Mask Refusal Backlash: The Politicization of Face Masks in the American Public Sphere During the Early Stages of the COVID-19 Pandemic

Caleb Scoville^{1,¶*}, Andrew McCumber^{2,¶}, Razvan Amironesei³, June Jeon⁴

¹ Department of Sociology, Tufts University, Medford, Massachusetts, United States of America

² Department of Earth and Environment, Boston University, Boston, Massachusetts, United States of America

³ Center for Applied Data Ethics, University of San Francisco, San Francisco, United States of America

⁴ Department of Sociology, Chungnam National University, Daejeon, Republic of Korea

*Corresponding author
E-mail: Caleb.Scoville@Tufts.edu (CS)

¶These authors contributed equally to this work.

This is a non-peer reviewed preprint, updated on March 10, 2022. The latest version of this preprint can be retrieved at SocArXiv: <https://osf.io/preprints/socarxiv/237x6>.

Abstract

This research shows how face masks took on discursive political significance during the early stages of the COVID-19 pandemic in the United States. We argue that political divisions over masks cannot be understood by looking to partisan differences in mask wearing behaviors alone. Instead, we show how the mask became a political symbol enrolled into patterns of affective polarization. This study relies on qualitative and computational analyses of opinion articles ($n = 7,970$), and supplemental analyses of Twitter data, the transcripts of major news networks, and longitudinal survey data. First, we show that anti-mask discourse was consistently marginal, and that backlash against mask refusal came to prominence and did not decline even as masking behaviors normalized and partly depolarized. Second, we show that backlash against mask refusal—rather than mask refusal itself—was the primary way that masks were discussed in relation to national electoral, governmental, and partisan themes.

The COVID-19 pandemic has been thoroughly politicized in the United States. Face masks, while effective at reducing the spread of the virus (Brooks and Butler 2021; Howard et al. 2021), have been particularly controversial. Survey evidence suggests that partisanship may be the most important explanation for differences in attitudes and behaviors associated with the pandemic, with conservatives reporting less concern about the pandemic and taking fewer behavioral precautions than liberals (Clinton et al. 2021; Gadarian et al. 2021; Kerr et al. 2021; Kahane 2021). Complementing this research with a large-scale and systematic media analysis, we argue that the political meaning of face masks cannot be reduced to partisan gaps in pandemic-related behavior. Instead, we show how the face mask became a political symbol enrolled in broader patterns of affective polarization that articulate shared meanings of “us” and “them,” even as mask wearing behaviors broadly normalized and partly depolarized across the political spectrum. In particular, we show that *backlash against mask refusal*—rather than mask refusal itself—was the primary way that masks took on political significance in opinion articles published in the United States from the beginning of the pandemic through the end of October 2020.

Political divisions around pandemic policy are at once old and new. On the one hand, they reflect a broader structural development in the nature of political identity in the United States. Since the 1970s, the Democratic and Republican party platforms have become more coherent and distinct, and voters have increasingly come to view those on the other side of the aisle with suspicion and hostility (Baldassarri and Gelman 2008; Iyengar et al. 2012; Hopkins 2017; McCarty 2019). On the other hand, the pandemic poses an opportunity for understanding how a new set of issues and objects became integrated into this political landscape in a short period of time. Face masks had little political significance in the United States before COVID-19. By June 2020, newspaper headlines like, “Masks Become a Flash Point in the Virus Culture Wars” had become commonplace (Rojas 2020).

Political polarization has been debated intensely in the social sciences for decades. Polarization was traditionally defined as increasingly extreme disagreement about policy issues (Fiorina et al 2005). More recently, “affective polarization”—a function of positive sentiment for one’s own group, and negative sentiment for opposing groups—is now widely understood as “a defining feature of twenty-first-century US politics” (Druckman et al. 2021). In the United States, warmth toward the opposing party (out-party) has significantly diminished since 1980, and out-party hate has proven to be a stronger force than in-party love (Finkel et al. 2020).¹

The face mask took on political significance in an intensely polarized environment. A county-level analysis of survey data from July 2020 shows that public mask wearing was significantly lower in counties where Donald Trump found strong

¹ A second important distinction that we do not discuss here is between polarization and partisan sorting. Partisan divergence can also result from voters with similar policy positions “sorting” into respective parties, even if voters’ views have not polarized overall (McCarty 2019).

support during the 2016 presidential election (Kahane’s 2021).² However, while divisions over masks persisted, survey evidence shows that partisan gaps in masking behaviors declined over the course of the pandemic. A Pew survey, for example, found that, by November 2020, 81% of Republicans and 91% of Democrats “said they wore a face mask all or most of the time in stores and businesses over the past month.” This 10% difference, while significant, is well under half of the 23% gap found in an earlier wave of the same survey in June 2020 (Schaeffer 2021).³ Although ascertaining the precise magnitude of the partisan gap in mask-wearing behavior may be sensitive to question wording and survey methodology, a broad pattern of the normalization of masking behaviors over the course of 2020 is corroborated across sources. The University of Southern California Center for Economic and Social Research’s Understanding Coronavirus in America tracking survey (which we rely on below and discuss in detail in Appendix C) shows a similar trend of a gradual normalization of masking that stabilizes at around 90% of the American population reporting wearing masks for protective purposes.

We suggest that a comprehensive understanding of the controversial nature of face masks cannot be gleaned from survey results on mask wearing behavior alone. As we show, the broad normalization and partial depolarization of masking *behaviors* does not necessarily correspond to the fizzling out of the mask’s role in America’s culture

² Survey and observational evidence also show persistent differences by gender, age, and location that are often associated with but not reducible to partisanship (Haischer et al. 2020; Cassino et al. 2020).

³ The February 2021 wave of the same survey found that 83% of Republicans and 93% of Democrats answered affirmatively to the same question, suggesting that mask wearing behaviors had (at least temporarily) stabilized across the political spectrum by November 2020. As we discuss below, our emphasis is less on the precise size of this gap, but instead on the diachronic pattern of stabilization and the broad normalization of mask-wearing society-wide, which we corroborate with the University of Southern California Center for Economic and Social Research’s Understanding Coronavirus in America tracking survey (see Appendix C). It is also important to acknowledge that our period of analysis does not include the many changes in state and local policies, and recommendations from the Centers for Disease Control and Prevention that occurred in 2021 following the introduction of vaccines and the delta and omicron waves of the pandemic.

wars. Instead, we argue that explaining the controversy surrounding face masks requires accounting for their role in patterns of affective polarization. While most research on the politics of the COVID-19 does not inquire into the relationship between affective polarization and pandemic protocols, one important counterexample draws on panel data with waves before and during the COVID-19 pandemic, to show “a strong association between citizens’ levels of partisan animosity and their attitudes about the pandemic, as well as the actions they take in response to it” (Druckman et al 2021). This finding helps explain why a common interest in surviving a potentially deadly virus did not facilitate increased solidarity and compassion among Americans, but instead gave them new ways to be divided.

We argue that while existing survey data are important for understanding pandemic politics, they offer an incomplete diagnosis of the divisiveness of face masks in the United States. The normalization and partial depolarization of individual masking behaviors occurred alongside a countervailing process: the face mask’s emergence as a *political symbol of partisan animosity*. Political symbols are entities which serve as vectors for collective meaning and sentiment (Edelman 1985). Their mobilization can potentially reduce or increase affective polarization. For instance, experimental research suggests that displaying symbols of shared national identity (e.g., the American flag) can undermine the sense of a partisan “us” and “them” (Levendusky 2018). On the other hand, political symbols that evoke out-group animosity may reinforce or even widen partisan divisions.

We understand political symbols of partisan animosity to derive their power by resonating with broader partisan identities and narratives. Sociological research on resonance focuses attention on “why certain discourses, messages, or other cultural objects have an advantage over others because they fit—or resonate with—prevailing

cultural worldviews of the audiences who receive them” (McDonnell et al. 2017). Thus, a political symbol of partisan animosity is effective to the extent it resonates with what Hochschild calls a “deep story,” an emotionally evocative account that defines *who we are* in opposition to a partisan *them* (Hochschild 2016). Many political symbols that divide Americans exist only through the divisive meanings that they encode (e.g., the Confederate flag). The mask was not designed as a political symbol, but only became one incidentally, a process that merits further explanation.⁴

The symbolic power of masks is perhaps most readily visible in mask refusal. In the first Presidential debate of 2020, Donald Trump mocked his electoral opponent Joe Biden’s mask wearing habits, stating, “I don’t wear a mask *like him*” (Collman 2020). Anti-mask protestors have been particularly evocative by variously referring to masks as “muzzles” and “the new symbol of tyranny” (Collinson and Hu 2020). Similarly, the spread of misinformation about COVID-19 online has been facilitated by linking masks to politically charged discourses that often invoke a nefarious “other,” whether an overreaching government, lying media organizations, or greedy pharmaceutical companies (Al-Ramahi et al. 2021).

The political symbolism of masks need not be discursive in order to resonate with consolidated senses of “us” versus “them.” This may be especially important when considering the role of elite cues. For example, Vice President Mike Pence’s maskless visit to the Mayo Clinic in 2020, and more recently Supreme Court Justice Neil

⁴ Experimental research finds “strong evidence of in-group favouritism among both mask and non-mask wearers” in a series of Prisoner’s Dilemma games (Powdthavee et al. 2021). Although framed in psychological terms without reference to affective polarization, the finding that the mask strongly signals group identification is broadly consistent with understanding masks as political symbols of partisan animosity.

Gorsuch's public refusal to wear a mask during the omicron surge in January 2022 encoded clear messages in nonverbal cues.⁵

The political symbolism of mask refusal as a source of meaning for conservatives has received some explicit scholarly attention to date (Wessel 2021; Kenworthy, Koon, and Mendenhall 2021; Lupton et al. 2021), but researchers have given less attention to the cultural-political meaning of pandemic protocols for liberals (although see Lang et al. 2021 for one exception). This, we argue, is a significant oversight that ignores how the specter of the “anti-masker” became meaningful for liberal proponents of masks by supporting out-group antipathy. Consider, for instance, the viral appeal of videos capturing individuals refusing to comply with mask mandates in stores and other shared spaces (e.g., KGW Staff 2021). Given that Americans now tend to view partisanship as the strongest source of conflict in society (Gramlich 2017), such episodes ostensibly resonate with existing patterns of affective polarization, offering Democrats the anti-masker as a stand-in for the broader partisan Republican *them*. Open-ended survey responses do provide some insights regarding this phenomenon. Pew found that Democrats accounted for 76% of those who expressed worries about others not wearing masks in a survey fielded between August 31st and September 7th of 2020 (Kessel and Quinn 2020).

We begin with the assumption that an important way that Americans make meaning about social events is by drawing on symbolic resources from the broader media environment. Even in normal times, media platforms relate to the formation and maintenance of partisan divisions in several ways. First, they operate as the stage on which partisan elites can disseminate political symbols of partisan animosity (Iyengar et

⁵ One limitation of our study is that because our data is discursive, we cannot directly account for the role of non-discursive cues such as these. We do, however, analyze the relationship between discursive political symbols and broad population-level trends in non-discursive mask wearing behavior.

al. 2012). Second, an increasingly high-choice media environment has resulted in the proliferation of extreme partisan media outlets (Iyengar et al. 2019). Third, an important pattern in recent decades has been the rise of “outrage” in the American media environment, a form of discourse that “is distinctly emotional, partial, antagonistic, and opinion-based” (Berry and Sobieraj 2014). And finally, social media adds fuel to the fire by amplifying extreme voices, muting moderates, and distorting our vision of our political opponents (Bail 2021). Americans’ reliance on media may have been particularly high in 2020 due to the physical and social isolation associated with pandemic protocols. Americans’ experience of the COVID-19 pandemic, which has dominated news cycles since early 2020, has thus been shaped by media reports in particularly deep ways.

By analyzing how the face mask took on meaning in the American public sphere, we do not claim a one-to-one correspondence between the mask as an object of media attention and mask wearing behavior. On the contrary, we exploit both similarities and disjunctures between survey evidence and media discourse about masks to show that the mask’s meaning as a symbol of partisan animosity in the American public sphere took on a life of its own, in important ways decoupled from broader trends in Americans’ individual mask wearing behaviors.

Our argument proceeds in two steps. First, we demonstrate the marginality of the anti-mask position in the American public sphere, with important qualifications. We show that, consonant with survey reports on masking behavior and attitudes, anti-mask discourse is a rare position in traditional print media. Also consistent with survey results on masking behavior and attitudes, anti-mask discourse is much more common, but still a minority position in online right-wing media outlets. Further, we show that the marginal nature of the anti-mask position in print media is not merely an effect of

editorial gatekeeping, because similar discursive patterns are observed on ungated social media (Twitter).

Second, we show that *backlash against* mask refusal is the primary means through which masks became discursively politicized in opinion articles published in the United States during the early stages of the pandemic. We find that mask refusal backlash was significantly more prevalent than anti-mask discourse, rapidly increased in the early phase of the pandemic, and did not experience a sustained decline as masking behaviors and attitudes normalized and partly depolarized. We verify the partisan appeal of the “anti-masker” discourse by showing differential usage of related terms across television networks. Finally, using topic modeling, a method of computational text analysis, we show that backlash against mask refusal is the primary vector through which face masks have been discussed in relation to national electoral, governmental, and partisan themes in opinion articles. Taken together, these findings suggest that masks became politically salient as a direct result of liberals’ antipathy for conservatives and conservative elites. Despite its marginality, mask refusal resonated with liberals’ pre-existing narratives about conservatives as ignorant, selfish, and dangerous.

Although the mask refusal backlash discourse did not mirror population-level trends in masking behaviors, we do not claim that backlash against mask refusal is grounded in a delusion. Mask refusal persisted through the pandemic as a fringe behavior and media specter, with a number of conservative elites including former President Trump actively discouraging widespread mask use implicitly or explicitly. We suggest that such refusals were refracted through existing patterns of affective polarization, representing a feedback loop between different forms of political division that is not observable in existing survey evidence when viewed in isolation.

ANALYTICAL APPROACH

To empirically analyze the meaning of masks during the COVID-19 pandemic, we built a large-scale dataset of relevant public discourse. The bulk of our argument rests on a combination of manual qualitative coding and computational text analysis of thousands of published opinion articles. We built our corpus of opinion articles by first collecting all US newspaper articles available in the Nexis Uni database that included the search term “mask” or “masks” as well as either “COVID” or “coronavirus.” We chose Nexis Uni as the basis for our convenience sample of mainstream news sources because it is likely the broadest database of news sources available to academic researchers, and includes many sources with a variety of national, regional, and local audiences. Our search covered the period of January 1 to October 31, 2020 to capture the emergence of the COVID-19 pandemic as a major political issue leading up to the 2020 US Presidential election. This initial search produced 89,089 articles.

We began our data collection under the assumption that mainstream news media would be less likely to publish articles expressing anti-mask sentiment than publications that target an audience on the right side of the political spectrum.⁶ While Nexis Uni’s sources span the political spectrum, the database primarily indexes traditional print media (although some sources it includes are only available online). When reviewing our Nexis Uni sample, we noticed that prominent right-wing websites were not included. Because we are aware of the importance of partisan echo-chambers, and in particular conservative’s increasing distrust of national and local news organizations in

⁶ We also acknowledge that there is ideological diversity among articles taking each position. This is especially likely for articles published during the first few months of 2020, before the United States Centers for Disease Control and Prevention first issued an advisory for the general public to wear masks. Speaking against mask wearing during this period of great uncertainty and mask shortages in the medical sector is not necessarily an endorsement of “mask refusal.” Our analysis does not rely on any such assumption. Yet, as we show in Appendix D, the vast majority of our coded articles appeared after this point. We include articles going back to the start of 2020 to capture the emergence and consolidation of the controversy.

recent years (Gottfried and Liedke 2021), we supplemented our Nexis Uni sample with a purposive sample of articles published on websites that explicitly target conservative audiences.⁷

To build our right-wing purposive sample, we used the same search terms and timeframe resulting in an additional 1,080 articles. To design this sampling frame, we selected the outlets with the highest Alexa Traffic Rankings of those appearing on a list of conservative news outlets on the crowd-sourced conservative website “Conservapedia” (Breitbart, Fox News, Newsmax, The Daily Wire, The Daily Caller, The Blaze, The Federalist, The New York Post, and The Washington Times).

Within the Nexis Uni (hereafter “mainstream”) and right-wing samples, we identified a total of 7,970 opinion articles including op-eds, opinion columns, editorials, and letters to the editor. (See Appendix A for a description of our data collection and cleaning process.) We supplemented our opinion article dataset with social media data from Twitter. As discussed in more detail in Appendix B, we analyzed a random sample of public English language Tweets that mentioned “mask” or “masks” within the same time frame as our published opinion data. We omitted the terms “COVID” and “coronavirus” from this sampling frame because of the short and informal nature of

⁷ We did not supplement our analysis with a symmetrical left-wing sample because no such symmetry exists in the partisan patterning of trust in news media organizations. According to the Pew Research Center, trust in national and local news organizations remains high and has not changed significantly among Democrats in recent years, while in “just five years, the percentage of Republicans with at least some trust in national news organizations has been cut in half – dropping from 70% in 2016 to 35%” in 2021 (Gottfried and Liedke 2021). Similarly, Gallup finds that Democrats’ trust in mass media has modestly increased in recent years, with 73% responding that they trust mass media a great deal in 2020. By comparison, in 2020 Republicans’ trust in mass media reached an all-time low since Gallup started asking the question in 1972 with only 10% reporting a great deal of trust (Brenan 2020). Because our main findings center on the meaning of mask refusal for a broadly liberal audience, it would be worthwhile to follow up our analysis by looking at how masks were discussed in more explicitly partisan liberal sources. However, given that our findings about mask refusal backlash comes through clearly without reliance on a purposive left-wing sample, we emphasize the broad resonance of these themes in the American public sphere outside of narrower hyper-partisan liberal echo-chambers. If our argument about the partisan appeal of the mask refusal backlash discourse holds, excluding a purposive left-wing sample likely attenuates the results we find here rather than biasing them in the direction of our conclusions. This said, we also acknowledge that some far-left publications (e.g., Jacobin) were less focused on COVID-19 culture wars than more partisan liberal publications (e.g., HuffPost).

Tweets and instead removed irrelevant items in our sample during analysis. Our analytical sample of Tweets after these omissions was 5,089. We also collected supplemental data on the occurrence of terms associated with masks, mask mandates, and anti-maskers using the Stanford Cable TV News Analyzer, which allows for “large-scale, data-driven analysis of the contents of cable TV news,” namely CNN, Fox News, and MSNBC (Stanford Cable TV News Analyzer 2021).

To compare the dual trajectories of public sphere treatment of masks with self-reported masking behavior, we rely on survey responses from the University of Southern California Center for Economic and Social Research’s Understanding Coronavirus in America tracking survey (USC 2021). This population-representative survey is updated daily, so it can be easily juxtaposed with the dynamics of our media analysis. Respondents were asked if they “wore a mask or face covering in the last 7 days” in order “to keep safe from coronavirus.” Because our central findings concern discourses about mask refusal, Figure 2 visualizes the weekly proportion of respondents who answered “no” to this question on each day of the survey until the end of October 2020. This allows us to track changes in societal norms on masks in the early months of the pandemic (since wearing masks in public was not required in most states/municipalities or even advised nationally for the general public during the first few months of 2020), as well as the degree of deviation from this norm once masking became dominant. We discuss the virtues and limitations of this particular survey response in Appendix C and triangulate our findings with Pew’s surveys on mask-related behaviors and partisanship discussed above.

As described in more detail in Appendix D, a group of research assistants coded every opinion article in our corpus ($n = 7,970$) and our analytical sample of Tweets ($n = 5,089$) along the following qualitative variables:

- (1) “Pro-Mask”: Articles/Tweets were coded “1” (“0”) if they elaborated a clear position in favor of mask-wearing as a public health measure.
- (2) “Anti-Mask”: Articles/Tweets were coded “1” (“0”) if they expressed a clear position in opposition to mask-wearing as a public health measure.
- (3) “Pro-Mask Mandate”: Articles/Tweets were coded “1” (“0”) if they expressed support for local, state, or national public policy requiring the wearing of masks in public places.
- (4) “Anti-Mask Mandate”: Articles/Tweets were coded “1” (“0”) if they expressed opposition to the above policy measures.
- (5) “Mask Refusal Backlash”: Articles/Tweets were coded “1” (“0”) if they expressed negative sentiment toward individuals who do not wear masks or discourage mask wearing.

These qualitative codes are not mutually exclusive, so individual articles/Tweets could be coded “1” on more than one variable. However, as discussed further in Appendix D, some variables logically preclude each other (e.g., “Pro-Mask” precludes “Anti-Mask”) or presuppose each other (e.g., “Pro-Mask Mandate” presupposes “Pro-Mask”). It was also possible to code articles/Tweets “0” across all variables if references to masks were incidental to the opinions expressed by the author. To avoid potential sources of bias, coders analyzed opinion articles in random order (with mainstream and right-wing samples pooled) and were blinded to meta-data including the articles’ sources and authors. Tweets were also randomized before distributing to coders.

Because Twitter's structure of engagement involves the ability to reply to others, we noticed a significant minority of Tweets that expressed negative sentiment toward expressions of "Mask Refusal Backlash," but did not necessarily express "Anti-Mask" sentiment. Therefore, in our Twitter analysis alone, we included an additional qualitative variable: "Backlash Against Mask Refusal Backlash." These Tweets tended to be accusations from conservatives that liberals were politicizing masks. As discussed in Appendix D, because we began coding Tweets after the opinion articles, the analysis of the latter does not include this variable, which our research assistants suggested would have only been coded affirmatively in a small minority of cases.

Finally, we used topic modeling to identify patterns in the substantive content of the opinion articles in our corpus and to examine how those patterns varied according to our coding categories. Topic modeling is a form of automated content analysis used to identify latent topics in a large corpus of text. The topics it identifies are not predetermined, but rather are arrived at through an algorithm that groups words together based on their co-occurrence in individual documents within the corpus. The underlying theory that drives topic modeling assumes that a "topic" is a cluster of words that tend to occur together. Analysts of cultural meaning have found that this computational tool is useful for measuring meaning on a large scale because its technical implementation proceeds from the core theoretical assumption that meaning is relational (or that words and concepts derive their meaning from their relationship to other words and concepts) (Mohr and Bogdanov 2013). The most well-known algorithm for topic modeling is called Latent Dirichlet Allocation (Blei et al. 2003), and it considers only word co-occurrences in assigning topics. More recently, scholars have increasingly used a different algorithm called Structural Topic Modeling (STM), which allows the researcher to include metadata that the algorithm considers when assigning topics

(Roberts et al. 2014). In our case we use STM and include metadata: each article's codes along the five qualitative variables described above; each article's date of publication; and the data source (our mainstream sample and our right-wing sample).

This approach allows us to identify the substantive topics in our corpus as a whole and generate metrics that represent how those topics are associated with our qualitative variables. While topics are not generated in advance, the researcher must specify the number of topics to be modeled. Although topics are computationally derived, their interpretation relies on a "deep human understanding of the corpus" (Farrell 2016). There is no objectively "correct" number of topics. Our goal was to arrive at a model that balanced granularity and interpretability. As elaborated further in Appendix E, we ran several specifications and converged on an STM model with six topics trained on our corpus of opinion articles. The results of our topic model are summarized in Figures 4 and 5 and discussed below.

RESULTS

The Marginality of Anti-Mask Discourse

First, as shown in Figure 1 our findings show that among published opinion articles in our mainstream sample, anti-mask discourse is rare (2.72%). The vast majority of opinion articles that we analyzed reflected pro-mask sentiment (75.9%). Articles expressing a negative opinion about governmental mask mandates were somewhat more common than articles expressing negative opinions about mask wearing, but this was still a small minority opinion (3.88%), and less common than articles expressing support for mask mandates (9.76%).

[Figure 1 About Here]

While these results are broadly consistent with polling on masks, which suggests that outright mask refusal rapidly became a fringe position in the course of the pandemic, it is possible that editorial authorities in traditional print media moderated the content of the views expressed in publications to suppress anti-mask discourse. Editors may have done so out of concern for public safety, to combat disinformation, or because of a liberal ideological slant. We address this possibility in two ways: by analyzing a purposive sample of articles from right-wing media outlets and a random sample of Tweets mentioning masks.

First, our analysis of our purposive sample of right-wing online media shows that anti-mask discourse was more common in our right-wing sample (12.3%) than our mainstream sample. Despite making up just 9.8% of our analytical sample of opinion articles, opinion articles from right-wing websites account for 44.2% of those coded anti-mask. However, this difference should not be overstated in substantive terms. Even within our right-wing sample, articles expressing anti-mask opinions remain a minority, with articles expressing pro-mask opinions representing a (slim) majority (50.8%). Among right-wing articles, those expressing negative sentiment toward governmental mask mandates were slightly more common than articles expressing negative opinions about mask wearing (14.9%), and more common than articles expressing positive sentiment about mask mandates (6.02%).

A second possibility is that the marginal nature of anti-mask discourse in opinion articles is a reflection of editorial gatekeeping practices in general, irrespective of partisan slant. To evaluate this proposition, we coded a random sample of Tweets mentioning masks along the same variables as the articles in our analysis. We note from the onset that our focus on revealing broad patterns is distinct but does not detract from

an approach that seeks to understand how disinformation spreads online (Johnson et al. 2020; Al-Ramahi et al. 2021). Twitter is an ungated social network with limited moderation. Since any individual or organization can create an account and post, content on Twitter does not reflect the editorial gatekeeping that we would expect in the articles we analyzed. Further, while Twitter’s policy is to remove COVID-19 related Tweets “that are claims of fact, demonstrably false or misleading, and likely to cause harm,” it is unlikely that most anti-mask discourse would rise to this level (Twitter 2020). (One of the first authors spoke directly with a Twitter employee with insider knowledge on the company’s content moderation practices who confirmed that most anti-mask discourse would not rise to the level of triggering removal from the platform.) The results of our analysis show that the discourse on Twitter about masks generally mirrors the discourse in opinion articles, albeit with anti-mask engagements closer in proportion to our right-wing sample. In our Twitter sample, anti-mask posts represent a minority of engagements (10.9%), with a strong majority expressing pro-mask sentiment (59.2%).⁸

The Politicization of Masks by Way of Backlash Against Mask Refusal

We have thus far focused on a way that the public sphere treatment of masks broadly aligns with the polling data on mask wearing. However, our analysis also shows that polling on mask behaviors fails to account for important ways that masks became politicized. While anti-mask discourse was marginal in the public sphere, backlash against mask refusal accounts for a significant proportion of the opinion articles published and Tweets posted about masks.

⁸ In Appendix D, we show that this finding is robust to an alternative specification of weighting Tweets in proportion to the number of times they were “liked” on the platform.

Figure 1 shows that mask refusal backlash is the second most prominent position taken in our mainstream opinion articles (29%) as well as posts in our Twitter sample (25.3%). (As discussed above, we also coded Tweets for backlash against mask refusal backlash, a position expressed in 3.07% of our sample.) Further, Figure 2 shows that as a proportion of all opinion articles (mainstream and right-wing samples pooled), mask refusal backlash discourse increased precipitously as mask wearing became normative, while anti-mask discourse rapidly declined from its initial peak at the beginning of our period of analysis and remained marginal thereafter.

[Figure 2 About Here]

The rise of mask refusal backlash in the public sphere coincided with increased agreement reflected in mask wearing behaviors. Rather than reflecting a widening behavioral divide, the trend suggests that (many) advocates of mask wearing in the public sphere increasingly came to frame their positions *against* those who refuse masks. Simply put, this finding suggests that an important way that masks become controversial was by way of antipathy toward the figure of the “anti-masker.”

Patterns of mask refusal backlash discourse provide important insights into the mask’s role as a symbol of partisan animosity. First, we consider the sources in which mask refusal backlash discourse appeared. Second, we consider how mask refusal backlash discourse resonated with broader national political narratives.

Figure 1 shows that mask refusal backlash is much less common among opinion articles from our right-wing sample (9.87%) than from our mainstream sample (29%) and posts on Twitter (25.3%). This suggests that mask refusal backlash is less appealing

to right-wing writers and audiences on average than those in the broader ideological spaces of mainstream print and social media.⁹

Reaching beyond our content analysis of news and social media, we conducted a supplemental analysis of the occurrence of terms associated with the mask refusal backlash discourse (“anti masker,” “anti maskers” and “anti mask”) on three major television networks: CNN, Fox News, and MSNBC. For our terms in question, we counted the number of unique days in which one or more of the terms associated with mask refusal was uttered on each of the networks. Figure 3 shows that terms associated with the mask refusal backlash discourse were uttered on 66 unique days on CNN, 87 days on MSNBC, and 0 (zero) days on Fox News. Given that Fox News targets a conservative audience, and CNN and MSNBC target liberal audiences, this finding provides additional evidence that the mask refusal backlash discourse is partisan in its appeal (Greiko 2020). Figure 3 also shows that this partisan association is not an artifact of how networks discussed masks overall. It does not hold for terms associated with masks (using the terms, “mask” and “masks”) or face masks (using the terms “face mask” or “face masks”) in general. Nor does it hold for terms associated with mask mandates (using the terms, “mask requirement,” “mask mandate,” and “mask law”). All of these terms were uttered a similar number of unique days across all networks during our period of analysis and the small observed differences are not associated with the partisan lean of TV networks.

⁹ See also Appendix D for a version of the analysis in Figure 2 that is disaggregated by sampling source. It shows that, although it was a marginal position overall, anti-mask discourse increased in relative terms over time in our right-wing sample. However, this change in right-wing sources appears to also be isolated to discourse (at least at the broad societal level), given that it did not correspond to a detectable behavioral change in masking behaviors. Indeed, in addition to our main source for diachronic behavioral trends (University of Southern California Center for Economic and Social Research’s Understanding Coronavirus in America tracking survey, which shows a pattern of the normalization of mask wearing during our sample period), Pew’s national surveys discussed above show a pattern of modest partisan depolarization with respect to masking behaviors during the same period (Schaeffer 2021).

[Figure 3 About Here]

The preceding evidence demonstrates that mask refusal backlash discourse cannot be understood as a simple reflection of the dynamics of population-level mask wearing behaviors, and that it appeals mostly to liberals. In the final step of our analysis, we turn to our computational text analysis to provide evidence that backlash against mask refusal is the primary vector through which face masks became discursively politicized in American opinion articles during the early stages of the pandemic. Here we understand discursive politicization as a function of being discussed in relation to national electoral, governmental, and partisan terms and topics. Masks are politicized in this sense when they are discussed in relation to national politicians, parties, and institutions. Taken together, our findings suggest that an important way that masks became politically salient was through the resonance of the mask refusal backlash discourse with liberals' antipathy for conservatives. As discussed above and in Appendix E, we use a six-topic STM model to analyze how latent topics vary across documents coded along our manually coded qualitative variables. Figure 4 is a corpus-level visualization of all topics along with the top 12 words in each topic, and the proportion of the corpus that belongs to each topic.¹⁰

[Figure 4 About Here]

We used STM's built-in "estimateEffect" function to visualize the results of our structural topic model that pertain to the "National Politics/Election" topic. This function conducts a linear regression, producing coefficients for each metadata

¹⁰ See Appendix E for word clouds that provide further interpretative detail to complement Figure 4, as well as the results of an alternative specification using eight topics instead of six.

variable's estimated effect on the prevalence of each topic identified in the topic model. Our regression contains each of the metadata variables included in the structural topic model itself (our qualitative coding categories, date of publication, and sampling source). The top panel of Figure 5 displays the estimates of all six topics from our model for the "Mask Refusal Backlash" coding category, along with 95% confidence intervals. This analysis shows that the estimate for the "National Politics/Election" topic is statistically significant, positive, and larger in magnitude than any other topic for the Mask Refusal Backlash coding category. The bottom panel of Figure 5 displays the estimates for the "National Politics/Election" topic for each of our six qualitative coding categories. Here, "Mask Refusal Backlash" has a statistically significant, positive, and by far the greatest estimate in terms of magnitude for the prevalence of the "National Politics/Election" among our coding categories.

[Figure 5 About Here]

It is important to note that backlash against mask refusal need not be connected to national politics. Backlash against mask refusal may be expressed as antipathy for others in the same community who do not wear masks, for elites who do not wear masks thus setting a bad example for the public, or in response to episodes of mask refusal captured or discussed in traditional or social media. We did not differentiate among or code for references to various types of actors or other qualitative themes when analyzing opinion articles. The strong association between national politics and backlash against mask refusal emerged from our computational text analysis alone.

DISCUSSION AND CONCLUSION

Face masks became controversial during the COVID-19 pandemic, but research on the controversy has focused primarily on how partisanship shapes mask-related behaviors. We complement this research with a large-scale systematic media analysis to better understand the political meaning of masks in the American public sphere. Anti-mask discourse was marginal in the public sphere during the first 10 months of 2020. Yet the results of our qualitative coding and computational analyses show that backlash against mask refusal was an important way that face masks took on political meaning during the pandemic. Because backlash against mask refusal came to prominence and did not decline as mask wearing behaviors broadly normalized and partly depolarized, we argue that understanding the mask controversy through the lens of partisan behavioral divides alone is insufficient. Instead, we argue that the mask became a symbol of partisan animosity enrolled into larger patterns of affective polarization. Specifically, we suggest that liberal antipathy for the figure of the “anti-masker” played an important and underappreciated role in shaping the public meaning of COVID-19.

While the historical period of this study provides a significant analytical opportunity, it also comes with weaknesses. 2020 was dominated, not only by a pandemic, but also by a particularly vitriolic presidential election. It is impossible to disentangle former President Trump’s frequent disregard and even disdain for mask protocols and the broader social meaning of masks in everyday life. Indeed, our analysis suggests that backlash against mask refusal was deeply partisan in its appeal and took on meaning in relation to national politics and the presidential election.

This qualification suggests important takeaways about the presidential election itself. Given the marginality of anti-mask discourse and the prominence of mask refusal backlash, it is likely that conservative elites who eschewed masks did more to motivate their liberal opponents than they did embolden their own base. While it is likely that

cues from conservative elites account for some of the polarization surrounding masks, our findings suggest that masks, either as an object of refusal or adornment, may have simply been less meaningful overall for conservatives. Thus, public anti-mask behaviors and discourses were likely to generate more antipathy from liberals than sympathy from conservatives, who appear to have been more likely to consider masks an inconvenience or source of discomfort than a source of political identity (Kessel & Quinn 2020). Consistent with this hypothesis, right-wing opinion articles were more likely not to have a clear position on masks and thus be coded “0” across all categories (30.2%) than were mainstream opinion articles (18.1%). Other aspects of the pandemic may have been more meaningful for conservatives than liberals (e.g., vaccines and various “lockdown” policies including school and business closures), but our data can only speak directly to the topic of masks.

While our findings are limited to the COVID-19 pandemic, our analysis has broader theoretical and methodological implications. First, this study provides a template for the analysis of political symbols in relation to polarization. More specifically, it demonstrates how to bring media data to bear on the prominence, appeal, and resonance of various positions as a new object of controversy enters the political landscape. This kind of analysis provides insights into political divisions that may not be evident in survey responses.

Conversely, these findings also provide insights into the media’s role in amplifying symbolic divisions that can potentially inform survey design itself. Especially in a context in which Americans’ personal networks have become increasingly politically homogeneous (Iyengar et al. 2018), political symbolism in the media may be an important driver of “false partisanship,” the extent to which partisans “believe the ideological divide to be far wider than it actually is” (Wilson et al. 2020).

In the case of the pandemic, the sustained and increasing resonance of backlash against mask refusal may have resulted in liberals underestimating the behavioral consensus on mask wearing across political parties. However, in order to test this hypothesis directly, survey researchers must design their questionnaires to elicit beliefs about the attitudes and behaviors of respondents' partisan opponents (Westfall 2019).

We conclude on a cautionary note. Although our findings suggest a disjuncture between the normalization and partial depolarization of masking behaviors and the discursive prominence, appeal, and resonance of the mask refusal backlash discourse, we want to clarify that we do not view the latter as somehow grounded in a delusion. On the contrary, it remains clear that a set of partisan elites and a marginal but significant group of ordinary Americans rejected the scientific evidence that wearing masks in public spaces reduces the risk of the transmission of COVID-19 to the detriment of public health and safety. Rather, our findings suggest that the actions of this group backfired politically due to their resonance with extant patterns of liberal antipathy for conservatives. The political divisiveness of masks may have thus taken shape through a feedback loop between ideological and affective polarization: modest partisan differences and highly visible elite cues were rendered meaningful in relation to consolidated senses of a partisan "us" and "them." Even a public increasingly united by masking behaviors could not quell the potency of masks as a divisive political symbol. The divisions in the American polity cannot be healed through scholarship alone, but our hope is that a more comprehensive understanding of how those divisions operate will inform responsible public interventions.

REFERENCES

- Al-Ramahi, Mohammad, et al. "Public discourse against masks in the COVID-19 era: Infodemiology study of Twitter data." *JMIR Public Health and Surveillance* 7.4 (2021): e26780.
- Bail, C. A. *Breaking the Social Media Prism: How to Make Our Platforms Less Polarizing*. Princeton University Press, 2021.
- Baldassarri, D. and A. Gelman, Partisans without constraint: Political polarization and trends in American public opinion. *American Journal of Sociology* 114.2, 408-446 (2008).
- Berry, J. M. S. Sobieraj. *The Outrage industry: Political Opinion Media and the New Incivility*. Oxford University Press, 2014.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research*, 3, 993-1022.
- Brenan, Megan. 2020. "Americans Remain Distrustful of Mass Media." Gallup. 30 September 2020. Available at: <https://news.gallup.com/poll/321116/americans-remain-distrustful-mass-media.aspx>
- Brooks, J. T. and J. C. Butler, Effectiveness of mask wearing to control community spread of sars-cov-2. *Jama* 325.10, 998-999 (2021).
- Cassino, Dan, and Yasemin Besen-Cassino. "Of masks and men? Gender, sex, and protective measures during COVID-19." *Politics & Gender* 16.4 (2020): 1052-1062.
- Clinton, J. et al., Partisan pandemic: How partisanship and public health concerns affect individuals' social mobility during COVID-19. *Science Advances* 7.2 (2021).
- Collinson, S. and C. Hu, America's mask resistance is just the latest example of a perennial struggle. CNN. 23 June 2020
<https://www.cnn.com/2020/06/23/world/meanwhile-in-america-june-23-intl/index.html>
- Collman, A. 2 days before his coronavirus diagnosis, Trump mocked Biden for wearing a face mask. *Business Insider*. 2 October 2020.
<https://www.businessinsider.com/trump-coronavirus-mocked-biden-face-mask-presidential-debate-2020-10>
- Druckman, J. N. et al., Affective polarization, local contexts and public opinion in America. *Nature human behaviour* 5.1, 28-38 (2021).
- Edelman, M. J. *The symbolic uses of politics*. University of Illinois Press (1985).
- Farrell, J. Corporate funding and ideological polarization about climate change. *Proc. Natl. Acad. Sci. U.S.A.* 113(1), 92-97 (2016).
- Finkel, E. J. et al., Political sectarianism in America. *Science* 370.6516, 533-536 (2020).
- Fiorina, M. S. Abrams and J. Pope, *Culture War? The Myth of a Polarized America*. New York: Pearson-Longman (2005).
- Gadarian, S. K., S. W. Goodman, T. B. Pepinsky, Partisanship, health behavior, and policy attitudes in the early stages of the COVID-19 pandemic. *PLOS One* 16.4 (2021).
- Gottfried, Jeffrey and Jacob Liedke. "Partisan divides in media trust widen, driven by a decline among Republicans." Pew Research Center. 30 August 2021. Available at: <https://www.pewresearch.org/fact-tank/2021/08/30/partisan-divides-in-media-trust-widen-driven-by-a-decline-among-republicans/>
- Gramlich, J. Far more Americans say there are strong conflicts between partisans than between other groups in society. 19 December 2017. Available at: <https://www.pewresearch.org/fact-tank/2017/12/19/far-more-americans-say-there-are-strong-conflicts-between-partisans-than-between-other-groups-in-society/>

- Greiko, E. Americans' main sources for political news vary by party and age. Pew Research Center. 1 April 2020. Available at: <https://www.pewresearch.org/fact-tank/2020/04/01/americans-main-sources-for-political-news-vary-by-party-and-age/>
- Haischer, Michael H., et al. "Who is wearing a mask? Gender-, age-, and location-related differences during the COVID-19 pandemic." *PloS one* 15.10 (2020): e0240785.
- Hochschild, A. R. *Strangers in Their Own Land: Anger and Mourning on the American Right*. The New Press (2016).
- Hopkins, D. A. *Red fighting blue: How geography and electoral rules polarize American politics*. Cambridge University Press (2017).
- Howard, J. et al. An evidence review of face masks against COVID-19. *Proceedings of the National Academy of Sciences* 118.4 (2021).
- Iyengar, S. G. Sood, and Y. Lelkes, Affect, not ideology: a social identity perspective on polarization. *Public Opinion Quarterly* 76.3, 405-431 (2012).
- Iyengar, S. T. Konitzer, K. Tedin. The home as a political fortress: family agreement in an era of polarization. *J. Politics* 80(4), 1326-38 (2018).
- Iyengar, S. et al., The origins and consequences of affective polarization in the United States. *Annual Review of Political Science* 22, 129-146 (2019).
- Johnson, N. F. et al., The online competition between pro-and anti-vaccination views. *Nature* 582.7811, 230-233 (2020).
- Kahane, Leo H. "Politicizing the mask: Political, economic and demographic factors affecting mask wearing behavior in the USA." *Eastern Economic Journal* 47.2 (2021): 163-183.
- Kenworthy, Nora, Adam D. Koon, and Emily Mendenhall. "On symbols and scripts: The politics of the American COVID-19 response." *Global public health* 16.8-9 (2021): 1424-1438.
- Kerr, J. C. Panagopoulos, and S. van der Linden, Political polarization on COVID-19 pandemic response in the United States. *Personality and Individual Differences* 179 (2021).
- Kessel, P. V. and D. Quinn. Both Republicans and Democrats cite masks as a negative effect of COVID-19, but for very different reasons. Pew Research Center. 29 October 2020. Available at: <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/>
- KGW Staff, Video of Oregon Trader Joe's manager taking on anti-mask shoppers goes viral. KGW8. 17 January 2021. <https://www.kgw.com/article/news/local/video-of-oregon-trader-joes-manager-taking-on-anti-mask-shoppers-goes-viral/283-41a77f89-8bbb-42a1-9b7f-9c563cd7c3c2>
- Lang, J., W. W. Erickson, and Z. Jing-Schmidt. "# MaskOn!# MaskOff! Digital polarization of mask-wearing in the United States during COVID-19." *PloS One* 16.4 (2021). <https://doi.org/10.1371/journal.pone.0250817>.
- Levendusky, M. S. Americans, not partisans: Can priming American national identity reduce affective polarization? *The Journal of Politics* 80.1, 59-70 (2018).
- Lupton, Deborah, Clare Southerton, Marianne Clark, and Ash Watson. *The Face Mask in COVID Times*. De Gruyter, 2021.
- McCarty, N. *Polarization: What everyone needs to know*. Oxford University Press (2019).
- McDonnell, T. E., C. A. Bail, and I. Tavory, A theory of resonance. *Sociological Theory* 35.1, 1-14 (2017).

- Mohr, J. W., P. Bogdanov. Introduction—topic models: what they are and why they matter. *Poetics*. 41(6), 545-569 (2013).
- Powdthavee, Nattavudh, et al. “When face masks signal social identity: Explaining the deep face-mask divide during the COVID-19 pandemic.” *PloS one* 16.6 (2021): e0253195.
- Reyes, V., E., Bogumil, and L. E. Welch, The Living Codebook: Documenting the Process of Qualitative Data Analysis. *Sociological Methods & Research*. Online First (2021). <https://journals.sagepub.com/doi/full/10.1177/0049124120986185>
- Roberts, M. E., B. M. Stewart, D. Tingley. stm: An R package for the structural topic model. R package. (2014).
- Rojas, R., Masks Become a Flash Point in the Virus Culture Wars. The New York Times. 11 June 2020. <https://www.nytimes.com/2020/05/03/us/coronavirus-masks-protests.html>
- Schaeffer, K. Despite wide partisan gaps in views of many aspects of the pandemic, some common ground exists. Pew Research Center. 24 March 2021. <https://www.pewresearch.org/fact-tank/2021/03/24/despite-wide-partisan-gaps-in-views-of-many-aspects-of-the-pandemic-some-common-ground-exists/>
- Stanford Cable TV News Analyzer. Frequently Asked Questions, Stanford Cable TV News Analyzer. 2021. <https://tvnews.stanford.edu/faq>
- Twitter. Coronavirus: Staying safe and informed on Twitter. Twitter, Inc. 12 January 2021. https://blog.twitter.com/en_us/topics/company/2020/covid-19.html
- USC. The Understanding Coronavirus in America tracking survey, University of Southern California. 2021. <https://covid19pulse.usc.edu>
- Vasi, I. B., et al., “No fracking way!” Documentary film, discursive opportunity, and local opposition against hydraulic fracturing in the United States, 2010 to 2013. *American Sociological Review* 80.5 (2015).
- Wessel, L. A., Messaging and the political symbolism of the face mask, *International Social Science Review* 97.1 (2021).
- Westfall, J. et al., Perceiving political polarization in the United States: Party identity strength and attitude extremity exacerbate the perceived partisan divide. *Perspectives on Psychological Science* 10.2, 145-158 (2015).
- Wilson, A. E., V. A. Parker, M. Feinberg. Polarization in the contemporary political and media landscape. *Current Opinion in Behavioral Sciences*. 34, 223-228 (2020).

FIGURES

Figure 1. Results of qualitative coding across all samples. These results are reported as within-sample proportions for comparability across data sources. Because coded positions are not mutually exclusive (except where logically contradictory), and items that mentioned masks only incidentally or did not express any of the positions in our coding scheme could be coded “0” across all variables, totals do not sum to 1.

Figure 2. Normalized proportions of opinion articles (mainstream and right-wing samples pooled) coded as containing anti mask refusal backlash (top panel) and anti-mask discourse (bottom panel), each juxtaposed with self-reported non-masking. The dotted vertical line indicates when the United States Centers for Disease Control and Prevention first issued an advisory to wear masks on April 3rd, 2020. These plots allow for a visual comparison of the dynamics of mask refusal backlash discourse (which increased over time without experiencing a sustained decline) and anti-mask discourse (which initially declined and remained marginal) with the survey results on masking behavior (which show a rapid decline in non-masking stabilizing around an estimated 10% of the American population by late summer). Because the measures of mask refusal backlash and anti-mask discourses are normalized, their magnitudes cannot be compared using this figure. See Figure 1 for a breakdown of all coding categories across all data sources. See Appendix D for a version of this analysis broken down by sampling source. Survey data source: University of Southern California Center for Economic and Social Research’s Understanding Coronavirus in America tracking survey.

Figure 3. Mask-related terms mentioned on cable TV News from January 1st to October 31st, 2020. The y axis indicates how many unique days terms in each category were mentioned. Source: Stanford Cable TV News Analyzer.

Figure 4. Structural topic model results from all opinion articles (mainstream and right-wing samples pooled). Topic labels are followed by the top 12 words associated with each topic. The topic proportions indicate the proportion of the corpus that belongs to each topic. See Appendix E for word clouds that provide further interpretative detail.

Figure 5. Structural topic modeling analysis for the mask refusal backlash discourse and the national politics/election topic. This figure visualizes the results of a linear regression which estimates the effect of covariates (our qualitative coding categories) on the prevalence of topics identified by the STM with 95% confidence intervals. The top panel displays estimates for the effect of the “Mask Refusal Backlash” coding category on all topics in our model, while the bottom panel displays estimates for the effects of all coding categories on the prevalence of the National Politics/Election topic.

APPENDICES

Appendix A. Opinion Articles: Data Collection and Cleaning

To build our sample of mainstream opinion articles, we collected every article in a US newspaper in the Nexis Uni database (formerly known as LexisNexis Academic) that contained the term “mask” or “masks” as well as either “COVID” or “coronavirus.” Our search covered the period of January 1 to October 31, 2020. This search returned a total of 89,089 articles, which were downloaded manually by research assistants. After automatically parsing the articles and metadata into spreadsheets, we isolated articles that appeared in opinion sections of newspapers. Research assistants then manually filtered the remaining corpus to identify opinion articles including op-eds, opinion columns, editorials, and letters to the editor that we could not classify automatically.

To collect our purposive sample of right-wing articles, we chose nine outlets with the highest Alexa Traffic Rankings of those appearing on a list of conservative news outlets on the crowd-sourced conservative website “Conservapedia.”¹¹ The outlets we sampled included: Breitbart, Fox News, Newsmax, The Daily Wire, The Daily Caller, The Blaze, The Federalist, The New York Post, and The Washington Times. Our search included the same parameters as our mainstream sample. We gathered the text of these articles as well as metadata (author, date, headline, and section). Some of the outlets we included did not have a dedicated opinion section, so manually cleaning this sample was particularly important. Research assistants included articles in the analytical sample of opinion articles if they articulated a clear editorial position, rather than simply reporting news.

¹¹ https://www.conservapedia.com/Top_Conservative_news_websites

Appendix B. Twitter Data: Data Collection and Cleaning

To build the analytical sample for our supplemental analysis of Twitter data, we used only the search terms “mask” and “masks,” reasoning that, given the short format of Twitter and informality of Tweets relative to articles, requiring that posts also mentioned “COVID” or “coronavirus” would eliminate many, if not most of the posts with a focus on COVID-related face coverings. (Indeed, when coding Tweets, we found that few mentioned the pandemic in such explicit terms.) This search returned approximately 29 million Tweets. We then extracted a random sample of these Tweets for analysis. Research assistants who conducted qualitative coding of Tweets were asked to identify Tweets that were not about face masks in relation to the pandemic at all. These were primarily Tweets that referred to other senses of the term “mask” (for example, Halloween masks). These Tweets were then removed from the analytical sample altogether. After the removal of these irrelevant Tweets, we were left with an analytical sample of 5,089 Tweets, on which we conducted qualitative coding as discussed below.

Our sample of articles only includes publications based in the United States of America, but Twitter is a global platform. While we omitted Tweets that were not in the English language, our sample does contain some Tweets from outside of the United States. Individual accounts are not required to indicate their accurate location, and the vast majority of Twitter users opt out of sharing the geolocation of their Tweets (Vasi et al. 2015). Therefore, omitting all accounts that did not indicate a location within the United States could introduce its own set of systemic biases that could threaten our analysis even more than including them. We consequently decided not to filter our sample in this way. With respect to the possibility of manually filtering Tweets for location, some Tweets could clearly be classified as originating outside of the United

States, but in many cases, ascertaining the location was either impossible or would require significant additional research for our coding team. We therefore decided not to code Tweets for perceived national origin. After the coding process was complete, we asked our coders if they believed that the vast majority of Tweets in our analytical sample originated in the United States, and they answered affirmatively. Although our data can more formally be said to give insights into the anglophone, rather than American “Twittersphere,” given the global nature of Twitter as a social media platform, which includes the possibility for significant interaction across national borders, and the relative dominance of the United States on the platform overall, we believe our sampling choices do not undermine the overall analytical goals of our supplementary analysis of Twitter data. This supplementary analysis was meant to test the hypothesis that the patterns observed in opinion articles was deeply shaped by editorial gatekeeping practices, not as a stand-alone analysis of the meaning of masks in the American public sphere.

Appendix C. Survey Data on Masking Behaviors

We additionally rely on survey data to draw comparisons between masking behavior and the public sphere treatment of masks. Specifically, we use responses from the University of Southern California Center for Economic and Social Research’s Understanding Coronavirus in America tracking survey.¹² This population-representative survey has the major benefit of being updated daily, allowing a clear juxtaposition between the dynamics of masks’ treatment in news media with self-reported masking behaviors. The question we rely on in Figure 2 in the manuscript asks

¹² <https://covid19pulse.usc.edu/>

respondents if they “wore a mask or face covering in the last 7 days” in order “to keep safe from coronavirus.” We visualize the proportion of respondents who answered “no” to this question to compare with the dynamics of anti-mask, and mask refusal backlash discourse.

One limitation of using this survey question is that it does not account for many variations in masking behavior, including how frequently respondents wore masks, and in what contexts. For instance, a respondent who always wore masks when in the presence of others outside of their household indoors and outdoors would not be differentiated from a respondent who only occasionally wore masks while shopping, but not when gathering with individuals outside of their household. A second major limitation of this survey is that it was not designed with the analysis of political divisions around COVID-19 behaviors and does not include questions on partisan affiliation or lean.

For these reasons, we triangulate our findings with the more in-depth but less frequent polling on mask-related behaviors and attitudes conducted by Pew Research (Schaeffer 2021). The Pew polling shows a similar overall trend as the University of Southern California poll but provides additional detail. The Pew data demonstrates, not only that mask wearing normalized over the course of our study period, but that mask wearing behaviors partly depolarized across political parties. This said, important differences across the political spectrum persisted throughout 2020 and into 2021, as we discuss in the main text.

Appendix D. Qualitative Coding

The bulk of the analysis presented in this article relies on qualitative coding conducted by a team of research assistants and overseen by one of the first authors of the study.

The analytical sample for the qualitative coding of opinion articles includes both the mainstream and right-wing opinion articles pooled together, for a total of 7,970 opinion articles. Figure D1 visualizes the weekly raw counts of these codes for all opinion articles in this sample.

Figure D1. Results of qualitative coding of opinion articles. These results are reported as weekly counts of articles coded affirmatively for each of our qualitative variables.

Figures D2 and D3 replicate Figure 2 in the main analysis but disaggregate opinion articles by sampling source (mainstream and right-wing samples). In line with Figure 1 in the main analysis, they show that there are systematic differences between the two samples. These figures show how those differences are patterned over time. In the mainstream sample (Figure D2), we observe a steady increase in mask refusal backlash discourse without a sustained decline thereafter. We also see a rapid decline in anti-mask discourse in the very beginning of the sample period without a sustained

increase thereafter. In the right-wing sample (Figure D3), the dynamics of the mask refusal backlash discourse tell no clear story. However, there is a clear, albeit modest, increase in anti-mask discourse in the final months of our sample period. As Figure 1 in the main analysis shows, “Anti-Mask” and “Mask Refusal Backlash” were small minority positions among opinion articles in our already smaller right-wing sample, so some caution is warranted when interpreting these results.

Together, these findings support the view that right-wing and mainstream opinion articles had distinct dynamics during our sample period. Mask refusal backlash discourse only clearly increased over time in our mainstream sample (with right-wing sample not showing a clear temporal pattern), and anti-mask discourse only increased over time in our right-wing sample (although later, and somewhat less dramatically than the mask refusal backlash discourse in our mainstream sample). These changes can be interpreted as an increasing polarization of mask discourse in the months leading up to the 2020 presidential election. Yet this polarization in mask discourse is not associated with a population-level change in masking behavior as reported in the survey results, which we juxtapose with the dynamics of published opinions in the same figure. This supports our claim that the mask’s role as a political symbol is irreducible to partisan gaps in masking behavior.

Figure D2. Normalized proportions of opinion articles (mainstream sample only) coded as containing anti mask refusal backlash (top panel) and anti-mask discourse (bottom panel), each juxtaposed with self-reported non-masking. The dotted vertical line indicates when the United States Centers for Disease Control and Prevention first issued an advisory to wear masks on April 3rd 2020. These plots replicate Figure 2 in the main analysis with the mainstream sample only. Survey data source: University of Southern California Center for Economic and Social Research’s Understanding Coronavirus in America tracking survey.

Figure D3. Normalized proportions of opinion articles (right-wing sample only) coded as containing anti mask refusal backlash (top panel) and anti-mask discourse (bottom panel), each juxtaposed with self-reported non-masking. The dotted vertical line indicates when the United States Centers for Disease Control and Prevention first issued an advisory to wear masks on April 3rd 2020. These plots replicate Figure 2 in the main analysis with the right-wing sample only. Survey data source: University of Southern California Center for Economic and Social Research’s Understanding Coronavirus in America tracking survey.

We prepared our data for qualitative coding by pooling mainstream and right-wing opinion articles and randomizing their order. As a result, coders did not know which articles originated in the mainstream or right-wing samples. We also blinded coders to other meta-data that might introduce bias, including the publication source and author. Randomization also ensured that any biases or errors would be uncorrelated with the publication date of articles, thus strengthening our confidence in the temporal patterns we found in our analysis.

The four authors of this study met several times to discuss potential coding schemes. We then each read a random sample of 100 opinion articles and coded them as either “Pro-Mask,” “Anti-Mask,” while also taking notes on various positions that this simple scheme did not account for well and other thematic elements in the articles that could inform a more detailed coding scheme. After comparing our results and deliberating about our findings, we converged on the coding scheme discussed below and in the main text.

The vast majority of the qualitative coding was conducted by a team of research assistants. Among the research assistants who worked on this project, a total of twelve research assistants conducted qualitative coding. As a first step, each research assistant was given the same 100 opinion articles to code. After they completed this test coding set, their completed sheets were quantitatively analyzed for coding agreement. One of the lead authors then met with the research assistants to discuss significant areas of coding disagreement with the aim of arriving at a shared understanding of how coding should be conducted in the future. Out of this process, one of the lead authors developed a codebook which included the explicit meanings of each code and was distributed to all coders. We treated this document as a “living codebook,” as it was periodically updated with more detailed coding decisions and rules that arose out of questions in the process

of coding, and subsequent team meetings that were held periodically to discuss the coding and issues that arose in the process (Reyes, Bogumil and Welch 2021).

The following is excerpted from our codebook that was distributed to the research assistants and reflects only light edits for readability and terminological consistency:

Coding Basics:

All of the variables are binary. A code of “1” (without quotes) is an affirmative code. When you are coding, you should get to the meaning of the statement with respect to masks in the context of the whole article. This said, you shouldn't try to read into the position of the person based on claims that aren't related to masks (for example, just because someone espouses generally conservative views, it doesn't make them anti-mask or anti-mask mandate). This is admittedly challenging. We also need to be careful about irony and the use of quotations. Feel free to read slowly if you aren't sure. It is possible to code all variables “0” if there is no opinion on masks and masks show up incidentally in the article. This is relevant information for us as well.

Notes On Coding Particular Variables

Pro-Mask

Does the author of the article express pro-mask sentiment in the article? For this variable in particular, we need to pay careful attention to implicit claims. Many will make implicitly pro-mask claims. The point, however, is to get to the

meaning of the text, not to make unsubstantiated guesses about the author's views.

Anti-Mask

Does the author of the article express anti-mask sentiment in the article? As with the above, we need to pay attention to implicit claims, and get to the meaning of the text.

A very good question that came up in a meeting was the following: if you are anti-mask, does that automatically make you anti-mask mandate? After deliberating, my collaborator and I decided that the answer to this question is no. Our goal is to get at the meaning of the text, so being anti-mask doesn't automatically make you anti-mask mandate. In order for a text to be coded as anti-mask mandate, the topic of a mask mandate has to be mentioned or invoked. This is because mandates emerged as a topic in particular time and context. We want to measure that. We don't want to undermine our analysis by making unnecessary assumptions. (The difference between this decision and what we're doing with the pro-mask positions is that in that situation, we are inferring a position on masks from a position on mandates. We don't want to infer a position on mandates from a position on masks. This makes sense because the concept of "mask" is included in "mask mandates," but the concept of "mandates" is not included in the concept of "masks").

Pro-Mask Mandate

Does the article express positive sentiment about a mask mandate? The only relevant mandates are government mandates (not mandates at places of work, schools, stores, etc.). It should be about city, county, state, or federal policy requiring people to wear masks in certain settings.

If you are pro-mask mandate, you are automatically pro mask (but being pro-mask doesn't automatically make you pro-mask mandate).

Anti-Mask Mandate

Does the article express negative sentiment about a mask mandate? The only relevant mandates are government mandates (not mandates at places of work, schools, stores, etc.). It should be about city, county, state, or federal policy requiring people to wear masks in certain settings. It is possible to be anti-mask mandate and pro-mask, or to have no expressed position on masks themselves.

Per above, we are not automatically coding anti-mask articles as anti-mask mandate. These two variables should be coded independently.

Backlash Against Mask Refusal

This is to indicate articles in which the author takes a stance against “anti-maskers,” those who do not wear masks, or those who discourage others from wearing masks. The targets of this criticism can be anyone, including ordinary people (generalized or specific), public figures, etc.

Articles coded in this category should also be coded pro-mask (but being pro-mask, does not necessarily imply backlash against mask refusal).

Move to News

If you encounter an article that is not an opinion article, code this as a “1” and move on without coding it along any other variables. Some are ambiguous, so you’ll have to make a judgment call. These will be removed from our analytical sample altogether.

Each coder received a sheet with a randomized set of opinion articles. Coders were initially distributed 600 articles to code and could request additional articles after they completed their first sheet. Each coder only had access to their own individual sheet. The two lead authors of the study reviewed all of the sheets throughout the coding process and noted trends. Because sheets comprised a random sample of articles, they were expected to have a similar distribution of codes. We expected some random variations in the samples, however. Further, although we sought to remove as much ambiguity as possible, the coding process inherently relies on subjective interpretations of text, introducing additional variation. The random distribution of individual sheets across the corpus (including publication sources and dates) ensured that any systematic difference across coders would also be distributed randomly with respect to our structural variables of interest. In a small minority of cases, individual sheets had unusual total quantities of codes along one or more of our coding categories. In such cases, those sheets were assigned to a second coder for an audit. In all such cases, the distributions conformed much more closely to the norm after the audit. Throughout the

entire process, the first authors monitored the coders' progress, including periodically checking samples of their work, and brought up any problems with the coders directly.

Coders also removed extraneous text (for example, copyright information) from articles to clean them in preparation for computational text analysis. Letters to the editor are often published in blocks of letters corresponding to a publication date. As a result, one or more relevant letters will be embedded with several irrelevant letters to the editor. Coders identified relevant letters by noting which mentioned the word "mask," and removed irrelevant letters from each entry. In cases where there were multiple letters in a single entry, a new entry was created for each relevant letter, so that each letter to the editor by a unique author would be treated as its own opinion article, to be coded separately.

The coders also coded our analytical sample of Tweets using the same coding scheme. All of the same rules applied about coding, including the possibility for Tweets to be coded "0" along all variables. However, given the structure of Twitter as a social network and the short length of Tweets, we learned early on in the coding process that the meaning of Tweets was more difficult to understand than articles. In order to assist with interpretation, coders were given access to the URL of each Tweet, which they could visit in order to better decipher its meaning.

Although the Twitter data generally lent itself well to coding using the scheme developed for articles, we noticed that the structure of Twitter's engagement, which involves the ability to reply to others, resulted in a significant minority of Tweets that expressed negative sentiment toward expressions of "Mask Refusal Backlash," but did not necessarily express "Anti-Mask" sentiment. Therefore, in our Twitter analysis alone, we included an additional qualitative code: "Backlash Against Mask Refusal Backlash." We did not include this option in our coding scheme for opinion articles

because we began coding Tweets when we were already well into the process of coding opinion articles. After consulting with all of the coders, it was clear that this was not a common position in the opinion articles, although most coders believed that a small minority would have been coded that way if it had been included in the original coding scheme.

Unlike opinion articles, each Tweet in our dataset is accompanied with metadata on engagements on the Twitter platform. This gave us the opportunity to ask if Tweet popularity was differently patterned in a distinct way from raw Tweet counts. We considered both “retweets” and “likes” as measures of engagement but settled on likes because the number of “retweets” includes “quote Tweets” with added comments that are frequently negative in sentiment with respect to the original Tweet. This means that while “likes” are straightforwardly reflective of positive engagement with a Tweet, “retweets” are fundamentally ambiguous.

To create a measure of Tweets adjusted for likes, we added the total number of Tweets to the total number of likes within each category. Our reasoning was that Tweets with no likes should still be included in this adjusted measure but weighted less than Tweets with a single like. The result is that in our measure, Tweets are weighted by $n+1$, in which n represents a Tweet’s number of likes. We found that the distribution of likes in our analytical sample of Tweets was extremely skewed, with just a few Tweets having several thousands of likes. Because of the modest size of our random analytical sample of Tweets, we ascertained that including these outliers would cause more analytical problems than including them. This is because a single outlier of this sort could dramatically alter the weighted totals for any category. Therefore, we excluded tweets above the 99th percentile in terms of likes when creating our adjusted measure.

Figure D4 compares the coding categories of raw Tweet counts and our measure of Tweets adjusted for likes. The generic increase in proportions of the adjusted measure compared to raw Tweet counts can be explained by the fact that Tweets could be coded “0” along all coding categories (expressing no discernable position), and Tweets that expressed a position on masks were more likely to receive likes than those that did not. The results shown in Figure D4 show that the proportions in each coding category remain fundamentally similar when adjusted for likes. No position receives a wildly disproportionate level of engagement. This supports the view that the raw counts of Tweets in each category is a valid measure of the shape of the overall engagement with the topic of masks on Twitter.

Figure D4. Twitter analysis adjusted for “likes.” This figure compares the raw proportions of Tweets in each coding category with the proportions of Tweets adjusted for likes.

Appendix E. Computational Text Analysis

We trained structural topic models with varying numbers of topics (k) before settling on the six-topic model presented in the main analysis. We judged this model to be the most

parsimonious option in terms of the semantic coherence of its topics and its minimal redundancy in terms of the themes these topics represented. The two lead authors of this study arrived at the topic labels for our model through deliberation. They considered the top terms in each topic by probability of occurrence (beta), referring to word clouds generated by STM. Because the most important topic for our analysis is the topic termed “National Politics/Election,” our primary concern during model specification and assigning topic names was making sure this topic was accurately labeled and that it was clearly distinct from the other topics. Figure E1 displays word clouds for each of our six topics, where larger words have higher probabilities of occurring within the given topic. This figure represents the same topics as Figure 4 in the main analysis but provides more interpretive detail regarding the lexical content of each topic.

Figure E1. Word clouds visualizing each topic in our six-topic model ($k = 6$) used in the main analysis. These word clouds display words highly associated with each topic, with higher probability words represented larger.

Figure E2 represents the topic proportions for an alternative specification of the topic model with more topics ($k = 8$). For the most part, the overall pattern is the same as in the six-topic model. This eight-topic model contains several topics that correspond neatly with topics from the six-topic model in the main analysis. For instance, Topic 7 is similar to the topic we termed “Economy,” although Topic 6 also contains terms from the “Economy” topic as well. Topic 4 corresponds with “Schools,” and Topic 5 corresponds with “National Politics/Presidential Election.” Topics 3, 6, 1, and 8 contain terms that are distributed among the topics from our six-topic model termed “Public Health,” “Social Interaction,” and to a lesser extent, “Economy.” While they are semantically coherent for the most part, we judged that, from an analytical perspective, little was gained from adding additional topics to our model.

Figure E2. Alternative specification of the structural topic model results from all opinion articles (mainstream and right-wing samples pooled). This figure reproduces Figure 4 in the main analysis with an eight-topic model ($k = 8$). This

alternative specification includes the top 12 words associated with each topic. The topic proportions indicate the proportion of the corpus that belongs to each topic.

Our other reason for preferring the simpler six-topic model was that our main finding from this portion of our analysis remained robust to both specifications. That is, both models contain a single topic clearly related to national politics and the 2020 presidential election, and both models demonstrate that this topic is by far most prevalent in articles coded as containing “Mask Refusal Backlash” discourse. Both also show that this national politics topic is the most prevalent topic *within* these “Mask Refusal Backlash” articles. Figure E3 reproduces Figure 5 in the main analysis with the eight-topic model. We have not assigned names to all of the topics in this model but have labeled only the “National Politics/Presidential Election” topic in order to demonstrate the robustness of our central results that rely on computational text analysis.

Figure E3. A reproduction of Figure 5 in the main analysis with an eight-topic model ($k = 8$) showing robustness of our results to an alternative specification. In this specification we have only assigned Topic 5 from Figure E2 a descriptive label because it clearly corresponds to the “National Politics/Presidential Election” topic in the six-topic model and is the primary topic of interest for the purposes of our argument. This figure visualizes the results of a linear regression which estimates the effect of covariates (our qualitative coding categories) on the prevalence of the topics identified by the STM with 95% confidence intervals. The top panel displays estimates for the effect of the “Mask Refusal Backlash” coding category on all topics in our model, while the bottom panel displays estimates for the effects of all coding categories on the prevalence of the National Politics/Election topic.

ACKNOWLEDGEMENTS

This research would not have been possible without excellent assistance collecting, cleaning, and analyzing data from Gabriella Cantor, Zachary Churchill, Margot Durfee, Lila Felser, Jamie Gareh, Grace Hagopian, Ilisa Lama, Amanda Joseph, Maya Lewis-Burgess, Ellen Murphy, Lauren Pollak, Abigail Pomeranz, Rachel Richman, and Elizabeth Sullivan (listed in alphabetical order). The authors also thank (without implicating) Adam Davis, Felipe Dias, Shai Dromi, Anjuli Fahlberg, Yuval Feinstein,

Tara Gonsalves, Laura Halcomb, Paul Joseph, Helen Marrow, Bo Yun Park, Deborah Seligsohn, Ioana Sendroiu, Mary Shi, Sarah Sobieraj, Adam Storer, Rosemary Taylor, Natasha Warikoo, and Jill Weinberg for helpful comments on this work. Earlier versions of this research were presented at the Tufts Department of Sociology and the Seminar of the Research Cluster on Comparative Inequality and Inclusion at Harvard University's Weatherhead Center for International Affairs.