

UC Irvine

UC Irvine Previously Published Works

Title

Understanding underlying moral values and language use of COVID-19 vaccine attitudes on twitter.

Permalink

<https://escholarship.org/uc/item/5t78v9mp>

Journal

PNAS Nexus, 2(3)

Authors

Borghouts, Judith

Huang, Yicong

Gibbs, Sydney

et al.

Publication Date

2023-03-01

DOI

10.1093/pnasnexus/pgad013

Peer reviewed

Understanding underlying moral values and language use of COVID-19 vaccine attitudes on twitter

Judith Borghouts¹^{a,*}, Yicong Huang^b, Sydney Gibbs^b, Suellen Hopfer^c, Chen Li^b and Gloria Mark^b

^aDepartment of Medicine, University of California Irvine, Irvine, CA 92617, USA

^bDepartment of Computer Science, University of California Irvine, Irvine, CA 92697, USA

^cDepartment of Health, Society & Behavior in the Program in Public Health, University of California Irvine, Irvine, CA 92697, USA

*To whom correspondence should be addressed: Email: jborghou@uci.edu

Edited By: Jay Van Bavel

Abstract

Public sentiment toward the COVID-19 vaccine as expressed on social media can interfere with communication by public health agencies on the importance of getting vaccinated. We investigated Twitter data to understand differences in sentiment, moral values, and language use between political ideologies on the COVID-19 vaccine. We estimated political ideology, conducted a sentiment analysis, and guided by the tenets of moral foundations theory (MFT), we analyzed 262,267 English language tweets from the United States containing COVID-19 vaccine-related keywords between May 2020 and October 2021. We applied the Moral Foundations Dictionary and used topic modeling and Word2Vec to understand moral values and the context of words central to the discussion of the vaccine debate. A quadratic trend showed that extreme ideologies of both Liberals and Conservatives expressed a higher negative sentiment than Moderates, with Conservatives expressing more negative sentiment than Liberals. Compared to Conservative tweets, we found the expression of Liberal tweets to be rooted in a wider set of moral values, associated with moral foundations of care (getting the vaccine for protection), fairness (having access to the vaccine), liberty (related to the vaccine mandate), and authority (trusting the vaccine mandate imposed by the government). Conservative tweets were found to be associated with harm (around safety of the vaccine) and oppression (around the government mandate). Furthermore, political ideology was associated with the expression of different meanings for the same words, e.g. “science” and “death.” Our results inform public health outreach communication strategies to best tailor vaccine information to different groups.

Keywords: moral foundations theory, social media, vaccination, COVID-19, political ideology, sentiment

Significance Statement

Social media has been frequently used as a platform for public discourse around vaccinations. We used moral foundations theory (MFT) to understand the moral underpinnings of how people with different political ideologies express attitudes toward the COVID-19 vaccine on Twitter. We find that people with more extreme ideologies share more similarities in vaccine attitudes—more negative sentiment—compared to people with a moderate ideology who exhibit more positive sentiment. Although the vaccine attitudes of Liberals stem from a wider set of moral values of care, authority, liberty, and fairness, the attitudes of Conservatives are predominantly rooted in the moral values of oppression and harm. This study is the first to examine political ideology, vaccine attitudes, and social media through the lens of MFT. These results have important implications for public health messaging, as Liberals and Conservatives may be more accepting of information around vaccination that align with their moral values. This knowledge can inform outreach communication strategies to best tailor vaccine information that resonates with different groups.

Introduction

As of September 2022, more than 616 million people have been diagnosed with COVID-19 worldwide, with more than 6.5 million reported deaths [1]. The infection rate and high contagion highlight the importance of an effective vaccine to prevent the further spread of the virus and mitigate effects for those infected.

However, hesitancy to get vaccinated can hamper the wide-scale adoption needed to achieve herd immunity, and vaccine hesitancy has contributed to disease outbreaks in the past [2]. A survey in May 2020 among 1,056 US respondents indicated that only half of respondents planned to get vaccinated against

COVID-19 if a vaccine became available to them [3], and a longitudinal study between May and September 2020 showed that people’s intention to vaccinate decreased over time [4]. A recent survey conducted in April 2022 in the United States found that one in four adults remains unvaccinated, including one in six who say they will definitely not get vaccinated [5].

A major contributor to vaccine hesitancy is the dissemination of antivaccination information [6]. In particular, social media platforms such as Twitter have become a frequently used and influential platform to express opinions and disseminate information and misinformation related to vaccination. Prior work found that

Competing Interest: The authors declare no competing interest.

Received: June 6, 2022. **Revised:** December 23, 2022. **Accepted:** January 9, 2023

© The Author(s) 2023. Published by Oxford University Press on behalf of National Academy of Sciences. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial reproduction and distribution of the work, in any medium, provided the original work is not altered or transformed in any way, and that the work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

around 30 to 60% of vaccine information on social media contained antivaccination content [7], and that tweets expressing negative sentiment toward the vaccine and/or misinformation attracted more public engagement [8]. Common factors influencing people's attitude toward the COVID-19 vaccine include the perceived severity of COVID-19 [9, 10], perceived safety of the vaccine [11, 12], and trust in regulatory authorities [10].

Political ideology

Multiple survey studies have found a link between an individual's political ideology and vaccine attitudes. Across studies in the United States, people who identified as Conservative were significantly less likely to report their intention to be vaccinated [3, 13, 14].

In the United States, Conservatives typically value tradition, whereas Liberals place more emphasis on equality [15]. While the US-based surveys mainly found a difference in COVID-19 vaccine attitudes between Liberal and Conservative individuals, a study in France found a different trend [10]. Rather than a Left-Right axis, respondents who identified with radical parties and those who did not identify with any party (and did not vote in the most recent presidential elections) were significantly more likely to refuse the vaccine.

Moral foundations theory to understand vaccination attitudes

To better understand differences between people's beliefs and opinions, previous studies have used moral foundations theory (MFT), which proposes five moral foundations upon which people base judgments and decisions [16]. Each of these foundations is framed as a polarity: (1) Care/Harm relates to caring for others and minimizing harm; (2) Fairness/Cheating relates to equality and avoiding selfishness; (3) Loyalty/Betrayal relates to loyalty among group members; (4) Authority/Subversion relates to respect and obedience to authorities and respecting hierarchy; and (5) Sanctity/Degradation relates to purity and avoiding contamination of the body or mind, such as unnatural toxins entering one's body. More recent work has started including a sixth domain, (6) Liberty/Oppression, which relates to a feeling of being dominated and restricted in liberty. Research linking moral foundations with political ideologies has found that political Liberals care most strongly about Harm and Fairness, whereas Conservatives care about Loyalty, Authority, and Purity in addition to Harm and Fairness [17, 15].

MFT was developed to explain differential reactions to societal topics that became politicized in the United States (e.g. abortion, gun control, stem cell research, immigration, death penalty, climate change, vaccination). More extreme conservative and liberal political ideologies have politicized vaccine policies, particularly mandates. Former President Trump set in motion politicizing COVID-19 vaccination during the pandemic when he did not publicly endorse vaccination initially and promoted adopting alternative unproven treatments [18]. Political rhetoric invokes moral judgments to reframe vaccination as the freedom to choose one's own health decisions (not the government or health authorities) including vaccination.

MFT has explained how deeply rooted intuitions or moral values influence and explain differential vaccine attitudes [19–21]. A survey study that used MFT to understand childhood vaccine attitudes among parents found that those rejecting vaccines endorsed moral foundations of Liberty, Harm, and Purity, whereas those accepting vaccines had a moral preference for Authority (indicating a trust for authorities that provide and recommend vaccination services) [22]. When thinking of common factors influencing vaccine

attitudes, perceived safety of the vaccine and severity of COVID-19 may be related to the foundations Harm and Purity, whereas trust in regulatory authorities is related to Authority. Although Rossen et al.'s study [22] considered political ideologies, they did not find an association between antivaccination attitudes and ideology.

MFT has also been used to understand moral values expressed on Facebook pages in support or opposition of (non-COVID-19) vaccination [23]. Pages that defended vaccines tended to focus on the value of the family (related to the moral foundation Care), whereas vaccine hesitancy pages were more focused on the value of freedom (related to the moral foundation Liberty). These findings suggest that differences in vaccine attitudes may be reflected in the type of topics and arguments used.

MFT has shown that people's world views and realities can be very different for different ideologies. In this article, we use MFT to help explain underlying moral values for different vaccine attitudes according to political ideology. Previously, MFT has been used to understand differences in moral values for political ideologies on the one hand [24], and the theory has been used to look at differences in moral values for provaccine and antivaccine attitudes on the other hand. However, to the best of our knowledge, a study is lacking bringing these elements together, to understand the link between political ideology and vaccine attitudes through the lens of MFT. This is important because Liberals and Conservatives may show a preference for different moral foundations in the specific context of the COVID-19 vaccine and accept information and beliefs that align with their moral values. In contrast, if information about taking the vaccine is at odds with the audience's cultural and moral values, attempts to disseminate this information, say by public health agencies, may not resonate and instead may make people even less likely to get vaccinated [25].

Understanding differences in language use

Attitudes toward certain societal issues can be further reflected in someone's language use [26]. The lexical hypothesis [27] is a theoretical framework that proposes that individual characteristics such as personality can influence language. For example, personality traits can predict one's language use both on social media [28] and in offline communication [29]. In addition, political ideology has been found to affect language use, shown with politicians [30, 31]: conservative politicians in US Congress have been found to use more language related to religion, tradition, power, and threat, whereas liberal politicians use more language related to social issues, benevolence, and achievements. We therefore further explore how people's political ideology and attitudes toward the COVID-19 vaccine are reflected in their word choice on Twitter.

Using Twitter to understand topics and sentiment

Twitter has frequently been used as a platform to discuss information related to vaccinations [32] and to share viewpoints on social issues [33]. As such, Twitter can be a useful source to gauge sentiment and topics discussed in relation to COVID-19 vaccines, along with the overall expression of emotions and opinions during the pandemic. For example, a study analyzing COVID-19 tweets between December 13, 2019, and March 9, 2020, found COVID-19 sentiment to be associated with significant events: a negative sentiment increased toward the start of 2020 with the official announcement of COVID-19 [34]. While overall sentiment became more positive over time with more information about protection against COVID-19, peaks in negative sentiment corresponded to national announcements of social distancing and mask-wearing policies. Topics on Twitter concerning COVID-19

related to news reports, complaints, misconceptions on how to control COVID-19, and attitudes toward COVID-19 and vaccinations [34, 35]. These studies, however, were not specific to the vaccine and were not linked to Twitter users' individual characteristics, such as their political ideology.

In summary, while survey studies have highlighted differences between political ideologies in terms of their moral foundations, and previous work has looked at how political ideology is linked to differing vaccine attitudes, little work has examined how these differences are then expressed in online discourse to publicly express opinions on COVID-19 vaccines, especially in a public sphere like Twitter. We examine whether groups with differing political ideologies discuss the COVID-19 vaccine differently, including whether the same words might convey different meanings. This understanding is crucial in developing effective outreach and communication strategies to inform people about the vaccine.

Our research questions investigate differences in sentiment, moral values, and language use between political ideologies on the COVID-19 vaccine, using Twitter data as follows:

RQ1: Is there a relation between political ideology and sentiment toward the COVID-19 vaccine?

Previous survey studies have indicated that people who identified as Conservative had a more negative sentiment toward the COVID-19 vaccine [3, 13, 14]. In this article, we test whether this relationship holds true online using Twitter data. Instead of asking people to self-report their political orientation, we use a previously validated algorithm that estimates users' political ideology based on the Twitter accounts they follow [36, 37]. This method allows us to derive ideologies for users who may not explicitly state their political preference.

RQ2: What moral values underlie attitudes toward the COVID-19 vaccine, as expressed by different political ideologies?

Exposure to positive and negative sentiment toward the vaccine, especially on social media, can affect people's intention to get vaccinated [6]. As a public forum, opinions expressed on Twitter can be far-reaching. People's moral values can influence the narrative they use to support their stance on various issues [23, 38]. Different attitudes toward the COVID-19 vaccine can be potentially traced to underlying moral values. The Moral Foundations Dictionary (MFD) [39] was developed and validated to identify specific words that are associated with a moral foundation. A person's political ideology can influence what moral values appeal to them, with political Liberals generally caring most strongly about Harm and Fairness, whereas Conservatives also care about Loyalty, Authority, and Purity in addition to Harm and Fairness [15, 17]. Specifically, we use MFT to understand what moral values underlie disparate attitudes toward the COVID-19 vaccine.

RQ3: How does language use reflect differences in political ideologies in discussing the COVID-19 vaccine?

Language use can be informed by individual characteristics (i.e. the lexical hypothesis [27]). We therefore further explore how people's political ideology and attitudes toward the COVID-19 vaccine are reflected in their word choice. Word usage may develop different meanings in separate groups, and arguments using selected words may resonate differently and appeal to different audiences (e.g. using scientific research as an argument to take the vaccine may appeal to some but not other audiences).

Results

The relation between sentiment and political ideology

Our first research question addressed whether a relationship exists between political ideology and sentiment toward the COVID-19 vaccine. In plotting the ideology and sentiment scores (see Fig. 1), we observed a quadratic relationship between sentiment score and political ideology. Sentiment scores were more

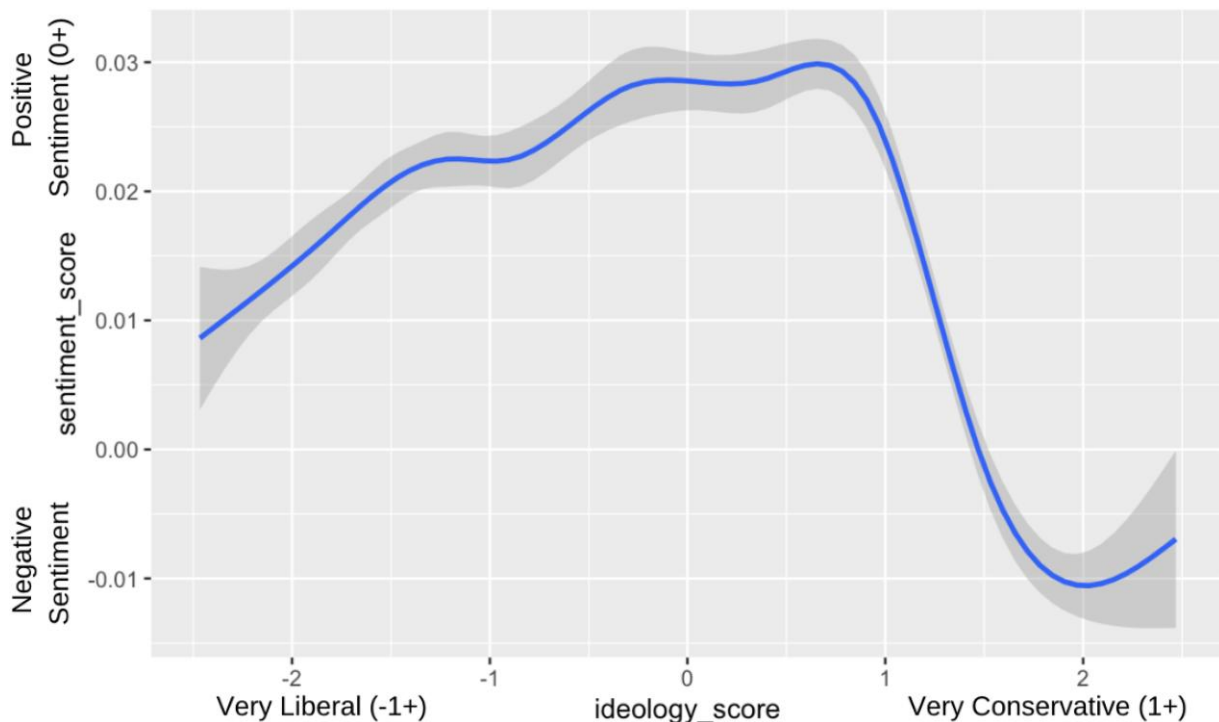


Fig. 1. A plot of ideology and sentiment scores.

negative for more extreme ideology scores (i.e. very Liberal and very Conservative users) in relation to more moderate ideology scores. However, those who are very Liberal (i.e. ideology scores < -1.5) showed a slightly positive sentiment, whereas those who are very Conservative (i.e. ideology scores > 1.5) showed sentiment that is negative toward the COVID-19 vaccine.

We therefore added a quadratic term to model the relationship between political ideology and sentiment. As sentiment toward the vaccine may differ by region [40], we controlled for the US regions (see Fig. S1 for a map of the regions we used). Table 1 shows the fit indices for the model with and without the quadratic term. Adding a quadratic term to model 2 resulted in a lower BIC, as well as a marginal R^2 increase, indicating that model 2 is a better fitting model.

Table 1 shows that both ideology score and region are significantly associated with sentiment score. Model 1 shows that the higher the ideology score (and thus, the more conservative the user), the more negative the sentiment of the tweet. Model 2 indicates a significant quadratic fit in the relationship between vaccine sentiment and political ideology.

As shown in Table 2, the average ideological score for each US region tends to be liberal as the mean score within each region is negative. An analysis of variance test showed ideologies to be significantly different across regions, $P < 0.001$. Also, on average, each region has a positive sentiment toward the vaccine.

Moral values underlying COVID-19 vaccine topics and attitudes in different political ideologies

For our second research question, we first examine moral values in the tweets using the MFD [39]. Fig. 2 shows the percentage of words that correspond to the different foundations, comparing very Liberal, Liberal, Moderate, Conservative, and very Conservative tweets. See Methods for determining subgroups. Using MFD, the most prevalent moral foundations across tweets are Care/Harm and Authority/Subversion, with no significant differences between the political ideological groups.

Although MFD has been a widely used technique to understand morality in texts, the dictionary is not domain specific and contains limited terms. It is, therefore, possible that important words in our dataset for understanding moral values about vaccination may be missed. Furthermore, the MFD does not show the context in which certain moral values are applied: e.g. the foundation Care/Harm may be used in the context of vaccinating others and avoiding harm, but also in the context of being harmed by the vaccine itself.

To further understand the specific terms and topics discussed based on political ideology, we next describe the emergent topics and themes identified using topic modeling [41]. Topic modeling is a method to understand hidden semantic structures that might exist in the different ideological groups. Past research shows clear patterns underlying moral foundations of Liberals and Conservatives [15]. Also, Twitter debates are predominantly led by individuals with strong and more extreme political views [37]. We thus focus on very Liberal and very Conservative ideological groups only in this analysis as this may be more likely to reveal disparate structures in attitudes.

We ran one Latent Dirichlet Allocation (LDA) model for very Liberal tweets and one for very Conservative tweets, and used five-topic models based on the highest coherence score, which indicates the degree of semantic similarity between words within a topic (see Methods section). Tables 3 and 4 show the topics discovered by the two models, with keywords identified for each topic. We added a descriptive label for each topic to aid in interpretability (e.g. “receiving the second dose or booster”). Words for each topic are shown in a descending order according to its beta value, indicating its contribution to the topic (i.e. the first word is the most important word for each topic). The beta values of each keyword per topic are presented in Tables S1 and S2.

To aid in interpretability, we grouped similar topics under larger themes. For example, two topics contained words related to the labels “receiving the second dose or booster” and “distribution and access to the vaccine.” These labels both relate broadly to the notion of getting the vaccine and were thus grouped together under a larger theme “Getting the vaccine.” We grouped the labeled topics in the following four themes: Getting the vaccine, Safety, Trust, and the Government mandate (which can relate to mandates such as the mask mandate, vaccine mandate, and the stay-at-home order).

We next applied MFT to understand the moral underpinnings of the different expressions of the two ideological groups based on the topics. Table 5 summarizes how moral foundations were represented in the expression of opinions on the COVID-19 vaccine for very Liberal and Conservative tweets. One author considered each topic and whether it related to one or more of the six moral foundations. Interpretations were then assessed, and reasons for associating topics with specific foundations were discussed among authors until agreement was reached.

Among Liberal tweets, we interpreted topics to be related to moral foundations of Liberty (related to the mandate), Care (getting the vaccine to protect against the virus), Authority (trusting

Table 1. Models predicting sentiment score.

Predictors	Model 1			Model 2		
	β	Std. error	t-value	β	Std. error	t-value
(Intercept)	0.0176	0.0010	16.807***	0.0281	0.0012	22.742***
Ideology score	-0.0029	0.0004	-6.603***	-0.0045	0.0005	-10.064***
Midwest	0.0025	0.0020	1.253	0.0028	0.0020	1.401
South	0.0011	0.0014	0.758	0.0015	0.0014	1.072
Northeast	0.0063	0.0015	4.234***	0.0059	0.0015	3.989***
Ideology score ²				-0.0062	0.0004	-15.698***
BIC		-105936.2			-106144.6	
N		19,191			19,191	
Observations		262,267			262,267	
Marginal R^2 /conditional R^2		0.0005/0.0617			0.0026/0.0615	

The dependent variable is vaccine sentiment, and the reference category for region is West. model 2 includes the quadratic term of the ideology score.
*** $P < 0.001$.

Table 2. Average sentiment and ideology scores by region.

Region	Mean sentiment score	Mean ideology score	Number of tweets
West	0.015	-0.433	80,136
Midwest	0.018	-0.135	26,915
South	0.016	-0.083	85,512
Northeast	0.022	-0.398	69,704

regulatory authorities regarding the vaccine), and Fairness (having access to the vaccine). Among Conservative tweets, we considered topics to be associated with moral foundations of Oppression (around the government mandate) and Harm (around effectiveness of the vaccine). We next describe our justification for deriving these moral foundations.

Liberal tweets discussed the government mandate (topics TL2 and TL3). Because a mandate can be seen as restricting people in their personal liberty and is often discussed in the context of civil rights issues in mandating vaccination [21], we consider these topics to be associated with the moral foundation of Liberty. The mandate is discussed using terms such as love, care, safe, and protect, suggesting that Liberal users discuss the mandate as being protective against COVID-19, and the moral foundation of Care may underlie these topics. It also suggests that they trust the government mandate, suggesting the topics may be related to the moral foundation of Authority. Finally, Liberal tweets write about distribution and access to the vaccine (topic TL5), which can be associated with fair access to the vaccine and the foundation of Fairness. Thus, the topic modeling suggests that those with very Liberal political orientations express attitudes of vaccines in terms of caring for others; viewing authority as promoting the good of vaccines for citizens; that public health may be more important than personal liberty; and fairness in enabling people to get the vaccines.

Conservative tweets also discussed the government mandate (topics TC1, TC2, and TC3). However, in contrast, they refer to the mandate using the terms of force, choice, body, and freedom. These terms suggest that the discussion is likely around limits on free choice and liberty, suggesting that the moral foundation of Oppression underlies it.

Conservative tweets further included a topic about trust in the effectiveness of the vaccine (topic TC4). The presumed effectiveness of the vaccine is often used as an argument to take the vaccine (to protect and care for others) or not take it (to mitigate harm from the vaccine) [22]; thus, we also associate this topic with the foundation of Harm.

Finally, topic TC5 (safety) refers to the spread of COVID-19 and the risk of death, which suggests that the moral foundation of Harm underlies it. Therefore, the topic modeling results suggest that those with very Conservative political ideologies express moral arguments in terms of how the vaccine might cause harm and in terms of how the vaccine initiative deprives individual liberties.

In summary, the MFD shows that the most prevalent moral foundations are Care/Harm and Authority/Subversion, without large differences between the political ideological groups. The topic modeling shows more nuance. Political ideologies express attitudes on the two ends of the Care/Harm spectrum: caring for others (Liberals) vs the vaccine being harmful (Conservatives). Also, very Liberal tweets discuss authority as promoting the good of vaccines for citizens and public health as more important than personal liberty; very Conservative tweets rely on the Liberty/Oppression foundation in terms of how the vaccine mandate impacts individual liberties.

Understanding language use of the COVID-19 vaccine

To answer our third research question on how language use reflects differences in political ideologies, we identified several themes. Two overarching themes present in both ideology groups related to (1) Safety (around COVID-19 and the vaccine), and (2) the Government mandate/Trust in authorities (on the vaccine). To explore the context in which these themes were discussed and to understand how the same words might be used differently based on political ideology, we employed Word2Vec [42]. Word2Vec is a natural language processing technique used to learn word associations and to discover which words are used in a similar way based on the context in which they are used. It can be used to elaborate on and distinguish meanings among words.

We trained two Word2Vec models to compare whether the same word might be used in different contexts for Liberals and Conservatives such as “death” or “science.” Within each corpus of tweets, the Word2Vec models show the similarity of words based on context and co-occurrence. From the topic models, we first considered all terms related to the two themes of safety and the government mandate/trust. Through discussion among the authors, the following words were considered to be most central to discussions on vaccine attitudes: safe, effective, death, side_effect, die, sick, kill, protect, love, risk, spread, test, prevent, trial, child, symptom (relating to the safety theme) and mandate, right, believe, trust, force, require, mandatory, news, doctor, science, potus, government, push, fauci, trump (relating to the government mandate and trust theme). We then used Word2Vec’s similarity feature to identify the most similar words for each word and explore the contexts in which each word is used by the different ideological groups.

The full results for each term are included in Tables S3–S6. Here, we describe the results of a subset of the aforementioned terms that in our view best illustrate the differences between ideologies: “death,” “effective,” “safe,” “science,” “force,” and “choice.” Table 6 shows the most similar words for each of these terms. The words are shown in the descending order where the first word is most similar to that term, then the next more similar, etc. The similarity values of each word per term are included in Tables S3–S6.

Overall, the results suggest that the diverse ideology groups discuss the two themes of safety and government mandate/trust from two different stances. Liberals tend to use the word “Death” in the context of vaccines preventing illness and death as indicated by the top words “prevent, serious-illness.” This is supported in our examination of specific Liberal tweets from which the Word2Vec results originate, e.g. “Thanks. I feel very lucky. I’m hoping everyone will get this vaccine so we can end 4,000+ deaths a day in this country.” On the other hand, Conservatives use the word “Death” in the context of reported cases of death, as indicated by the top words “Report, high case rate.” The tweet content backs this up, suggesting “death” refers to reporting death rates due to the vaccine (“Why is no one reporting the deaths and other adverse reactions due to the vaccine?”; “The media are fools, they are not reporting on the thousands of deaths caused by taking the vaccine”).

Second, for the word “Science,” the closest words in Word2Vec for Liberals are “Believe,” “Trust,” “Pro.” For example, a Liberal tweet reads: “I believe in science and vaccines.” For Conservatives, the closest word to “Science” is “Lie,” and Conservative tweets show a more negative sentiment toward




Fig. 2. The percentage of words associated with moral foundations across different ideology groups.

science: “Yeah that’s a no for me, dawg. I don’t believe the science suggests that young people should be vaccinated.” These results suggest distinct beliefs: Liberals trust science of the vaccine; Conservatives associate science about the vaccine with lies. Conservative tweets use “choice” and “force” similarly to “body,” “freedom,” and “right” (“Why I’m NEVER TAKING THE ‘VACCINE’!!! #MYBODYMYCHOICE I WILL NOT BE FORCED BY SOCIAL SHAMING”), whereas Liberal tweets use words like “choice” and “force” similar to “selfish” (“Refusing the vaccine is not a ‘personal choice’, it is a collective one. Your selfishness will likely cause harm (and potentially death) to not only yourself, but also others.”). This suggests that Conservatives may express that the government mandate is limiting free choice over their bodies, while Liberals may regard this argument as being selfish.

Discussion

Previous studies have used MFT to understand differences between political ideologies [24] and differences between provaccine and antivaccine attitudes [22]. Our study brings these elements together to understand the link between political ideology and vaccine attitudes on social media through the lens of MFT. By examining these linkages, we find that Liberals use a broader range of moral values than Conservatives within the specific context of COVID-19 vaccine discussions on Twitter. This finding highlights the importance of studying unsolicited opinions as expressed on social media, an important mode of communication related to expressing vaccine attitudes [32]. It can be a useful source to gauge people’s sentiment on social issues [33].

Political ideology and vaccine sentiment

Our results confirm previous work showing Conservatives expressed more negative sentiment toward the COVID-19 vaccine [3, 13, 14]. We build on this work showing a quadratic relationship: moderate political ideologies have the most positive sentiment toward the vaccine. Other studies have clustered individuals into categories of Republican and Democratic [4]. Our results suggest that it is not just the orientation of one’s ideology but also the intensity of ideology, which matters in vaccine attitudes. Our study raises a question for future research: Why do extreme political

ideologies—both on the right and left—tend to have less positive views toward vaccines than moderates? Some research suggests that degree of conservatism is linked with susceptibility and spread of misinformation [43], which could account for the more negative sentiment of very Conservative individuals.

Using MFT to understand vaccine attitudes

MFT identifies “intuitive ethics” that guide people’s behavior and may be informative for informing effective vaccine messaging. COVID-19 vaccine uptake has presented considerable policy and behavior adoption challenges during the pandemic [8]. Results from our study show first that morality dimensions characterize COVID-19 vaccine discussions on Twitter. Results that evaluated the frequency of moral foundation terms used, rather than the context in which these were used, showed that all six moral foundations were expressed in some capacity in Twitter discussions, and that Care/Harm and Authority/Subversion were expressed most frequently among very Liberal and Conservative political ideologies. These findings are consistent with a Twitter COVID-19 vaccine study that found that vaccine hesitancy profiles predominantly expressed Care/Harm and Authority/Subversion morality dimensions [21].

LDA topic modeling contextualizes COVID-19 vaccine discussions to a greater extent. In contrast to the prior literature that has shown Conservatives rely on an expanded range of moral dimensions when interpreting politicized social topics beyond justice and care in the US [16], for COVID-19 vaccination discussion on Twitter, this is not the case: very Liberals rely on four morality dimensions, whereas very Conservatives rely on two, namely, Care/Harm and Liberty/Oppression. Moreover, the Word2Vec results indicate that when these moral dimensions are expressed, the same words are used (e.g. science, death) to express diametrically opposed moral dimensions by very Liberals and Conservatives.

An analysis by the Dutch Government’s dissemination of a public health vaccine communication campaign before the pandemic (2011–2019) found that Authority/Subversion and Liberty/Oppression moral framings used in brochure vaccine messaging distributed to households by the government had a significant association with vaccine uptake, while Care/Harm and Loyalty/Betrayal did not [20]. This is one of the few studies to show a significant effect of public health vaccine messaging using moral

Table 3. The topics and themes in very Liberal tweets about the COVID-19 vaccine.

ID	Theme	Label	Most common words	Contr. (%)
TL1	Getting the vaccine	Receiving the second dose or booster	Work good dose second need first trump tell fully shot mask year give shoot way booster receive thank help coronavirus hope available new day sure never mandate republican passport wait	24.0%
TL2	Safety	The mask mandate to care for Americans	Make mask trump mandate require think day virus feel johnson free shot today first biden put care american use love much need start flu work life anti give wait mean	21.8%
TL3	Safety	Safety and effectiveness of the mandate at schools	Need think mandate first school day mask thing fully thank make effective feel good require safe friend shot rate back child state low soon give test work hear receive virus	19.2%
TL4	Trust	Trust in COVID-19 prevention behavior	Today fully die dose mask appointment work think trump shot right refuse need yet family try wear america stop second give science death rate believe actually protect new week kid	18.5%
TL5	Getting the vaccine	Distribution and access to the vaccine	First need right refuse state dose ask fully week today card appointment eligible happen county able person day think show kid life tell give point already mandate effective new site	16.5%

The percentage column indicates the percentage of words in the dataset that are grouped under this topic.

Table 4. The topics and themes in very Conservative tweets about the COVID-19 vaccine.

ID	Theme	Label	Most common words	Contr. (%)
TC1	Government mandate	Mask mandate and protection, individual choice over body	Trump mandate mask need never work make test flu thank protect biden passport give die state person tell first maybe choice require president thing back refuse issue news johnson body	23.1%
TC2	Government mandate	Force of government mandate to prevent the spread of COVID-19 from China	Need virus give biden year force make mandate good call fully work flu first trump think business china right death joe plan day tell feel require cdc dose spread never	21.1%
TC3	Government mandate	Force of government mandate and trusting government advice	Mandate work virus mask make passport government think tell joe biden need fake fully free point actually ever force right new good lie rate world fauci hear last biden today wear_mask	20.1%
TC4	Trust	Trust in government mandate and vaccine effectiveness	Trump biden child work stop effective mask need die make think good way mandate believe receive american use trust tell coronavirus already fully death president year require spread develop anti	18.4%
TC5	Safety	Spread of COVID-19 and risk of death	Stop make think virus biden die spread show need mask death right cdc trump shot use approve potus fully push mandate sure work try trial kid lie exactly care vaccinate	17.3%

The percentage column indicates the percentage of words in the dataset that are grouped under this topic.

framing as a driving contributor to vaccine adoption. Political ideology, however, was not taken into account and is important as the US context has uniquely politicized COVID-19 vaccination under the Trump administration [18, 44, 45].

Another Twitter study on the COVID-19 vaccine in 2022 identified nine COVID-19 vaccine hesitancy profiles in which stance toward Authority/Subversion and Care/Harm moral framings shaped vaccine ambivalent profiles [21]. This study's findings empirically showed the important role that moral framing plays in contributing toward vaccine confidence in a social media context. Political ideology however, was not taken into account as we did.

Research on vaccine messaging before the pandemic found that moderately hesitant parents were twice as likely to emphasize Purity/Degradation morals and very vaccine hesitant parents were twice as likely to emphasize Purity/Degradation and Liberty/Oppression [19]. In public health vaccine messaging, care and fairness morality dimensions often implicitly underlie vaccine messages despite little evidence about whether such moral frames effectively motivate subgroups to vaccinate. Two studies examining the effectiveness of moral framing in motivating vaccine adoption found favorable effects among vaccine hesitant audiences when using Authority/Subversion, Liberty/Oppression, and Purity/Degradation [19, 20].

Implications for public health messaging

Public health vaccine communication may benefit from tailoring vaccine messaging not only by political ideology but also by moral foundations. To more effectively reach Conservatives who have not yet vaccinated, public health messaging may need to use a broader repertoire of moral vaccine messaging. Vaccine messaging framed as granting the freedom to choose and make one's personal health decisions may lead to greater message acceptance among Conservative subgroups who respond to oppression/liberty moral values. Reimer et al.'s [46] research showed that subgroups who had stronger endorsements of fairness and loyalty messaging were more likely vaccinated. Messaging that opens with "Vaccinating is your own choice" before emphasizing other aspects of vaccination may function as a "foot in the door" communication strategy [47, 48] to lessen the possibility of reactance against vaccine messaging among Conservative subgroups. Traditional public health messages that evoke care and fairness moral dimensions may simply not offer an effective strategy nor provide the motivation for behavior change in politicized contexts in the United States. Future research may broaden our understanding for why and in what contexts evoking certain moral foundations may reduce vaccine hesitancy.

Table 5. Moral foundations found in topic modeling results, in expressing attitudes toward the COVID-19 vaccine in Liberal and Conservative tweets.

Foundation	Liberal tweets	Conservative tweets
Authority/ Subversion	Liberal tweets discussed the government mandate using terms such as “required,” “safe,” and “good,” suggesting they agreed with adherence to restrictions imposed by regulatory authorities.	
Care/Harm	Liberals discussed getting the vaccine, and the vaccine mandate was discussed using terms such as “love,” “care,” “safe,” and “protect,” suggesting the vaccine is considered to protect against COVID-19.	Conservative tweets included topics related to harm such as risk of death and trust in the effectiveness of the vaccine, using words such as “die,” “fake,” and “lie.”
Fairness/Cheating	Liberal tweets included one topic related to distribution and fair access to the vaccine.	
Liberty/Oppression	Liberal tweets included topics related to the government mandate.	Conservative tweets relied on this foundation by discussing the government mandate being forced and limiting free choice.
Loyalty/Betrayal Sanctity/Degradation	Not present in the topic modeling results. Not present in the topic modeling results.	

Table 6. The most similar words to the terms “death,” “effective,” “safe,” “science,” “force,” and “choice,” obtained from the Word2Vec models for Liberal and Conservative tweets.

Theme	Term	Word2Vec results (liberal)	Word2Vec results (conservative)
Safety	Death	Prevent serious_illness infection spread hospitalization virus variant risk transmission readily	Report case high rate cdc die cause fully study injury
Safety	Effective	Variant study thoroughly mrna studies_suggest scientifically protection efficacy safe thus_far	Safe antibodies wuhan_chinese month mrna immunity flu year study side_effect
Safety	Safe	Effective prove protect thus_far keep rigorously variant symptom delta cld	Prove less survival_rate deadly illness transmission symptom delta variant mutate
Govt. mandate/ Trust	Science	Believe trust pro understand true trump_supporter expert fact change politic	Lie believe actually understand stupid fuck keep stop work worry
Govt. mandate/ Trust	Force	Fox_news fox foxnews cop selfish moron block democrat freedom fire	Freedom right demand government illegal protest refuse company comply governor
Govt. mandate/ Trust	Choice	Freedom selfish foxnews fucking choose point reason force moron fox_news	Body without_risk risk choose life abortion freedom right care pro

Difference in language use will be important to consider when developing effective outreach and communication strategies to inform people about the vaccine knowing that certain words and arguments may resonate differently and appeal to different audiences (e.g. using scientific research as an argument to take the vaccine may appeal to some but not other audiences). Echo chambers might be using language common within groups but different across groups [49]. Literature has shown this to be the case for vaccine messaging by political ideology [50]. Universal vaccine messaging may in fact have difficulty using language that resonates with everyone.

Limitations

Twitter data may not represent the broader US opinion, as many users are young (ages 18 to 29 comprise 42% of users), are more likely to identify as Democrat, have high education and incomes, are men (62%), and are concentrated in urban and suburban areas [51]. Tweets are also not equally distributed among users: a small group of users tweet more than others.

Although the MFD did not show a clear difference between political groups, the topic modeling showed that Liberals seem to rely on a broader range of moral foundations than Conservatives. This difference in results may be because our dataset contains moral terms not included in the MFD that were apparent in the topic

modeling, such as “mandate,” “choice” (related to Liberty/Oppression), and “eligible” (related to Fairness/Cheating). To understand what moral values underlie attitudes toward the COVID-19 vaccine, we used two techniques: we used the MFD to extract moral values from tweets and then analyzed the moral values from the topic modeling qualitatively. By using this combination of techniques, we were able to get a more holistic view of moral values. The MFD is intended to be applicable to a wide variety of texts and therefore only includes a limited set of terms. In response to this limitation of using a standardized lexicon, some studies have developed their own modified and/or expanded versions of the dictionary [20, 52]—see Fig. S2 for a comparison with the MoralStrength [53] lexicon. It would be worthwhile to conduct follow-up studies to develop an extended dictionary including words that are more common in the vaccine debate.

The issue of vaccination attitudes is complex, and there are other factors besides political ideology such as social determinants of health and barriers to vaccine access that influence vaccine uptake [54]. In addition to targeted messaging, it is important to highlight and further study the need for easy access to health services.

Finally, our study considered data from May 15, 2020, to October 31, 2021. Our results can generalize to earlier time periods of a pandemic (in this study, 17 months, but future work needs to

examine vaccine attitudes reflected in Twitter for longer time periods).

Conclusion

Our research contributes to a growing body of literature on the effective use of moral framing to motivate COVID-19 vaccine behavior. This study advances a more nuanced understanding of the differential role that political ideology and underlying moral arguments play for understanding vaccine attitudes and reactions to public health vaccine messaging especially in the United States. Extreme ideologies of both Conservatives and Liberals expressed a more negative sentiment than Moderates, with Conservatives expressing more negative sentiment than Liberals. We found the expression of Liberal tweets to be rooted in a wider set of moral values (care, authority, liberty, fairness) than Conservative tweets (harm and oppression). Finally, specific words, such as death and science, used differently among Liberals and Conservatives may make it difficult to develop universal messaging that will resonate with everyone. Public health vaccine communication may instead benefit from tailoring vaccine messaging by moral values depending on the intended audience, such as emphasizing the notion of free choice to Conservatives or framing vaccination to care for others to a Liberal audience.

Materials and methods

Data collection and filtering

We utilized the Twitter COVID-19 Stream^a to retrieve the initial dataset used in this study. The stream was filtered with a list of COVID-19-related terms defined by Twitter.^b Tweets were collected from the stream from May 15, 2020, to October 31, 2021. The initial dataset contained 3.4 billion tweets with original tweets and retweets. After filtering on US location and removing duplicate tweets and retweets, the dataset was reduced to 539,652 tweets. After filtering the dataset to only include users for whom we were able to estimate their political ideology, the final dataset included 262,267 tweets.

Geo-filtering and Geo-tagging

A coarse granularity geo-filtering was applied to the collected filtered stream dataset based on the “country_code” field tagged by Twitter. As the study was concerned with COVID-19-related tweets in the United States, we included tweets with “US” as the “country_code” tag. Since only the original tweets have geo-information from Twitter, we also filtered out any retweets in this step.

Although the collected tweets already have geo-tagging from Twitter, some of the tags are county name, state name, or country name, which are not accurate enough. Thus, further geo-tagging was applied to each tweet, which converted all place names into approximated coordinates.

US regions

We used the regions determined by the US Census [55] to analyze the average ideology score based on the user’s geo-location. The US Census categorizes the United States into four main regions: northeast, midwest, south, and west. Fig. S1 includes a map showing US states within each region.

Data preprocessing

Before data analysis, words were converted to lower case, and the Python package gensim [56] was used to remove white space,

URLs, emoticons, symbols, and punctuation. Standard English stopwords included in the Python package SpaCy [57], such as “I” and “the,” were removed (additional custom stopwords are included in the SI Appendix).

The package gensim was also used to detect common bigrams and trigrams. Bigrams and trigrams that occurred 100 times or more in the dataset were included as tokens and recognized as key words by the LDA model during topic modeling.

Calculating sentiment and ideology scores

The R library sentimentr [58] was used to assign a continuous sentiment score for each tweet ranging from -2 (negative sentiment) to $+2$ (positive sentiment).

We adopted Barbera et al.’s methodology [36] to estimate users’ political ideology. The estimation was based on the assumption that the more liberal or conservative accounts a user follows, the more likely that user identifies with the same ideology. The algorithm has been validated in prior work by estimating the ideology of known democrats on Twitter (e.g. Barack Obama) and known conservatives (e.g. George Bush) [37]. To further check reliability with our dataset, we randomly sampled six Twitter users from our dataset. Three members of our research team independently coded the political ideology of each user and compared this with the estimated ideology, further validating the algorithm (see the SI Appendix for inter-rater reliability assessment). We used the Python package tweepy [59] to scrape the Twitter accounts that each user follows, and the R library tweetscores to estimate users’ ideology score [36]. The tweetscores package included pre-estimated ideologies for well-known political accounts and media outlets (e.g. CNN or Fox News) and calculated a user’s ideology score based on the accounts they followed. There were no cut-off points for the ideology scores: the more negative or positive a score, the more liberal or conservative was the user, respectively. In the analyses, we only included users and tweets for which an ideology score could be calculated. This resulted in a dataset of 262,267 tweets from 19,191 unique users.

Previous research has found that Twitter debates are predominantly led by individuals with strong political views [37]. For the topic modeling and Word2Vec analyses, we therefore compared two subsets of the dataset with more extreme ideologies: those with scores lower than -1.5 (labeled as “very Liberal”) and those with scores higher than 1.5 (labeled as “very Conservative”). We chose these cut-off points based on prior work, which found ideologies to become more extreme beyond -1 and 1 [37] as well as the distribution of sentiment scores in our dataset, which showed a trend of becoming more negative for users with ideology scores beyond -1.5 and 1.5 (see Fig. S3 for a comparison of different cut-off points). The very Liberal group contained 59,570 tweets from 4,344 unique users with mean ideology score of -1.87 . The very Conservative group contained 31,535 tweets from 2,073 unique users with mean ideology score of 1.84 .

Modeling sentiment and ideology LMER model

We used Linear Mixed Effects Regression (LMER) [60] in R to run a linear mixed-effects model to investigate the relation between political ideology and sentiment, and compared the two models. Both models used a tweet’s sentiment score as a dependent variable.

Moral Foundations Dictionary

We used the MFD [39], a validated dictionary which contained stemmed words related to each moral foundation. For the Liberty/

Oppression foundation, we used the list of words used in the study by Teemstra et al. [61]. After preprocessing the data, we searched the tweets for terms listed in the MFD. If a term in a tweet matched with an MFD term, the word was counted. For each tweet, we divided the occurrences of words from the dictionary by the total number of words in the original tweet, to get a percentage for each moral foundation. We took the average values of very Liberal (with ideology scores of -1.5 or lower), Liberal (scores between -1.5 – 0.5), Moderate (scores between -0.5 – 0.5), Conservative (scores between 0.5 – 1.5) and very Conservative (scores of 1.5 or higher) tweets to compare the presence of moral values for each group.

Topic modeling

We used LDA [41] for topic modeling, a widely used method in other COVID-19-related Twitter work [34, 62]. LDA tends to have better performance on large datasets than alternatives and provides easy to interpret results [63]. To compare LDA performance, we ran two alternative modeling techniques on our dataset (GSDMM and NMF); GSDMM did not converge and NMF produced a lower coherence score using the UMass coherence score [64] from the gensim package for comparison. LDA was therefore selected as a method.

Choosing the number of topics was based on the coherence score, which evaluates topics by measuring the degree of semantic similarity between words within a topic. A topic model was considered good if words in a topic are similar. Following prior work [35, 65], we used a combination of four metrics specific for LDA models to identify the best number of topics (the SI Appendix includes more detail on the metrics, and Figs. S4 and S5 plot coherence scores by number of topics). The Python package gensim and R package ldatuning [66] were used to compute the scores for models ranging from 2 to 40 topics. Models with four to eight topics produced the highest coherence scores; after a manual comparison, we used five as the optimal number of topics (i.e. this model produced the most distinct topics with minimal overlap). Two models were trained, one for “very Liberal” tweets and one for “very Conservative” tweets. These models allowed us to compare differences in topics between ideology groups. We used MFT to interpret commonly discussed topics and related these to moral foundations.

Although unsupervised topic modeling is advantageous in extracting topics in a nonsubjective manner, researchers are needed to interpret the topics and identify larger themes [34]. To contextualize findings from unsupervised learning, we therefore employed a qualitative approach to extract themes. We started with a single researcher to identify themes that allowed for consistency and interpretation of topics. To check reliability, topic labels and themes were assessed among the rest of the authors. As an understanding of the data was developed, the first author iteratively reviewed and refined themes in discussion with the other authors. This process is common in topic modeling and qualitative research [34, 67].

Word2Vec

We used the Word2Vec implementation [42] in the gensim package to look at the context of specific words related to themes found by both political ideologies. This allowed us to understand how people’s attitudes toward the vaccine were reflected in their language use. We used the Continuous Bag of Words (CBOW) implementation, in which the surrounding words are used to predict a specific word. The Word2Vec model mapped words into a high-dimensional vector space so that words used in a similar way were

close to each other. Words used in a similar context presented similar vectors. Two Word2Vec models were trained: one for “very Conservative” tweets, and one for “very Liberal” tweets. These two models allowed us to compare whether the same word was used differently depending on a user’s political ideology.

Notes

- <https://developer.twitter.com/en/docs/twitter-api/tweets/covid-19-stream/overview>
- <https://developer.twitter.com/en/docs/twitter-api/tweets/covid-19-stream/filtering-rules>

Acknowledgments

Part of this research was conducted on the NSF-funded Texera platform and supported by a grant from the Orange County Health Care Agency. The authors declare no competing interest.

Supplementary Material

Supplementary material is available at PNAS Nexus online.

Funding

The work was supported by the National Science Foundation (NSF) Awards 2027254 and 2107150, and a University of California Irvine Information and Computer Sciences (UCI ICS) research award.

Authors’ Contributions

J.B., Y.H., S.G., S.H., C.L., and G.M. designed the research. Y.H. collected the data. J.B., Y.H. and S.G. wrote the code to estimate sentiment and ideology scores and the code to analyze the data. J.B., Y.H., S.G., S.H., C.L., and G.M. were involved in analysis of the data, interpretation of the results, and wrote and reviewed the manuscript.

Data Availability

We provide data according to Twitter’s Terms and Services at <https://github.com/ISG-ICS/COVID-19-Vaccine-Attitudes-dataset>. The code used in analysis was prepared to run on a custom platform and can be shared upon request.

References

- Ritchie H, et al. 2020. Coronavirus pandemic (COVID-19). Our world in data. <https://ourworldindata.org/coronavirus>
- Shetty P. 2010. Experts concerned about vaccination backlash. *Lancet*. 375(9719):970–971.
- Latkin CA, Dayton L, Yi G, Colon B, Kong X. 2021. Mask usage, social distancing, racial, and gender correlates of COVID-19 vaccine intentions among adults in the US. *PLoS ONE*. 16(2): e0246970.
- Fridman A, Gershon R, Gneezy A. 2021. COVID-19 and vaccine hesitancy: a longitudinal study. *PLoS ONE*. 16(4):e0250123.
- Sparks G, Lopes L, Montero A, Hamel L, Brodie M. 2022. KFF COVID-19 Vaccine Monitor, April 2022. [accessed 2023 Feb 07]. <https://www.kff.org/coronavirus-covid-19/dashboard/kff-covid-19-vaccine-monitor-dashboard/>

- 6 Betsch C, Renkewitz F, Betsch T, Ulshöfer C. 2010. The influence of vaccine-critical websites on perceiving vaccination risks. *J. Health Psychol.* 15(3):446–455.
- 7 Hussain A, Ali S, Ahmed M, Hussain S. 2018. The anti-vaccination movement: a regression in modern medicine. *Cureus.* 10(7):e2919.
- 8 Hou Z, et al. 2021. Assessing COVID-19 vaccine hesitancy, confidence, and public engagement: a global social listening study. *JMIR.* 23(6):e27632.
- 9 Attwell K, et al. 2021. Converting the maybes: crucial for a successful COVID-19 vaccination strategy. *PLoS ONE.* 16:e0245907.
- 10 Ward JK, et al. 2020. The French public's attitudes to a future COVID-19 vaccine: the politicization of a public health issue. *Soc Sci Med.* 265:113414.
- 11 Griffith J, Marani H, Monkman H. 2021. COVID-19 vaccine hesitancy in Canada: content analysis of tweets using the theoretical domains framework. *JMIR.* 23(4):e26874.
- 12 Thelwall M, Kousha K, Thelwall S. 2021. COVID-19 vaccine hesitancy on English-language Twitter. *Prof de la Inf.* 30(2):e300212.
- 13 Killgore WDS, Cloonan SA, Taylor EC, Dailey NS, Diclemente J. 2021. The COVID-19 vaccine is here—now who is willing to get it? *Vaccines.* 9:339.
- 14 Pennycook G, McPhetres J, Bago B, Rand DG. 2022. Beliefs about COVID-19 in Canada, the United Kingdom, and the United States: a novel test of political polarization and motivated reasoning. *Pers Soc Psychol Bull.* 48(5):750–765.
- 15 Stewart BD, Morris DSM. 2021. Moving morality beyond the ingroup: liberals and conservatives show differences on group-framed moral foundations and these differences mediate the relationships to perceived bias and threat. *Front Psychol.* 12:912.
- 16 Haidt J, Joseph C. 2008. The moral mind: how five sets of innate intuitions guide the development of many culture-specific virtues, and perhaps even modules. In: Carruthers P, Laurence S, Stich S, editors. *The innate mind.* Oxford (UK): Oxford University.
- 17 Haidt J. 2012. *The righteous mind: why good people are divided by politics and religion.* London (UK): Penguin.
- 18 Niburski K, Niburski O. 2020. Impact of Trump's promotion of unproven COVID-19 treatments and subsequent internet trends: observational study. *JMIR.* 22(11):e20044.
- 19 Amin AB, et al. 2017. Association of moral values with vaccine hesitancy. *Nat Human Behav.* 1(12):873–880.
- 20 Heine F, Wolters E. 2021. Using moral foundations in government communication to reduce vaccine hesitancy. *PLoS ONE.* 16(11):1–20.
- 21 Weinzierl M, Harabagiu S. 2022. Identifying the adoption or rejection of misinformation targeting COVID-19 vaccines in Twitter discourse. In: Laforest F, Troncy R, Simperl E, Agarwal D, Gionis A, Herman I, Médini L, editors, WWW' 22. Vol. 10. Lyon (France): ACM. p. 3196–3205.
- 22 Rossen I, Hurlstone MJ, Dunlop PD, Lawrence C. 2019. Accepters, fence sitters, or rejecters: moral profiles of vaccination attitudes. *Soc Sci Med.* 224:23–27.
- 23 Kalimeri K, et al. 2019. Human values and attitudes towards vaccination in social media. In: Liu L, White R, editors. WWW'19, San Francisco (CA): ACM. p. 248–254.
- 24 Johnson K, Goldwasser D. 2018. Classification of moral foundations in microblog political discourse. In: Gurevych I, Miyao Y, editors. ACL'2018. Melbourne, Australia: Association for Computational Linguistics. p. 720–730.
- 25 Nyhan B, Reifler J. 2015. Does correcting myths about the flu vaccine work? An experimental evaluation of the effects of corrective information. *Vaccine.* 33:459–464.
- 26 Abbas AH. 2022. Politicizing COVID-19 vaccines in the press: a critical discourse analysis. *Int J Semiot Law.* 35:1167–1185.
- 27 Galton F. 1884. Measurement of character. *Fortn Rev.* 36:179–185.
- 28 Park G, et al. 2015. Automatic personality assessment through social media language. *J Pers Soc Psychol.* 108(6):934–952.
- 29 Uher J. 2013. Personality psychology: lexical approaches, assessment methods, and trait concepts reveal only half of the story—why it is time for a paradigm shift. *Integr Psych Behav.* 47(1):1–55.
- 30 Jost JT, Sterling J. 2020. The language of politics: ideological differences in congressional communication on social media and the floor of congress. *Soc Infl.* 15(2–4):80–103.
- 31 Raubach EE. 2019. Does political ideology influence how politicians speak? A linguistic analysis of one-minute speeches from the US House of Representatives of the 115th Congress. *Document Numér.* 22:127–138.
- 32 Bahk CY, et al. 2016. Publicly available online tool facilitates real-time monitoring of vaccine conversations and sentiments. *Health Aff.* 35(2):341–347.
- 33 de Zuniga HG, Jung N, Valenzuela S. 2012. Social media use for news and individuals' social capital, civic engagement and political participation. *J Comput-Med Commun.* 17:319–336.
- 34 Boon-Itt S, Skunkan Y. 2020. Public perception of the COVID-19 pandemic on Twitter: sentiment analysis and topic modeling study. *JPHS.* 6(4):e21978.
- 35 Kwok SWH, Vadde SK, Wang G. 2021. Tweet topics and sentiments relating to COVID-19 vaccination among Australian twitter users: machine learning analysis. *JMIR.* 23(5):e26953.
- 36 Barbera P. 2015. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Polit Anal.* 23:76–91.
- 37 Barberá P, Jost JT, Nagler J, Tucker JA, Bonneau R. 2015. Tweeting from left to right: is online political communication more than an echo chamber? *Psychol Sci.* 26(10):1531–1542.
- 38 Rymes B. 1995. The construction of moral agency in the narratives of high-school drop-outs. *Disc Soc.* 6(4):495–516.
- 39 Graham J, Haidt J. 2012. Moral Foundations Dictionary. [accessed 2023 Feb 07]. https://moralfoundations.org/wp-content/uploads/files/downloads/moral%20foundations%20dictionary_dic
- 40 Nguyen KH, Nguyen K, Corlin L, Allen JD, Chung M. 2021. Changes in COVID-19 vaccination receipt and intention to vaccinate by socioeconomic characteristics and geographic area, United States, January 6–March 29, 2021. *Ann Med.* 53(1):1419–1428.
- 41 Blei DM, Ng AY, Jordan MI. 2003. Latent dirichlet allocation. *J Mach Learn Res.* 3:993–1022.
- 42 Mikolov T, Chen K, Corrado G, Dean J. 2013. Distributed representations of words and phrases and their compositionality. In: Burges CJ, Bottou L, Welling M, Ghahramani Z, Weinberger KQ, editors. *Adv in NIPS.* Stateline, NV: Curran Associates, Inc. p. 3111–3119.
- 43 Bilewicz M, Soral W. 2021. The politics of vaccine hesitancy: an ideological dual-process approach. *Soc Psychol Pers Sci.* 13(6):1–10.
- 44 Milosh M, Painter M, Sonin K, Van Dijcke D, Wright A. 2020. Unmasking partisanship: polarization undermines public response to collective risk. Chicago, IL: Technical report, Becker Friedman Institute.
- 45 Milosh M, Painter M, Sonin K, Van Dijcke D, Wright AL. 2020. Political polarisation impedes the public policy response to COVID-19. London, UK: Technical report, Vox EUCEPR.
- 46 Reimer NK, et al. 2022. Moral values predict county-level COVID-19 vaccination rates in the United States. *Am Psychol.* 77:743–759.

- 47 Carducci BJ, Deuser PS, Bauer A, Large M, Ramaekers M. 1989. An application of the foot in the door technique to organ donation. *J Bus Psychol.* 4(2):245–249.
- 48 Dolinski D. 2015. *Techniques of social influence: the psychology of gaining compliance.* 1st Edition. London, UK: Routledge.
- 49 Cinelli M, De Francisci Morales G, Galeazzi A, Quattrociocchi W, Starnini M. 2020. The echo chamber effect on social media. *PNAS.* 118(9):e2023301118.
- 50 Baumgaertner B, Carlisle JE, Justwan F. 2018. The influence of political ideology and trust on willingness to vaccinate. *PLoS ONE.* 13(1):e0191728.
- 51 Odabaş M. 2022. Ten facts about Americans and Twitter. [accessed 2023 Feb 07]. <https://www.pewresearch.org/fact-tank/2022/05/05/10-facts-about-americans-and-twitter/>
- 52 Matsuoid A, Sasaharaid K, Taguchi Y, Karasawa M. 2019. Development and validation of the Japanese Moral Foundations Dictionary. *PLoS ONE.* 14(3):e0213343.
- 53 Araque O, Gatti L, Kalimeri K. 2020. Moralstrength: exploiting a moral lexicon and embedding similarity for moral foundations prediction. *Knowl Based Syst.* 191:105184.
- 54 Atwell K, Hannah A, Leask J. 2022. COVID-19: talk of “vaccine hesitancy” lets governments off the hook. *Nature.* 602:574–577.
- 55 U. S. C. Bureau. Geographic levels. [accessed 2023 Feb 07]. <https://www.census.gov/programs-surveys/economic-census/geographies.html>
- 56 Řehůřek R, Sojka P. 2010. Software framework for topic modeling with large corpora. In: Calzolari N, Choukri K, Maegaard B, Mariani J, Odijk J, Piperidis S, Rosner M, Tapias D, editors. *LREC’2010 Valletta, Malta:University of Malta.* p. 46–50.
- 57 Honnibal M, Montani I, Van Landeghem S, Boyd A. 2020. spaCy: industrial-strength natural language processing in Python. [accessed 2023 Feb 07] <https://spacy.io/>
- 58 Rinker TW. 2021. *sentimentr: calculate text polarity sentiment.* New York: Buffalo. Version 2.9.0.
- 59 Roesslein J. 2020. Tweepy: Twitter for python! <https://github.com/tweepy/tweepy>
- 60 Bates D, Mächler M, Bolker B, Walker S. 2015. Fitting linear mixed-effects models using lme4. *J Stat Softw.* 67(1):1–48.
- 61 Teernstra L, van der Putten P, Noordegraaf-Eelens L, Verbeek F. 2016. The morality machine: tracking moral values in tweets. In: Boström H, Knobbe A, Soares C, Papapetrou P, editors. *Adv in Intell Data Anal XV.* New York City: Springer International Publishing. p. 26–37.
- 62 Bogdanowicz A, Guan C. 2022. Dynamic topic modeling of Twitter data during the COVID-19 pandemic. *PLoS ONE.* 17(5): 1–22.
- 63 Albalawi R, Yeap TH, Benyoucef M. 2020. Using topic modeling methods for short-text data: a comparative analysis. *Front Artif Intell.* 3:42.
- 64 Röder M, Both A, Hinneburg A. 2015. Exploring the space of topic coherence measures. In: Cheng X, Li H, Gabrilovich E, Tang J, editors. *WSDM’15.* New York (NY). p. 399–408.
- 65 Li Y, Rapkin B, Atkinson TM, Schofield E, Bochner BH. 2019. Leveraging latent Dirichlet allocation in processing free-text personal goals among patients undergoing bladder cancer surgery. *Qual Life Res.* 6:1441–1455.
- 66 Nikita M, Chaney N. 2020. ldatuning: tuning of the latent Dirichlet allocation models parameters. [accessed 2023 Feb 07]. <https://github.com/nikita-moor/ldatuning>
- 67 Guest G, MacQueen KM, Namey EE. 2011. *Applied thematic analysis.* Thousand Oaks, CA: SAGE Publications, Inc.