



Moral and Emotional Influences on Attitude Stability Towards COVID-19 Vaccines on Social Media

Samantha C. Phillips¹, Lynnette Hui Xian Ng¹, Wenqi Zhou²,
and Kathleen M. Carley¹

¹ CASOS Center, Software and Societal Systems Department, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA

{samanthp,huixiann,carley}@andrew.cmu.edu

² Palumbo-Donahue School of Business, Duquesne University, 600 Forbes Ave, Pittsburgh, PA 15282, USA

zhouw@duq.edu

Abstract. Effective public health messaging benefits from understanding antecedents to unstable attitudes that are more likely to be influenced. This work investigates the relationship between moral and emotional bases for attitudes towards COVID-19 vaccines and variance in stance. Evaluating nearly 1 million X users over a two month period, we find that emotional language in tweets about COVID-19 vaccines is largely associated with more variation in stance of the posting user, except anger and surprise. The strength of COVID-19 vaccine attitudes associated with moral values varies across foundations. Most notably, liberty is consistently used by users with no or less variation in stance, while fairness and sanctity are used by users with more variation. Our work has implications for designing constructive pro-vaccine messaging and identifying receptive audiences.

Keywords: Attitude strength · Emotion · Moral values · Social media

1 Introduction

Moral values and emotions are ubiquitous experiences that shape how we process information and form judgments. Efforts to curb public opinion, such as

This work was supported in part by the Knight Foundation and the Office of Naval Research grant MURI: Persuasion, Identity, & Morality in Social-Cyber Environments, N00014-21-12749. Additional support was provided by the Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Knight Foundation, Office of Naval Research, or the U.S. Government.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2024
R. Thomson et al. (Eds.): SBP-BRIMS 2024, LNCS 14972, pp. 216–225, 2024.

https://doi.org/10.1007/978-3-031-72241-7_21

promoting the acceptance of COVID-19 vaccines, can leverage moral and emotional reasoning to appeal to target audiences [4, 12]. This practice can be more effective when the moral values or emotions are relevant to members of the target audience, that is, anti-vaccine belief holders [1, 15, 17, 20]. Furthermore, moral values and emotions associated with an attitude can affect the stability of and openness to updating it. For example, attitudes with a moral basis have been associated with more attitude-consistent behavior and resistance to persuasion [2, 9]. Similarly, emotional reactions to attitude-relevant information have been associated with greater attachment to that belief [22]. Understanding antecedents of attitude strength can reveal which individuals are most receptive to public health messaging.

Social media provides insight into how people’s public-facing views change over time, and the corresponding change in justification for those views. In this work, we examine the relationship between the use of moral and emotional language in tweets about COVID-19 vaccines and the variance of stance towards COVID-19 vaccines expressed by the posting user. The variance of stance across two months is used as a proxy for attitude strength, which is characterized by stability over time, resistance to persuasion, and influence on behavior and information processing [7]. Through an empirical analysis the tweets of nearly 1 million X (formerly Twitter) users who tweeted about COVID-19 vaccines at least twice in a two month period, we aim to address the following:

RQ1: How is moral and emotional language used differently in pro- and anti-COVID-19 vaccine tweets?

RQ2: How is moral and emotional language used differently by X users who indicate variation in expressed stance versus those who have no variation in expressed stance?

RQ3: Do X users who include more moral and emotional language in their tweets about COVID-19 vaccines have more or less variation in stance towards COVID-19 vaccines?

The selected time period between March and May of 2021 coincides with some of the highest vaccination rates in the United States throughout the COVID-19 pandemic¹. For each tweet, we applied a lexicon-based methodology to detect moral foundations and emotions. In addition, we developed and applied a classifier to label the stance of each tweet towards COVID-19 vaccines as {pro, anti, neutral}. We then calculated the average use of each moral value and emotion, as well as the standard deviation of stances, across tweets for each user.

In *RQ1*, we provide an overview of which moral foundations and emotions are used in pro- and anti-COVID-19 vaccine tweets. In *RQ2*, we compare the moral and emotional language used by the 310,258 users who expressed a single stance about COVID-19 vaccines with the 603,158 users with any variation in stance represented in their tweets. In addition, in *RQ3*, we identify associations between moral and emotional language and stance variance among users with any variation in stance, excluding users with zero variation.

¹ <https://usafacts.org/visualizations/covid-vaccine-tracker-states/>.

2 Related Works

Extensive previous work has examined associations between moral values and judgments. Moral foundation theory proposes six automatic intuitions (i.e., foundation or value) that affect judgments [6]. These foundations include: care (well-being of others), fairness (justice), sanctity (purity), liberty (freedoms), loyalty (in-group/out-group relations), and authority (following rules and traditions). Moral foundations most associated with vaccine hesitancy are purity and liberty, followed by authority [1]. In addition, vaccine hesitancy is associated with less need for care [1, 17]. However, these studies simply ask people about their automatic moral intuitions and attitudes towards vaccines, so there is no assessment of changes in stance over time or how moral reasoning is used to justify judgments about vaccines specifically. Nonetheless, they demonstrate existing interest in how moral values affect attitudes towards vaccines.

Like moral intuitions, emotions can affect information processing and attitude formation. While there are a multitude of ways to conceptualize emotions, we use Plutchik’s six-emotion model: happy, sad, anger, fear, disgust, surprise [14]. This model has been used to study emotions on social media [8, 13]. Certain emotional reactions to vaccine-relevant events can predict vaccine hesitancy, like anger and anxiety [20]. We build on this work to assess which emotions are used in pro- and anti-vaccine tweets.

Scholars have evaluated moral aspects of an attitude as an antecedent to attitude strength, finding attitudes with a moral basis are more resistant to influence [2, 9] and predict consistent behavior [9]. Furthermore, emotional responses to attitude-relevant information typically indicate a high degree of importance has been placed on that view, implying attitude strength [11, 22]. We investigate if using emotional language to discuss a stance is also associated with stability over time. Unlike previous analyses of antecedents of attitude stability, we use stance variability in social media posts over time as a measure of attitude strength, extending these works to consider stance over time outside a laboratory environment. Taking these studies together, we expect that moral and emotional language are associated with *less* variation in stance towards COVID-19 vaccines.

3 Data and Methods

3.1 Dataset

Our dataset was initially collected using a streaming keyword search via Twitter v1 API between March 12, 2021 thru May 11, 2021. We collected tweets that contained at least one of the following terms: coronavirus, Wuhan virus, Wuhanvirus, 2019nCoV, NCoV, NCoV2019, covid-19, covid19, covid 19. We then further filtered the tweets to select those about vaccines using the following keywords: vaccine, vax, mrna, autoimmuneencephalitis, vaccination, getvaccinated, covidisjustacold, autism, covidshotcount, dose1, dose2, VAERS, GBS, believe-mothers, mybodymychoice, thisisourshot, killthevirus, proscience, immunization,

gotmyshot, igottheshot, covidvaccinated, beatcovid19, moderna, astrazeneca, pfizer, johnson & johnson, j&j, johnson and johnson, jandj.

In total, we extracted 2,283,281 users, 1,369,865 of which only tweeted one time about COVID-19 vaccines during this time period. In this work, we analyze the remaining 913,416 users who tweeted at least twice in our time period for a total of 6,811,854 tweets.

3.2 Stance Detection

To evaluate stance consistency, we first developed a classifier to label the stance (pro, anti, neutral) of each tweet towards COVID-19 vaccines. To generate a training dataset, two annotators labelled 1034 tweets with Cohen’s Kappa of 0.741 and agreement of 84.2%. A third annotator labelled the 163 tweets the initial two annotators disagreed on. We removed the 6 tweets that all three disagreed on. We then fine-tuned a BERTweet model² for the stance detection task to obtain an accuracy of 72.8% on the evaluation data. The stance classifier returns a stance label and confidence score for each tweet.

We applied this stance detection classifier to original tweets, retweets, replies, and quote tweets in our dataset. This gives us 3,547,645 pro-vaccine, 2,928,820 neutral, and 335,389 anti-vaccine tweets. To obtain a measure of variation in stance per user, we calculated the standard deviation of stance expressed across their tweets. Of the 913,416 users in our dataset, 310,258 express a consistent stance. Removing the 310,258 users with a standard deviation of 0, the average standard deviation is 0.62 (std. dev. 0.18) and median is 0.58.

3.3 Moral Value and Emotion Extraction

We use the Netmapper software to extract words and phrases associated with each moral value (care, fairness, liberty, sanctity, loyalty, authority) and emotion (happiness, sadness, fear, anger, disgust, surprise) [3, 21]. We recorded the number of concepts representing each moral/emotion variable in each tweet.

4 Results

4.1 RQ1: Moral Values and Emotions in Pro and Anti-COVID-19 Vaccine Tweets

Figure 1 displays the average and 95% confidence interval of the number of concepts associated with each moral value and emotion in pro- and anti-COVID-19 vaccine tweets. Overall, anti-vaccine tweets contain more moral and emotional language than pro-vaccine tweets for most types of moral values and emotions. That is, anti-vaccine tweets contain more references to care, fairness, authority, sanctity, and liberty foundations than pro-vaccine tweets. On the other hand, pro-vaccine tweets contain more loyalty terms. Anti-vaccine tweets are also more

² https://huggingface.co/docs/transformers/model_doc/bertweet.

likely to contain sad, fear, anger, disgust and surprise emotions, while pro-vaccine tweets contain more happiness. Table 1 contains examples of tweets containing select moral values and stances.

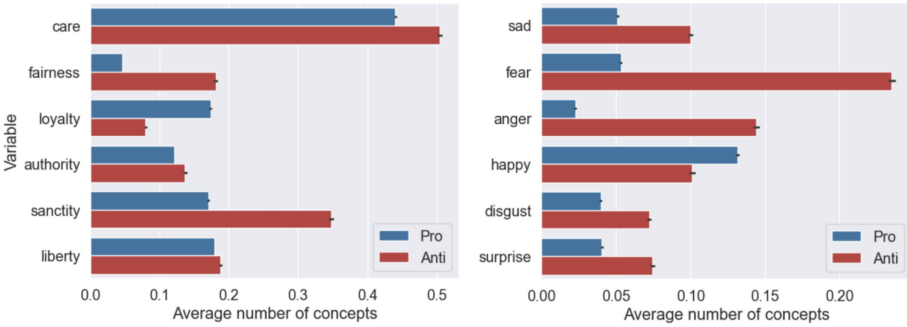


Fig. 1. Average and 95% CI number of concepts associated with each moral foundation (left) and emotion (right) in pro- and anti-COVID-19 vaccine tweets.

Table 1. Example tweets

Stance	Moral foundation	Example tweet
pro	care	If you can't see the urgency of saving humanity by granting a #TRIPSWaiver to ensure #VaccineForAll, I don't see what else you see. Time to stop putting Health in the Marketplace
anti	care	"Vaccinating" babies now with the mRNA jabs? How can this even be ethical?
pro	loyalty	His son has taken both doses of the COVID-19 vaccine, and hopefully was wearing a mask. They are both putting action behind their push towards herd immunity. I hope you'll get the vaccine too!
anti	sanctity	I am not taking the vaccine. I take 5000 units of vitamin D3. If God wants me to die from COVID-19 no vaccine will stop God's Will

4.2 RQ2: Moral Values and Emotions in Tweets by Users Expressing no Versus any Variation in Stance

Table 2 contains t-statistics from t-tests comparing the use of moral and emotional language by users with no versus any variance in expressed stance. Users

that are completely consistent in their expressed stance tend to use anger, loyalty, care, and liberty in their tweets more than users that express variation in their stance. Users who express any variation in stance are more likely to use happy, disgust, surprise, fear, or sad emotional language. Furthermore, they tend to refer to sanctity, authority and fairness foundations more in their tweets.

Table 2. T-statistics from t-tests comparing the average number of concepts associated with each moral value and emotion in tweets by users who expressed no variation versus any variation in stance towards COVID-19 vaccines. Negative values indicate the variable is used more in tweets by users that express no variation in stance. All t-tests are significant ($p < 0.0001$)

Emotion		Moral value	
Variable	T-statistic	Variable	T-statistic
Anger	-44.22	Loyalty	-37.95
Happy	6.36	Care	-34.36
Disgust	17.26	Liberty	-22.57
Surprise	30.18	Sanctity	21.90
Fear	32.3	Authority	27.42
Sad	37.86	Fairness	44.43

4.3 RQ3: Moral Values and Emotions in Tweets by Users Expressing Some Degree of Variation in Stance

Figure 2 displays the coefficients for an OLS regression model predicting the standard deviation of expressed stance across tweets for each user given the use of moral and emotional language in their tweets about COVID-19 vaccines. All coefficients are significant ($p < 0.0001$) except for anger ($p > 0.1$).

More emotional language is largely associated with more variation in stance, except surprise. Liberty and loyalty are associated with less variation in stance, while the remaining moral foundations (fairness, sanctity, authority, care) are associated with more variation.

5 Discussion

The aim of this work is to investigate how moral foundations and emotions associated with an attitude affect the stability of that attitude in a social media context. To do so, we analyzed tweets expressing a stance toward COVID-19 vaccines.

First, we examined which moral foundations and emotions are used in tweets expressing pro- and anti-COVID-19 vaccine views to provide an overview of the context in which each feature is typically used. The findings in RQ1 indicate that anti-vaccine tweets tend to contain more moral and emotional language than




Fig. 2. Regression coefficients for the standard deviation of expressed stance across tweets. Negative values indicate the variable is associated with less variation in stance. Error bars indicate the 95% confidence intervals. All coefficients are significant ($p < 0.0001$) except for anger ($p > 0.1$) and care ($p = 0.049$)

pro-vaccine tweets except for loyalty and happiness. Many pro-vaccine tweets that contain language associated with loyalty contain information about where a community can access vaccines or celebrate that vaccines were successfully given to a community (e.g., Americans, indigenous tribes in Oklahoma). As one may expect, vaccine hesitant messaging tends to use more negative emotional language than those expressing support.

Our findings echo previous work showing vaccine hesitancy is associated with moral values like sanctity and liberty [1]. Anti-vaccine tweets that use language associated with sanctity include concerns about the efficacy and safety of COVID-19 vaccines [1], and use more negative emotional language [5]. However, while vaccine hesitancy has been associated with lack of need for care [1, 17], we find care is the most used moral foundation in anti-vaccine tweets. These tweets include concern about how negative side effects from vaccines are affecting or may affect others.

In addressing *RQ2* and *RQ3*, we find emotional language is associated with more variation in stance, except for anger and surprise. That is, certain emotions (fear, sadness, happiness, disgust) indicate unstable attitudes. Users may be less tied to their view once their emotional reaction is less salient. Indeed, studies show that both positive and negative emotions can override analytical thinking processes, increasing the effect of biased information processing on truth discernment [16, 18].

Conversely, anger is associated with consistent stance, but the association between anger and stance variance among those who do vary to some degree is not significant. This highlights the different ways different emotions can impact judgments depending on the individual and context. Surprise is most associated with users that express some, but small, variance in stance. Surprise drives belief updating by attracting attention to unexpected information, motivating people to reconsider and potentially rectify their understanding [10].

The role of moral language varies across stance foundations. Liberty is consistently used by users with no or less variation in stance, while fairness and sanctity are used by users with more variation in stance towards COVID-19 vaccines. Care and loyalty are associated with no or high variation, while authority is associated with some but low variation. The same moral value can play a very different role depending on the context and individual differences. For example, moral identity is how important the individual considers being a “moral” person to their self-concept. Prioritizing “binding” values (loyalty, authority, sanctity) represent a group-oriented view of morality. Adherence to binding foundations is associated with out-group derogation if moral identity is weak and willingness to help out-group members if moral identity is strong [19]. More work is needed to identify variables influencing the effect of moral values on stance consistency.

The results of this work inform how moral values and emotions associated with a belief affect the stability of that belief. Our work has implications for constructing effective campaigns to promote vaccines and debunk vaccine misinformation, as well as directing these messages to receptive audiences. Specifically, promoting and tailoring pro-vaccine messaging to people expressing moral and emotional language associated with less attitude stability (i.e., sanctity, fairness, fear, sadness, happiness, disgust) can help to optimize public health outreach.

5.1 Future Work and Limitations

Our study is limited to the platform X. Future work should assess the use of moral values and emotions towards vaccine attitudes on other social media platforms. These analyses could provide vital insight into how organizations designing public health messaging should anticipate shifts in motivations as vaccines are developed and approved for the public. Furthermore, we did not attempt to identify automated accounts (“bots”) in our analysis [13]. Future work should examine the extent to which removing users displaying bot-like behavior affects the observed relationship between moral-emotional bases for an attitude and the consistency of that attitude.

Rather than separately examining which moral values and emotions predict more or less vaccine hesitancy over time, we looked at any change in stance regardless of the direction. Disentangling the effects on increases and decreases in anti-vaccine attitudes can directly inform public health messaging aiming to increase vaccinations [17].

Lastly, attitudes towards vaccines are multi-faceted, complex judgments that often include far more nuance than we include in this work. Indeed, vaccine hesitancy can be motivated by a wide range of factors, such as concerns over vaccine

testing protocols, institutional distrust, and belief in alternatives [23]. Considering more fine-grained stance categories may reveal important distinctions in the role of moral values and emotions in stance (in)consistency.

6 Conclusion

We studied the influence of moral and emotional associations with an attitude on the stability of that attitude over time. Specifically, we examined the moral values, emotions, and stance expressed in tweets by nearly one million users on X over a two month period. We find that (1) moral and emotional language is used more in anti-vaccine tweets, except for happiness and loyalty. (2) Emotional language is associated with larger variation with stance, except for anger and surprise, which are associated with consistent stance. (3) The role of moral language and stance variability differs across values, e.g. liberty is associated with consistent stances while fairness and sanctity are associated with larger stance variation. Our work shows that moral and emotional reasoning in attitudes can be used to predict receptive audiences, which can be especially useful in counter-misinformation campaigns.

References

1. Amin, A.B., et al.: Association of moral values with vaccine hesitancy. *Nat. Hum. Behav.* **1**(12), 873–880 (2017)
2. Aramovich, N.P., Lytle, B.L., Skitka, L.J.: Opposing torture: moral conviction and resistance to majority influence. *Soc. Infl.* **7**(1), 21–34 (2012)
3. Carley, L.R., Reminga, J., Carley, K.M.: Ora & netmapper. In: Thomson, R., Dancy, C., Hyder, A., Bisgin, H. (eds.) *Social, Cultural, and Behavioral Modeling. Demo at SBP-BRiMS 2018*. Springer, Cham (2018)
4. Chou, W.Y.S., Budenz, A.: Considering emotion in COVID-19 vaccine communication: addressing vaccine hesitancy and fostering vaccine confidence. *Health Commun.* **35**(14), 1718–1722 (2020)
5. Faasse, K., Chatman, C.J., Martin, L.R.: A comparison of language use in pro-and anti-vaccination comments in response to a high profile Facebook post. *Vaccine* **34**(47), 5808–5814 (2016)
6. Graham, J., Nosek, B.A., Haidt, J., Iyer, R., Koleva, S., Ditto, P.H.: Mapping the moral domain. *J. Pers. Soc. Psychol.* **101**(2), 366 (2011)
7. Krosnick, J.A., Petty, R.E.: Attitude strength: an overview. In: *Attitude Strength*, pp. 1–24 (2014)
8. Liu, Z., Zhang, T., Yang, K., Thompson, P., Yu, Z., Ananiadou, S.: Emotion detection for misinformation: a review. *Inf. Fusion* **107**, 102300 (2024)
9. Luttrell, A., Petty, R.E., Briñol, P., Wagner, B.C.: Making it moral: merely labeling an attitude as moral increases its strength. *J. Exp. Soc. Psychol.* **65**, 82–93 (2016)
10. Maguire, R., Maguire, P., Keane, M.T.: Making sense of surprise: an investigation of the factors influencing surprise judgments. *J. Exp. Psychol. Learn. Mem. Cogn.* **37**(1), 176 (2011)
11. Miller, J.M., Krosnick, J.A., Holbrook, A., Tahk, A., Dionne, L.: The impact of policy change threat on financial contributions to interest groups. In: *Political Psychology: New Explorations*, pp. 172–202 (2016)

12. Nan, X., Iles, I.A., Yang, B., Ma, Z.: Public health messaging during the COVID-19 pandemic and beyond: lessons from communication science. *Health Commun.* **37**(1), 1–19 (2022)
13. Ng, L.H.X., Carley, K.M.: Bot-based emotion behavior differences in images during Kashmir black day event. In: Thomson, R., Hussain, M.N., Dancy, C., Pyke, A. (eds.) *Social, Cultural, and Behavioral Modeling*. SBP-BRiMS 2021. LNCS, vol. 12720, pp. 184–194. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-80387-2_18
14. Plutchik, R.: The nature of emotions: human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *Am. Sci.* **89**(4), 344–350 (2001)
15. Rossen, I., Hurlstone, M.J., Dunlop, P.D., Lawrence, C.: Accepters, fence sitters, or rejecters: moral profiles of vaccination attitudes. *Soc. Sci. Med.* **224**, 23–27 (2019)
16. Scheufele, D.A., Krause, N.M.: Science audiences, misinformation, and fake news. *Proc. Natl. Acad. Sci.* **116**(16), 7662–7669 (2019)
17. Schmidtke, K.A., et al.: Evaluating the relationship between moral values and vaccine hesitancy in Great Britain during the COVID-19 pandemic: a cross-sectional survey. *Soc. Sci. Med.* **308**, 115218 (2022)
18. Sinclair, M., Ashkanasy, N.M., Chattopadhyay, P.: Affective antecedents of intuitive decision making. *J. Manag. Organ.* **16**(3), 382–398 (2010)
19. Smith, I.H., Aquino, K., Koleva, S., Graham, J.: The moral ties that bind... even to out-groups: the interactive effect of moral identity and the binding moral foundations. *Psychol. Sci.* **25**(8), 1554–1562 (2014)
20. Sun, R., Wang, X., Lin, L., Zhang, N., Li, L., Zhou, X.: The impact of negative emotional reactions on parental vaccine hesitancy after the 2018 vaccine event in China: a cross-sectional survey. *Hum. Vaccines Immunotherapeutics* **17**(9), 3042–3051 (2021)
21. Tausczik, Y.R., Pennebaker, J.W.: The psychological meaning of words: LIWC and computerized text analysis methods. *J. Lang. Soc. Psychol.* **29**(1), 24–54 (2010)
22. Visser, P.S., Krosnick, J.A., Norris, C.J.: Attitude importance and attitude-relevant knowledge: motivator and enabler. In: *Political Psychology*, pp. 217–259. Psychology Press (2016)
23. Weinzierl, M.A., Harabagiu, S.M.: From hesitancy framings to vaccine hesitancy profiles: a journey of stance, ontological commitments and moral foundations. In: *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 16, pp. 1087–1097 (2022)