Authority without Care:
Moral Values behind the Mask Mandate Response

Yelena Mejova[1], Kyriaki Kalimeri[1], Gianmarco De Francisci Morales[2]

[1]ISI Foundation, V. Chisola 5, Turin, Italy
[2]CENTAI, C. Inghilterra 3, Turin, Italy
yelenamejova@acm.org, kyriaki.kalimeri@isi.it, gdfm@acm.org

Abstract

Face masks are one of the cheapest and most effective non-pharmaceutical interventions available against airborne diseases such as COVID-19. Unfortunately, they have been met with resistance by a substantial fraction of the populace, especially in the U.S. In this study, we uncover the latent moral values that underpin the response to the mask mandate, and paint them against the country's political backdrop. We monitor the discussion about masks on Twitter, which involves almost 600k users in a time span of 7 months. By using a combination of graph mining, natural language processing, topic modeling, content analysis, and time series analysis, we characterize the responses to the mask mandate of both those in favor and against them. We base our analysis on the theoretical frameworks of Moral Foundation Theory and Hofstede's cultural dimensions.

Our results show that, while the anti-mask stance is associated with a conservative political leaning, the moral values expressed by its adherents diverge from the ones typically used by conservatives. In particular, the expected emphasis on the values of authority and purity is accompanied by an atypical dearth of in-group loyalty. We find that after the mandate, both pro- and anti-mask sides decrease their emphasis on care about others, and increase their attention on authority and fairness, further politicizing the issue. In addition, the mask mandate reverses the expression of Individualism-Collectivism between the two sides, with an increase of individualism in the anti-mask narrative, and a decrease in the pro-mask one. We argue that monitoring the dynamics of moral positioning is crucial for designing effective public health campaigns that are sensitive to the underlying values of the target audience.

Introduction

The COVID-19 pandemic has been an unprecedented event, which has also brought about an infodemic that makes public health response difficult. COVID denial, anti-vaccine sentiment, and other flavors of theories (from doubts to full-blown conspiracies) have been documented in social media (Patwa et al. 2021). Among the many controversies, surprisingly, wearing a mask has become extremely politicized and contentious. Inconsistent messaging by public health organizations has seriously undermined public compliance to this simple measure. Initially, the WHO recommended that masks should be worn only by professionals or those taking care of a sick person (WHO 2020), and in the U.S., the Center for Disease Control and Prevention (CDC) recommended the use of masks only for healthcare workers and people who were sick. Only on April 3rd, 2020, the CDC officially recommended wearing non-medical cloth face coverings when in public places (Landsverk 2020). Soon after, speculations that masks do more harm than good (Gillespie 2020) and that they foster a false sense of security (Khan 2020) started to proliferate. Public health experts have acknowledged that it has been a challenge to communicate with groups holding strong values of freedom (Hartsoe 2020). Indeed, globally, research has found that the value of individualism is associated with higher COVID-19 mortality rates (Maaravi et al. 2021) and less adherence to preventive measures (Biddlestone, Green, and Douglas 2020). However, how personal moral values evolve around governmental messaging concerning these preventive measures is not yet understood.

Mask-wearing behavior and attitudes can be influenced by multiple factors. There is extensive evidence that political ideology and identity influences attitudes, judgments, and behaviors (Van Bavel and Pereira 2018). At the same time, political identity is deeply intertwined with moral values (Graham, Haidt, and Nosek 2009; Hatemi, Crabtree, and Smith 2019). Finally, and perhaps obviously, moral values are directly linked to moral decision making (Karandikar et al. 2019), which completes the triad. In particular, there is an altruistic component in mask-wearing, as their main purpose is to protect others, which is related to solidarity towards the in-group in the face of an out-group threat (Campbell 1965). For this reason, among moral values we pay particular attention to the individualist-collectivist angle. Finally, there is ample literature on the division and polarization of news and information sources in general.[1] As users increasingly rely on social media to satisfy their information needs,[2] the lines between professional reporting and personal opinions begin to blur. For this reason, we contextualize the information environment in which the mask debate takes place in terms of information sources from peers—within the Twitter

---

[1]https://www.pewresearch.org/journalism/2020/03/04/about-one-fifth-of-democrats-and-republicans-get-political-news-in-a-kind-of-media-bubble

[2]https://www.pewresearch.org/journalism/2021/11/15/news-on-twitter-consumed-by-most-users-and-trusted-by-many

community—and from external sources.

This work provides a fine-grained analysis of the moral values of those expressing opinions around masking by applying the Moral Foundations Theory to a dataset of Twitter posts spanning the beginning of the pandemic, from January to July 2020. In particular, we ask: *What is the anatomy of the collective discussion on mask wearing around the mask mandate on Twitter?* In particular, we analyze different facets of this discussion in the U.S.:

1. How does the users' stance relate to their political leaning?

2. What moral values do adherents to pro- or anti-masking stances hold?

3. What is the information environment around their arguments?

Our findings confirm the known political divisions, with liberals supporting the pro-mask and conservatives the anti-mask stance. The moral narrative of the two sides also differs: those on the anti-mask side invoke the values of *authority* and *purity*, while those promoting mask-wearing emphasize *care* and *loyalty*. The introduction of the mask mandate shifts these values from an emphasis on *care* to an increased attention on *authority* and *fairness*, thus escalating the politicization of the issue. Content analysis shows that the increasingly individualistic views espoused by the anti-mask side are accompanied by the allegations of the dangers and ineffectiveness of mask-wearing, supported by resources from social media and sometimes even governmental and scientific sources. We argue that the mask-wearing debate has a dynamic moral landscape that should be carefully monitored to design effective public messaging campaigns that reflect the shifting values of their target audience. Such studies are especially important, given the rising prominence of social media as a platform for public policy discussion and influence (Shapiro and Hemphill 2017).

Related Work

Masking Behavior.   Mask mandates have been shown to be effective in decreasing COVID infection rates in the U.S., both in urbanized and rural areas (Krishnamachari et al. 2021). Early studies have tested the effect of mask adoption, with simulations showing that even relatively ineffective face coverings could meaningfully reduce community transmission (Eikenberry et al. 2020). Unfortunately, the public response to the masking mandates has been partial, with only a 12% increase in mask-wearing immediately after the CDC masking guideline on April 3, 2020 (Goldberg et al. 2020). A survey of residents of 10 states in the U.S. relates mask-wearing to COVID cases in the state, political party in power, and individual measures of "social capital", as well as some demographics (Hao, Shao, and Huang 2021). They find that respondents are more likely to wear a mask if there are more COVID-related deaths in the state (Odds Ratio of 1.26), the Democratic party is in power (OR = 2.0), and if the respondents more often speak to their friends and family (OR = 1.16). Although misinformation has been blamed for the resistance to mask-wearing and social distancing, a nationally representative survey has found that the beliefs about

the consequences of these behaviors are more predictive of people's compliance (Hornik et al. 2021). Especially in the U.S., beliefs and trust in authority are often strongly related to the people's political affiliations. For instance, Republicans, conservatives, and nationalists are less likely to believe that the World Health Organization (WHO) can effectively manage the pandemic (Bayram and Shields 2021).

Social Media & Masking.   A study of tweets during the early days of the COVID pandemic (Feb-March 2020) identified methods for decreasing the spread of COVID as one of the main themes and the wearing of masks as one of those most associated with positive sentiment (Abd-Alrazaq et al. 2020). However, authors of a later study spanning January to October 2020 showed that the output of automated sentiment analysis tools corresponds poorly with the mask-related sentiment expressed in the tweets due to the richness of the language used (He et al. 2021). Instead, they perform manual coding of a sample of the tweets to identify several major categories of concerns around masking, including physical discomfort, effectiveness, appropriateness, and political beliefs. A network analysis of the mask-related tweets has shown the pro-mask activists exist in a kind of "echo chamber" (Cinelli et al. 2021; Garimella et al. 2018a), and that they tend to ignore the subversive rhetoric of the anti-mask fringe (Lang, Erickson, and Jing-Schmidt 2021). Instead, a recent Twitter study found a focus on ongoing news (Cotfas et al. 2021). This fringe is more likely to use toxic language, including insults and profanity, than the pro-mask ones (Pascual-Ferrá et al. 2021). The authors link this behavior to either the vociferous protestations of a minority group (Miller and Morrison 2009) or potential signaling as an in-group behavior and a marker of personal identification. Alongside this toxicity, other studies find widespread misinformation and misunderstandings in the social media discussions around mask use. These include the beliefs that COVID19 is over-hyped by the media, that masks are ineffective, and that they do more harm than good (Keller, Honea, and Ollivant 2021). These beliefs were then shown to impact the mask-wearing (and social distancing) behaviors (Hornik et al. 2021). In this study, we introduce a dimension of *moral values*, which we argue underlies some of the disagreements on the appropriate use of masks during the epidemic, and which may shed some light on the moral reasoning behind the rhetoric.

Theoretical Framework.   Our study of the rhetoric around masks during COVID-19 is grounded in the Moral Foundations Theory (MFT), its manifestation in interpersonal and inter-group communication, and its reflection at the societal level as an individualist or collectivist cultural dimension. According to social identity theory, members of an in-group will look for negative aspects of an out-group, thus enhancing their self-image (Tajfel et al. 1979). Strong in-group and out-group reasoning at a societal level determines whether a society is individualistic or collectivist (IC), as defined by Hofstede (2001). The IC dimension considers the degree to which societies are integrated into groups and their perceived obligations and dependence on groups. Individualism indicates there is a greater importance placed on attaining personal goals. Collectivism indicates there is a greater im-

portance placed on the goals and well-being of the group. People in collectivist societies generally distinguish sharply between in- and out-groups, while people in individualistic societies treat everyone as a potential in-group member and thus apply universal values to everyone. Cross-cultural research has demonstrated that the United States is the prototypical individualist culture based on the IC dimension (Hofstede 1984; Kim et al. 1994). At the same time, the U.S. show a measurable variation on this dimension (Vandello and Cohen 1999) at state level.

The societal dimensions of collectivism and individualism can be related to the individuals' adherence to specific moral dimensions, as postulated by the Moral Foundations Theory (Graham et al. 2012; Samuel et al. 2016), which include: *care/harm*, fundamental concerns for the suffering of others, including virtues of caring and compassion; *fairness/cheating*, concerns about unfair treatment, inequality, and more abstract notions of justice; *loyalty/betrayal*, concerns related to obligations of group membership; *authority/subversion*, concerns related to social order and the obligations of hierarchical relationships such as obedience, respect, and proper role fulfillment; and *purity/degradation*, with concerns about physical and spiritual contagion, including virtues of chastity, wholesomeness, and control of desires. These foundations are shown to underlie human judgements and decision-making (Weber et al. 2015) on societal topics ranging from vaccine hesitancy (Kalimeri et al. 2019; Beiró et al. 2023), to politics (Iyer et al. 2012), religion, and social cooperation (Haidt 2012; Curry 2016). Here, we place the focal point on the linguistic analysis of values expressed via Twitter, and aim to clarify peoples' dispositions and attitudes towards interpersonal and inter-group processes related to persuasion and communication narratives.

The moral dimensions of MFT have also been linked to political ideologies (Day et al. 2014; Kugler, Jost, and Noorbaloochi 2014), with conservatives emphasizing the in-group relationships and tradition, and liberals endorsing fairness and equal opportunity. In the United States, the mask-wearing measure has also been strongly associated with the partisan divide. Social identity is a primary reason behind people's decision whether to wear a mask during the pandemic (Powdthavee et al. 2021), and surveys show that faith in President Trump is a strong predictor of refusal to social distance, and its effect is largest among individuals high in binding foundations (Graham et al. 2020). Indeed, the U.S. counties that showed strong support for Trump in 2016 practiced significantly lower mask-wearing in 2020 (Kahane 2021). This work illustrates the strong connection between the attitudes expressed towards mask-wearing on Twitter and the the political leaning of those expressing them, and shows the shifts in moral emphasis after the mask mandates are introduced.

Data

We begin by collecting tweets mentioning the keywords "mask", "facemask", "ffp3", and "n95" (the latter two refer to popular kinds of masks), spanning the dates of January 1st to July 30th, 2020, using the GOT3 library (Mottl 2021). These keywords were chosen by considering the spe-

cial Twitter Covid-19, stream[3] and picking the most common English keywords related to masks. This collection results in $18\,245\,298$ tweets from $5\,935\,103$ users. Following recommendations from existing literature, we then perform several filtering steps in order to ensure that the tweets can be used to assess the stance of the user on masking, and that the account is likely to belong to a human living in U.S.:

1. Relevance classifier (details next).
2. Exclude users whose location cannot be mapped to one of U.S. states or Washington DC.
3. Exclude those not having at least 1 English tweet (Garimella et al. 2016).
4. Exclude users with only 1 tweet (Garimella et al. 2016).
5. Exclude top 0.1% of users by the number of tweets (having higher posting rate) (des Mesnards et al. 2022).
6. Exclude users whose friends to followers ratio is >10 (Aggarwal, Rajadesingan, and Kumaraguru 2012).

We proceed by making sure the tweets we collect are indeed about mask-wearing due to the COVID-19 epidemic. Upon manual examination, we find several other topics captured, such as advice on masking during protests, sports-related wear, and beauty products. To remove such content, we use a set of *distant* labels to identify the non-relevant content.[4] We manually annotate a balanced random selection of 300 tweets associated with a distant label and find 96% accuracy in distinguishing between relevant and non-relevant tweets (expected accuracy of a random baseline is 50%). We use $430\,568$ tweets with the respective distant labels to train the relevance classifier. The model is a logistic regression, trained on tf-idf-weighted unigram counts extracted from the tweets and with inverse-proportional class weighting to mitigate class imbalance. As text pre-processing, we remove URLs, numbers, handles, hashtag (the # symbols), retweet indicators, and stopwords, keep whole words, and perform snowball stemming. The accuracy in 5-fold cross-validation is 99.7%. We draw another balanced sample of 300 tweets thus classified and find the accuracy of the classifier to be 82.3% (inter-annotator agreement overlap of 90%, Cohen's kappa of 0.76 between 3 annotators).

Next, we geolocate the tweets identified as relevant by mapping the user location strings to the Geonames ID by using custom string matching,[5] and to U.S. zip codes by using the 'uszipcode' library.[6] We apply basic pre-processing (i.e., stopword and non-ASCII character removal). First, we filter locations assigned to the U.S. territory by 'uszipcode' and then the remaining ones by the Geonames library, aiming to

---

[3]https://developer.twitter.com/en/docs/twitter-api/tweets/covid-19-stream/filtering-rules

[4]We use lists compiled by finding top bigrams (considering the 300 most frequent) in the dataset and then manually labeling them for non-relevance. The irrelevant list is {'hair mask','gas mask','sleeping mask', 'majora's mask','ski mask','eye mask','clay mask','tear gas'} and the relevant list is {'covid','coronavirus','sars-cov-2','sars-cov2','social distanc','socialdistanc'}.

[5]Code available at https://sites.google.com/site/yelenamejova/resources

[6]https://pypi.org/project/uszipcode/

Dataset stage	N tweets	N users
Keyword-based collection	18 245 298	5 935 103
Geo-location filter	5 685 866	1 383 729
Engagement & network filter	3 614 343	598 792

Table 1: Number of tweets and associated users at different stages of data collection.
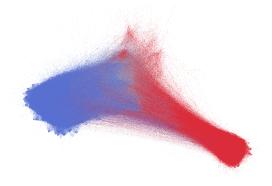

Figure 1: GCC of follower network, colored by METIS score.

recover the GPS coordinates of the smallest possible administrative area. We are able to locate 1 383 729 of the users within a U.S. state or Washington DC. Our sample is representative of the U.S. population distribution (2019 Census estimates) with a Pearson correlation of 0.96 (n=51), which suggests that there is little bias in the sampling of the state in terms of number of users.

The latter four filtering steps aim at excluding users with either too little engagement on the topic or those who post so much that they are likely to be business or automated accounts. We do not use the popular tool Botometer, as a recent study on mask-related tweets shows it to mostly find active human Twitter users (He et al. 2021). At the end of this process, we are left with 647 730 users. Finally, we use the Twitter API Friends call to collect the information about whom these users follow ("followees" or "friends"), thus resulting in the coverage of 598 792 users.

Results

Stance Classification. Following previous work on identifying controversial topics on social media (Garimella et al. 2016, 2018b), we look for a bi-partitioning of the network that would indicate polarization. We look for two sides of the debate because, ultimately, the decision whether to wear a mask is binary. We start our analysis by considering the follower network induced by the users captured in our data, shown in Figure 1. In total, we are able to retrieve 509 755 293 follower-followee relationships, such that followers are all users whose tweets relate to masks. When we constrain the followees to the set of users in our tweet dataset, the network contains $|V| = 598\,792$ users and $|E| = 35\,763\,336$ edges. We use the graph partitioning algorithm METIS (Karypis and Kumar 1998) to partition the network into two groups, repeatedly $N = 100$ times with different random seeds, so to

get an ensemble of partition assignments for each node, and use the average partition assignment for each node across the N repetitions as a polarity score $p \in [0, 1]$ (Mejova et al. 2022).[7] We tune the relative size of the partitions by maximizing the number of users within 95% confidence interval of either extreme, and find the optimal proportion to be 1.1:1 of pro-mask to anti-mask users. As the algorithm does not indicate the actual stances of the users (in fact, it uses no content information whatsoever), each partition is assigned its stance manually by examining 10 sampled users on each side. All users in the same partition expressed the same stance, which suggests a good separation. The sampled users' stance is then propagated to all users within their partition, 0 for pro-mask and 1 for anti-mask.

We also verify the robustness of our result by looking at the community structure of the network. The top-4 communities found by the Louvain algorithm (Blondel et al. 2008), which include more than 90% of the nodes of the network, have median polarity scores of $0.2, 1.0, 0.0, 1.0$, which indicate clear separation and polarization in their composition. The largest community, which includes 18 708 users and is mainly pro-mask, has a slightly larger spread of polarity scores, which may indicate a larger divergence of opinions (and possibly doubts) in the population. This result can also be seen in Figure 4a, the histogram of polarity scores of the network, which shows a smoother decline on the left side than on the right side. Finally, we assign a discrete label to each node according to its polarity score. We heuristically choose the threshold 0.05, which allows us to assign a label to 56.4% of users: 28.8% with 0 and 27.6% with 1, thus leaving 43.5% of users with an unknown label.[8] The roughly equal shares of pro- and anti-mask users are in line with recent other manual annotation work on this topic (Cotfas et al. 2021).

The two sides are colored as blue and red, and unknown as grey in Figure 1. To assess the accuracy of the approach described so far, we sample 60 users from those assigned a stance label by METIS and annotate them manually. To determine the manual label, we use all captured tweets by each user, thus considering in aggregate 1866 tweets. The annotation was done by three authors, all knowledgeable in the U.S. mask debate. Twelve percent of the data was used for inter-annotator agreement, which shows perfect agreement. The exercise revealed an overall precision of 86.4% (95% confidence interval [0.76, 0.94]), with perfect precision for pro-mask class, but only 72.4% for anti-mask case, with several users incorrectly labeled as anti-mask by the algorithm. The network structure suggests some connection of people who express doubts but are not clearly anti-mask with more extreme positions. Indeed, it is likely that the anti-mask stance is more diverse than the pro-mask one, and cannot be easily represented with a single stance, similarly to the vaccine hesitancy case (Cossard et al. 2020).

---

[7]Out of the two possible relabeling of each partition assignment, we choose the one that minimizes the distance from the current node labeling to maintain the partition identity fixed.

[8]The 95% Agresti-Coull confidence interval upper bound for binomial proportions of a polarity score of 0 is 0.0444, which we round to 0.05.
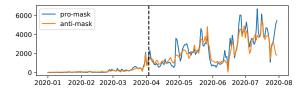
Figure 2: Daily time series of number of tweets by users classified by their mask stance. Vertical line on mandate date.
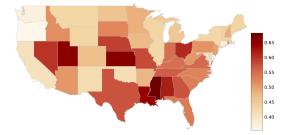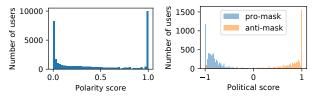

Figure 3: Fraction of anti-mask users over the total geolocated users in the 48 U.S. states.

(a) Polarity score distribution (b) Pol. leaning per mask stance

Figure 4: Distributions of (a) polarity scores of users (computed via METIS), and (b) political leaning of users (inferred via follower relationships on Twitter) grouped by inferred stance on mask-wearing.

Figure 2 shows the daily volume of tweets by users classified as either pro- or anti-mask. The two time series are highly correlated, with a Pearson correlation of $r = 0.938$. The engagement begins roughly at the time of the first CDC recommendation to wear masks on 2020-04-03. The increase in volume confirms a previous study which finds increased "appetite" to share opinions after major mask-related news (Cotfas et al. 2021). The subsequent peaks often revolve around major news stories involving masks, such as one when the U.S. Vice-President Pence visited a hospital without wearing a mask towards the end of April,[9] continuous comparisons of the masking behavior of the two contenders for the U.S. Presidency in late May,[10] and subsequent adjustments to the guidelines by the public health officials who were trying to "correct" their previous messaging in mid-July.[11] Given the high correlation of the two time series, we ask whether this effect is endogenous, i.e., if there is a causal feedback loop whereby one of the two stances answers the other. A Granger causality test on these two time series for time lags from 1 to 14 days finds no strong relationship in either direction. The likely implication of this negative result is that the volume of both stances has a common cause that is exogenous: external events that get discussed on Twitter. Although we cannot exclude an effect with lag shorter than one day, the fact that the networks of two stances are well separated makes this hypothesis less likely.

Figure 3 shows the fraction of anti-mask users in each

state. We find a particularly high concentration in the south, as well as some western states. These findings largely agree with contemporaneous studies of mask-wearing around the U.S. (Katz, Sanger-Katz, and Quealy 2020).

Political Leaning. As outlined in the Introduction, masking regulations have been strongly politicized. According to PEW, when surveyed in 2020, Republicans were more likely to call masks unnecessary, ineffective, oppressive, or unfair and were much less likely than Democrats to be concerned about unknowingly spreading COVID-19 to others.[12]

We compare the mask stance of the users we are able to classify to their political affiliation, which can be glimpsed via their Twitter social network. In particular, we consider the Twitter accounts with known political leaning that the users follows. In accordance to previous work, we assume that users mostly follow accounts that are in agreement with their political views (Golbeck and Hansen 2011), and create a list of prominent political accounts in order to propagate their leaning to their followers. The list, which we make available to the research community,[13] includes 501 accounts of members of the U.S. Congress, 79 governors, 70 party entities, and 67 Attorney Generals, as well as 157 media accounts from allsides.com,[14] and 67 journalists from politico.com.[15] In total, it has a rather balanced set of 487 right-leaning and 454 left-leaning accounts. Note that we do not use prominent politicians such as the U.S. President, as such accounts may be followed purely out of their popularity or importance. For each user in our mask-related dataset, we count the number of accounts they follow in our list: right-leaning N_R and left-leaning N_L. We consider only users who follow at least 5 accounts in our list from either side, and calculate the aggregated political leaning score as $S_{PL} = (N_R - N_L)/(N_R + N_L)$, which results in $S_{PL} \in [-1, 1]$ with 1 the most right-leaning score. Thus, we are able to identify the political leaning of 18 422 users.

---

[9]https://www.washingtonpost.com/politics/pence-meets-with-mayo-clinic-patients-staff-while-not-wearing-face-mask/2020/04/28/57c4200c-897e-11ea-9dfd-990f9dcc71fc_story.html

[10]https://edition.cnn.com/2020/05/26/opinions/biden-mask-trump-ghitis/index.html

[11]https://edition.cnn.com/2020/07/12/politics/jerome-adams-surgeon-general-mask-mandate/index.html

[12]https://www.pewresearch.org/fact-tank/2020/10/29/both-republicans-and-democrats-cite-masks-as-a-negative-effect-of-covid-19-but-for-very-different-reasons/

[13]Available at https://tinyurl.com/poliaccounts

[14]https://www.allsides.com/media-bias/media-bias-ratings

[15]https://www.politico.com/blogs/media/2015/04/twitters-most-influential-political-journalists-205510

Figure 4b presents the distribution of political leaning scores for the two categories of users based on their stances on masks. Pro-mask users are more likely to be following left-leaning accounts, and anti-mask ones the right-leaning ones, with almost no users existing in the middle political ground. In fact, we find a strong polarization at the extremes of the political spectrum, especially for anti-mask users.

Moral Values. Moral values are directly linked to decision making (Karandikar et al. 2019). Since there is an altruistic component to mask wearing, we ask what are the moral values expressed by the holders of the two stances. We assess the moral narratives by employing the MoralStrength lexicon (Araque, Gatti, and Kalimeri 2020), which holds the state-of-the-art performance in moral text prediction. MoralStrength lexicon provides, along with each lemma, the *Moral Valence score*, a numeric assessment that indicates both the polarity and the intensity of the lemma in each moral foundation. According to this lexicon, the Moral Valence is expressed in a Likert-scale from 1 to 9, with 5 to be considered as neutral. Scores lower than 5 reflect notions closer to Harm, Cheating, Betrayal, Subversion, and Degradation, while values higher than 5 indicate Care, Fairness, Loyalty, Authority, and Purity. For each lemma in a tweet and for each foundation, we obtain a moral valence score which is then averaged for each tweet. Negation correction was not applied, as foundation polarities do not directly translate as opposites (e.g., ânot careâ is not the same as âharmâ). The MoralStrength lexicon has a limited linguistic coverage; as a result only the 41.5% of the tweets were found to express a moral foundation. For all the rest, we assigned the value 5, neutral point of the Likert scale. This approach pushes the observed mean towards the center of the scale, but captures the variability of the value across all documents (we discuss the implications of this methodological step in Discussion). To assess statistical significance, we use a Student's t-test (with the Benjamini-Hochberg correction for multiple hypothesis testing) on the scores obtained from tweets written before and after the mandate date (2020-04-03) for each moral dimension.

Figure 5 shows the mean moral value scores of each side in the periods before and after the mandate. Before, the two sides display comparably similar values, except for *care*, which is by far higher for the pro-mask side (significant at $p < 0.001$). However, there is a clear shift in the moral narratives expressed after the mandate by both sides of the debate. To understand the context of these morally-charged expressions, we examine a sample of tweets for each side and value.

First, we find an increase in the valence of *authority* for the anti-mask side ($p < 0.001$), which is mostly accompanied by criticism and mistrust of the decisions made by the authorities. For instance, the anti-mask side associates wearing a mask with weakness and lack of leadership (*"@JoeBiden Real leadership? With that thing on you look feeble"*). Conversely, the pro-mask side sees a lack of leadership in former President Trump's refusal to wear a mask (*"Real leaders won't mind when the mask smears your absurd orange makeup"*). Most of the examples we find in this moral category are indeed

Figure 5: Moral valence in narratives expressed by pro-mask and anti-mask users in the periods before (lighter points) and after (darker points) the mandate. Dot represents the median value while the whiskers represent 5-95% quantiles.

criticism of the authorities, which nevertheless signifies that authority is held in high importance, especially for the anti-mask side (it is the value with the highest valence). The fact that post-mandate the authority-related keywords have higher valence on the anti-mask side suggests stronger criticism of the authorities than the pro-mask side (for whom the increase is significant only at $p = 0.004$ before the correction).

In terms of *care*, both sides have a downwards shift after the mandate. For the pro-mask side, this shift is accompanied by an increase in *fairness* and *loyalty*, which can be interpreted as a shift in focus from personal choice based on caring for others to complying with the mandate. Also, in this case, the spotlight is often on the opposite side (*"Masks protect others, which Trump doesn't care about. He cares only about himself."*). Conversely, anti-mask supporters express themselves by prioritizing much less the notion of *care*, explicitly showing disregard for the protection of others, or simply stating that they do not care about being criticized for not wearing a mask (*"I'm tired of being called a murderer because I don't wear a mask. Cannot understand how people are such thoughtless idiots"*).

In addition, pro-mask supporters express significantly more *loyalty* in their messaging after the mandate ($p < 0.0001$). Upon examining a sample of posts, we find the increase is primarily related to criticism of those not wearing masks as loyalists to a political affiliation or to Trump personally (*"People don't wear masks because of Trump. Dear leader doesn't wear one, loyal followers go along"*).

After the mandate, the *fairness* value increases for both sides (both at $p < 0.0001$). Narratives about the fair treatment of individuals appear to be present equally on both sides, focusing on negotiation when wearing masks is reasonable (pro-mask: *"What if a person not wearing it understands the risks?"*), comparison with other rights violations (pro-mask: *"The only "rights" that are violated is just not wearing a mask? They don't appreciate the nice life they have."*), or whether the criticism is fair (anti-mask: *"Colorado Governor Says People Who Refuse to Wear Mask are "Selfish Bastards" OK, and anybody who wants others to harm their immune system is a POS Marxist propagandist"*). Instead, the valence of *purity* does not change substantially. We find that the notion of purity is not tied to religious views, but to what people
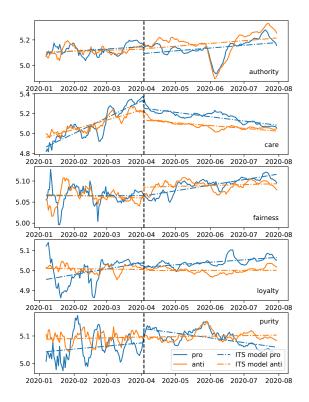
Figure 6: Time series of moral value scores of pro-mask and anti-mask users, along with an interrupted time series analysis model.

Pronoun	Singular		Plural	
Mandate	before	after	before	after
Pro-mask	0.55	0.48	0.11	0.10
Anti-mask	0.53	0.60	0.09	0.10

Table 2: Use of singular and plural personal pronouns in a tweet by side, before and after the mask mandate.

consider 'natural' and 'healthy' (e.g. *"My 8 y.o. niece will be heading back to school but will have to wear a mask! How is it physically or emotionally healthy?"*).

We further examine the evolution of moral values expressed by the two sides in time and model it via an Interrupted Time Series (ITS) linear model. Figure 6 depicts the evolution of the average moral foundation score per side, with a vertical dashed black line indicating the date of the official mask mandate. Complementing our previous analysis, the interrupted time series model shows that for all the moral dimensions, after the mandate, there is an evident change in behavior by both sides. Perhaps the most interesting moral dimension is *loyalty*, whose signal is evidently diverging for two sides exactly after the mandate date and continues the same trend until the end of our data collection. We also observe that not only does the value of *care* decreases, the trend is downward over time, signaling a progressive shift in the debate. Similarly, the value of *purity* has a progressively negative trend for pro-mask side over time. Thus, we find that the temporal dimension of the data can be instructive about the evolution of the rhetoric in terms of divergence between the two sides of conversation and changes in emphasis.

Collectivism vs Individualism. One of the main purposes of mask wearing is the protection of others, an expression of solidarity within the in-group against an external threat. Thus, we turn to the Individualism-Collectivism (IC) dimension (Hofstede 2001), which captures the standing of individu-

als as interdependent members of a collective. We operationalize it via the personal pronouns used in the tweets, mainly first-person singular ("I", "me","mine" etc.) and first-person plural ("we", "us", "ours" etc), following existing literature (Twenge, Campbell, and Gentile 2013).

Table 2 shows the average usage of the two sets of pronouns in the tweets posted by the two sides of the debate, separately before and after the mask mandate. In the period before mask mandate, anti-mask supporters use *I* and other singular pronouns at 0.53, and after the prevalence increases to 0.60 ($+0.07$, $p < 0.0001$), which points to an increased focus on the individual. Instead, pro-mask supporters decrease their usage of singular pronouns (-0.07, $p < 0.0001$). The mention of *We* and other plural pronouns by both pro- and anti-mask supporters remain at the same levels (± 0.01), thus indicating no significant shifts in the focus on the self-identified group. Thus, although having comparatively similar usages of singular pronouns before the mandate, the debate after the government's messaging becomes more individualistic for anti-mask side and less so for advocates of masking.

Information Environment. Finally, we turn to the information context where the debate takes place, in form of links to external sources and information from peers. We begin by performing basic normalization steps on the original tweet text, including removing punctuation, accents, contractions, and stopwords, substituting emojis with a text description,[16] and finally by performing lemmatization. We then identify the most distinguishing words used by each side by calculating the odds ratio of using particular words (lemmas) by one side compared to the other before and after the mask mandate. We filter these words by their total frequency since otherwise rare terms would emerge as most distinguishing. Because the periods before and after the mandate have significant differences in volumes, we set the word frequency threshold to 100 and 800 for the two periods, respectively. Politician names appeared in the top terms and were grouped together, hence resulting in two-word terms. The resulting top terms are:

- Pro-mask: *trumpvirus, Louie Gohmert, Putin, Herman Cain, wearadamnmask, penny, Mayo, GOP, DeSantis, covidiots*

- Anti-mask: *riot, micron, virtue, leftish, loot, sheeple, unhealthy, bacterium, MSM, antifa*

The most distinctive words by each side are highly politicized: aside from references to politicians (similar to other recent findings around masking (Sanders et al. 2021)), we find

---

[16]Using `emoji` library https://pypi.org/project/emoji/.

	Topic	Words
pro	T1 (.35)	distance, wear, social, day, mandate, today, friend, hope, state, stay
	T2 (.33)	wear, people, virus, covid, bad, protect, medical, infect, spread, science
	T3 (.31)	wear, Trump, Cain, rally, covid, Herman, die, Tulsa, death, kill
anti	T1 (.46)	wear, people, covid, virus, work, distance, die, thing, protect, spread
	T2 (.31)	wear, mandate, people, state, store, vote, business, leave, today, order
	T3 (.22)	man, Trump, rally, school, eye, fuck, kid, hope, stupid, big

Table 3: LDA topics with the most representative words, extracted separately from tweets for each mask debate side. The topic prevalence is reported in parenthesis.

several words used to attack the other side. For instance, pro-maskers often use the term 'trumpvirus' to refer to COVID-19, as a political response against the term 'chinavirus' used by President Trump: *"@realDonaldTrump #HermanCain-RIP one more death due to #TrumpVirus . Just think. If only he wore a mask and NOT attended the Tulsa #Coronavirus-Rally"*. The focus on individuals is worth of notice, as in the case of Luie Gohmert, who strongly supported the use of hydroxychloroquine and attended a House Judiciary Committee hearing without wearing a mask, or in the case of Herman Cain, a Republican politician who opposed the mask mandate and later died of COVID-19.[17] Similarly, the derogative term 'covidiots' is used to describe anti-mask supporters.

The opposite side instead focuses on the 'riots' that would happen if lockdown and mask mandates are enforced, and on the claim that masks can stop a 'bacterium' but not a 'micron'-size virus. There are also claims that masks are 'unhealthy' as they impede breathing, and are instead just 'virtue signaling' devices: *"Virtue signaling: mask and gloves. People lived with coronaviruses for 100+ years. Turn off the TV."* A right-wing, anti-establishment sentiment can be inferred from the reference to 'sheeple': *"And the CDC telling everyone to mask up is just another test to see how long the sheeple will obey. #Globalists"*, and from references to supposedly derogative terms such as 'leftish' and 'antifa'. Nonetheless, what is labeled as the anti-mask side is also more varied in its opinions, ranging from just hesitant to fully conspiratorial, somewhat similar to the spread of opinions around vaccines (Cossard et al. 2020).

Moving to higher-level constructs, we aim to uncover common patterns in the argumentation proposed by both sides by applying a topic modeling approach base on latent Dirichlet allocation (LDA) (Blei, Ng, and Jordan 2003). Limitations of LDA clustering of short text are known, still, it offers a good compromise between clustering performance and computational cost (Qiang et al. 2020).[18] To derive the optimum number of topics k, we optimize the topic coherency (C_v

---

Pro-mask		Anti-mask	
rawstory.com	2317	*youtube.com*	3485
cnn.com	2000	thegatewaypundit.com	1341
youtube.com	1751	**etsy.me**	1210
washingtonpost.com	1393	*instagram.com*	912
a.msn.com	935	foxnews.com	903
apple.news	872	**zazzle.com**	796
huffpost.com	758	breitbart.com	781
news.yahoo.com	686	nypost.com	472
flip.it	630	**fineartamerica.com**	453
nytimes.com	587	fxn.ws	393
nbcnews.com	573	westernjournal.com	362
thehill.com	527	*dlvr.it*	344
dailykos.com	486	**pixels.com**	317
instagram.com	458	*buff.ly*	288
businessinsider.com	449	theblaze.com	282
thedailybeast.com	402	bizpacreview.com	268
newsweek.com	371	infowars.com	250
theguardian.com	343	**etsy.com**	246
usatoday.com	337	*ift.tt*	236
yahoo.com	333	cnn.com	231
cnbc.com	307	twitchy.com	217
politico.com	301	**ebay.us**	202
newsbreakapp.com	295	`ncbi.nlm.nih.gov`	196
buff.ly	260	*facebook.com*	192
npr.org	252	`nejm.org`	191
politicususa.com	236	*newsbreakapp.com*	175
apnews.com	231	*a.msn.com*	173
nypost.com	223	google.com	162

Table 4: Counts of the top 28 URL domains posted by pro- and anti-mask users. Domains colored by class: news and news aggregators, social media and social media automator/aggregators (italic font), business platforms (bold font), medical organization (typewriter font).

metric (Röder, Both, and Hinneburg 2015)) models with $k \in [2, 10]$. For both the pro-mask and anti-mask sides, the model with $k = 3$ is the best fit. Table 3 presents the salient keywords that form the corresponding topics. The topics are ranked according to their prevalence, with T1 the most prevalent one, and similarly, the terms are ranked by descending importance for the specific topic.

From the emergent topics, the most prominent one on the pro-mask side concerns the various interventions, including *social distancing* and *wearing* a mask. (*"COVID is not a flu! Everyone needs to wear a mask to protect others from these germs."*), echoing our earlier finding of higher care value of this side. The second topic includes references to *medicine* and *science*, while the third centers around (Republican) political figures. On the anti-mask side, the most prominent topic also concerns the interventions, but instead focuses on whether interventions *work* against the *spread*. The second one puts the mandate in the context of the *businesses* and *stores*, in contrast with T1 from the pro-mask side which speaks about other interventions such as social *distancing* and *staying* at home. The third topic is about political *rallies*, and argues that masks are not useful (since the *eyes* are exposed to the virus: *"Still, even with a mask, you are uncovering the mucous membranes in your eyes"*).

Finally, the two sides of the debate have about the same

proportion of tweets with URLs (around 23%). Earlier works have shown a polarization in terms of news sources accompanying political polarization (Garimella et al. 2021), however here we find differences in behavior beyond sharing news media. Table 4 shows the top domains of the URLs posted by pro- and anti-mask users, along with the counts. Pro-mask users overwhelmingly post URLs pointing to news websites or aggregators. YouTube and Instagram feature prominently in both lists, though anti-mask users favor YouTube more than twice the second most popular domain. Anti-mask tweets also link to a variety of business platforms, including Etsy and Ebay, and lesser-known ones such as Zazzle, a platform for custom-designed products. In addition, anti-mask users link to the governmental agency National Center for Biotechnology Information (NCBI), the New England Journal of Medicine (NEJM), and the Association of American Physicians and Surgeons (AAPS). These findings stand in contrast to a smaller recent study of geolocated-only tweets (He et al. 2021) that finds that anti-mask tweets were less likely to share external information from public health authorities. An explicit comparison of the captured tweet sets would be necessary to resolve these findings.

Discussion & Conclusions

Our analysis reveals that the government messaging about mask wearing provoked—instead of the intended focus on the benefits to the communities and society at large—a marked shift in the moral values towards higher politicization of the issue, with an increased focus on authority and fairness. We argue that interventions targeted to those resistant to mask-wearing should center around these values, instead of appealing to those valued less as the debate goes on.

It is no surprise that we find the mask debate captured on Twitter to be highly politically polarized. It is well-known that polarization in the U.S. political scene has been growing in the last decades,[19] and this growing polarization has important effects on several areas, including public health. As polarization and health behaviors intersect, it is crucial to understand their interaction to design effective health policies. For instance, the ideological and moral facets of health intervention perceptions are closely related to the compliance in the population, and the resulting effectiveness, as shown for instance in recent statistics relating vaccination uptake to political leaning.[20] However, shifting established associations between health-related beliefs and political leaning may be challenging. For instance, the attendees of Donald Trump's rally in August 2021 booed him after he voiced support for the COVID-19 vaccination drive,[21] thus illustrating that even the forerunners of the Republican party may encounter challenges in connecting with their constituents. As the pandemic develops, and more mandates are issued by the governments, constant polling and monitoring is essential in establishing

the public response to these measures (as of mid-2021, the attitudes toward masking are still highly polarized).[22]

We find that different moral values underpin the reasoning emphasized by the two camps. Pro-mask arguments highlight loyalty and fairness while criticizing the opposing leadership. The anti-mask ones, on the other hand, focus on the authoritarian and oppressive aspects of the mandate and show a lack of concern for the effects of their actions on others. In accordance with these results, when examining this phenomenon through the lenses of the Collectivism-Individualism theory, we notice a decisive shift of the anti-mask community towards individualism, with more intense use of first-person personal pronouns. We note the lack of loyalty among the values emphasized by the anti-mask side, which tends to hold a conservative political view, and differs from the commonly observed ones associated with conservatism: authority, loyalty, and purity (Kivikangas et al. 2021; Fulgoni et al. 2016). This may be a response to the pro-social framing of the mask intervention, thus leading its opponents to de-emphasize the in-group narrative usually common to their side. This interpretation may point to motivated reasoning, wherein the desired conclusion modifies the worldview usually taken. Interestingly, the ITS analysis shows divergence over time between the two groups on the value of loyalty, which is alarming since such disagreement, if not adequately addressed, can lead to severe societal polarization.

This polarization seems to have been accompanied by a lively commercial activity. When examining the links posted by the two sides of the issue, we find a prominent existence of commerce-related platforms including *etsy.com*, *ebay.us*, as well as platforms for custom creation of merchandise such as *zazzle.com*. Indeed, a brief search on these websites reveals an assortment of t-shirts, coffee mugs, face masks, and baby bibs with political messages from both sides of the debate. Sometimes historical symbols were used to make a stance, such as the use of the yellow star—like those forced on Jews by Nazi Germany—which was sold by a Nashville store protesting against the vaccination campaign. After community criticism, the item was removed[23]. Our findings suggest that there is an active development of symbolism and aesthetics of the resistance movement, and it would be a fascinating subject of research to uncover the non-verbal representations of the self and the group, expressions of values, and calls to action (McGarry et al. 2019; Awad and Wagoner 2020). Awareness of such symbolism and self-conceptualization is vital for crafting appropriate messages and fostering communication between the two sides.

Despite the volume and span of the data analyzed here, this study has some limitations as most social media studies. Especially when concerning politically charged topics, often a minority of vocal users dominates the conversation (Mustafaraj et al. 2011), thus making the opinions of the "silent majority" challenging to discern. The observed relationships between the stances on masking, moral values, and political stance are limited to those who choose to express themselves

---

[19]https://www.pewresearch.org/politics/2021/11/09/beyond-red-vs-blue-the-political-typology-2/

[20]https://acasignups.net/21/07/20/update-us-covid19-vaccination-levels-county

[21]https://www.theguardian.com/us-news/2021/aug/22/donald-trump-rally-alabama-covid-vaccine

[22]https://www.politico.com/news/2021/08/02/poll-americans-back-return-of-masking-502144

[23]https://www.bbc.com/news/world-us-canada-57297902

vocally. In addition, we use a high-precision vocabulary to measure moral values, which has the side effect of pushing the average valence in our results close to the neutral point, due to a large number of tweets for which we cannot extract moral valence. Traditional surveys are necessary to reach those not as comfortable expressing their opinions online; however, even those have selection biases. The automated tools utilized in this study are not perfect: after manual examination, we find that the network-based classifier achieved an accuracy of 72.4% for the anti-mask class, thus introducing noise in the subsequent analysis. However, from the experience of manually labeling the users, we postulate that the complexities of human expression, including humor and sarcasm, may limit the best possible performance of such a classifier. Similarly, the inference of moral values from the text may struggle with vocabulary mismatch, sarcasm, and self-censorship. Finally, the generalizability of this study is limited by the unique circumstances of an unprecedented global pandemic happening in a hyper-connected world during a politically polarized social environment. These considerations must be taken into account when comparing our findings to new scenarios.

Ethical Statement

The dataset presented here contains only tweets which were publicly available at the time of the collection. We make the dataset publicly available to the research community in compliance with the Twitter Terms of Service,[24] that is, sharing only the tweet IDs of the collected posts, which will have to be re-collected. In this paper, we have rephrased all quoted tweets to prevent re-identification of their authors. This practice ensures that the tweets which have been removed (either by the user or the platform) will not be available. Although large, this dataset does not include users with particular disabilities which may disallow them to interact with the platform, as well as minors and those blocked by Twitter. On the other hand, the content collected here affects not only those who have posted it, but also those who viewed or interacted with it, which may be orders of magnitude more users, since most users are "lurkers" who consume social media content without posting (Van Mierlo 2014). Ultimately, the masking decisions made by people engaging in this deliberation may directly affect the health and life of vulnerable people, such as those with autoimmune disorders or other conditions making them especially vulnerable to COVID-19. Additionally, the sometimes aggressive rhetoric in this material may not be suitable for young Twitter users, or those dealing with mental health issues. Also, although the focus of this study is the moral dimension of the debate, we caution public health communicators not to overemphasize moral or emotional dimensions of their message (or attempt to emotionally manipulate their audience), but rather provide the clearest and most informative messaging possible. Further, we would like to discourage the tools used in this study to be used for targeting individuals espousing particular opinions for harassment or undue surveillance, and to follow the AAAI

---

[24]https://developer.twitter.com/en/developer-terms/

ethical guidelines[25] in the application of these findings.

References

Abd-Alrazaq, A.; Alhuwail, D.; Househ, M.; Hamdi, M.; and Shah, Z. 2020. Top concerns of tweeters during the COVID-19 pandemic: infoveillance study. *Journal of medical Internet research*, 22(4): e19016.

Aggarwal, A.; Rajadesingan, A.; and Kumaraguru, P. 2012. PhishAri: Automatic realtime phishing detection on twitter. In *2012 eCrime Researchers Summit*, 1–12. IEEE.

Araque, O.; Gatti, L.; and Kalimeri, K. 2020. MoralStrength: Exploiting a moral lexicon and embedding similarity for moral foundations prediction. *Knowledge-based systems*, 191: 105184.

Awad, S. H.; and Wagoner, B. 2020. Protest symbols. *Current Opinion in Psychology*, 35: 98–102.

Bayram, A. B.; and Shields, T. 2021. Who Trusts the WHO? Heuristics and Americans' Trust in the World Health Organization During the COVID-19 Pandemic. *Social Science Quarterly*.

Beiró, M.; DìIgnazi, J.; Prado, M.; Bustos, V.; and Kalimeri, K. 2023. Moral Narratives Around the Vaccination Debate on Facebook.

Biddlestone, M.; Green, R.; and Douglas, K. M. 2020. Cultural orientation, power, belief in conspiracy theories, and intentions to reduce the spread of COVID-19. *British Journal of Social Psychology*, 59(3): 663–673.

Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *JMLR*, 3: 993–1022.

Blondel, V. D.; Guillaume, J.-L.; Lambiotte, R.; and Lefebvre, E. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10): P10008.

Campbell, D. T. 1965. Ethnocentric and other altruistic motives. In *Nebraska symposium on motivation*, volume 13, 283–311.

Cinelli, M.; De Francisci Morales, G.; Galeazzi, A.; Quattrociocchi, W.; and Starnini, M. 2021. The Echo Chamber Effect on Social Media. *Proceedings of the National Academy of Sciences*, 118(9): e2023301118.

Cossard, A.; De Francisci Morales, G.; Kalimeri, K.; Mejova, Y.; Paolotti, D.; and Starnini, M. 2020. Falling into the Echo Chamber: the Italian Vaccination Debate on Twitter. In *ICWSM*.

Cotfas, L.-A.; Delcea, C.; Gherai, R.; and Roxin, I. 2021. Unmasking People's Opinions behind Mask-Wearing during COVID-19 Pandemic—A Twitter Stance Analysis. *Symmetry*, 13(11): 1995.

---

[25]https://www.aaai.org/Conferences/code-of-ethics-and-conduct.php

Curry, O. S. 2016. Morality as cooperation: A problem-centred approach. In *The evolution of morality*, 27–51. Springer.

Day, M. V.; Fiske, S. T.; Downing, E. L.; and Trail, T. E. 2014. Shifting liberal and conservative attitudes using moral foundations theory. *Personality and Social Psychology Bulletin*, 40(12): 1559–1573.

des Mesnards, N. G.; Hunter, D. S.; el Hjouji, Z.; and Zaman, T. 2022. Detecting bots and assessing their impact in social networks. *Operations Research*, 70(1): 1–22.

Eikenberry, S. E.; Mancuso, M.; Iboi, E.; Phan, T.; Eikenberry, K.; Kuang, Y.; Kostelich, E.; and Gumel, A. B. 2020. To mask or not to mask: Modeling the potential for face mask use by the general public to curtail the COVID-19 pandemic. *Infectious Disease Modelling*.

Fulgoni, D.; Carpenter, J.; Ungar, L.; and Preoţiuc-Pietro, D. 2016. An empirical exploration of moral foundations theory in partisan news sources. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, 3730–3736.

Garimella, K.; De Francisci Morales, G.; Gionis, A.; and Mathioudakis, M. 2016. Quantifying Controversy in Social Media. In *WSDM*, 33–42.

Garimella, K.; De Francisci Morales, G.; Gionis, A.; and Mathioudakis, M. 2018a. Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship. In *The Web Conference*, WWW, 913–922.

Garimella, K.; De Francisci Morales, G.; Gionis, A.; and Mathioudakis, M. 2018b. Quantifying Controversy on Social Media. *ACM Transactions on Social Computing*, 1(1): 3.

Garimella, K.; Smith, T.; Weiss, R.; and West, R. 2021. Political polarization in online news consumption. *ICWSM*, 152–162.

Gillespie, C. 2020. Does wearing a face mask reduce oxygen— and can it increase CO2 levels? Here's what experts say. https://www.health.com/condition/ infectious-diseases/coronavirus/does-wearing-face-maskincrease-co2-levels. Health. (Accessed on April 28, 2021).

Golbeck, J.; and Hansen, D. 2011. Computing political preference among twitter followers. In *SIGCHI conference on human factors in computing systems*, 1105–1108.

Goldberg, M. H.; Gustafson, A.; Maibach, E. W.; Ballew, M. T.; Bergquist, P.; Kotcher, J. E.; Marlon, J. R.; Rosenthal, S. A.; and Leiserowitz, A. 2020. Mask-wearing increased after a government recommendation: A natural experiment in the US during the COVID-19 pandemic. *Frontiers in Communication*, 5: 44.

Graham, A.; Cullen, F.; Pickett, J.; Jonson, C.; Haner, M.; and Sloan, M. 2020. Faith in Trump, Moral Foundations, and Social Distancing Defiance During the Coronavirus Pandemic. *Moral Foundations, and Social Distancing Defiance During the Coronavirus Pandemic (April 22, 2020)*.

Graham, J.; Haidt, J.; Koleva, S.; Motyl, M.; Iyer, R.; Wojcik, S. P.; and Ditto, P. H. 2012. Moral foundations theory: The pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology, Forthcoming*.

Graham, J.; Haidt, J.; and Nosek, B. A. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96(5): 1029.

Haidt, J. 2012. *The righteous mind: Why good people are divided by politics and religion*. Vintage.

Hao, F.; Shao, W.; and Huang, W. 2021. Understanding the influence of contextual factors and individual social capital on American public mask wearing in response to COVID–19. *Health & Place*, 68: 102537.

Hartsoe, S. 2020. Experts urge empathetic approach when talking with anti-maskers. https://globalhealth.duke.edu/news/experts-urge-empathetic-approach-when-talking-anti-maskers. (Accessed on Dec 2, 2021).

Hatemi, P. K.; Crabtree, C.; and Smith, K. B. 2019. Ideology justifies morality: Political beliefs predict moral foundations. *American Journal of Political Science*, 63(4): 788–806.

He, L.; He, C.; Reynolds, T. L.; Bai, Q.; Huang, Y.; Li, C.; Zheng, K.; and Chen, Y. 2021. Why do people oppose mask wearing? A comprehensive analysis of US tweets during the COVID-19 pandemic. *Journal of the American Medical Informatics Association*, 28(7): 1564–1573.

Hofstede, G. 1984. *Culture's consequences: International differences in work-related values*, volume 5. Sage publications.

Hofstede, G. 2001. *Culture's consequences: Comparing values, behaviors, institutions and organizations across nations*. Sage publications.

Hornik, R.; Kikut, A.; Jesch, E.; Woko, C.; Siegel, L.; and Kim, K. 2021. Association of COVID-19 misinformation with face mask wearing and social distancing in a nationally representative US sample. *Health communication*, 36(1): 6–14.

Iyer, R.; Koleva, S.; Graham, J.; Ditto, P.; and Haidt, J. 2012. Understanding libertarian morality: The psychological dispositions of self-identified libertarians. *PloS one*, 7(8): e42366.

Kahane, L. H. 2021. Politicizing the mask: Political, economic and demographic factors affecting mask wearing behavior in the usa. *Eastern economic journal*, 47(2): 163–183.

Kalimeri, K.; G. Beiró, M.; Urbinati, A.; Bonanomi, A.; Rosina, A.; and Cattuto, C. 2019. Human values and attitudes towards vaccination in social media. In *Companion Proceedings of the World Wide Web Conference*, 248–254.

Karandikar, S.; Kapoor, H.; Fernandes, S.; and Jonason, P. K. 2019. Predicting moral decision-making with dark personalities and moral values. *Personality and Individual Differences*, 140: 70–75.

Karypis, G.; and Kumar, V. 1998. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing*, 20(1): 359–392.

Katz, J.; Sanger-Katz, M.; and Quealy, K. 2020. A Detailed Map of Who Is Wearing Masks in the US. https://www.nytimes.com/interactive/2020/07/17/upshot/coronavirus-face-mask-map.html. (Accessed on Dec 1, 2021).

Keller, S. N.; Honea, J. C.; and Ollivant, R. 2021. How Social Media Comments Inform the Promotion of Mask-Wearing and Other COVID-19 Prevention Strategies. *Int. J. of Environmental Research and Public Health*, 18(11): 5624.

Khan, A. 2020. 'Wearing a mask can give a false sense of security': What the experts say about wearing face masks. https://news.yahoo.com/facemask-proper-use-necessary221717984.html. (Yahoo! News, 30 March. Accessed on April 28, 2021).

Kim, U. E.; Triandis, H. C.; Kâğitçibaşi, Ç. E.; Choi, S.-C. E.; and Yoon, G. E. 1994. *Individualism and collectivism: Theory, method, and applications.* Sage Publications, Inc.

Kivikangas, J. M.; Fernández-Castilla, B.; Järvelä, S.; Ravaja, N.; and Lönnqvist, J.-E. 2021. Moral foundations and political orientation: Systematic review and meta-analysis. *Psychological Bulletin*, 147(1): 55.

Krishnamachari, B.; Morris, A.; Zastrow, D.; Dsida, A.; Harper, B.; and Santella, A. J. 2021. The role of mask mandates, stay at home orders and school closure in curbing the COVID-19 pandemic prior to vaccination. *American journal of infection control*.

Kugler, M.; Jost, J. T.; and Noorbaloochi, S. 2014. Another look at moral foundations theory: Do authoritarianism and social dominance orientation explain liberal-conservative differences in "moral" intuitions? *Social Justice Research*, 27(4): 413–431.

Landsverk, G. 2020. The CDC says healthcare workers who can't get a mask should use a bandana or scarf as a 'last resort' as supplies run short. https://www.businessinsider.com/cdc-recommends-health-workers-use-bandanas-face-masks-crisis-2020-3. (Accessed on April 28, 2021).

Lang, J.; Erickson, W. W.; and Jing-Schmidt, Z. 2021. # MaskOn!# MaskOff! Digital polarization of mask-wearing in the United States during COVID-19. *PloS one*, 16(4): e0250817.

Maaravi, Y.; Levy, A.; Gur, T.; Confino, D.; and Segal, S. 2021. "The tragedy of the commons": How individualism and collectivism affected the spread of the COVID-19 pandemic. *Frontiers in Public Health*, 9: 37.

McGarry, A.; Erhart, I.; Eslen-Ziya, H.; Jenzen, O.; and Korkut, U. 2019. *The aesthetics of global protest: Visual culture and communication.* Amsterdam University Press.

Mejova, Y.; An, J.; Morales, G. D. F.; and Kwak, H. 2022. Modeling Political Activism around Gun Debate via Social Media. *ACM Transactions on Social Computing*, 3532102.

Miller, D. T.; and Morrison, K. R. 2009. Expressing deviant opinions: Believing you are in the majority helps. *Journal of Experimental Social Psychology*, 45(4): 740–747.

Mottl, D. 2021. GetOldTweets3. https://pypi.org/project/GetOldTweets3/. (Accessed on April 1, 2021).

Mustafaraj, E.; Finn, S.; Whitlock, C.; and Metaxas, P. T. 2011. Vocal Minority versus Silent Majority: Discovering the Opinions of the Long Tail. In *International Conference on Social Computing*, 103–110. IEEE.

Pascual-Ferrá, P.; Alperstein, N.; Barnett, D. J.; and Rimal, R. N. 2021. Toxicity and verbal aggression on social media: Polarized discourse on wearing face masks during the COVID-19 pandemic. *Big Data & Society*, 8(1).

Patwa, P.; Sharma, S.; Pykl, S.; Guptha, V.; Kumari, G.; Akhtar, M. S.; Ekbal, A.; Das, A.; and Chakraborty, T. 2021. Fighting an Infodemic: COVID-19 Fake News Dataset. In Chakraborty, T.; Shu, K.; Bernard, H. R.; Liu, H.; and Akhtar, M. S., eds., *Combating Online Hostile Posts in Regional Languages during Emergency Situation*, 21–29. Cham: Springer International Publishing. ISBN 978-3-030-73696-5.

Powdthavee, N.; Riyanto, Y. E.; Wong, E. C.; Yeo, J. X.; and Chan, Q. Y. 2021. When face masks signal social identity: Explaining the deep face-mask divide during the COVID-19 pandemic. *PloS one*, 16(6): e0253195.

Qiang, J.; Qian, Z.; Li, Y.; Yuan, Y.; and Wu, X. 2020. Short text topic modeling techniques, applications, and performance: a survey. *IEEE Transactions on Knowledge and Data Engineering*.

Röder, M.; Both, A.; and Hinneburg, A. 2015. Exploring the space of topic coherence measures. In *WSDM*, 399–408.

Samuel, J.; et al. 2016. Authoritariansim and Collectivism: Antecedents and Consequences Among College Students. *Honors Undergraduate Theses*, (12).

Sanders, A. C.; White, R. C.; Severson, L. S.; Ma, R.; McQueen, R.; Paulo, H. C. A.; Zhang, Y.; Erickson, J. S.; and Bennett, K. P. 2021. Unmasking the conversation on masks: Natural language processing for topical sentiment analysis of COVID-19 Twitter discourse. In *AMIA Annual Symposium Proceedings*, volume 2021, 555. American Medical Informatics Association.

Shapiro, M. A.; and Hemphill, L. 2017. Politicians and the policy agenda: Does use of Twitter by the US Congress direct New York Times content? *Policy & internet*, 9(1): 109–132.

Tajfel, H.; Turner, J. C.; Austin, W. G.; and Worchel, S. 1979. An integrative theory of intergroup conflict. *Organizational identity: A reader*, 56(65): 9780203505984–16.

Twenge, J. M.; Campbell, W. K.; and Gentile, B. 2013. Changes in pronoun use in American books and the rise of individualism, 1960-2008. *Journal of cross-cultural psychology*, 44(3): 406–415.

Van Bavel, J. J.; and Pereira, A. 2018. The partisan brain: An identity-based model of political belief. *Trends in cognitive sciences*, 22(3): 213–224.

Van Mierlo, T. 2014. The 1% rule in four digital health social networks: an observational study. *J Med Internet Res*, 16(2): e2966.

Vandello, J. A.; and Cohen, D. 1999. Patterns of individualism and collectivism across the United States. *Journal of personality and social psychology*, 77(2): 279.

Weber, R.; Eden, A.; Huskey, R.; Mangus, J. M.; and Falk, E. 2015. Bridging Media Psychology and Cognitive Neuroscience. *Journal of Media Psychology*.

WHO. 2020. Coronavirus disease (COVID-19) advice for the public: When and how to use masks. https://archive.vn/cwkJ3. (Archived on March 31, 2020, Accessed on April 28, 2021).