

Authority without Care: Moral Values behind the Mask Mandate Response

Yelena Mejova¹, Kyriaki Kalimeri¹, Gianmarco De Francisci Morales²

¹ISI Foundation, Turin, Italy; ²Centai, Turin, Italy

Keywords: *Twitter, mask wearing, moral values, topic modeling*

Face masks are one of the cheapest and most effective non-pharmaceutical interventions available against airborne diseases such as COVID-19. In the U.S., masks have been met with resistance by a substantial fraction of the populace. Being a prosocial behavior, mask-wearing is influenced by our political ideology [1] and moral values [2, 3] which are directly linked to moral decision making [4]. In this work, we provide a fine-grained analysis of the moral values of those expressing opinions around masking in the U.S. 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, focusing on the U.S., we analyze different facets of this discussion:

1. **RQ1.** How does their stance relate to their political leaning;
2. **RQ2.** What moral values do adherents to pro- or anti-masking stances hold;
3. **RQ3.** What is the information environment around their arguments?

Data Collection. 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 [5]. These keywords were chosen by considering the special Twitter Covid-19, stream¹ and picking the most common English keywords related to masks. This collection results in 18 245 298 tweets from 5 935 103 users. After performing basic pre-processing steps, we employed 430 568 tweets to train a relevance classifier, maintaining only the tweets that were related to the pandemic. We geolocated the relevant tweets by direct string matching of the declared location of the user in the Location strings to the Geonames ID. Finally, we used 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.

Stance Classification. Following previous work on identifying controversial topics on social media [6, 7], we look for a bi-partitioning of the network that would indicate polarization. Constraining 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 [8] to partition the network into two groups,




Fig. 1. GCC of follower network, colored by METIS score.

allowing 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. The two sides are colored as blue and red, and unknown as grey in Figure 1. Manual annotation of a sample of users from each stance revealed an overall precision of 86.4%, 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, although are not clearly anti-mask, with more extreme positions.

Political Leaning. From the mask stance of the users we are able to classify to their political affiliation, which can be glimpsed via their Twitter social network. Previous literature suggests that users mostly follow accounts that are in agreement with their political views [9]. We created a list of prominent political accounts in order to propagate their leaning to their followers² 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.com³ and 67 journalists from politico.com⁴.

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. 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.

²Available at <https://tinyurl.com/poliaccounts>

³<https://www.allsides.com/media-bias/media-bias-ratings>

⁴<https://www.politico.com/blogs/media/2015/04/twitters-most-influential-political-journalists-205510>

¹<https://developer.twitter.com/en/docs/twitter-api/tweets/covid-19-stream/filtering-rules>




Fig. 2. 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.

Moral Values. Applying the MoralStrength lexicon [10] on all tweets we obtained an average moral score per foundation for each tweet. 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. Figure 2 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 $p < 0.001$). There is a clear shift in the moral narratives expressed after the mandate by both sides of the debate. 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. The pro-mask side sees a lack of leadership in former President Trump’s refusal to wear a mask. 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 Benjamini-Hochberg correction for multiple hypothesis testing).

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. 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. In addition, pro-mask supporters express significantly more *loyalty* in their messaging after the mandate ($p < 0.0001$). After the mandate, the *fairness* value increases for both sides (both at $p < 0.0001$). The expected emphasis on the values of authority and purity is accompanied by an atypical dearth of in-group loyalty.

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. 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




Fig. 3. Time series of moral value scores of pro-mask and anti-mask users, along with an interrupted time series analysis model.

Table 1. Use of singular and plural personal pronouns in a tweet by side, before and after the mask mandate.

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

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 [11], which captures the standing of individuals 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 [12]. Table 1 shows that 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. 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.

Table 2. Counts of the top 30 URL domains posted by pro- and anti-mask users. Domains colored by class: news and news aggregators (black), social media and social media automator/aggregators (red), business platforms (blue), medical organization (green).

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	twitch.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
latimes.com	209	dennismichaelylnch.com	152
mol.im	208	aapsonline.org	142

Information Environment. Finally, we apply LDA to the argumentation obtained by each side to find major topics mentioned by either side. The most prominent one on the pro-mask side concerns the various interventions, including *social distancing* and *wearing a mask*. On the anti-mask side, the most prominent topic also concerns the interventions, but instead focuses on whether interventions *work* against the *spread*.

Finally, we turn to the information sources used by the two sides. Table 2 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 conclusion, 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. This interpretation may point to motivated reasoning, wherein the desired conclusion modifies the worldview usually taken. Our findings suggest that there is an active development of symbolism and aesthetics of the

resistance movement. Awareness of such symbolism and self-conceptualization is vital for crafting appropriate messages and fostering communication between the two sides. 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.

-
- [1] Jay J Van Bavel and Andrea Pereira. The partisan brain: An identity-based model of political belief. *Trends in cognitive sciences*, 22(3):213–224, 2018.
 - [2] Jesse Graham, Jonathan Haidt, and Brian A Nosek. Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96(5):1029, 2009.
 - [3] Peter K Hatemi, Charles Crabtree, and Kevin B Smith. Ideology justifies morality: Political beliefs predict moral foundations. *American Journal of Political Science*, 63(4):788–806, 2019.
 - [4] Sampada Karandikar, Hansika Kapoor, Sharlene Fernandes, and Peter K Jonason. Predicting moral decision-making with dark personalities and moral values. *Personality and Individual Differences*, 140:70–75, 2019.
 - [5] Dmitry Mottl. Getoldtweets3. <https://pypi.org/project/GetOldTweets3/>, 2021. (Accessed on April 1, 2021).
 - [6] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. Quantifying Controversy in Social Media. In *WSDM*, pages 33–42, 2016.
 - [7] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. Quantifying Controversy on Social Media. *ACM Transactions on Social Computing*, 1(1):3, 2018.
 - [8] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing*, 20(1):359–392, 1998.
 - [9] Jennifer Golbeck and Derek Hansen. Computing political preference among twitter followers. In *SIGCHI conference on human factors in computing systems*, pages 1105–1108, 2011.
 - [10] Oscar Araque, Lorenzo Gatti, and Kyriaki Kalimeri. Moralstrength: Exploiting a moral lexicon and embedding similarity for moral foundations prediction. *Knowledge-based systems*, 191:105184, 2020.
 - [11] Geert Hofstede. *Culture’s consequences: Comparing values, behaviors, institutions and organizations across nations*. Sage publications, 2001.
 - [12] Jean M Twenge, W Keith Campbell, and Brittany Gentile. Changes in pronoun use in american books and the rise of individualism, 1960-2008. *Journal of cross-cultural psychology*, 44(3):406–415, 2013.