptible impact of just 2 more pro–vaccine tweets over the 20–day period following his post).

"Negative" influencers

For each influencer in this category, all models predicted a significant impact (i.e., $p < 0.05$) — furthermore, each impact was fairly large in magnitude (note that not all of these impacts were in the positive direction, meaning that some influencers' anti–vaccine posts led to fewer anti–vaccine tweets over the next 20 days – possible explanations are explored in Discussions). Here, the weighted ensemble performed near the best in every case, as would be expected given that it is the most robust of all the models built. The summary results (which are the results from the weighted ensemble described above) for the five "negative" influencers are shown below, with the model results for each individual influencer displayed in Appendix A.

Influencer	RMSE	Average Impact	Total Impact	p–value
@evangelinelillyofficial	1.256	352.94	7058.87	0.001
@chethanx	0.905	–162.603	–3252.053	0.001
@RobSchneider	0.742	48.99	979.85	0.001
@NICKIMINAJ	0.487	–123.84	–2476.88	0.001
@doutzen	0.447	–89.974	–1799.484	0.001

Table 2. Weighted ensemble estimates for influencers sharing anti–vaccination messages.

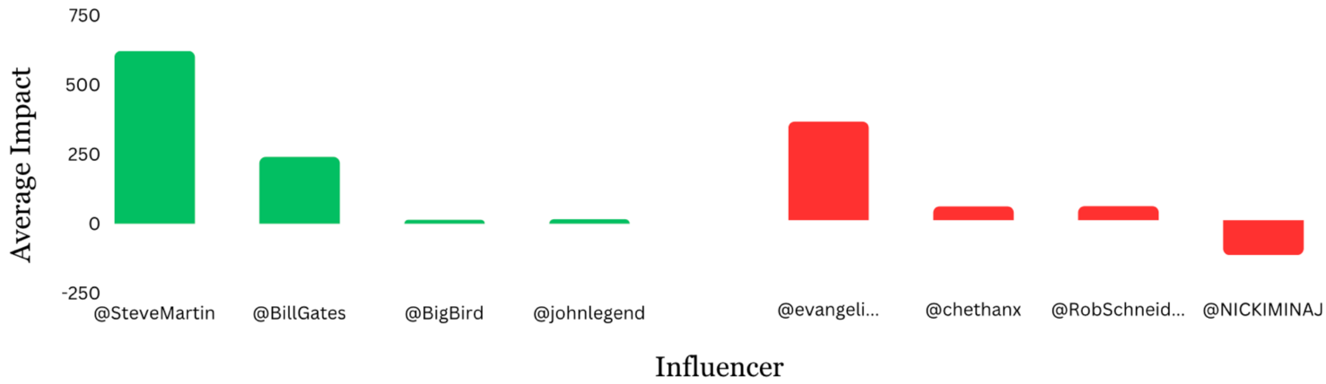

Figure 5. Plot showing the weighted ensemble's estimated average impacts for pro–vaccine influencers (left and in green) and anti–vaccine influencers (right and in red). The anti–vaccine influencer estimates are more volatile with @NICKIMINAJ even having a negative impact (i.e., there were roughly 123 fewer anti–vaccine tweets over the next 20 days as a result of her post).

DISCUSSION

Both @SteveMartinToGo and @BillGates were estimated to have statistically significant, positive impacts unanimously by all models. Based on these results, it seems as though at this stage of the vaccination campaign, influencer positive statements about the vaccines may have large, positive effects. It is interesting to note that, in both cases, the linear model predicted significantly larger values — on the order of 20 or more times larger — than the other three standalone models (the ensemble is excluded as it factors in the predictions of the linear model). This may be due to the fact that the positive time series displays a sharp downward trend towards the first quarter of the dataset, shortly before both influencers shared their posts. This trend has a heavier effect on the linear model than the others, leading to a more pessimistic counterfactual (far lower than other models), and consequently a larger difference. The linear models also had relatively high (worse) RMSEs in both incidents, with only the base Prophet model performing worse. The XGBoost model performed surprisingly well, surpassing even the Weighted Ensemble in both instances.

The results from @BigBird and @johnlegend were subject to slightly more variation among the models. Here, interestingly, the XGBoost models performed considerably worse than others, while the far simpler linear models had significantly better performances. In both instances, the weighted ensemble performed near the best; the ensemble is typically the most reliable for predictions based on longer–term data, as is the case here. However, in both instances, the impact estimates, regardless of the chosen model, were far lower than for @SteveMartinToGo and @BillGates. There may be many possible reasons for this difference, including that the latter two posts were shared toward the end of 2021, by which point the global focus on vaccines had decreased compared to the earlier months. Indeed, the estimated impacts only seem to decrease as one moves through the "positive" influencers chronologically.

@DollyParton and @americaferrera had results between those of the other four "positive" influencers discussed thus far. While their impacts were not as high as @SteveMartinToGo's or @BillGates', they were certainly higher and more consistent than those of @BigBird and @johnlegend. @DollyParton, in particular, had posted one of the most popular pro-vaccine tweets during the window of the study, with over 145,000 likes and nearly 30,000 reposts, which may have contributed to a result that is on neither extreme (since more popular tweets are likely to be seen by more people who agree and more people who disagree).

Overall, the results for the "negative" influencers were more volatile than the "positive" influencers, with the different models providing contrasting results for many of the influencers. However, the patterns found amongst these influencers imply that "negative" posts made by influencers did indeed affect the conversations Twitter users had regarding the vaccines. In fact, in almost every model in the "negative" section, the estimated cumulative impacts were on the order of 103 or greater. However, not all of these impacts were positive. In fact, in multiple instances, a "negative" post by an influencer was able to significantly increase the number of "positive" tweets in the following days. This phenomenon can be directly observed by reading selected "negative" posts by influencers on a selected platform — many users, in response to the statements in the original post, begin to provide their own counterarguments in favor of the vaccine. This generates a "counter–effect", where a "negative" post leads to a higher influx of "positive" tweets. Some influencers (specifically, @NICKIMINAJ and @doutzen) fell victim to this counter–effect according to all models, whereas others (e.g., @evangelinelillyofficial) had conflicting results between models. This counter–effect is, in fact, also present (to a far lesser degree) in "positive" influencers' posts as well (where users against the vaccines share their thoughts in response to the original post). Another point of interest is that in each category, the posts by the two influencers who performed worst (@BigBird and @NICKIMINAJ) had significantly more involvement from political figures than did the others (for instance, the account @POTUS replied to @BigBird's tweet) — from this, it seems that tweets with more political interference have a lessened impact.

CONCLUSIONS AND FUTURE WORK

This study aimed to evaluate the effect influencers had on the public sentiment of COVID–19 vaccines during the vaccine roll–out period. The public sentiment was determined using the VADER lexicon, and the impact was evaluated using Bayesian structural analysis with the help of the CausalImpact library. Based on the results, pro–vaccination messages sent out by influencers tend to increase the number of pro–vaccination tweets in the following days by measurable, statistically significant amounts. On the other hand, influencers that send out anti–vaccination or anti–mandate messages have large, statistically significant impacts, although sometimes in the opposite direction as intended.

Future studies in this area may attempt to quantify the diffusion rate of a statement on social media networks by using benchmarks such as the retweet/like count, and then use this value as an additional feature for each model. After all, a post that diffuses through networks rapidly enough to reach considerable viewership levels within a few hours will likely have a larger impact than a post that takes days. Of course, one limitation of this work is that we only examined eleven influencers, and future work should add more influencers to the analysis, to develop a more comprehensive picture of the impact. Finally, this study did not establish any form of causality — although the CausalImpact library was used, the counterfactuals were not built using synthetic controls (as the name of the library suggests), but by using forecasted predictions based solely on the prior data. Designing a randomized experiment to determine whether or not a post actually causes the observed differences should be explored in the future. For instance, an experiment may show individuals tweets that either support, or don't support, COVID-19 vaccination and then measure changes in knowledge and attitudes.

This paper's findings have a significant impact on the information systems (IS) community and the world at large. Firstly, these findings could provide valuable insights into social media's dynamics (i.e., how information spreads and how sentiment is shaped on online social networks); Also, the analysis of public sentiment from a large database is an important problem in IS (Yue et al. 2018), and the techniques used and presented here could be valuable to IS researchers. The results of this study also provide a clearer picture to IS practitioners and the world at large of the impact prominent figures have on the public perception of vaccines. These findings will also help inform IS practitioners who are communicating public health strategy findings to health officials in better understanding vaccine hesitancy; this will be crucial as information — and misinformation — spread even faster with the growth of social media.

REFERENCES

1. Adams, S. H., Schaub, J. P., Nagata, J. M., Park, M. J., Brindis, C. D., & Irwin Jr, C. E. (2021). Young adult perspectives on COVID–19 vaccinations, *Journal of Adolescent Health*, 69, 3, 511–514. https://doi.org/10.1016/j.jadohealth.2021.06.003
2. Adhikari, R., Verma, G. Khandelwal, I. (2015). A Model Ranking Based Selective Ensemble Approach for Time Series Forecasting, *Procedia Computer Science*, 48, 2015, 14–21. https://doi.org/10.1016/j.procs.2015.04.104
3. Ajzen, I. (1991). The Theory of Planned Behavior. Organizational Behavior and Human Decision Processes, 50, 179–211.
4. Afifi, T. O., Salmon, S., Taillieu, T., Stewart–Tufescu, A., Fortier, J., & Driedger, S. M. (2021). Older adolescents and young adults willingness to receive the COVID–19 vaccine: Implications for informing public health strategies. *Vaccine*, 39, 26, 3473–3479. https://doi.org/10.1016/j.vaccine.2021.05.026
5. Alatas, V., Chandrasekhar, A. G., Mobious, M., Olken, B. A. & Paladines, C. (2019). When Celebrities Speak: A Nationwide Twitter Experiment Promoting Vaccination in Indonesia. National Bureau of Economic Research, Cambridge, Massachusetts, *NBER Working Paper Series*, 25589. https://doi.org/10.3386/w25589
6. Alim, M., Ye, G. H., Guan, P, Huang, D. S., Zhou, B. S. & Wu, W. (2020). Comparison of ARIMA model and XGBoost model for prediction of human brucellosis in mainland China: a time– series study, *BMJ Open*, 10, e039676. https://doi.org/ 10.1136/ bmjopen–2020–039676
7. Asul, S. & Huberman, B. A. (2010). Predicting the Future With Social Media. *2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, 1, 2010, 492–499. https://doi.org/10.1109/WI–IAT.2010.63
8. Bandura, A. (1977). Self–Efficacy: Toward a Unifying Theory of Behavioral Change. *Psychological Review*. 84, 2, 191–215. https://doi.org/10.1037/0033–295X.84.2.191
9. Broderson, K. H., Galluser, F., Koehler, J., Remy, N. & Scott, S. L. (2015). Inferring causal impact using Bayesian structural time–series models. *The Annals of Applied Statistics*, 9, 1, 247–274. https://doi.org/10.1214/14–AOAS788
10. Bloom, D. E., Cadarette, D., Ferranna, M., Hyer, R. N., & Tortorice, D. L. (2021). How New Models Of Vaccine Development For COVID–19 Have Helped Address An Epic Public Health Crisis, *Health Affairs*, 40, 3, 410–418. https://doi.org/ 10.1377/hlthaff.2020.02012
11. Bonnevie, E., Gellegos-Jeffrey, A., Byrd, B. & Smyser, J. (2020a). Quantifying the rise of vaccine opposition on Twitter during the COVID-19 pandemic, *Journal of Communication in Healthcare*, 14, 1, 12-19.
12. Bonnevie, E., Goldbarg, J., Gallegos-Jeffery, A. K., Rosenberg, S. D., Wartella, E. & Smyser, J. (2020b). Content Themes and Influential Voices Within Vaccine Opposition on Twitter, 2019. *American Journal of Public Health*, 110, S3, S326-S330. https://doi.org/10.2105/AJPH.2020.305901
13. Calac, A. J., Haupt, M. R., Li, Z. & Mackey, T. (2022). Spread of COVID–19 Vaccine Misinformation in the Ninth Inning: Retrospective Observational Infodemic Study. *JMIR Infodemiology*, 2, 1, e33587. https://doi.org/10.2196/33587
14. Carvalho, T., Krammer, F. & Iwasaki, A. (2021). The first 12 months of COVID–19: a timeline of immunological insights, *Nature Reviews Immunology*, 21, 245–256. https://doi.org/10.1038/s41577–021–00522–1
15. Chansolme, N. M. (2022). Pivoting from Tradition: An Analysis of Influencer Communication in COVID–19 Vaccination Campaigns, *TCU Digital* Repository, 211–58

16. Chen, T. & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1, 785–794. https://doi.org/10.1145/2939672.2939785
17. Fishbein, M. A theory of reasoned action: some applications and implications. *Nebraska Symposium on Motivation*, 27, 65-116.
18. Fishbein, M. & Ajzen, I. (1975). Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research. Reading, MA: Addison–Wesley
19. Gupta, R., Yadav, A. K., Jha, S. K., & Pathak, P. K. (2022). Time Series Forecasting of Solar Power Generation using Facebook Prophet and XGBoost. *2022 IEEE Delhi Section Conference (DELCON)*, 2022, 1–5. https://doi.org/10.1109/DELCON54057.2022.9752916
20. Hoffman, S. J. & Tan, C. (2013). Following celebrities' medical advice: meta-narrative analysis, *BMJ*, 2013, 347. https://doi.org/10.1136/bmj.f7151
21. – & –. (2015). Biological, psychological and social practices that explain celebrities' influence on patients' health-related behaviors, *Archives of Public Health*, 73, 1, 3. https://doi.org/10.1186/2049-3258-73-3
22. Hovland, C. I. & Weiss, W. (1952). The Influence of Source Credibility on Communication Effectiveness. *The Public Opinion Quarterly*, 15, 4, 635-650.
23. Jomnonkwao, S., Uttra, S. & Ratanavaraha, V. (2020). Forecasting Road Traffic Deaths in Thailand: Applications of Time–Series, Curve Estimation, Multiple Linear Regression, and Path Analysis Models, *Sustainability*, 2020, 12, 395. https://doi.org/10.3390/su12010395
24. MacDonald, N. E. & SAGE Working Group on Vaccine Hesitancy (2015). Vaccine hesitancy: Definition, scope and determinants. *Vaccine*, 33, 34, 4161–4164. https://doi.org/10.1016/j.vaccine.2015.04.036
25. Morton, F. (2020). Influencer marketing: An exploratory study on the motivations of young adults to follow social media influencers, *Journal of Digital & Social Media Marketing*, 8, 2, 156–165.
26. Mubang, F. & Hall, L. Simulating New and Old Twitter User Activity with XGBoost and Probabilistic Hybrid Models. *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2022, 132-139. https://doi.org/10.1109/ICMLA55696.2022.00026.
27. Ohanian R. Construction and validation of a scale to measure celebrity endorsers' perceived expertise, trustworthiness, and attractiveness. *Journal of Advertising*. 1990, 19, 39–52. https://doi.org/10.1080/00913367.1990.10673191
28. Park, J. W., Lagniton, P.N.P., Liu, Y. & Xu, R.H. (2021). mRNA vaccines for COVID–19: what, why & how, *International Journal for Biological Sciences*, 17, 6, 1446–1460.
29. Randolph, H. E. & Barreiro, L. B. (2020). Herd Immunity: Understanding COVID–19, *Immunity*, 52, 5, 737–741. https://doi.org/10.7150/ijbs.59233
30. Robinson, W. A. (2003). AIDS in the African-American community: celebrities as an ineffective communication strategy, *Theses and Dissertations*, 1367. https://rdw.rowan.edu/etd/1367
31. Rzymski, P., Borkowski, L., Drag, M., Flisiak, R., Jemielity, J., Krajewski, J., Mastalerz–Migas, A., Matyja, A., Pyr´c, K., Simon, K., Sutkowski, M., Wysocki, J., Zajkowska, J. & Fal, A. (2021), The strategies to support the COVID–19 vaccination with evidence–based communication and tackling misinformation. *Vaccines*, 9, 2, 109. https://doi.org/10.3390/vaccines9020109
32. Salathé, M. & Khandelwal, S. (2011). "Assessing Vaccination Sentiments with Online Social Media: Implications for Infectious Disease Dynamics and Control." *PLOS Computational* Biology, 7, 10, e1002199. https://doi.org/10.1371/journal.pcbi.1002199
33. Scannell, D., Desens, L., Guadagno, M., Tra, Y., Acker, E., Sheridan, K., Rosner, M., Mathieu, J. & Fulk, M. (2021). COVID-19 Vaccine Discourse on Twitter: A Content Analysis of Persuasion Techniques, Sentiment and Mis/Disinformation. *Journal of Health Communication*, 26, 7, 443-459. https://doi.org/10.1080/10810730.2021.1955050
34. Selkie, E. (2022). Influence at the Intersection of Social Media and Celebrity. *JAMA Network Open*, 5, 1, e2143096. https://doi.org/10.1001/jamanetworkopen.2021.43096

35. Ssanyu, J. N., Kiguba, R., Olum, R., Kiguli, J., Kitutu, F. E. (2022). Using community influencer groups to address COVID–19 misinformation and vaccine hesitancy in Uganada: a protocol for a prospective quasi–experimental study. *BMJ Open*, 12, 8, e057994.
36. Sutrave, K., Godasu, R. & Liu, J. (2021). Understanding the Public Sentiment and Discourse on COVID-19 Vaccine. *AMCIS 2021 Proceedings*, 24.
37. Sweileh, W. M. (2020). Bibliometric analysis of global scientific literature on vaccine hesitancy in peer–reviewed journals (1990–2019), *BMC Public Health*, 20, 1–15. https://doi.org/10.1186/s12889–020–09368–z
38. Taylor, S. J. & Letham, B. (2018). Forecasting at scale. *The American Statistician*, 72, 1, 37–45. https://doi.org/10.1080/00031305.2017.1380080
39. Wang, J., Du, X. & Qi, X. (2022). Strain prediction for historical timber buildings with a hybrid Prophet–XGBoost model, *Mechanical Systems and Signal Processing*, 179, 1, 109316. https://doi.org/10.1016/j.ymssp.2022.109316
40. Webster, P. (2021). COVID–19 timeline of events. *Nat. Med.*, 27, 12, 2054–2055. https://doi.org/10.1038/s41591–021–01618–w
41. White, B. M., Melton, C., Zareie, P., Davis, R. L., Bednarczyk, R. A., & Shaban–Nejad, A. (2023). Exploring celebrity influence on public attitude towards the COVID–19 pandemic: social media shared sentiment analysis. BMJ Health & Care Informatics, 30, 1, e100665. https://doi.org/10.1136/bmjhci–2022–100665
42. Yannopoulos, P. (2012). Celebrity advertising: Literature review and propositions, *World Review of Business Research*, 2, 4, 24–36.
43. Yasaman, E., Amin, S. H., Zhang, G. & Shah, B. (2022). Time–series forecasting of seasonal items sales using machine learning – A comparative analysis, *International Journal of Information Management and Data Insights*, 2, 1, 100058. https://doi.org/10.1016/j.jjimei.2022.100058
44. Yue, L., Chen, W., Li, X., Zuo, W. & Yin, M. (2018). A survey of sentiment analysis in social media, *Knowledge and Information Systems*, 60, 617–663. https://doi.org/10.1007/s10115–018–1236–4
45. Zhang, L., Bian, W., Qu, W., Tuo, L. & Wang, Y. (2021). Time series forecast of sales volume based on XGBoost, *Journal of Physics: Conference Series, 2021 2nd International Workshop on Electronic communication and Artificial Intelligence*, 1873. https://doi.org/10.1088/1742–6596/1873/1/012067
46. Zhu, N., Zhang, D., Wang, W., Li, X., Yang, B., Song, J., Zhao, X., Huang, B., Shi, W., Lu, R., Niu, P., Zhan, F., Ma, X., Wang, D., Xu, W., Wu, G., Gao, G., Tan, W. & China Novel Coronavirus Investigating & Research Team (2020). A Novel Coronavirus from Patients with Pneumonia in China, 2019, *New England Journal of Medicine*, 382, 8, 727–733. https://doi.org/ 10.1056/NEJMoa2001017

APPENDIX A

The following eleven tables display more detailed results than those discussed above, with each celebrity's model results displayed individually.

"Positive Influencers"

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	0.891	1242.55	24850.95	0.001
Prophet with Regressors	1.005	64.17	1283.49	0.001
XGBoost	0.752	44.39	887.80	0.001
Prophet with Boosted Errors	0.845	22.36	447.30	0.022
Rank Ensemble	0.785	620.02	12400.32	0.013

Table 3. Model results for tweet by @SteveMartinToGo

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	1.111	864.52	17290.35	0.001
Prophet with Regressors	1.183	43.92	878.34	0.001
XGBoost	0.520	38.34	766.74	0.001
Prophet with Boosted Errors	0.838	53.51	1066.18	0.022
Rank Ensemble	0.690	283.59	5671.80	0.001

Table 4. Model results for tweet by @BillGates

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	0.649	6.85	136.99	0.096
Prophet with Regressors	0.694	20.04	400.71	0.001
XGBoost	0.819	–17.72	–354.35	0.001
Prophet with Boosted Errors	0.796	19.88	397.56	0.001
Rank Ensemble	0.665	12.10	241.97	0.013

Table 5. Model results for tweet by @BigBird

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	0.675	9.19	183.87	0.001
Prophet with Regressors	0.609	24.08	481.53	0.001
XGBoost	0.786	0.105	2.106	0.488

Prophet with Boosted Errors	0.713	7.50	150.09	0.085
Rank Ensemble	0.612	14.204	284.079	0.003

Table 6. Model results for Instagram post by @johnlegend.

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	0.675	71.076	1421.517	0.001
Prophet with Regressors	0.609	37.586	751.734	0.001
XGBoost	0.786	42.166	843.322	0.001
Prophet with Boosted Errors	0.713	23.952	479.043	0.007
Rank Ensemble	0.612	14.204	284.079	0.001

Table 7. Model results for tweet by @DollyParton

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	0.727	42.461	849.210	0.001
Prophet with Regressors	0.581	64.379	1287.581	0.001
XGBoost	0.792	–11.191	–223.815	0.015
Prophet with Boosted Errors	0.737	40.187	803.738	0.001
Rank Ensemble	0.614	45.685	913.692	0.001

Table 8. Model results for Instagram post by @americaferrera

"Negative Influencers"

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	1.294	782.43	15648.65	0.001
Prophet with Regressors	1.364	–622.61	–12452.21	0.001
XGBoost	1.714	–40.29	805.86	0.001
Prophet with Boosted Errors	1.620	–104.02	–2080.51	0.001
Rank Ensemble	1.256	352.94	7058.87	0.001

Table 9. Model results for Instagram post by @evangelinelillyofficial.

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	0.783	–63.852	–1277.048	0.001
Prophet with Regressors	1.106	–221.564	–4431.271	0.001
XGBoost	1.112	–109.019	–2180.387	0.001

Prophet with Boosted Errors	1.472	–181.980	–3639.604	0.001
Rank Ensemble	0.905	–162.603	–3252.053	0.001

Table 10. Model results for (now deleted) Instagram post by @chethanx.

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	0.803	–90.15	1802.92	0.001
Prophet with Regressors	1.034	–104.58	–2091.64	0.001
XGBoost	1.188	142.44	2848.74	0.001
Prophet with Boosted Errors	1.404	119.96	2399.19	0.001
Rank Ensemble	0.742	48.99	979.85	0.001

Table 11. Model results for tweet by @RobSchneider.

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	0.482	–100.15	–2003.04	0.001
Prophet with Regressors	0.727	–264.97	–5299.39	0.001
XGBoost	0.802	–24.87	–497.42	0.004
Prophet with Boosted Errors	0.712	–104.53	–2090.60	0.001
Rank Ensemble	0.487	–123.84	–2476.88	0.001

Table 12. Model results for tweet by @NICKIMINAJ.

Model	RMSE	Average Impact	Total Impact	p–value
Multiple Linear Regression	0.519	–79.264	–1585.278	0.001
Prophet with Regressors	0.954	–258.359	–5167.185	0.001
XGBoost	0.599	–16.938	–338.761	0.017
Prophet with Boosted Errors	0.561	–102.15	–2043.14	0.001
Rank Ensemble	0.447	–89.974	–-1799.484	0.001

Table 13. Model results for Instagram post by @doutzen.