Pandemic Culture Wars: Partisan Differences in the Moral Language of COVID-19 Discussions

Ashwin Rao*^{1,2},Siyi Guo*^{1,2}, Sze Yuh Nina Wang³, Fred Morstatter¹, and Kristina Lerman¹ Information Sciences Institute, ²University of Southern California, ³University of Regina {mohanrao,siyiguo}@usc.edu, syw332@uregina.ca,{fredmors, lerman}@isi.edu

Abstract-Effective response to pandemics requires coordinated adoption of mitigation measures, like masking and quarantines, to curb a virus's spread. However, as the COVID-19 pandemic demonstrated, political divisions can hinder consensus on the appropriate response. To better understand these divisions, our study examines a vast collection of COVID-19related tweets. We focus on five contentious issues: coronavirus origins, lockdowns, masking, education, and vaccines. We describe a weakly supervised method to identify issue-relevant tweets and employ state-of-the-art computational methods to analyze moral language and infer political ideology. We explore how partisanship and moral language shape conversations about these issues. Our findings reveal ideological differences in issue salience and moral language used by different groups. We find that conservatives use more negatively-valenced moral language than liberals and that political elites use moral rhetoric to a greater extent than non-elites across most issues. Examining the evolution and moralization on divisive issues can provide valuable insights into the dynamics of COVID-19 discussions and assist policymakers in better understanding the emergence of ideological divisions.

Index Terms—social networks, Twitter, moral foundations, polarization, Covid-19

I. INTRODUCTION

The COVID-19 pandemic presented a major challenge to society, testing its resilience and capacity to respond to a crisis. Slowing the rapidly spreading illness required coordinated adoption of non-pharmaceutical interventions recommended by public health experts, such as wearing a mask, staying home, social distancing, and getting vaccines once they became available. However, the response to the pandemic became politically polarized, with Republicans and Democrats disagreeing both on the gravity of the crisis and the appropriate measures to address it [1]. Some Republican leaders downplayed the severity of the virus and expressed skepticism about the effectiveness of measures like maskwearing and social distancing, while some Democratic leaders advocated for more stringent measures like closing schools, businesses, and even parks and beaches. Policy differences in pandemic mitigation emerged at the state level as well. Republican-led states lifted restrictions on businesses and public gatherings earlier than Democratic-led states. When the COVID-19 vaccine was approved, many Democratic-led states introduced vaccine and mask mandates for returning to work and school, while some Republican-led states passed bills to ban such mandates. These political divisions hampered effective pandemic response, leading to more than a million deaths in US, one of the highest mortality rates among advanced economies [2].

To understand how the pandemic response became so polarized, we study a large corpus of pandemic-related tweets [3], comprising over 270M tweets from 2.1M users around United States. We focus on five highly politicized and contentious wedge issues: coronavirus origins, lockdowns and business closures, masking, online education, and vaccines. These issues have been previously found to be salient during the pandemic [4]–[10]. We classify user ideology based on the partisanship of information sources they share and also describe a method for identifying issues raised in a tweet. These models enable us to label millions of tweets and users, thereby allowing us to study polarized discussions at scale.

We study online conversations about the pandemic from the lens of Moral Foundations Theory (MFT) [11]. This theory provides a framework for understanding how moral values shape people's political attitudes and behaviors. MFT proposes that individuals' values and judgments can be described by five moral foundations: care/harm, fairness/cheating, lovalty/betrayal, authority/subversion, and sanctity/degradation. Individuals and groups vary in how much emphasis they place on each foundation. While both liberals and conservatives consider the care and fairness to be important, conservatives are thought to endorse the latter three foundations to a greater extent than liberals [12]. Studying online conversations from the perspective of MFT can provide insights into the ways in which partisans use moral appeals to articulate and defend their political positions. It can also help us understand how to frame messages so as to appeal to the values of each group.

We organize our investigations around the following research questions:

- 1) What COVID-19-related issues did liberals and conservatives discuss online?
- 2) How did attention to issues vary by ideology? Which issues were more salient to different groups and when?
- 3) How did liberals and conservatives differ in their use of moral language on different issues? How did the groups react to COVID-19-related events?
- 4) How does moral language use differ between political elites and non-elites?

We propose and evaluate a weakly-supervised method to identify issue-relevant tweets. This novel method extracts relevant phrases from publicly available Wikipedia pages and then

^{*}These authors contributed equally to this work

TABLE I Issue Detection: Wikipedia Articles

Issue	Wikipedia Articles
Origins	Plandemic, Investigations into the origin of COVID-19, COVID-19 lab leak theory
Lockdowns	COVID-19 lockdowns, COVID-19 protests in the United States, U.S. federal government response to the COVID-19 pandemic, Protests against responses to the COVID-19 pandemic, Social impact of the COVID-19 pandemic in the United States, Stay-athome order, Social distancing measures related to the COVID-19 pandemic, Quarantine
Masking	Face masks during the COVID-19 pandemic, Face masks during the COVID-19 pandemic in the United States, Maskne, Anti-mask sentiment
Education	Impact of the COVID-19 pandemic on education, Homeschooling during the COVID-19 pandemic, Impact of the COVID-19 pandemic on education in the United States
Vaccines	COVID-19 vaccination in the United States, COVID-19 vaccine, COVID-19 vaccine hesitancy in the United States, COVID-19 vaccine misinformation and hesitancy, Deaths of anti-vaccine advocates from COVID-19, Herman Cain Award
Baseline	COVID-19 pandemic, COVID-19 pandemic in the United States, Politics of the United States

uses them to detect issues in the text of tweets. In addition, we use state-of-the-art computational methods to detect moral foundations in tweets [13] and to infer the political ideology of users [14]. These models enable us to track partisan attention to polarized issues at scale from the start of the COVID-19 pandemic to November 2021.

We uncover important differences in how liberals and conservatives discuss issues over time, as well as their use of moral language. In particular, we find that discussions of issues and reactions to events were negatively valenced, especially among conservatives. Liberals showed more sustained interest in the topic of education, while conservatives focused on the origins of COVID-19. We also unveil differences in the moral appeals of political elites and non-elites across the ideological spectrum. Understanding the moralization on divisive issues can provide valuable insights into the emergence of polarization and help public health experts better tailor messaging to the values and concerns of each group.

II. RELATED WORK

1) Polarization: American society has been divided by "culture wars" [15], with liberals and conservatives disagreeing on controversial issues (also known as wedge issues) like abortion, women's and LGBTQ+ rights [16]–[18]. This polarization was further exacerbated by the COVID-19 pandemic. Partisans disagreed on the virus's existence and severity, its origins, mitigation strategies like social distancing and masking, and vaccine mandates. Political ideology was found to explain differences in compliance with health guidelines and

mandates [19], [20]. A study of elite messaging on Twitter at the onset of the pandemic [21] found that Democratic elites—discussed the pandemic more frequently and emphasized its—threats to public health, while Republican elites focused on China and the pandemic's impact on businesses.

These divisions extend beyond political elites: nearly 3 out of 10 Americans believed that COVID-19 was created in a lab, with 4 in 10 Republicans endorsing this claim [4]. Government efforts to restrict the spread of COVID-19 through lockdowns were met with resistance in the form of protests and demonstrations, which were largely driven by conservatives: nearly 52% of conservatives in the United States preferred to have fewer COVID-19 restrictions as opposed to only 7% of liberals [5]. These divisions can have important behavioural consequences: following and interacting with right-leaning political elites, influencers and hyper-partisan media sources was found to be associated with vaccine hesitancy in the US [9], [10].

2) Moral Foundations: Differences in moral values may underlie these partisan divisions. We analyze these through the framework of moral foundations theory (MFT), which proposes five moral foundations of care, fairness, in-group loyalty, respect for authority, and purity/sanctity [11]. These foundations can be further categorized into virtues (care, authority, fairness, loyalty & purity) and vices (harm, subversion, cheating, betrayal & degradation). A number of language analysis techniques have been used to quantify the moral foundations in text data (e.g. [13], [22], [23]), with past work suggesting that moral appeals may increase message diffusion [24], [25].

To comprehend and mitigate issue-based political divisions, it is crucial to understand how individuals attach moral significance to their viewpoints [26]. Some past work has done this in the context of COVID-19. For example, [27] found that valuing the care foundation was correlated with compliance with COVID-19 public health recommendations. Similarly, [8] found that endorsing the fairness and care moral foundations predicted compliance with calls for staying-at-home, wearing masks, and social distancing, while purity predicted compliance with face masks and social distancing. Others have found that framing masking-related health messaging using ideologymatched moral arguments was effective for liberals, but not conservatives [7]. Moral values can also be used to predict county-level vaccination rates: counties with residents who prioritize moral concerns about purity had lower vaccination rates while counties whose residents prioritized fairness and loyalty had higher vaccination rates [28].

Some initial work has assessed COVID-19 vaccine discourse on Twitter. [29] found that liberal tweets about the COVID-19 vaccine expressed care, fairness, liberty¹ and authority moral foundations more than conservative tweets, while oppression and harm were referenced more by conservatives. Similarly, [30] found that care/harm was associated with provaccine sentiment whereas, liberty/oppression was correlated

¹Note that some researchers include a sixth foundation, liberty/oppression.

TABLE II
SAMPLE TWEETS HIGHLIGHTING WEDGE ISSUES IN COVID-19 ONLINE DISCUSSIONS.

Issue	Tweets	
Origins	No matter what the Chinese Communist Party says, given the mounting evidence, the most likely origins for the China virus are the Wuhan labs studying bats and coronavirus.	
Lockdowns	This is a GREAT idea. We're all in this together. Take care of each other. #StayHome #TakeItSeriously #FlattenTheCurve #COVID19	
Masking	We're in the middle of a pandemic and y'all are still coughing and sneezing without covering your mouths? Come on now.	
Education	More glimmers of hope as we "safely" move forward and open up Texas A&M University while containing #COVID19.	
Vaccines	You are joking right? Zero sympathy for anti- vaxxers who quit their jobs rather than get vac- cinated. They put us all at risk and make the pandemic prolonged for the world.	

with anti-vaccine attitudes. While vaccinations are a critical polarizing issue in the discussion of COVID-19, no work as of yet has explored differences in moral appeals across a broader range of contentious COVID-19 issues.

III. METHODS

A. Data

We use a publicly available dataset [3] consisting of 1.4B tweets about COVID-19 posted between January 21, 2020 and November 4, 2021. These tweets contained one or more COVID-19-related keywords, such as coronavirus, pandemic, and Wuhan, among others. We utilized Carmen [31], a geolocation identification technique for Twitter data to assign tweets to locations within US. This method leverages tweet metadata, including "place" and "coordinates" objects that encode location information, such as country, city, and geographic coordinates. Additionally, the technique uses mentions of locations in a user's bio on Twitter to infer their location. A manual review confirmed that this approach was effective in identifying a user's home state. As a result, we were left with 270 million tweets generated by 2.1 million geo-located users in the United States.

B. Issue Detection

As the pandemic progressed, a set of polarizing issues such as origins of the coronavirus, lockdowns and masking mandates, school closures and online schooling and vaccines, emerged and grew to dominate conversations about COVID-19 [4]–[6], [9], [10]. Previous studies [24], [32] relied on a small set of manually-selected keywords and phrases to harvest issue-relevant tweets. However, there is a lack of data on the reliability of these methods, and a systematic framework to identify issue-relevant conversations is largely missing. While

issue detection has analogues to topic-detection [33], [34], one cannot control for the stochasticity of topics identified in traditional topic-detection approaches. Although [35] presented a corpus of several thousand news articles annotated with frames on traditional policy issues such as smoking, same-sex marriages, and immigration, there has not been an attempt to propose a framework for identifying COVID-19 issues in social media discourse. The lack of expertly annotated data on the issues we explore in this work prevents us from employing supervised learning techniques.

We define the origins issue to encompass discussions surrounding the origins of the pandemic, including topics such as pangolins, gain of function research, wet-markets, and bats. The lockdown issue comprises content pertaining to early state and federal mitigation efforts, such as quarantines, stay-athome orders, business closures, reopening, and calls for social distancing. The masking issue is defined by discussions on the use of face coverings, mask mandates, mask shortages, and anti-mask sentiment. Education-related content involves tweets about school closures, reopening of educational institutions, homeschooling, and online learning during the pandemic. The vaccines issue pertains to discussions about COVID-19 vaccines, vaccine mandates, anti-vaccine sentiment, and vaccine hesitancy in the US.

We describe a weakly-supervised method to harvest relevant keywords from Wikipedia pages discussing these issues (see Table I). We use SAGE [36] to identify distinctive keywords and phrases that are relevant to a specific issue. SAGE calculates the deviation in log-frequencies of words from a baseline lexical distribution. As a baseline, we use Wikipedia articles discussing general aspects of the pandemic and politics in the US (see I). We concatenate these pages for each issue. Depending on how we tokenize the corpus, we can use SAGE to identify issue-relevant n-grams. In this study, we restrict analysis to unigrams, bigrams and trigrams. Using SAGE we then compare tokens from issue-relevant articles to baseline tokens to identify keywords and phrases that uniquely define each issue. We manually verify the keywords identified by SAGE to ensure precision and relevance. Word clouds in Fig. 1 show the top 50 keywords/phrases extracted from Wikipedia articles for each issue. We assess the relevance of a tweet to a issue by the presence of these keywords and phrases. For example, terms such as 'Wuhan labs' or 'wet markets' would indicate a tweet is about the origins of the pandemic, while 'cover your mouth', 'N-95 masks' would indicate a maskingrelated tweet. Examples of tweets discussing these issues are shown in Table II.

a) Validation: We evaluate issue detection on a subset of tweets, which we draw at random with the constraint that there are at least five tweets for each combination of issues and moral foundation. This gave us a test set of 784 tweets, which were then labeled by five trained annotators. Table III shows that the mean Cohen's Kappa between each pair of annotators for different issues were all above 0.7. The Multiannotator Fleiss' Kappa values are lower because it relies on agreement among all five annotators rather than pairs of



Fig. 1. Wordclouds showing phrases identified from Wikipedia articles using SAGE for each issue. We show the top-50 most frequently occurring phrases in tweets. Bigrams and trigrams are connected by __ for better visualization.

TABLE III
INTER-RATER AGREEMENT AND EVALUATION OF ISSUE DETECTION.

Issue	Pairwise Cohen's κ	Multi- annotator Fleiss' κ	F1- Score	Support
Origins	0.73 ± 0.06	0.45	0.51	25
Lockdowns	0.79 ± 0.04	0.47	0.90	108
Masking	0.92 ± 0.04	0.55	0.83	101
Education	0.71 ± 0.07	0.50	0.74	54
Vaccines	0.87 ± 0.05	0.53	0.92	84

annotators. Finally, we evaluate performance of issue detection method on this test set and show that F1 scores for most issues are above 0.7 (Table III), indicating good model performance.

C. Morality Detection

We capture the moral sentiments in tweets based on the Moral Foundations Theory. We train our morality detection model on top of the transformer-based pretrained language model BERT [37]. We train BERT with three Twitter datasets, including a manually annotated COVID dataset [38], the Moral Foundation Twitter Corpus dataset with six different topics [39], and a dataset of political tweets published by US congress members [40]. Fusing an in-domain training set that is also about COVID-19, along with other datasets consisting of various topics, we are improving the model generalizability when applied to the target data [13].

D. Ideology Detection

A number of methods identify the political ideology of social media users, with recent works [14], [41]–[43] leveraging URLs that users share in tweets. We use the method discussed in [14] to estimate the ideology of individual users. This method relies on Media Bias-Fact Check (MBFC) [44] to compute ideology scores for individuals. MBFC provides ideological leaning for over 6K Pay-Level Domains (PLDs) with leaning categorized as Left/Hardline Liberal (0), Left-Center (0.25), Least-Biased/Center (0.5), Right-Center (0.75), Right/Hardline Conservative (1). Each individual's ideological leaning is estimated to be the weighted average of the leanings of the URLs they share.

These scores are then binarized with a threshold of ≤ 0.4 for liberals (0) and ≥ 0.6 for conservatives (+1). The model

described in [14] then leverages a fastText embedding model pre-trained on Twitter data [45], to generate embeddings for user's tweets over time. These embeddings serve as features to train a Logistic Regression classifier to predict the binarized ideological leanings for all users, including users who have not shared URLs from domains in the MBFC. We were able to identify ideology for 2.1M US users comprising of 1.6M liberal users and 500K conservatives.

a) Validation: We assess the level of agreement between the ideology scores estimated by the methods proposed in [14] and [46], which uses the follower graph to estimate user ideology. We discretize the continuous ideal point estimates in [46] using a threshold of < 0 for liberals and > 0 for conservatives, which we respectively label as 0 and +1. We identify approximately 35K users who appear in both datasets and calculate the F1-score between the two models. The resulting F1-score of 0.75 and Jaccard overlap of 0.87 indicates a high level of agreement between the two methods.

IV. RESULTS

Figure 2 shows the daily share of tweets, retweets and replies discussing each issue, smoothed using a 7-day rolling window, with vertical lines marking major issue-related events. The five issues we monitored accounted for small share of all messages at the beginning of the pandemic, but grew to dominate all discussions. Lockdowns were among the first mitigation efforts carried out by state governments, with California the first state in the mainland US to issue a stayat-home order on March 19, 2020. This date marks the start of an increase in lockdown related discussions. Although the peak in the masking discussions did not occur until the Black Lives Matter protests in June 2020, Fig. 2 indicates an early upward trend on April 3rd, 2020, when the CDC recommended face coverings to prevent infections. On July 8, 2020 President Trump promised to cut funding to schools that did not reopen for in-person instruction, which is associated with a spike in education related discussions. Similarly, Pfizer-BioNTech's announcement [47] on November 9, 2020 that their COVID-19 vaccine was highly effective in Phase 3 clinical trials created a spike in vaccine related tweets. A spike in the discussions of COVID-19 origins is seen on May 23, 2021 when the Wall Street Journal published an article citing an US Intelligence Report that three researchers from the Wuhan Institute of Virology had sought hospital care in November 2019 shortly before the outbreak [48]. The story marked the first major

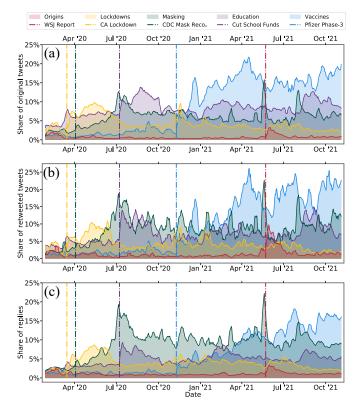


Fig. 2. Dynamics of attention to issues. Daily fraction of (a) original tweets, (b) retweets and (c) replies related to each issue. We use 7-day rolling average to reduce noise. Major events are marked with vertical lines of the same color as the issue and placed along the timeline.

report on the origins of the pandemic by a major news organization. These results provide a qualitative check of robustness of our issue detection, as well as insights into importance of issues at different times. Fig. 2 also shows differences in the engagement across issues. For example, tweets about masking seemed to garner more replies, suggesting that they may have generated more discussion or controversy.

A. Partisan Differences in Attention to Issues

We examine the relative prevalence of issues in the tweets of conservatives and liberals. The user population discussing the pandemic in our sample is ideologically unbalanced, with 5.143 times as many tweets by liberals as conservatives, on average; therefore, we cannot directly compare the volume of messages on an issue posted by each side. Moreover, the average imbalance hides a high degree of variability: the number of active liberals or conservatives may vary over time in response to events, making the issues more or less salient to each group. Using a measure like share (percentage) of tweets on an issue fails to capture the increased interest from a group. To address these challenges, we propose to use *the ratio of the daily number of original tweets* to capture differences in partisan attention to issues.

Figure 3 shows the time series of the daily tweet ratio $r(i,d) = \tau_l(i,d)/\tau_c(i,d)$, where $\tau_l(i,d)$ and $\tau_c(i,d)$ represent the number of original tweets posted by liberals and

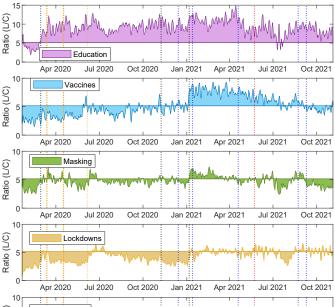


Fig. 3 Partisancies ymmetries in attention to issues. Each plot shows the ratio of the daily number of original tweets on a given issue posted by liberals and conservatives (smoothed using a three day window). Vertical lines correspond to major events during the time period: first US death 02/29/20, national emergency/ude/daration 202/13/20,202Liberate/02/II" Jux2021s 04/17/20, BLM protests 06/6/20, Biden defeats Trump 11/8/20, Covid vaccine rollout 12/14/20, Jan 6 2021, Covid vaccine available to 65+ in CA 01/13/21, Covid vaccine available to 16+ 04/19/21, WSJ Covid lab leak story 05/23/21, FDA approves first covid vaccine 08/23/21, federal vaccine mandate 09/9/21.

conservatives respectively about an issue i on day d. Each plot shows the deviation of the ratio from its mean value. Overall, the education issue is of greater interest to liberals, with the mean ratio of liberal to conservative tweets being 8.951. In contrast, the issue of the origins of COVID-19 is of most interest to conservatives, with the mean ratio of 2.126. Masking and lockdowns are of slightly more interest to conservatives, with the mean ratios of 4.743 and 4.282 respectively. Different events make the issues more salient to liberals or conservatives, as captured by the positive or negative deviations from the mean. Specifically, the first US death due to COVID-19 was a turning point for liberals: their share of issue discussions increased. This is also confirmed by the event analysis, which shows more liberal reactions to this event. With respect to specific issues, although education is a liberal issue, conservatives show comparatively more sustained interest in it during the summer. The availability of COVID-19 vaccines at the end of 2020 (vertical lines show important events, including vaccine roll-out on December 14, 2020) led to more tweets from liberals about the vaccine compared to conservatives. Likewise, national emergency declaration in March 2020, and anti-lockdown protests the following month brought sustained interest from conservatives to the lockdowns issue. Following vaccine mandates for federal workers on September 9, 2021, the vaccines issue became a greater focus for conservatives, likely due to their outrage about the mandates. After Wall Street Journal's May 2021 report [48] claiming that COVID-19 originated at the Wuhan Institute of



Fig. 4. Partisan differences in framing each issue. Values show the log odds ratio of phrases used by liberals (blue) and conservatives (red) in issue-relevant tweets across the time period.



Fig. 5. Partisan differences in moral language use along each issue. Values indicate percentage of tweets shared by liberals (blue) and conservatives (red) containing the specific moral foundation. Values for Care and Harm are downscaled by a factor of two for better visualization.

Virology increased conservative attention to the COVID-19 origins issue. While we limit our discussion here to original tweets due to space constraints, we see similar trends in our analysis of retweets and replies.

To characterize differences in liberal and conservative tweets about issues, we identify adjectives that were used to modify issue-relevant phrases (anchors) extracted earlier. We leverage SpaCy's Dependency Parser to extract dependency relations of the form: \mathbf{XX} - $amod \rightarrow ANC$ and \mathbf{XX} - $amod \rightarrow YY \leftarrow amod$ -ANC, where XX refers to the extracted adjective and ANC refers to an issue-relevant phrase. We then pair the extracted adjective with its corresponding anchor as a phrase of the form XX-ANC and compute the log-odds ratio of these phrases being expressed in liberal and conservative tweets. Wordclouds in Fig. 4 depict the top 10 most-likely used phrases in liberal (blue) and conservative (red) issue-discourse.

Conservatives use conspiratorial phrases such as "batshit conspiracy", "fake plandemic", and "military bioweapon" to discuss origins of the pandemic. In contrast, liberals talk about "likely bats" and "medical virology". When it comes to lockdowns, conservatives express negative sentiment, framing them as "deadly", "unconstitutional", and "stringent". Liberals, however, approach the topic cautiously, using phrases such as "premature reopening", "safe reopening", and "soft lockdowns". Similar patterns emerge in the discourse on masking, with conservatives deeming mask mandates as "unconstitutional" and "useless", while liberals focus on "reusable masks" and "mandatory masking". Concerning education, conservatives express distrust in mitigation strategies employed by educational institutions, using terms like "fascist education" and "unmasked students", while liberals highlight "hybrid learning", "vulnerable students", and "early education." On the issue of vaccines, conservatives exhibit skepticism, referring to

"experimental mRNA" and "dangerous vaccines". In contrast, liberals emphasize the benefits, discussing "free vaccinations" and "mobile vaccines".

B. Partisan Differences in Moral Language

To study partisan differences in the use of moral language, we measure the share of issue-related tweets posted by each group that expresses each moral foundation. We examine moral valence by separately calculating the share for vices and virtues. The proportion $\rho(m,i)$ of original tweets related to issue i that express moral foundation vice or virtue m is: $\rho(m,i) = \tau(m,i)\tau(i)$, where, $\tau(i)$ returns the number of original tweets on issue i.

Figure 5 shows partisan differences in the proportions ρ . Across issues, we find that care and harm the most prevalent moral virtues and vices, followed by authority and subversion, with fewer tweets expressing purity and degradation, despite the relevance of that moral foundation to the infection concerns. Conservatives express more moral vices than virtues in their tweets. There are other partisan differences in the use of moral language. For example, while both appealed to subversion, betrayal, and cheating when discussing the origins of COVID-19, liberals focus more on harm (see Fig. 5(a)). This could reflect greater concern among liberals about the potential harms of these speculations, as they led to a rise in anti-Asian hate crimes across the nation [49].

Across lockdowns, masking, education, and vaccine issues (Figs. 5(b)-(e)), conservatives expressed greater levels of harm and subversion than liberals, while liberals expressed more care. This may indicate greater support for mandates (e.g., school shutdowns and vaccine mandates) among liberals, and opposition by conservatives, consistent with partisan divides on shutdowns, social distancing, and masks [50].

Overall, we find similarities and important differences in the moral language use by liberals and conservatives. Conservatives appear to make more negatively valenced moral appeals when it comes to lockdowns, masking, education, and vaccines, while liberals use greater care language.

C. Partisan Differences in Reactions to Events

To understand how liberals and conservatives react to different events during the pandemic, we perform interrupted time series analysis. Our goal is to quantify both the changes in the use of moral language and the amount of discussions about different issues. First, we compute the time series of the partisan tweet ratios for each concern and moral foundation, $r(i,m,h) = \tau_l(i,m,h)/\tau_c(i,m,h)$, where $\tau_l(i,m,h)$ and $\tau_c(i,m,h)$ represent the number of original tweets posted by liberals and conservatives respectively about an issue i and expressing a moral foundation m in an hour h. We choose to use hourly data here to get more data points and therefore more robust regression results.

We select four major events during the pandemic: (1) the first US death due to COVID-19 on 02/29/20, (2) the Black Lives Matter (BLM) protests on 06/06/20, (3) FDA approving the first COVID vaccine on 08/23/21, and (4) the federal vaccine mandate on 09/09/21. For each event, we select the time window from seven days before the event to four days after, as discussions on Twitter usually die down within that time period [51]. We perform linear regression within this time window on each time series of the hourly tweet ratios:

$$r(i, m, h) = \beta_0 + \beta_1 h + \beta_2 \mathbb{1}_{after\ event} + \beta_3 h * \mathbb{1}_{after\ event}$$

where h is hour, and $\mathbb{1}_{after\ event}$ is 0 before the event and 1 after the event. To represent the change associated with an event, we use $\frac{\beta_3}{mean(r(i,m,h))}$, the coefficient for the interaction term normalized by the mean of this time series. Different issues have different mean ratios, so this normalization allows us to fairly compare changes across issues. Figure 6 shows an example of such interrupted time series analysis with linear regression.

Figure 7 shows how liberals and conservatives reacted differently to different events. The color of the cell relatively change in the discussions of the specific issue and moral framing: $\frac{\beta_3}{mean(r(i,m,h))}$. A positive value, colored blue, means that there was relatively more engagement from liberals in response to this event, and a negative value, in red, means that there was more conservative engagement. We see that liberals and conservatives had distinct moral reactions to different events. To begin with, the first US death attributed to COVID-19 on 02/29/20 elicited more reactions from liberals. This implies that liberals were paying more attention to the COVID-19 disease at this early stage of the pandemic. During the Black Lives Matter (BLM) protests on 06/06/20, conservative engagement on the issues of education and lockdowns increased overall. However, reliance of the cheating, loyalty/betrayal and subversion moral foundations on the issues of masking and lockdowns significantly increases among liberals. This suggests that liberals expressed more concerns about

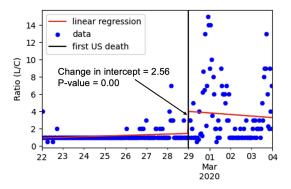


Fig. 6. Interrupted time series analysis of response to the first US death (02/29/20). We carry out linear regression from seven days before the event to four days after it, using the time series of hourly ratio between the number of original tweets posted by liberals and conservatives respectively about an issue (specifically, *education*) and expressing a moral concern (specifically, *subversion*). The discontinuity, i.e., change in the intercept, shows more discussions expressing subversion about education from liberals than conservatives in response to the first US death.

COVID safety during mass gatherings. More interestingly, the FDA approval of the first COVID vaccine exhibited diverging moral reactions. While, liberals focused on fairness more, the conservative reactions to this event were more negatively valenced, expressing more cheating, betrayal and subversion about different issues, suggesting a more pessimistic view of the COVID vaccines. On another vaccine-related event, the federal vaccine mandate, liberals consistently had more discussions about fairness, especially when talking about the issues of education and vaccines. These results show us the different moral reactions in the two partisan groups towards important events and agendas along the pandemic timeline.

D. Differences in Moral Language of Elites and Non-Elites

Past research has shown that moral language used by political elites [52], [53] in the US can influence the opinions and moral reasoning of non-elites [54]. In this work, we compare the moral language used by elites and non-elites. Political elites in US include prominent members of Congress, politicians, journalists and media pundits. Previous studies [55], [56] have curated extensive lists of Twitter handles for over 4K US political elites. To ensure a fair comparison with non-elites, we randomly sample non-elite original tweets such that the number matches the number of tweets from elites on a given day. The random sampling also ensures that tweets from non-elites with higher overall activity are selected. We then calculate the daily proportion of issue-related tweets that express a certain moral foundation. Specifically, we calculate the proportion ρ of original tweets relevant to issue i that express moral foundation m on day d: $\rho(m,i,d) = \frac{\tau(m,i,d)}{\tau(i,d)}$ where, τ returns the number of original tweets. We calculate this separately for elites and non-elites. To account for variation, we bootstrap non-elite tweets 100 times.

Figure 8((a)-(e)) compares moral language use by elites and non-elites across issues. Boxplots show the distribution of the daily share of tweets in each category. We do not

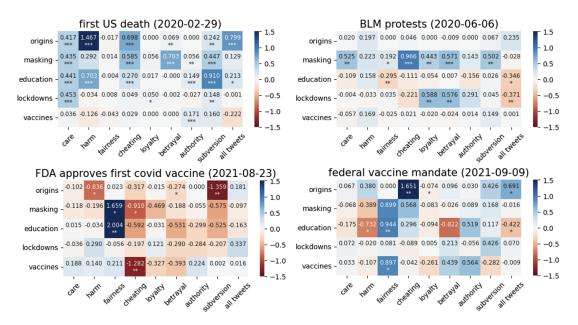


Fig. 7. Different responses from Liberals and Conservatives to COVID-related events. The blue color (positive changes in L/C ratio after event) implies that there was a bigger reaction in liberals, and the red color implies that conservatives reacted more. Asterisks indicate significance values: * (p-value < 0.05), ** (p-value < 0.01), *** (p-value < 0.001), and no asterisk indicates p-value ≥ 0.05 .

show loyalty/betrayal and purity/degradation foundations, as they were not substantially expressed by elites. The lower variance for non-elites is an attribute of bootstrapping. We rely on the non-parametric Mann-Whitney U Test to test the significance of our comparisons. All differences are assessed for significance at p < 0.001 unless otherwise specified.

On the issue of origins (Fig. 8(a)), non-elites express more care, subversion and cheating than elites. Conservative non-elites use more harm language than conservative elites, a difference that is not significant for liberals. Lockdowns, masking, and education tweets all show a pattern of elites using more moral language than non-elites across most moral foundations, with liberal elites and non-elites tending to refer more to care (Fig. 8(b-d)).

Discussions of vaccines deviate somewhat from these trends (Fig. 8(e)), with conservative non-elites expressing more harm than both liberals and conservative elites. Conservatives (elites and non-elites) also express more subversion than liberals.

V. Conclusions

The COVID-19 pandemic presented significant challenges to society both within and beyond the sphere of public health. The ensuing response to the pandemic quickly became polarized, with liberals and conservatives disagreeing on the severity of the pandemic and appropriate measures to address it. Sharp differences in issue positions were observed in discussions relating to the origins of the pandemic, lockdowns and business closures, masking mandates, disruptions to education, and vaccines. Previous studies [4]–[10] attest to the salience of these issues during the pandemic. We use a massive corpus of over 200M tweets [3] to assess these ideological divisions. Our findings suggest that not only do conservatives and liberals

differ in which issues they discuss and how they frame these issues, but also how they moralize these issues. We also demonstrate that a weakly-supervised classifier can be used to accurately identify issue relevant tweets.

While liberals dominated the discussion on education, tweets about COVID-19 origins and lockdowns were more prevalent among conservative. Tweets about the vaccine were predominantly made by liberals once the vaccine became available, but switched to being more important to conservatives after federal vaccine mandate were announced. We also found strong differences in moral language among liberals and conservatives. We found that tweets expressing the care/harm and authority/subversion foundations were most persistent, suggesting that these values resonated with both groups. This contrasts with past results suggesting that authority/subversion is primarily considered an important moral dimension by conservatives [12], and supports the idea endorsing the moral foundation of authority may be largely dependent on which authorities are being discussed, and what ideological stereotypes are associated with those authorities [57].

Our analysis revealed further differences in the moral language of partisans. While conservatives relied on negatively-valenced moral language, like harm and subversion, liberals primarily invoked care. The endorsement of negatively-valenced morality in conservatives is especially reflected in their moral reactions to COVID-related events, such as the FDA approving the first vaccine. This has important implications for message tailoring and public policy: establishing trust in public health institutions is critical to increasing compliance with public health measures, as is addressing moral concerns about these measures. We also found that political elites in the US tended to use more moral language compared to non-elites,

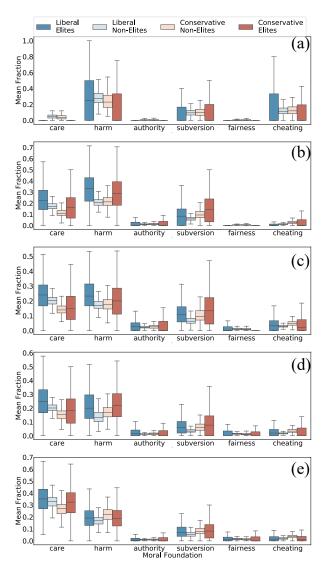


Fig. 8. Moral language differences between elites and non-elites. Boxplot shows the differences in the share of daily tweets made by each group on the issues: (a) origins, (b) lockdowns,(c) masking, (d) education and (e) vaccines.

regardless of ideology. Given the greater influence that elites have, this may lead to outsize perceptions of the extent of political polarization on these issues, particularly as negative tweets by elites tend to spread particularly widely [58]. In sum, our findings suggest that differences in the expressions of moral values drove ideological divisions on different issues, in response to different events, fueling polarization. Overall, our findings suggest that polarization is reflected in 1) the issues that partisans focus on, 2) the way that they frame these issues, and 3) the moral values they apply to issues.

However, it is worth noting that our methods have several limitations. We infer geo-location of tweets based on tweet metadata and user bios' which may not be accurate. Expression of ideological preferences and moral appeals can be both subtle and highly subjective. Methods leveraged to identify ideology and moral foundations, while state-of-the-art, may

suffer from inconsistencies and biases given the task's inherent complexity. While our method to identify issue relevant tweets is highly scalable, less explicit references to issues in tweets may be missed by this method. Future work could test a supervised classifier that may more accurately identify all relevant tweets. While we made an effort to combine available lists of US political elites on Twitter, this list is not exhaustive.

Ethical considerations

All data used in this study is publicly available [3], [44]. The study was deemed exempt from review by the Institutional Review Board (IRB), as it relied solely on publicly available data. Our study also adheres to Twitter's terms of service [59]. Tweet objects contain user information and this brings user anonymity into question. Users can restrict their tweets by switching to a private account or by deleting their tweets. Additionally, we preserve user anonymity during analysis by relying on user IDs instead of screen names. However, we acknowledge that many Twitter users may not be aware their data is used for research purposes [60]. In this article, we present aggregated statistics to address further mitigate this concern. The authors declare that there are no competing interests.

REFERENCES

- [1] C. Funk and A. Tyson, "Partisan differences over the pandemic response are growing," Pew Research Center, 2020, accessed: 2023-04-29.
- [2] M. Mordecai and A. Connaughton, "Public opinion about coronavirus is more politically divided in u.s. than in other advanced economies," Pew Research Center, 2020, accessed: 2023-04-29.
- [3] E. Chen, K. Lerman, E. Ferrara et al., "Tracking social media discourse about the covid-19 pandemic: Development of a public coronavirus twitter data set," *JMIR public health and surveillance*, vol. 6, no. 2, p. e19273, 2020.
- [4] K. Schaeffer, "Nearly three-in-ten americans believe covid-19 was made in a lab," Pew Research Center, 2020, accessed: 2023-04-29.
- [5] A. Connaughton, "Those on ideological right favor fewer covid-19 restrictions in most advanced economies," Pew Research Center, 2021, accessed: 2023-04-29.
- [6] R. Rojas, "Masks become a flash point in the virus culture wars," New York Times, 2020, accessed: 2023-04-29.
- [7] A. Luttrell and J. T. Trentadue, "Advocating for mask-wearing across the aisle: Applying moral reframing in health communication," *Health Communication*, pp. 1–13, 2023.
- [8] E. Y. Chan, "Moral foundations underlying behavioral compliance during the covid-19 pandemic," *Personality and individual differences*, vol. 171, p. 110463, 2021.
- [9] F. Pierri, B. L. Perry, M. R. DeVerna, K.-C. Yang, A. Flammini, F. Menczer, and J. Bryden, "Online misinformation is linked to early covid-19 vaccination hesitancy and refusal," *Scientific reports*, vol. 12, no. 1, pp. 1–7, 2022.
- [10] S. Rathje, J. K. He, J. Roozenbeek, J. J. Van Bavel, and S. van der Linden, "Social media behavior is associated with vaccine hesitancy," *PNAS Nexus*, vol. 1, no. 4, p. pgac207, 2022.
- [11] J. Haidt and J. Graham, "When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize," *Social justice* research, vol. 20, no. 1, pp. 98–116, 2007.
- [12] J. Graham, J. Haidt, and B. A. Nosek, "Liberals and conservatives rely on different sets of moral foundations." *Journal of personality and social* psychology, vol. 96, no. 5, p. 1029, 2009.
- [13] S. Guo, N. Mokhberian, and K. Lerman, "A data fusion framework for multi-domain morality learning," arXiv preprint arXiv:2304.02144, 2023.
- [14] A. Rao, F. Morstatter, M. Hu, E. Chen, K. Burghardt, E. Ferrara, and K. Lerman, "Political partisanship and antiscience attitudes in online discussions about covid-19: Twitter content analysis," *Journal of medical Internet research*, vol. 23, no. 6, p. e26692, 2021.

- [15] J. D. Hunter, Culture wars: The struggle to control the family, art, education, law, and politics in America. Avalon Publishing, 1992.
- [16] P. DiMaggio, J. Evans, and B. Bryson, "Have american's social attitudes become more polarized?" *American journal of Sociology*, vol. 102, no. 3, pp. 690–755, 1996.
- [17] J. H. Evans, "Have americans' attitudes become more polarized?—an update," *Social Science Quarterly*, vol. 84, no. 1, pp. 71–90, 2003.
- [18] A. I. Abramowitz and K. L. Saunders, "Is polarization a myth?" The Journal of Politics, vol. 70, no. 2, pp. 542–555, 2008.
- [19] A. Gollwitzer, C. Martel, W. J. Brady, P. Pärnamets, I. G. Freedman, E. D. Knowles, and J. J. Van Bavel, "Partisan differences in physical distancing are linked to health outcomes during the covid-19 pandemic," *Nature Human Behaviour*, pp. 1–12, 2020.
- [20] G. Grossman, S. Kim, J. M. Rexer, and H. Thirumurthy, "Political partisanship influences behavioral responses to governors' recommendations for covid-19 prevention in the united states," *PNAS*, vol. 117, no. 39, pp. 24144–24153, 2020.
- [21] J. Green, J. Edgerton, D. Naftel, K. Shoub, and S. J. Cranmer, "Elusive consensus: Polarization in elite communication on the covid-19 pandemic," *Science advances*, vol. 6, no. 28, p. eabc2717, 2020.
- [22] J. Graham and J. Haidt, "Moral foundation dictionary," https://moralfoundations.org, 2012, accessed: 2023-04-29.
- [23] J. Garten, R. Boghrati, J. Hoover, K. M. Johnson, and M. Dehghani, "Morality between the lines: Detecting moral sentiment in text," in *IJCAI* 2016 workshop on Computational Modeling of Attitudes, 2016.
- [24] W. J. Brady, J. A. Wills, J. T. Jost, J. A. Tucker, and J. J. Van Bavel, "Emotion shapes the diffusion of moralized content in social networks," *PNAS*, vol. 114, no. 28, pp. 7313–7318, 2017.
- [25] S. Y. N. Wang and Y. Inbar, "Re-examining the spread of moralized rhetoric from political elites: Effects of valence and ideology," *Journal* of Experimental Psychology: General, vol. 151, pp. 3292–3303, 2022.
- [26] S. P. Koleva, J. Graham, R. Iyer, P. H. Ditto, and J. Haidt, "Tracing the threads: How five moral concerns (especially purity) help explain culture war attitudes," *Journal of research in personality*, vol. 46, no. 2, pp. 184–194, 2012.
- [27] R. Díaz and F. Cova, "Reactance, morality, and disgust: the relationship between affective dispositions and compliance with official health recommendations during the covid-19 pandemic," *Cognition and Emotion*, vol. 36, no. 1, pp. 120–136, 2022.
- [28] N. K. Reimer, M. Atari, F. Karimi-Malekabadi, J. Trager, B. Kennedy, J. Graham, and M. Dehghani, "Moral values predict county-level covid-19 vaccination rates in the united states." *American Psychologist*, vol. 77, no. 6, p. 743, 2022.
- [29] J. Borghouts, Y. Huang, S. Gibbs, S. Hopfer, C. Li, and G. Mark, "Understanding underlying moral values and language use of covid-19 vaccine attitudes on twitter," *PNAS nexus*, vol. 2, no. 3, p. pgad013, 2023.
- [30] M. L. Pacheco, T. Islam, M. Mahajan, A. Shor, M. Yin, L. Ungar, and D. Goldwasser, "A holistic framework for analyzing the covid-19 vaccine debate," arXiv preprint arXiv:2205.01817, 2022.
- [31] M. Dredze, M. J. Paul, S. Bergsma, and H. Tran, "Carmen: A twitter geolocation system with applications to public health," in AAAI workshop on HIAI, vol. 23. Citeseer, 2013, p. 45.
- [32] K. Garimella, G. De Francisci Morales, A. Gionis, and M. Mathioudakis, "Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship," in WWW-18, 2018, pp. 913–922.
- [33] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," JMLR, vol. 3, no. Jan, pp. 993–1022, 2003.
- [34] M. Grootendorst, "Bertopic: Neural topic modeling with a class-based tf-idf procedure," arXiv preprint arXiv:2203.05794, 2022.
- [35] D. Card, A. Boydstun, J. H. Gross, P. Resnik, and N. A. Smith, "The media frames corpus: Annotations of frames across issues," in ACL-IJCNLP 2015, 2015, pp. 438–444.
- [36] J. Eisenstein, A. Ahmed, and E. P. Xing, "Sparse additive generative models of text," in *ICML-11*, 2011, pp. 1041–1048.
- [37] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," arXiv preprint arXiv:1810.04805, 2018.
- [38] A. Rojecki, E. Zheleva, and L. Levine, "The moral imperatives of self-quarantining," *Annual meeting of the American Political Science Association*, 2021.
- [39] J. Hoover, G. Portillo-Wightman, L. Yeh, S. Havaldar, A. M. Davani, Y. Lin, B. Kennedy, M. Atari, Z. Kamel, M. Mendlen, G. Moreno, C. Park, T. E. Chang, J. Chin, C. Leong, J. Y. Leung, A. Mirinjian,

- and M. Dehghani, "Moral foundations twitter corpus: A collection of 35k tweets annotated for moral sentiment," *Social Psychological and Personality Science*, vol. 11, no. 8, pp. 1057–1071, 2020.
- [40] K. Johnson and D. Goldwasser, "Classification of moral foundations in microblog political discourse," in ACL. Melbourne, Australia: Association for Computational Linguistics, Jul. 2018, pp. 720–730.
- [41] H. Le, R. Maragh, B. Ekdale, A. High, T. Havens, and Z. Shafiq, "Measuring political personalization of google news search," in WWW-19, 2019, pp. 2957–2963.
- [42] D. Nikolov, A. Flammini, and F. Menczer, "Right and left, partisanship predicts (asymmetric) vulnerability to misinformation," arXiv preprint arXiv:2010.01462, 2020.
- [43] M. Cinelli, G. De Francisci Morales, A. Galeazzi, W. Quattrociocchi, and M. Starnini, "The echo chamber effect on social media," *PNAS*, vol. 118, no. 9, p. e2023301118, 2021.
- [44] M. B.-F. Check, "The media bias chart," https://mediabiasfactcheck.com, 2023, accessed: 2023-05-06.
- [45] M. Pagliardini, P. Gupta, and M. Jaggi, "Unsupervised Learning of Sentence Embeddings using Compositional n-Gram Features," in NAACL 2018, 2018.
- [46] P. Barberá, "Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data," *Political analysis*, vol. 23, no. 1, pp. 76–91, 2015.
- [47] Pfizer, "Pfizer and biontech announce vaccine candidate against covid-19," https://www.pfizer.com, 2020, accessed: 2023-05-06.
- [48] M. Gordon, W. Strobel, and D. Hinshaw, "Intelligence on sick staff at wuhan lab fuels debate on covid-19 origin," Wall Street Journal, 2021, accessed: 2023-05-06.
- [49] A. R. Gover, S. B. Harper, and L. Langton, "Anti-asian hate crime during the covid-19 pandemic: Exploring the reproduction of inequality," *American journal of criminal justice*, vol. 45, pp. 647–667, 2020.
- [50] C. Deane, K. Parker, and J. Gramlich, "U.s. public opinion on the coronavirus pandemic," Pew Research Center, 2021, accessed: 2023-05-11
- [51] J. Leskovec, L. Backstrom, and J. Kleinberg, "Meme-tracking and the dynamics of the news cycle," in *Proceedings of the 15th ACM SIGKDD* international conference on Knowledge discovery and data mining, 2009, pp. 497–506.
- [52] S. Clifford and J. Jerit, "How words do the work of politics: Moral foundations theory and the debate over stem cell research," *The Journal* of Politics, vol. 75, no. 3, pp. 659–671, 2013.
- [53] S.-Y. N. Wang and Y. Inbar, "Moral-language use by us political elites," Psychological Science, vol. 32, no. 1, pp. 14–26, 2021.
- [54] K. Clayton, N. T. Davis, B. Nyhan, E. Porter, T. J. Ryan, and T. J. Wood, "Elite rhetoric can undermine democratic norms," *PNAS*, vol. 118, no. 23, p. e2024125118, 2021.
- [55] M. Wojcieszak, A. Casas, X. Yu, J. Nagler, and J. A. Tucker, "Most users do not follow political elites on twitter; those who do show overwhelming preferences for ideological congruity," *Science advances*, vol. 8, no. 39, p. eabn9418, 2022.
- [56] S. Shugars, A. Gitomer, S. McCabe, R. J. Gallagher, K. Joseph, N. Grinberg, L. Doroshenko, B. F. Welles, and D. Lazer, "Pandemics, protests, and publics: Demographic activity and engagement on twitter in 2020," *Journal of Quantitative Description: Digital Media*, vol. 1, 2021.
- [57] J. G. Voelkel and M. J. Brandt, "The effect of ideological identification on the endorsement of moral values depends on the target group," *Personality and Social Psychology Bulletin*, vol. 45, pp. 851–863, 2019.
- [58] H. Hwang, Y. Kim, and C. U. Huh, "Seeing is believing: Effects of uncivil online debate on political polarization and expectations of deliberation," *Journal of Broadcasting and Electronic Media*, vol. 54, pp. 621–633, 2014.
- [59] Twitter, "Developer agreement and policy," https://developer.twitter.com, 2022. accessed: 2023-05-06.
- [60] C. Fiesler and N. Proferes, ""participant" perceptions of twitter research ethics," *Social Media + Society*, pp. 1–14, 2018.