Uyheng et al. BMC Digital Health (2023) 1:27
https://doi.org/10.1186/s44247-023-00026-z

BMC Digital Health

RESEARCH

Open Access

Bridging online and offline dynamics of the face mask infodemic

Joshua Uyheng[1*], Dawn C. Robertson[1] and Kathleen M. Carley[1]

Abstract

Background Online infodemics have represented a major obstacle to the offline success of public health interventions during the COVID-19 pandemic. Offline contexts have likewise fueled public susceptibility to online infodemics. We combine a large-scale dataset of Twitter conversations about face masks with high-performance machine learning tools to detect low-credibility information, bot activity, and stance toward face masks in online conversations. We match these digital analytics with offline data regarding mask-wearing and COVID-19 cases to investigate the bidirectional online-offline dynamics of the face mask infodemic in the United States.

Results Online prevalence of anti-mask over pro-mask stance predicts decreased offline mask-wearing behavior and subsequently increased COVID-19 infections. These effects are partially influenced by low-credibility information and automated bot activity, which consistently feature greater anti-mask stance online. Despite their purported controversy, mask mandates generally decrease anti-mask stance online and increase mask-wearing offline, thus reducing future COVID-19 infections. Notable asymmetries are observed, however, between states run by Democratic and Republican governors: the latter tend to see higher levels of low-credibility information and anti-mask stance online, and thus lower mask-wearing and higher infection rates offline.

Conclusions These findings contribute new insights around collective vulnerabilities to online infodemics and their links to evolving offline crises. We highlight the need to synergize and sustain targeted online campaigns from legitimate information sources alongside offline interventions in and beyond the pandemic.

Keywords COVID-19 pandemic, Face masks, Infodemic, Low-credibility information, Bots

Background

Over two years since the beginning the COVID-19 pandemic, broad scientific consensus has highlighted the need for both top-down systemic interventions and bottom-up collective changes in public behavior [1, 2]. Face masks have been crucial to such initiatives, given their effective disruption of disease transmission while posing minimal risk to users [3, 4]. However, many countries have seen public health mandates entangled with ideological conflicts: over and above personality, demographics, or local risk levels [5, 6]. These dynamics have obstructed pandemic responses worldwide, extending beyond masks and into later interventions like vaccines [7, 8].

Rich scholarship in digital health points to the overwhelming wave of online information linked to the pandemic, also known as the 'infodemic', as a major driver of these societal divides [9–11]. Social media platforms feature low-credibility information undermining expert

*Correspondence:
Joshua Uyheng
juyheng@cs.cmu.edu
[1] CASOS Center, Software and Societal Systems Department, Carnegie Mellon University, 5000 Forbes Ave, 15213 Pittsburgh, PA, USA

communications and public policy [8, 12, 13]. Evidence likewise continues to accrue around information campaigns involving automated bots to spread misinformation, hate speech, and conspiracy theories during the pandemic [14, 15].

Yet despite the rich literature on online and offline conflicts around face masks, surprisingly little is known that connects these two sets of processes. Prior scholarship in the social sciences suggests that offline events trigger beliefs and behavior reflected in the online world [15, 16]. Yet digital platforms also reshape conversations and communities that impact the offline world [17, 18]. From this perspective, we ask: How have offline events during the COVID-19 pandemic affected the online face mask infodemic? And conversely, what impacts has the online face mask infodemic had on offline public health outcomes?

Focusing on the United States, this paper quantifies the online-offline dynamics of the face mask infodemic. The United States serves as a crucial site for this work given the sharply politicized nature of face masks during the COVID-19 pandemic [5, 6]. Mirroring similar contexts in other countries, messaging from political leaders around face masks was hotly divided, and these rifts persisted among the general public in a form of 'culture war' against repression of 'individual freedoms'[1]. Nonetheless, state-wide mask mandates were implemented to bolster public health responses against the virus, even if in uneven ways that broke according to political lines [19]. Such societal dynamics make for fertile ground for triggering infodemics, as well as for infodemics to subsequently exacerbate.

We undertake this analysis by integrating large-scale online Twitter data with various offline records including official COVID-19 case reports, mask-wearing surveys, and mask mandate announcements. We take advantage of a retrospective and ecological design to closely examine a global period during which mask mandates served as a core intervention prior to the emergence of vaccines and challenges stemming from viral variants. Hence, though the crisis has since evolved over time, these findings offer unique, focused insights that may inform synergized measures to tackle "two fronts" of health and social cohesion in and beyond the pandemic [20, 21].

Data and methods

Twitter data

We collected a large sample of tweets from the United States between January 29, 2020 and November 30, 2020 ($N = 22,143,552$), filtered for tweets about the COVID-19 pandemic with any terms that referenced masks and mask-wearing (e.g., 'mask', 'shield', 'cover'). We combined explicit platform-enabled geolocation with a machine learning model for location prediction to limit our data corpus only to tweets that were likely to have originated from the United States and to identify their specific state of origin [15, 22]. This technique has previously produced fine-grained classifications at a 92.1% accuracy [22].

Offline records

Offline records of *daily state-level COVID-19 cases* were obtained from the Centers for Disease Control and Prevention's (CDC) COVID Data Tracker[2]. This was linked to the *schedule of mask mandates* taken from the United States government's open data catalog, indicating the start and end dates for various orders[3]. We additionally noted the *official party affiliation of each state's governor prior* to the November 2020 elections[4].

To capture *mask-wearing*, we used datasets from COVIDcast by Carnegie Mellon University's Delphi group for trend information around reported mask-wearing in each state[5]. This data is based on seven-day smoothed self-reports of mask-wearing behavior in public, and records only began to be recorded on September 8, 2020. Note that smoothing may introduce issues in lagged regressions and path modelling, which we address in the Supplementary Material. We additionally recognize that because of its reliance on self-reports and potentially biased samples, these measures are also highly coarse-grained estimates [23].

Collectively, our analyses thus focused only on the state level. Although mask mandates could indeed be enacted on finer-grained scales—cities and counties—such microscopic location information is rarely available to link to social media or mask-wearing activities. The state level was thus chosen as an optimal point of comparison given trade-offs in precision and information accessibility.

Computational predictions

To measure the sharing of *low-credibility information*, we used an existing list of low-credibility websites [13]. For a given state on a given day, we calculated the percentage of tweets that contained a link to any of these

[1] https://www.theguardian.com/world/2020/jun/29/face-masks-us-politics-coronavirus

[2] https://data.cdc.gov/Case-Surveillance/United-States-COVID-19-Cases-and-Deaths-by-State-o/9mfq-cb36/data

[3] https://catalog.data.gov/dataset/u-s-state-and-territorial-public-mask-mandates-from-april-10-2020-through-july-20-2021-by--7e5b8

[4] https://www.nga.org/governors/

[5] https://delphi.cmu.edu/covidcast/?mode=indicator+&%3Bsensor=fb-survey-smoothed_wwearing_mask_7d

websites. Because new low-credibility information sites are dynamically evolving, this measure provides a conservative estimate of how much low-credibility information was shared on Twitter in relation to masks and mask-wearing. Note this list includes satirical sites as we do not assume exposure to satire is well-discerned by the public.

Next, to assess the activity of *bots*, we used the BotHunter tool [24]. BotHunter is a machine learning model that achieves high accuracy and scalability in bot prediction on Twitter. In our dataset, BotHunter produced a probability that the account which sent each tweet was a bot. This accounts for both individual account features as well as some features of the social network in which the bot engages in its activities. While based on a well-established dataset of known automated accounts, the algorithm achieves over 90% accuracy but may not perfectly classify all accounts. In addition, it may pick up on accounts that behave in a bot-like fashion but are not fully automated (e.g., cyborg accounts). Taking the average BotHunter probability provided us with state-level daily estimates of the distribution of likely automated accounts in online masking conversations during the pandemic.

Finally, to classify *stance* toward face masks, we used a state-of-the-art machine learning approach with label propagation and co-training to identify pro-mask and anti-mask stance [25]. This method is a semi-supervised method, meaning that it uses the social network to identify pro- and anti-mask hashtags based on how they are used in similar contexts. For this work, some of the most common pro-mask hashtags were validated to include: #WearAMask, #WearADamnMask, and #YourActionsSaveLives. Anti-mask hashtags included: #MasksOffAmerica, #masksdontwork, and #RefuseMasks. More information around the top hashtags identified by the algorithm is available in the Supplementary Material. The top hashtags for each stance polarity validate the predictions of the methodology.

Given individual stance classifications, we computed the daily percentage of pro-mask and anti-mask stance on a collective level per state. To jointly consider the effects of anti-mask relative to pro-mask stance, we computed an aggregate measure represented by the natural logarithm of their quotient, i.e., $\ln \frac{\#\text{anti}-\text{masktweets}}{\#\text{pro}-\text{masktweets}}$.

Note that because bots may intend to skew the stance of online conversations toward face masks, we retain the original measure without filtering for bots in subsequent analysis. This is because we are interested here in how the social media conversation overall orients toward face masks, which automated tweets may also have an impact on. Hence, by measuring both bot activity and stance toward face masks, we can use their co-variation over time to account for the extent to which the former impacts the latter.

Multilevel mediation modeling

Multilevel models have been used in longitudinal studies where individuals (or in this study, states) are treated as clusters that change over time [26]. Two multilevel mediation models are implemented in this study: (a) one on the overall data corpus without mask-wearing behavior (Model 1); and (b) one on a subset of the dataset coinciding with available mask-wearing behavior (Model 2). These models quantified how offline factors like mask mandates, current COVID-19 cases, and the state governor's political party predicted online prevalence of anti-mask stance, low-credibility information, and bots. Conversely, we then used these online factors to predict offline outcomes such as future COVID-19 cases, partially mediated by actual mask-wearing behavior. Note that because we are studying these relationships at a state-level over time, the unit of analysis is the state-day. This means that while we have collected 22 million tweets featuring 1.7 million accounts in total, we are interested in measures that apply per state per day (e.g., cases per state per day). For Model 1, this results in 11266 data points; and for Model 2, this results in 4010 data points.

All measures were mean-centered and scaled to have unit variance. Coefficients may thus be interpreted as standard deviation changes in the response variable as a result of standard deviation changes in the predictor. Time variables, however, were simply divided by the total number of days in the dataset, while a binary variable was retained to indicate whether a state had a Republican governor. In the special case of mask mandates, we used interrupted time series techniques [27] to model three types of effects: (a) a shock variable, coded as a binary measure for before and after the implementation of mask mandates; (b) a pre-shock trend, which captures the slope of the outcome variable over time before mask mandates; and (c) a post-shock slope change, which captures the modified slope of the outcome variable over time after mask mandates are implemented.

All predictors were treated as fixed effects except for the governor's party affiliation, which was treated as a state-level random effect. Mediation models were evaluated with: (a) two fit indices, namely, the Comparative Fit Index (CFI) and the Tucker-Lewis Index (TLI); and (b) two error measures, namely, the root mean square error of approximation (RMSEA) and the standardized root mean square residual (SRMR) [28]. All models achieving high values in the fit indices (> 0.90) and low values on the error measures (< 0.10) were considered adequate models for describing the data.

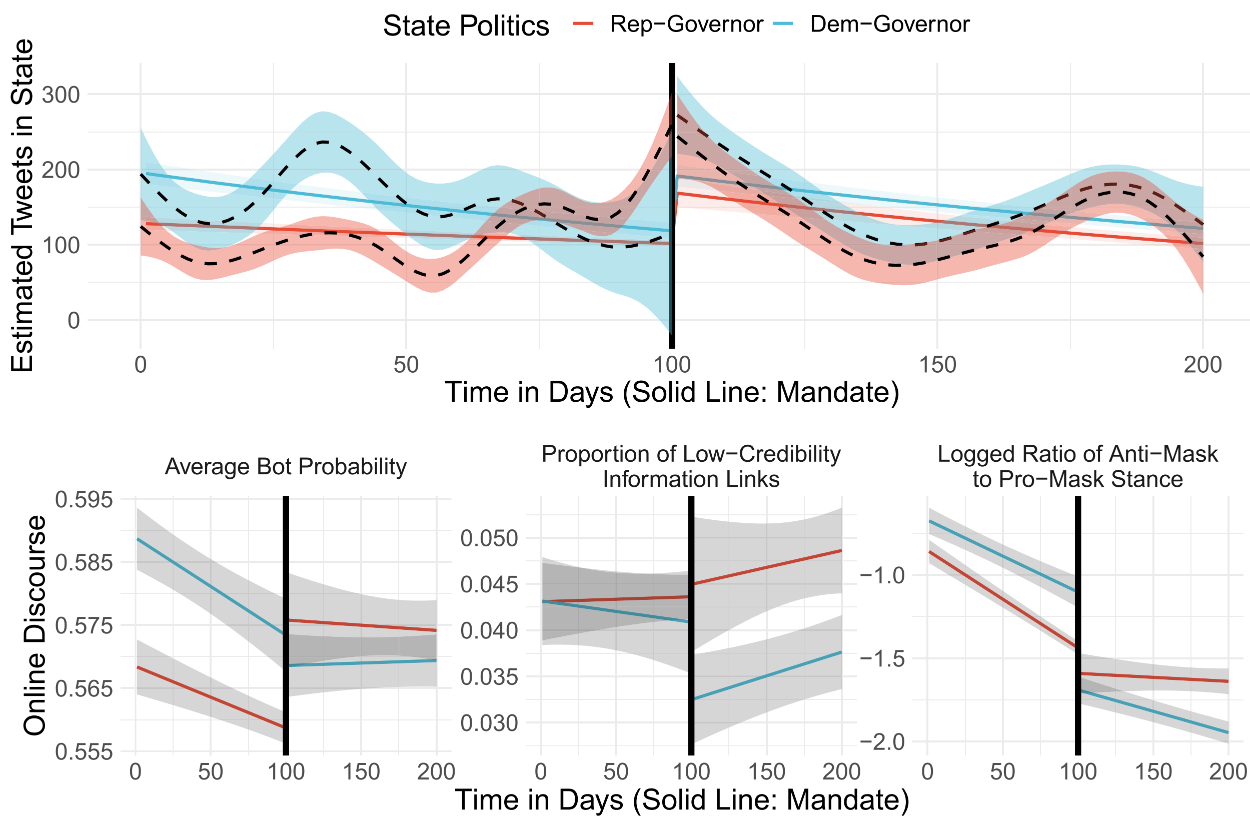

Fig. 1 Volume and features of social media conversations in states with Republican (red) and Democratic governors (blue), with mask mandates modelled as a shock in an interrupted time series setup. Top: Average daily volume of tweets over time. Loess trends (dashed) are presented alongside the predictions of negative binomial regressions with 95% confidence intervals. Bottom: Average bot probabillity, proportion of low-credibility information links, and logged ratio of anti-mask to pro-mask stance over time. Linear predictions are presented with 95% confidence intervals

Results

Offline contexts shift online discourse

Figure 1 depicts significant shifts in online discourse following the implementation of mask mandates. Initially, online talk in Republican-run states discussed face masks 34% less than in Democrat-run states ($p < .001$). Yet the face mask conversation saw a notable spike upon the announcement of mask mandates, approximately increasing 1.62 times in size across all states ($p < .001$). Further, over time, although the conversation generally waned in volume, it did so more quickly in Democrat-run states at a rate of -0.50% ($p < .001$) every day. In Republican-run states, the conversation decayed by only -0.23% a day ($p < .001$), eventually closing the gap in tweet volume.

Beyond the volume of tweets, even prior to mask mandates, bots predominantly targeted Democrat-run states ($p < .001$), such that there was no clear difference in the level of bot activity after mandates were put in place ($p > .05$). Conversely, bots swarmed to Republican-run states with the enforcement of mask mandates ($p < .001$). Before the mandates, bot activity decreased steadily over time in all states ($p < .001$), but the level of bot activity held steady after mask mandates were implemented ($p < .001$). Bots thus strategically targeted Republican-run states with mask mandates in place, and thereafter sustained their participation throughout online conversations everywhere.

Meanwhile, states generally saw their share of low-credibility information links reduced upon the announcement of mask mandates ($p < .001$), though there was weak evidence for the opposite effect in Republican-run states ($p < .10$). Hence, across states in general, legitimate information sources became relatively more prevalent on the day mask mandates began. However, both Democrat-run and Republican-run states eventually saw marked increases in low-credibility information links ($p < .05$). Low-credibility information thus demonstrates notable resilience against more legitimate sources of information, especially if the latter become dominant only in the initial stages of policy implementation.

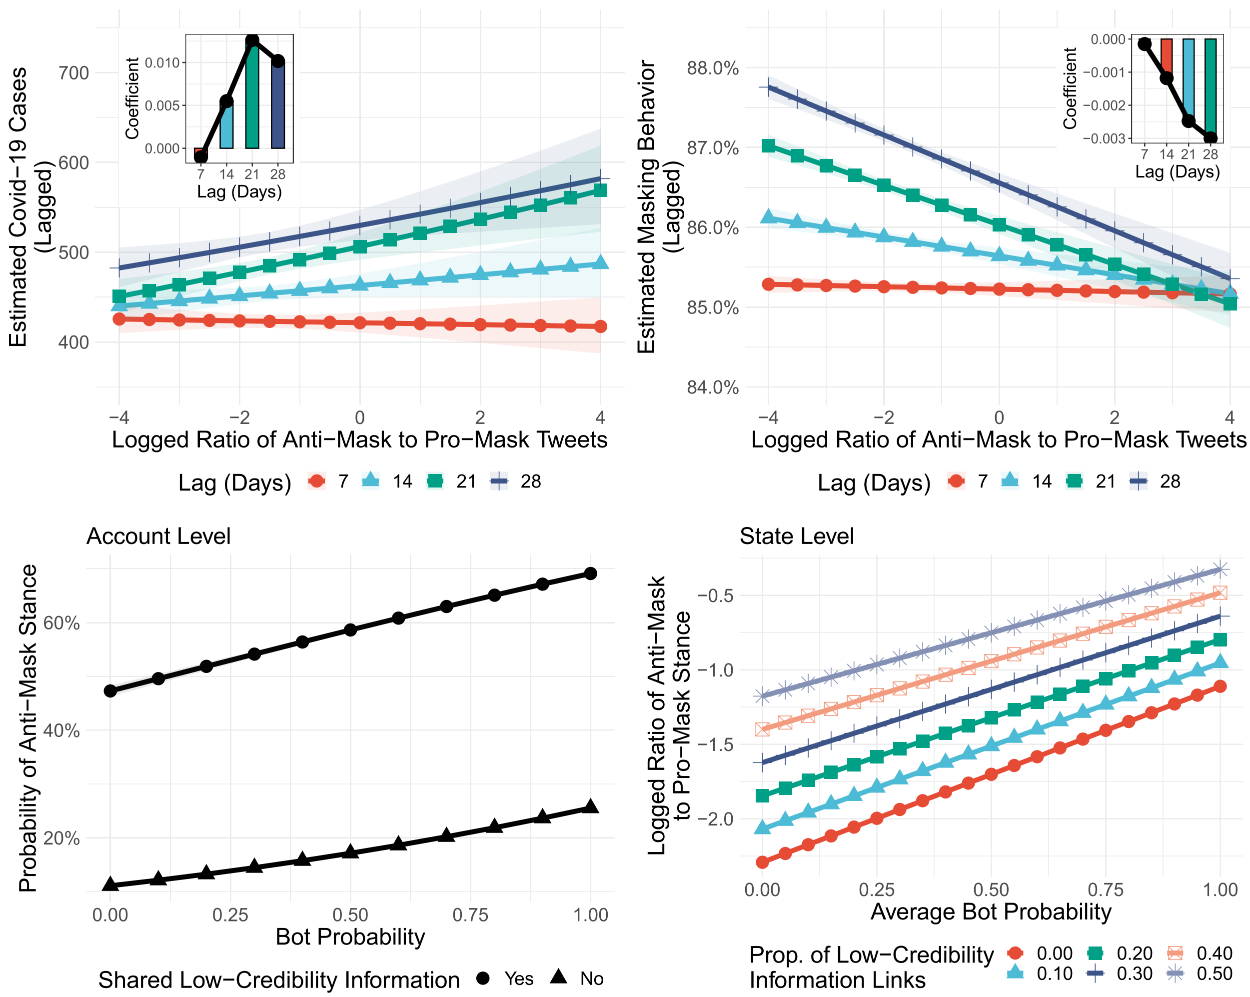

Fig. 2 Future COVID-19 cases and mask-wearing are associated with relative prevalence of anti-mask to pro-mask stance online, which are in turn shaped by bot activity and low-credibility information. Top-Left: Increased share of anti-mask stance predicts higher levels of lagged new daily COVID-19 infections in a given state ($N = 11266$). Top-Right: Increased share of anti-mask stance predicts lower levels of actual masking behavior in a given state ($N = 4010$). Top-Insets: Coefficient estimates across time lags, with higher absolute values indicating stronger associations between online discourse and offline impacts. Bottom-Left: Account-level estimates of anti-mask stance, given bot probability and low-credibility information-sharing ($N = 1787963$). Bottom-Right: State-level estimates of the logged ratio of anti-mask to pro-mask stance based on the average bot probability and proportion of low-credibility information-sharing ($N = 11266$)

Finally, we observed that pro-mask tweets generally rose significantly relative to their anti-mask counterparts with mask mandates put in place ($p < .001$). However, while pro-mask stance grew more steadily in Democrat-run states, relative levels of anti-mask stance held more rigidly in Republican-run states ($p < .001$).

Online discourse predicts offline impacts

Figure 2 shows that in states where anti-mask stance was more prevalent online relative to pro-mask stance, higher rates of COVID-19 infection were more likely to occur weeks later. The strongest relationship was apparent with a three-week lag ($p < .001$). Conversely, we saw a consistent negative relationship between the ratio of anti-mask to pro-mask stance, and actual mask-wearing behavior, controlling for present levels of mask-wearing ($p < .001$). In general, for both future COVID-19 cases and mask-wearing behavior, shorter-term lags yield weaker or non-significant associations, affirming that indeed digital influences may not link to offline behaviors—and crucially, pandemic outcomes—in an instantaneous fashion [3].

Importantly, in the context of an infodemic, online discourse did not evolve without influence from bots and low-credibility information. Figure 2 additionally shows that, on the account level, bots were more likely than humans to express anti-mask stance ($p < .001$). Sharing of low-credibility information also indicated that an account was likely to be anti-mask ($p < .001$). It was estimated that the most bot-like accounts which spread low-credibility information had about a 69.11% probability of being anti-mask, while the most human-like accounts who did not spread low-credibility information had only about an 11.08% probability of being anti-mask. A Welch two-sample t-test further revealed that sharers of low-credibility information had higher bot probabilities than accounts which did not share low-credibility information ($t(90524) = 57.053, p < .001$). Hence, not only were bots more likely to be anti-mask; they also shared more low-credibility information.

These same patterns hold on the state level. On any given day, a 10% increase in low-credibility information was specifically linked to about a 2.23 unit increase in the logged ratio of anti-mask to pro-mask stance ($p < .05$). Further, a 10% increase in a state's average bot probability was linked with a 1.18 unit increase in the logged ratio of anti-mask to pro-mask stance ($p < .001$). The greatest prevalence of anti-mask stance relative to pro-mask stance was thus detected in states where online conversations about face masks were inundated with the most low-credibility information links and social bots.

Bidirectional dynamics of the face mask infodemic

Piecing this information together, offline contexts thus trigger shifts in online discourse, and in turn, online discourse links to offline impacts in the context of the COVID-19 pandemic in the United States. Figure 3 displays the estimates of two multilevel mediation models. The first model reflects the full data corpus without mask-wearing information and predicts future COVID-19 cases with a 21-day lag. The second model predicts future COVID-19 cases with a 21-day lag, mediated by future mask-wearing with a 14-day lag. We obtain estimates of direct, indirect, and total effects of these factors on the trajectory of the COVID-19 pandemic in the United States, as summarized in Table 1.

Both models affirm that even after controlling for a variety of offline factors—such as mask mandates, the state governor's party, and current COVID-19 cases—the ratio of anti-mask to pro-mask stance was significantly linked to future COVID-19 cases. This was evident in both direct effects (Model 1: $p < .001$), and indirect effects as mediated by actual mask-wearing (Model 2: $p < .01$). Widespread online conversations about face masks therefore had credible public health impacts, especially via changes to mask-wearing behaviors to the benefit or detriment of pandemic outcomes.

Crucial to note is that greater prevalence of anti-mask over pro-mask stances was associated with higher levels of low-credibility information sharing (Model 1: $p < .001$, Model 2: $p < .001$) and greater bot activity (Model 1: $p < .001$, Model 2: $p < .001$). This highlights the susceptibility of public opinions online to undue influence from misinformation and disinformation, and underscores the public health costs associated with these online dangers.

These online vulnerabilities were in turn shaped by their evolving offline contexts. Sharply counter to public health interests, times with higher current infections invited more anti-mask stance (Model 1: $p < .001$). State politics were paramount in this regard, as Republican-run states generally saw more anti-mask stance online (Model 1: $p < .10$, Model 2: $p < .05$) amidst lower mask-wearing among their populations (Model 2: $p < .001$) and more low-credibility information spread by humans (Model 1: $p < .01$, Model 2: $p < .10$). Meanwhile, bots targeted states with the highest mask-wearing rates (Model 2: $p < .01$), which were more likely to be Democrat-run (Model 2: $p < .001$). In addition, both bots (Model 1: $p < .01$, Model 2: $p < .001$) and low-credibility information (Model 1: $p < .05$) targeted states at times when they had lower infection rates.

Finally, we observe the key impacts of mask mandates. For all states, mask mandates predicted lower levels of future COVID-19 infections (Model 1: $p < .001$), mediated by greater pro-mask stance over anti-mask stance online (Model 1: $p < .001$). These effects were mediated by offline mask-wearing behaviors, indicating a coherent causal link between online discourse about face masks and its offline impacts (Model 2: $p < .001$).

Discussion

Amidst a surge of research on infodemics coinciding with COVID-19 outbreaks around the world [9, 10], crucial unanswered questions have persisted around their effects on public health outcomes. We unite a wealth of disparate research separately examining online and offline processes [1, 17] to unearth a complex understanding of their interplay in the face mask infodemic. These insights suggest practical directions for forward-looking policies in and beyond the COVID-19 pandemic that has ravaged the globe for over two years now.

This research emphasizes three online-offline links uncovered in relation to the face mask infodemic. First, politically bifurcated vulnerabilities to online harms—that is, to low-credibility information in Republican-run states and social bots in Democrat-run states—expose unique informational challenges faced in a polarized democracy [8, 13]. Past scholarship highlights the need for robust interventions against pandemic misinformation [7, 21]. By linking online and offline processes, our work encourages that these responses be sensitively

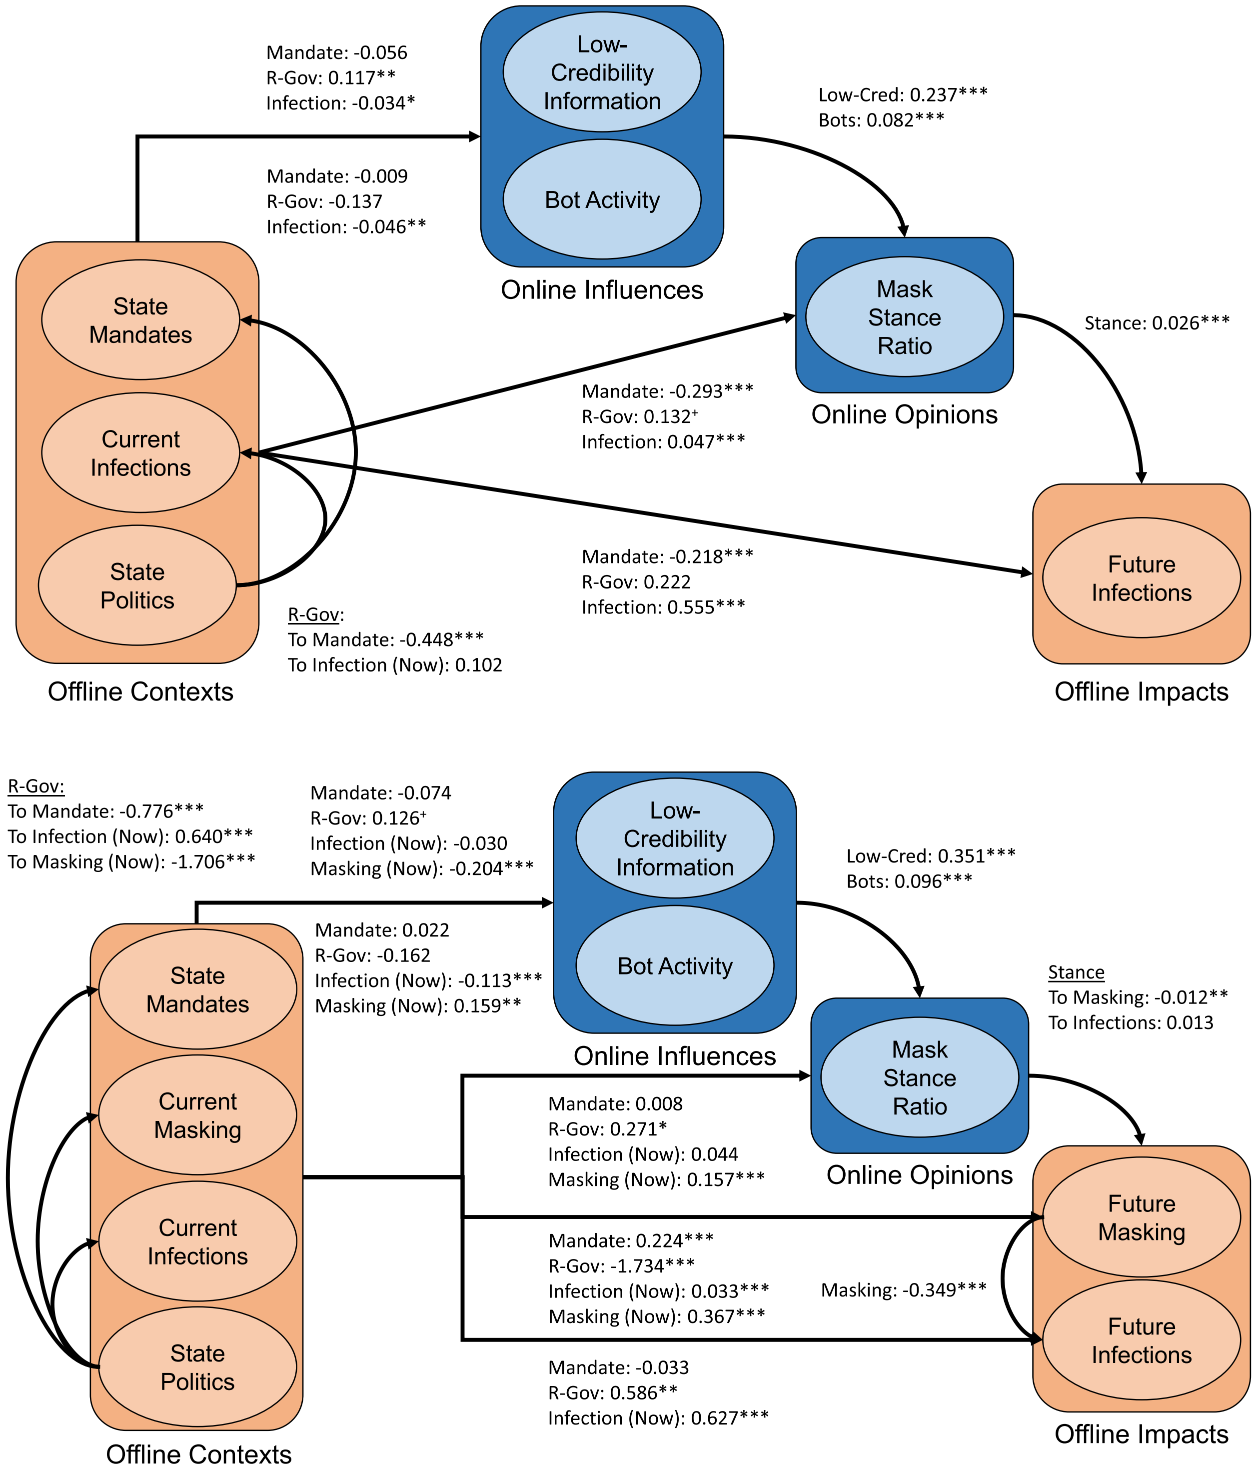

Fig. 3 Mediation models of online-offline dynamics in relation to the face mask infodemic. The effect of state politics and mask mandates on future mask-wearing (14-day lag) and COVID-19 cases (21-day lag) is partially mediated by low-credibility information and bot activity, which in turn shape online anti-mask and pro-mask stance. All coefficients are standardized, with statistical significance reported as: $^{+}p < .10$, $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$. Top: Model 1 includes the full data corpus ($N = 11266$) without mask-wearing data. Bottom: Model 2 uses a subset of the data ($N = 4010$) coinciding with mask-wearing data

Table 1 Direct, indirect, and total effects of offline contexts, online influences, and online discourse on future COVID-19 infections (21-day lag). Model 1 includes the full data corpus without mask-wearing data ($N = 11266$). Model 2 uses a subset of the data coinciding with mask-wearing data ($N = 4010$). Estimates are standardized except binary variable for Republican governor

Model	Variable	Direct	Indirect	Total
Model 1	Present Cases	0.555***	0.001*	0.556***
	Mask Mandates	-0.218***	-0.008***	-0.226***
	Republican Governor	0.222	0.162	0.383
	Stance Ratio	0.026***	-	0.026***
	Bot Activity	-	0.002***	0.002***
	Low-Credibility Information	-	0.006***	0.006***
Model 2	Present Cases	0.627***	-0.011***	0.616***
	Mask Mandates	-0.033	-0.078***	-0.111+
	Republican Governor	0.586**	1.306***	1.892***
	Mask Behavior	-0.349***	-	-0.349***
	Stance Ratio	0.013	0.004**	0.017*
	Bot Activity	-	0.002+	0.002+
	Low-Credibility Information	-	0.006+	0.006+

Note: $^{+}p < .10$, $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$

targeted to their social contexts. Organic misinformation in Republican-run states may be better addressed with accuracy nudges to correct organic misinformation [29]. On the other hand, more mask-compliant Democrat-run states may require additional investment in exposing and countering coordinated influence campaigns seeking to inorganically contaminate conversations that are otherwise informed by more legitimate sources. Future work drawing on these insights may delve deeper into the specific types of messages propagated by bots or low-credibility information sources across the political divide, in order to design more pointed strategies at countering their influence.

However, it is not sufficient to merely root out pandemic-related falsehoods without addressing the underlying conditions which enable their proliferation [30]. This brings us to our second insight: although mask mandates have been subject to widespread controversy [5, 6], the conflicts they precipitate have less to do with the policies themselves, and instead reflect already existing political cleavages. Our findings reaffirm that mandates improve mask-wearing and reduce infections [3, 4]. But we echo that communication around their implementation should aim to minimize partisan divides to avoid nullifying their positive impacts [12]. From this perspective, more in-depth analysis could be performed in future work that tackles specifically how pro-mask and anti-mask stances respond to the implementation of mask mandates, to identify effective messaging around these policies which mitigate political flare-ups.

These observations culminate in our third, overarching finding linking online stances to future infections, an association mediated by actual mask-wearing. Our results reaffirm that online talk does not remain isolated to the digital sphere. Instead, it explains offline outcomes beyond what one would expect purely from exclusively offline factors. This showcases social media talk's offline consequences for collective behavioral change, and its overall importance in the public health arena [7, 8]. Future work may benefit from drilling down on the specific social processes involved in this transition from online to offline dynamics. This affirms the need for sustained and targeted public health communication, both online and offline, with messages that mobilize identity-aligned narratives and originating from politically trustworthy sources of information [1, 7].

The interoperable methodology we employ likewise constitutes valuable contributions to future study and practical monitoring of infodemics during crisis [9, 31]. Although we focus on the face mask conversation in the United States, many of the tools we employ are readily applicable across novel contexts. From a practical standpoint, this may guide efforts to design informational counter-measures and enhance digital resilience, especially to misinformation and disinformation [1, 29].

As with any research on social media platforms like Twitter, nuance is needed in interpreting findings for generalizability. Our findings reflect a reliance on several external databases, and are thus also dependent on their quality and provenance. This includes, for instance, the reliance on self-reports around mask-wearing data, which may reflect social desirability biases. This likewise extends to the use of predictive algorithms which represent practical and scalable labeling tools for large-scale datasets based on prior knowledge, but may not always capture the ways these phenomena evolve in the future [13, 24, 25]. Finally, as we restrict our analysis to the state level, our work does not speak to more fine-grained settings (e.g., counties) which previous work has examined [6]. Relatedly, nor do we assume that our results are straightforwardly applicable outside the United States or the specific time period upon which we focus. However, these gaps open up valuable avenues for future work, especially to investigate the global impacts of diverse infodemics as they evolve alongside the pandemic around the world.

Conclusions

Considerable research examines the offline social conflicts surrounding face masks as a public health intervention during the COVID-19 pandemic. Digital health scholarship likewise characterizes the online infodemic that has muddied public discourse with a wave of low-credibility information and increased polarization throughout the global crisis. However, surprisingly little work links the online infodemic with its offline antecedents and outcomes. Our findings richly add to prior work which separately examines online and offline processes by demonstrating novel, powerful interconnections between them. We show that as states enforce mask mandates, online conversations shift in terms of anti-mask stance, bots, and low-credibility information. These factors in turn predict trajectories in COVID-19 infections, beyond the predictive capacity of offline factors and present COVID-19 infections. This is partially mediated by actual mask-wearing behavior, suggesting a coherent causal link between online discourse and future offline behavior. Taken together, these insights guide forward-looking action in and beyond the pandemic and inform measures for enhanced resilience to future crises and conflicts. Offline implementations of public health interventions require complementary online campaigns that minimize group division and promote legitimate information sources on digital platforms. Likewise, digital measures to counter online infodemics need to be sensitively tailored to their offline contexts of social division.

Supplementary Information

The online version contains supplementary material available at https://doi.org/10.1186/s44247-023-00026-z.

Additional file 1.

Acknowledgements
Not applicable.

Authors' contributions
J.U., D.C.R., and K.M.C. contributed to conceptualization, methodology, software, and validation. J.U. and D.C.R. performed data curation, formal analysis, and visualization. All authors contributed to writing the original draft, reviewing, and editing. All authors approved the submitted version of the manuscript.

Funding
This work was supported in part by the Knight Foundation and the Office of Naval Research grants N000141812106 and N000141812108. Additional support was provided by the Center for Computational Analysis of Social and Organizational Systems (CASOS) and the Center for Informed Democracy and Social Cybersecurity (IDeaS). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Knight Foundation, Office of Naval Research or the U.S. government. The funding bodies did not play any role in the design of the study, the collection, analysis, and interpretation of data, or in the writing of the manuscript.

Availability of data and materials
The Twitter data that support the findings of this study were collected using the Twitter academic API. Restrictions apply to the availability of these data but parts may be shared upon reasonable request. Offline data that support the findings of this study are openly available in the repositories listed in the method section. Contact J.U. for data access.

Declarations

Ethics approval and consent to participate
Not applicable.

Consent for publication
Not applicable.

Competing interests
The authors declare no competing interests.

Received: 17 February 2023 Accepted: 15 June 2023
Published online: 27 July 2023

References

1. Van Bavel JJ, Baicker K, Boggio PS, Capraro V, Cichocka A, Cikara M, et al. Using social and behavioural science to support COVID-19 pandemic response. Nat Hum Behav. 2020;4(5). https://doi.org/10.1038/s41562-020-0884-z.
2. Wu F, Zhao S, Yu B, Chen YM, Wang W, Song ZG, et al. A new coronavirus associated with human respiratory disease in China. Nature. 2020;579(7798). https://doi.org/10.1038/s41586-020-2008-3.
3. Clapham HE, Cook AR. Face masks help control transmission of COVID-19. Lancet Digit Health. 2021;3(3):e136–7.
4. Howard J, Huang A, Li Z, Tufekci Z, Zdimal V, van der Westhuizen HM, et al. An evidence review of face masks against COVID-19. Proc Natl Acad Sci. 2021;118(4):1–12.
5. Ruisch BC, Moore C, Granados Samayoa J, Boggs S, Ladanyi J, Fazio R. Examining the left-right divide through the lens of a global crisis: Ideological differences and their implications for responses to the COVID-19 pandemic. Polit Psychol. 2021;42(5):795–816.
6. Stroebe W, vanDellen MR, Abakoumkin G, Lemay EP Jr, Schiavone WM, Agostini M, et al. Politicization of COVID-19 health-protective behaviors in the United States: Longitudinal and cross-national evidence. PLoS ONE. 2021;16(10):e0256740.
7. Sylvia Chou WY, Gaysynsky A, Cappella JN. Where we go from here: health misinformation on social media. Am J Public Health. 2020;110(S3):S273–5.
8. Jamison AM, Broniatowski DA, Quinn SC. Malicious actors on Twitter: A guide for public health researchers. Am J Public Health. 2019;109(5):688–92. https://doi.org/10.2105/AJPH.2019.304969.
9. Gallotti R, Valle F, Castaldo N, Sacco P, De Domenico M. Assessing the risks of 'infodemics' in response to COVID-19 epidemics. Nat Hum Behav. 2020;4(12). https://doi.org/10.1038/s41562-020-00994-6.
10. Himelein-Wachowiak M, Giorgi S, Devoto A, Rahman M, Ungar L, Schwartz HA, et al. Bots and misinformation spread on social media: Implications for COVID-19. J Med Internet Res. 2021;23(5). https://doi.org/10.2196/26933.
11. Horton R. Offline: Managing the COVID-19 vaccine infodemic. Lancet. 2020;396(10261):1474.
12. De Freitas L, Basdeo D, Wang HI. Public trust, information sources and vaccine willingness related to the COVID-19 pandemic in Trinidad and Tobago: an online cross-sectional survey. Lancet Reg Health-Am. 2021;3:100051.
13. Shao C, Ciampaglia GL, Varol O, Yang KC, Flammini A, Menczer F. The spread of low-credibility content by social bots. Nat Commun. 2018;9(1):4787. https://doi.org/10.1038/s41467-018-06930-7.
14. Tsao SF, Chen H, Tisseverasinghe T, Yang Y, Li L, Butt ZA. What social media told us in the time of COVID-19: a scoping review. Lancet Digit Health. 2021;3(3):e175–94.
15. Uyheng J, Carley KM. Bots and online hate during the COVID-19 pandemic: Case studies in the United States and the Philippines. J Comput Soc Sci. 2020;3(2):445–68. https://doi.org/10.1007/s42001-020-00087-4.
16. Lin YR, Chung WT. The dynamics of Twitter users' gun narratives across major mass shooting events. Humanit Soc Sci Commun. 2020;7(1):1–16.

17. Greijdanus H, de Matos Fernandes CA, Turner-Zwinkels F, Honari A, Roos CA, Rosenbusch H, et al. The psychology of online activism and social movements: Relations between online and offline collective action. Curr Opin Psychol. 2020;35:49–54.
18. Kim DH, Ellison NB. From observation on social media to offline political participation: The social media affordances approach. N Media Soc. 2021;0(0):1461444821998346.
19. Holtz D, Zhao M, Benzell SG, Cao CY, Rahimian MA, Yang J, et al. Interdependence and the cost of uncoordinated responses to COVID-19. Proc Natl Acad Sci. 2020;117(33):19837–43. https://doi.org/10.1073/pnas.2009522117.
20. Health The Lancet Digital. Pandemic versus pandemonium: Fighting on two fronts. Lancet Digit Health. 2020;2(6):e268.
21. Limaye RJ, Sauer M, Ali J, Bernstein J, Wahl B, Barnhill A, et al. Building trust while influencing online COVID-19 content in the social media world. Lancet Digit Health. 2020;2(6):e277–8.
22. Huang B, Carley K. A hierarchical location prediction neural network for Twitter user geolocation. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). Hong Kong, China: Association for Computational Linguistics; 2019. p. 4732–4742. https://doi.org/10.18653/v1/D19-1480.
23. Salomon JA, Reinhart A, Bilinski A, Chua EJ, La Motte-Kerr W, Rönn MM, et al. The US COVID-19 Trends and Impact Survey: Continuous real-time measurement of COVID-19 symptoms, risks, protective behaviors, testing, and vaccination. Proc Natl Acad Sci. 2021;118(51):e2111454118.
24. Beskow DM, Carley KM. Bot conversations are different: Leveraging network metrics for bot detection in Twitter. In: 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), 2018. p. 825–832. ISSN: 2473-991X. https://doi.org/10.1109/ASONAM.2018.8508322.
25. Kumar S, Carley KM. Tree LSTMs with convolution units to predict stance and rumor veracity in social media conversations. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Florence: Association for Computational Linguistics; 2019. p. 5047–5058.
26. Preacher KJ, Zyphur MJ, Zhang Z. A general multilevel SEM framework for assessing multilevel mediation. Psychol Methods. 2010;15(3):209.
27. Ferron J, Rendina-Gobioff G. Interrupted time series design. In: Encyclopedia of Statistics in Behavioral Science. Wiley; 2005. https://doi.org/10.1002/0470013192.bsa312.
28. Bentler PM, Bonett DG. Significance tests and goodness of fit in the analysis of covariance structures. Psychol Bull. 1980;88(3):588.
29. Pennycook G, McPhetres J, Zhang Y, Lu JG, Rand DG. Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention. Psychol Sci. 2020;31(7):770–80.
30. Cotter K, DeCook JR, Kanthawala S. Fact-Checking the Crisis: COVID-19, Infodemics, and the Platformization of Truth. Soc Media+ Soc. 2022;8(1):20563051211069048.
31. Uyheng J, Magelinski T, Villa-Cox R, Sowa C, Carley KM. Interoperable pipelines for social cyber-security: Assessing Twitter information operations during NATO Trident Juncture 2018. Comput Math Organ Theory. 2020;26(4):465–83.

Publisher's Note