Can a Trusted Messenger Change Behavior when Information is Plentiful? Evidence from the First Months of the COVID-19 Pandemic in West Bengal

By ABHIJIT BANERJEE, MARCELLA ALSAN, EMILY BREZA, ARUN G. CHANDRASEKHAR, ABHIJIT CHOWDHURY, ESTHER DUFLO, PAUL GOLDSMITH-PINKHAM, AND BENJAMIN A. OLKEN*

Can information from a credible messenger shift behavior in an information-saturated environment? In a randomized controlled trial involving twenty-eight million individuals in West Bengal, we find that SMS-delivered video messages containing information about COVID-19 symptoms and health-preserving behaviors recorded by a credible messenger increased adherence to targeted and non-targeted preventive behaviors, measured by two objective measures (symptoms reported to a health worker, and phone usage at home), as well as self-reported behaviors. We find large spillovers onto non-targeted recipients. Credible light-touch messaging can play an important role in crisis response, even when similar information is widely available.

* Banerjee: Department of Economics, Massachusetts Institute of Technology and the National Bureau of Economic Research (email: banerjee@mit.edu); Alsan: John F. Kennedy School of Government, Harvard University and the National Bureau of Economic Research (email: marcella_alsan@hks.harvard.edu); Breza: Department of Economics, Harvard University and the National Bureau of Economic Research (email: ebreza@fas.harvard.edu); Chandrasekhar: Department of Economics, Stanford University and the National Bureau of Economic Research (email: arungc@stanford.edu); Chowdhury: John C. Martin Centre for Liver Research and Innovations (email: abhijit@liverfoundation.in); Duflo: Massachusetts Institute of Technology and the National Bureau of Economic Research (email: eduflo@mit.edu); Goldsmith-Pinkham: School of Management, Yale University (email: paul.goldsmith-pinkham@yale.edu); Olken: Department of Economics, Massachusetts Institute of Technology and the National Bureau of Economic Research (email: bolken@mit.edu). This study was approved by Massachusetts Institute of Technology's IRB (IRB #2003000118). We gratefully acknowledge funding from J-PAL/IGI, NSF RAPID #2029880 and the Weiss Foundation. Ben Golub and Alessandra Peter provided helpful comments. We thank J-PAL SA, and particularly Tithee Mukhopadhyay, Shreya Chaturvedi, Vasu Chaudhary, Shobitha Cherian, Arnesh Chowdhury, Anoop Singh Rawat, Advik Shreekumar, Meghna Yadav, Louis-Maël Jean, Nikhil Shankar, and Lukas Leister for exceptional research support.

Policymakers frequently disseminate information designed to influence behavior and improve the well-being of their constituents. Governments worldwide, for example, publicize statistics on the harms of smoking to discourage tobacco use, encourage participation in public programs ranging from health insurance to college financial assistance, and relay information on agricultural best practices to stimulate productive investments (Bindlish and Evenson 1997; MacMonegle et al. 2018; Bird et al. 2021; Goldin, Lurie, and McCubbin 2021).

These policy efforts, however, face an uphill battle. First, policymakers are competing with a deluge of information and disinformation that can hinder their messaging efforts. Second, and relatedly, even if policymakers successfully disseminate messages, they still need recipients to internalize the content of the message. As one indicator of this challenge, average mobile data consumption has increased dramatically in recent years, with India – the setting for our study – at the leading edge, averaging 18.4 gigabytes of data consumption per smartphone per month in 2021 compared to 14.5 gigabytes per month in North America (Statista 2022). During times of crisis, the flow of both accurate and mistaken information generally increases, thus ensuring that engagement with accurate information takes on even greater importance during these periods (Bursztyn et al. 2020; Miller 2020; Zarocostas 2020). A survey we conducted in the Indian state of West Bengal in mid-2020 indicates that such an information deluge occurred during the COVID-19 pandemic: our respondents reported that in the two previous days they heard about the importance of distancing on average 20.2 times, washing hands 16.9 times, and masking 17.2 times (N=408) (Appendix Table A1). Despite this profusion of information, however, compliance with government emergency health policies was incomplete: although our survey took place in a nationwide lockdown, 37 percent of respondents in our study's control group (N = 242) left their village at least once every two days, washing their hands only 68 percent (N = 233) of the time upon returning home.

Can policymakers design messages that motivate behavior change in such information-rich environments? Using a large-scale randomized experiment in which 25 million people in West Bengal were sent information on health-preserving behaviors and detailed symptoms of COVID-19, and 3 million people formed the control group, we study whether a credible, influential messenger can shift behavior in a population inundated with messages. We then investigate three hypotheses based on our messaging strategy. First, in this saturated information environment, can these short messages nudge people to act on information they may already know? Second, is it important that everyone is reached directly, or does network diffusion take care of spreading the message (Sutton et al. 2015; VanderWeele and Christakis 2019)? Finally, does the specific content of the message matter, or does the campaign act as a more general reminder, regardless of the motivation given and the specific behavior targeted?

We study these questions in the context of West Bengal's information campaign during the onset of COVID in May 2020. The Government of West Bengal, working with Reliance Jio, one of the largest telecom operators in India, disseminated SMS messages in May 2020 to all West Bengali subscribers of Reliance Jio. One of the study authors, Abhijit Banerjee, had been covered widely by the West Bengal media since he had won the 2019 Sveriges Riksbank Prize in Economic Sciences in memory of Alfred Nobel just a few months prior, and he had recently been named chair of the West Bengal Government's Global Advisory Board COVID-19 Response Policy. Given this, subscribers in treated Postal Index Number (PIN) codes received short messages advocating preventive health behavior with a link to a video delivered by Banerjee with information on preventing COVID-19, while subscribers in the control PIN codes received an analogous message with a link to a government website that also contained COVID-19 information but without Banerjee's message. To analyze the role of the treatment message content we worked with the Government to vary the video script over eight arms, each emphasizing a combination of a practice, a rationale for action, and a social problem.

This early in the pandemic, data on COVID-19 prevalence was not available. We thus elicited data on symptom reporting and COVID-19 related behavior and knowledge through two telephone survey instruments designed for this study, as well as the user location data made available by Facebook. For our first survey on symptom reporting, we contacted Accredited Social Health Activists (ASHAs) – community health workers who were re-purposed as frontline health workers in West Bengal's COVID-19 response. Having potential COVID-19 patients consult ASHAs as a part of the triage process was a key government objective and all our messages encouraged individuals to contact their ASHA if exhibiting symptoms of fever or cough. We then conducted a second survey on COVID-19 related knowledge and behaviors on the phone, sampling from a list of present and former village council members. In India, village councils are the only grassroots-level of formalized local self-governance at the village or small-town level. Finally, we merged at the PIN code level publicly available Facebook user counts data in treated and control areas on the days following the intervention.

We find that the intervention significantly increased the reporting of symptoms to ASHAs in the first days after the Banerjee messaging. Since these symptoms must have been due to infections contracted before the intervention, they indicate an increase in reporting. Three months after the treatment, we find no long-term effect on the reporting of respiratory systems, but a persistent increase in reporting of fever. Since many common diseases cause fever without respiratory illness while COVID-19 is commonly associated with both respiratory symptoms *and* fever, this is consistent with the hypothesis that villagers continued to report symptoms more often, but that COVID-19 may have decreased. The surveys did not mention the SMS messaging campaign and answers were mainly based on ASHAs' administrative records; hence it is unlikely that demand effects were prominent.

The intervention also exhibited a significant impact on COVID-19 related health behaviors as reported by the village elected leaders. Surveyed individuals in treated PIN codes reported less travel outside of their own village, a higher rate of handwashing upon returning home, and higher mask usage. We also find a decline in the number of COVID-19 related conversations in the treated PINs. Effects on the number of interactions and knowledge about COVID-19 were less pronounced. We find a significant spillover effect of these messages to non-targeted individuals and on protective behaviors that were not specifically mentioned; moreover, the specific content of the messages did not matter.

Finally, in the large-scale administrative data from Facebook, we find significantly higher Facebook user counts in treated areas on the days following the intervention. In the US, this is a proxy for user location (and this is how Facebook describes it). In India, the interpretation of this data is not as obvious, since users often switch off their phones when not in use. The data suggest that the intervention reduced mobility, led to more time online, or both. But regardless of the interpretation, this is strong objective evidence that the intervention was impactful at the population level.

Our work contributes to several literatures. We show that a light-touch behavioral intervention disseminated to over 25 million people successfully nudged individuals to improve both targeted and non-targeted health behaviors, adding to the emerging literature on information campaigns and nudge architecture at-scale (Athey et al. 2022; DellaVigna and Linos 2022). We demonstrate that messaging from credible sources can shift important behaviors even in a highly information-saturated environment, expanding on previous work on celebrity messaging (Alatas et al. 2019; Abu-Akel, Spitz, and West 2021) and sender identity (Alsan and Eichmeyer 2022). Our study documenting strong spillovers onto non-targeted individuals adds to research on seeding as a dissemination strategy (Banerjee et al. 2013; Banerjee et al. 2019; Beaman et al. 2021; Banerjee et al. 2022). The lessons learned from this experiment may also aid future messaging efforts by policymakers, such as by informing strategies to combat misinformation about vaccines, COVID-19 variants, or emerging infectious diseases in the developing world.

The paper proceeds as follows. Section I describes the intervention and data collection, Section II details the empirical approach and the main results, and Section III concludes.

I. Experimental Design and Data Collection

A. Experimental design

Informing citizens about COVID-19, the importance of symptom reporting, and desirable health behaviors was a key objective of the West Bengal government at the start of the pandemic. West Bengal has a population of 91.3 million, with 62.2 million living in rural areas. The literacy rate was 74 percent in 2018 and as of 2019, 57.7 million people had access to mobile devices (Department of Telecommunications 2019; World Bank 2022). To contribute to this effort, working with the West Bengal government, we designed eight 2.5-minute-long video clips delivered by Abhijit Banerjee. Since receiving the 2019 Sveriges Riksbank Prize in Economic Sciences in memory of Alfred Nobel, Banerjee has been covered widely by the West Bengal media and was named chair of the West Bengal Government's Global Advisory Board COVID-19 Response Policy. Thus, he was arguably considered a credible and well-known voice on COVID-19.

We partnered with Reliance Jio, one of the largest telecom operators in India. They agreed to send SMS messages in Bangla to all their subscribers in West Bengal, across the 1,264 PIN (Postal Index Number) codes that divide the state. For our analysis, we considered all PIN codes which correspond to either an urban or rural area, excluding 50 PIN codes that could not be classified.[1] Our resulting sample consists of 28 million subscribers. We assigned treatment at the PIN code level using stratified randomization at the district x urban PIN code level. In total, 1,085 PIN codes were assigned to one of the eight treatments and 129 PIN codes were assigned to the control, where users were sent messages directing them to a government web site. For ethical reasons, no PIN codes were excluded from receiving any health information.

All treatment SMSs stated: “Nobel laureate Abhijit Vinayak Banerjee's appeal on the subject of Coronavirus” and contained a unique YouTube link to that treatment's message. The control group had a near-identical SMS (“An appeal on the subject of Coronavirus”) but linked to a government website with COVID-19 information, similar to typical government messaging. Scripts were crafted under the guidance of a physician co-author: Abhijit Chowdhury, also a member of West Bengal Government's Global Advisory Board — COVID-19 Response Policy (see Appendix Methods).

Messages were assigned using a 2 x 2 x 2 cross-randomized design, with each emphasizing (i) a practice (social distancing or handwashing), (ii) a rationale for action (cost to self or everyone), and (iii) a social problem (ostracizing COVID-19 victims is unacceptable and should be reported to authorities or no mention of the issue). Messages were sent on May 4 and 5, 2020, during India's nationwide lockdown (Appendix Figure A1 and Appendix Materials). View counts of YouTube videos are shown in Appendix Figure A2. All treatments had similar average viewing rates (N = 327,860) of 1.1 percent, consistent with the literature on click-through-rates that finds rates of 0.3 percent-2.6 percent to be standard (Richardson, Dominowska, and Ragno 2007; Kanich et al. 2009).

[1] The Appendix Methods section describes our geocoding procedure to identify urban PIN codes.

B. Data sources

To evaluate the effect of our intervention, we conducted two phone-based surveys, detailed below, and merged in the Facebook administrative user count data after it became publicly available.

Survey of Health Workers – Symptom Reporting.—Community health workers (ASHAs) are frontline health workers in West Bengal's COVID-19 response. During the pandemic, ASHAs were re-purposed from their focus on maternal and child welfare to connect communities with the government's health initiatives and the formal healthcare system (Ministry of Health and Family Welfare 2020). Engagement with the healthcare system is low in rural India (Das et al. 2016; Das et al. 2020). For our survey of frontline health workers, we built a database of ASHAs' contact information using a publicly available directory of 44,312 current and former village council members (*Gram Panchayat*) from 3,340 village councils across 19 districts in West Bengal. In India, village councils are the only grassroots-level of formalized local self-governance at the village or small-town level and are comprised of elected representatives from each administrative unit (known as wards) in the cluster of villages. They are charged with providing basic amenities in the cluster of villages, implementing government schemes, executing other functions related to planning, and maintaining welfare during their five-year terms. We contacted 5,253 former and current local village council members randomly chosen from this database and requested the contact details of ASHAs in their village. The sampling was performed randomly across all the districts.

Our final sample consists of 759 ASHAs, who were interviewed on a rolling basis between April 21st and September 15th and who entered the rolling survey no later than May 10th. On contacting a health worker for the first time, we elicited information on demographics, the area and number of households they oversaw, the number of households they visited, the number of recent migrants that had returned to their area, and the number of cases of fever, cough, and shortness of breath in the previous three days. We also asked if the ASHAs were experiencing any symptoms themselves. We then obtained permission to conduct follow-up surveys with them at three-day intervals. In a random order, we followed-up with ASHAs every 3-5 days, subject to consent and availability, contacting each an average of 17.2 times through September 15, 2020, in five phases. The first phase is the window within 5 days of broadcasting (i.e., May 7-10th) and forms the basis for the main analysis. To examine effects over time, we include responses from post intervention – Phases 1 to 5 (see Appendix Figure A1).[2]

Survey of Village Council Members – Behavioral Outcomes.—Second, we conducted a survey of present and former village council members. Our sampling frame consisted of the same publicly available directory of 44,312 current and former village council members also used for the first survey. The directory included village council members from 2017, and 19.6 percent of our respondents were active village council members. Postal code information was extracted where possible for each council and mapped onto the list of PIN codes that received treatment and control SMS messages. We were able to map 11,614 council members onto our PIN code database. This list of respondents was separate from the 5,253 members that were sampled to get ASHA contact information. Following attempts to contact these councilors by enumerators, we successfully contacted 2,440 for a 26 percent response rate. Our final sample consists of 1,883 unique village council members who completed the survey from May 8-19, 2020 and had

[2] Additional details on our field procedures for the health worker survey is in the Appendix Methods section.

key variables including Jio status completed.[3] This approach, while focused on a subsample of the village community, enabled us to rapidly deploy the survey.

C. Data collection

Survey of Health Workers – Symptom Reporting.—In our frontline health worker survey, our primary outcomes include episodes of fever, cough, and shortness of breath (the latter two combined to form respiratory symptoms) reported to the worker over the previous three days. We also constructed variables for the total number of symptoms reported and an indicator for any symptom reported.[4]

Survey of Village Council Members – Behavioral Outcomes.— We collected several behavioral outcomes for our survey of current and former village council members. We asked respondents if they traveled outside their own village in the last two days and how many interactions (within two arms' distance of themselves, excluding members of their household) they had with other people in their village. Those who visited other villages, towns, or cities were also asked about the number of people they came within two arms' distance of in each of these locations. We asked respondents if they had received information or advice on COVID-19 from anyone in their village, if they had given information or advice on COVID-19 to anyone in their village (face-to-face, via phone, or via chat client), and the number of individuals for each. To measure levels of handwashing, respondents were asked to think about the behavior of a typical

[3] Those who provided Jio status are not statistically different from those who did not for a series of demographic variables (results available upon request).
[4] 57.5 percent of all surveyed individuals were Reliance Jio users (Jio users) and the remainder were not messaged (non-Jio users). For additional information on respondent characteristics, including balance and attrition tables, see Appendix Tables A2-A11.

person in their village, to limit any social desirability that may arise as a result of asking them about their own behavior.[5]

We elicited various measures of mask-wearing and attitudes about masks from respondents. We asked: “out of 100 people in their village, how many are wearing masks,” whether they themselves wear a mask or anything else such as a handkerchief to cover their face when they go out, whether they agree or disagree with the statement “if you wear a mask, you can meet and interact with people as you like," and whether they agree or disagree with the statement “if I wear a mask and go out in my location, I will not feel judged or people will not look at me differently." The last two items were asked to only a subset of respondents (randomized to appear in the digital survey 25 percent of the time) and answers could be given on a five-point Likert scale, ranging from “Agree Strongly" to “Disagree Strongly," with a neutral midpoint.

Lastly, we measured the respondents’ knowledge of symptoms and precautions, eliciting beliefs both about correct and incorrect symptoms and precautions (more details are provided in Appendix Methods). We then created a “knowledge index" for each respondent, using the following formula:

> Knowledge index = (total number of correct symptoms - total number of incorrect symptoms) + (total number of correct precautions – total number of incorrect precautions).

Administrative Population Count Data from Facebook “Data for Good” .— In addition to the survey data, we leverage data from Facebook’s “Data for Good” platform, which we merged with the data set by creating a crosswalk between the “Bing tile” at which the data is provided (which covers an area of approximately 10km x 10km at the equator), and our treatment/control PIN codes (Meta 2020).

[5] Note that at the time the survey was fielding, handwashing was a key component of the Covid-19 prevention strategy.

A key advantage of these data is that they are large-scale, capture population movements at a relatively fine scale, and are free of demand effects. A potential limitation is the representativity. About 25 percent of individuals in West Bengal are Facebook users.

Specifically, we use Facebook's population data, which tracks Facebook users who have given Facebook access to their device's location data. The population counts measure the average number of users in each geographical region, during each of three eight-hour blocks over the course of the day. For an individual to be counted, their phone needs to be switched on and transmit location information to Facebook. Thus, the population count gives a combined measure of a) the number of Facebook users located in a geographical area and b) the intensity of phone use. In the US context, the data is typically used to proxy for location, assuming the phones remain used. However, in India, users often switch off their phones when not in use, especially at night, and the data may therefore capture a combination of phone use and location choice, both of which could have been affected by the intervention.

II. Empirical Framework and Results

A. Reporting Symptoms to Health Workers

We analyze reports to ASHAs within five days of the intervention – the typical incubation period for COVID-19 (Chan et al. 2020; Guan et al. 2020; Lauer et al. 2020; Li et al. 2020). Our primary estimating equation for the outcome reporting of symptoms to health workers is as follows:

$$y_j \;=\; \beta \cdot Treatment_j + \delta' X_j + \epsilon_j \tag{1}$$

where y_j is, for ASHA worker j, the number of reports of fever or respiratory symptoms. Throughout the analysis, standard errors are clustered at the PIN code level.

Reports include all ways through which an ASHA might find out about a case: through regular home visits, over the phone, or via the patient or their household members approaching the ASHA in person. X_j is a vector of controls including smartphone access, the number of households supervised by the ASHA, and fixed effects including district, survey date, and total rounds of surveys done with the ASHA within a five-day horizon of the intervention. This specification therefore identifies a *reporting* effect uncontaminated by any impact on disease transmission.

These intent-to-treat estimates reflect the aggregate reduced-form impact of the seeding to Jio users (which may have generated multiple conversations as well as social signals), both on Jio and non-Jio users. Large spillovers from seeds to the rest of the population have been seen in other settings, such as in the diffusion of microfinance participation or agricultural technology (Banerjee et al. 2013; Beaman et al. 2021).

Table 1 shows reporting of fevers increases by 81.4 percent relative to the control mean (0.228, p = 0.004, q = 0.012, N = 675) and reporting of the number of respiratory symptoms increases by 90.2 percent (0.176, p = 0.045, q = 0.089, N = 675), with similar increases for total symptoms and smaller ones for any symptoms. Appendix Tables A2-A3 present results using alternative time windows of two and three days respectively. The estimated effects are generally higher using the shorter windows, but in the same range as the results reported here.

To evaluate persistence, we extend our sample through September 2020, allowing the treatment effect coefficient to vary by survey phase and controlling for survey rounds done with the ASHA. We estimate treatment effects on reporting over time using equation (2) for ASHA worker j, survey phase k and survey date t:

$$y_{jkt} = \beta_k \cdot Treatment_j + \delta_k{}' X_{jkt} + \epsilon_{jkt} \tag{2}$$

where y_{jkt} is the number of reports of fever or respiratory symptoms. As in equation (1), we consider reports from all modes of communication which might inform an ASHA worker about a case. X_{jkt} is a vector of fixed effects including district, survey date, smartphone access, total number of survey rounds done with each individual ASHA worker, and a control for the number of households supervised by the ASHA in each survey phase k.

We find significant increases with similar magnitudes for fevers (e.g., increases of 44.7 percent, 0.174, p = 0.026, q = 0.052 in June and 43.5 percent, 0.481, p = 0.018, q = 0.052 in August), though we cannot reject the null for respiratory symptoms (Figure 1 and Appendix Table A4). This is consistent with a net increase in reporting combined with a change in the composition of symptoms reported to ASHAs. While this is speculative, this could even be due to a decline in COVID-19, combined with an increase in reporting, since fever is a symptom associated with many diseases other than COVID-19 (Bush 2020). Appendix Table A5 shows that results are very similar for ASHAs who never attrited.

B. Health Behaviors and Social Interactions

To estimate treatment effects on health behaviors and social interactions, we turn to the village council member survey, and estimate the following for respondent *i* in PIN code *p* and on survey date *t*:

$$y_{ipt} \;=\; \beta \cdot Treatment_p + \delta' X_{ipt} + \epsilon_{ipt} \tag{3}$$

where y_{ipt} is one of the following outcomes: traveling outside the village, interactions within two arms' length, percent of times washing hands upon returning home, mask usage when leaving home, conversations about COVID-19, and knowledge about COVID-19. The variable $Treatment_p$ indicates whether the PIN code p was assigned to any of the eight treatment messages, and X_{ipt} is a vector of controls including a district fixed effect, a survey date fixed effect, a dummy for individual *i* being a Reliance Jio user, and controls for age, gender, and smartphone access.

Table 2 presents treatment effects on health behaviors and interactions for these specifications. To assess the extent of spillover, we present treatment effects for the full, Jio-only, and non-Jio-only samples and adjust for false discovery rates. Being in a treated PIN code decreased travel outside one's village by 20.0 percent (7.4 pp, p = 0.026, q = 0.089, N = 1,883). Treatment had no detectable effect on reductions in close interactions in the pooled sample, but there is a significant effect on non-Jio users (31.7 percent decline in number of interactions within two arms' length, 3.869, p = 0.063, q = 0.149, N = 799). Treatment increased the rate of handwashing upon returning home by 7.0 percent (4.7 pp, p = 0.044, q = 0.089, N = 1,821) relative to the control mean of 67.5 percent. The effect is indistinguishable from zero for Jio users (1.7 percent, 1.2 pp, p = 0.694, q = 0.783, N = 1,046) but large for non-Jio users (13.8 percent, 8.8 pp, p = 0.015, q = 0.089, N = 775). While masking was not explicitly mentioned in the message, mask usage increased by 1.9 percent (1.9 pp, p = 0.042, q = 0.089, N = 1,883, we used the question about the respondents' own mask-wearing described above). The number of conversations regarding COVID-19 declined slightly by 18 percent (2.099, p = 0.108, q = 0.162, N = 1881) and more so in the non-Jio group, where we also noted fewer interactions overall. The decline in conversations is consistent with reduced interactions, but of course requires that there is no substitution towards conversations through social media, SMS, or calls. Yet knowledge about COVID-19 improved marginally (4.5 percent) in the Jio treated group (0.227, $p = 0.106$, $q = 0.311$, $N = 1{,}082$), but not in the non-Jio group (the knowledge index described above was used as a measure of knowledge).

We next turn to the Facebook population count data. Figure 2 presents the impact of the intervention on the Facebook-estimated population count within a given geographic area. Each dot in the figure shows day-by-day estimates of the effect of treatment on Facebook population counts. In the days prior to the intervention, differences between treatment and control are close to zero and statistically insignificant.[6] The point estimates of Facebook population counts, however, steadily rise following the second day after the intervention, attaining statistical significance on the sixth day, and remaining significantly elevated for about a week before receding at the two-week mark. Overall, we find statistically significant and positive effects ($p = 0.074$) on cell phone populations in the treated areas in the two weeks after the intervention (Appendix Table A6). As population counts were similar before the intervention, our findings suggest that more people in treatment areas were localized with their phones on in their home location, since it is not plausible that people in control areas would have been attracted to the treatment locations by an intervention they were not exposed to. In our setting, this does not necessarily indicate that individuals were more likely to stay home: it could reflect people were more likely to use their phones while at home. The results from our Facebook analysis are consistent with the intervention reducing mobility, leading to more time online, or a combination of both. However, a reduction in mobility is also supported by our survey data where people reported being less likely to travel outside the village. Regardless of the specific interpretation, the

[6] Facebook population counts are not significantly different in the pre-period ($p = 0.902$).

Facebook findings provide clear evidence from administrative data that the intervention affected behavior.

C. Effects by Content

To estimate effects by content, we fit a separate regression for each of the three topics of interest (behavior, motivation, and ostracism), including indicator variables for the two message variants within a topic and omitting an indicator for the control group. We adjust for the same covariates as in equation (3) and use the specification

$$(4) \qquad y_{ipt} = \alpha + \beta_{m1} \cdot V_{pm1} + \beta_{m2} \cdot V_{pm2} + \delta' X_{ipt} + \epsilon_{ipt}$$

Here m indexes each of the three regressions conducted, one for each of behavior, motivation, and ostracism. V_{pm1} and V_{pm2} are indicators for the message variants (e.g., for the behavior topic, these are indicators for social distancing and hygiene, respectively), and β_{m1} and β_{m2} are the respective effects.

Figure 3 reports the estimates from this analysis and Appendix Table A7 presents *F*-tests across all outcomes. Both exhibits demonstrate there is no major difference in effectiveness across treatment arms. Appendix Table A7 also shows that the null of equality between pairs of message content across all outcomes cannot be rejected for any of the combinations in the pooled sample. Appendix Tables A8-A9 further disaggregate the Jio and non-Jio samples (equation 4), and Appendix Table A10 reports the effects scaled by the standard deviation of the outcome for the control group. Impacts are indistinguishable in all cases. Thus, we fail to reject the hypothesis that the specific message content does not make a difference to its effectiveness, although this may also reflect noisily estimated effects. Point estimates for effects ranged between -0.17 standard deviations to 0.16 standard deviations across all treatment arms. In Appendix Table A11, we report the maximal differences between effect sizes that we can reject at the 10 percent level between each arm. Across all arms, we can reject differences that are 0.16 standard deviations at the 10 percent level.

III. Conclusion

In a crowded information landscape, a message from a well-known individual nudged individuals to increase symptom reporting to the healthcare system, increased phone usage (of Facebook users) at home, and improved COVID-19 related health behaviors, with spillovers onto non-targeted individuals and non-covered content, relative to similar content from the government but without the personal appeal. Increased engagement with ASHAs continued for months post-treatment (reassuringly, treated ASHAs were not more likely to experience symptoms over this period - see Appendix Table A12).

It is unlikely that these measured changes were mostly driven by demand effects. While behavioral survey outcomes could reflect social desirability, the surveys did not mention the intervention. Meanwhile, the ASHA survey relied on their administrative records. And most strikingly, the large-scale administrative dataset of Facebook users' locations is consistent with the survey results indicating reduced travel outside their own village. Facebook population counts significantly increased in treated areas.

The evidence suggests that despite the low click-through rate, a process of seeding and relying on diffusion to reach a much larger population can be effective. Diffusion occurred even though interactions were discouraged and conversations about COVID-19 declined. However, it is possible that the content of conversations changed.

Our results have several policy implications. First, in an information-rich environment, messaging by a credible, well-known figure is valuable, perhaps especially when events are rapidly shifting (Abaluck et al. 2021). We did not explore whether less credible sources would have been effective because our governmental partner insisted that we deploy our most credible figure. Second, even temporary nudges can have durable effects, such as by creating a connection to the formal healthcare system. Third, the presence of spillovers, while noisy, indicates that seeding is an effective dissemination strategy. Fourth, the exact content may be less pivotal in saturated information environments, though our null effects across arms are imprecise and additional experimentation with more power is warranted. This is consistent with work we conducted among disadvantaged populations in the United States (Alsan et al. 2021; Torres et al. 2021), where we also find that the specific content of COVID-19 related messages did not seem to affect their effectiveness. These lessons may aid public health messaging efforts in other contexts, such as strategies to improve vaccine take-up and reduce the emergence and spread of variants of COVID-19 or other emerging diseases in the developing world.

REFERENCES

Abaluck, Jason, Laura H Kwong, Ashley Styczynski, Ashraful Haque, Md. Alamgir Kabir, Ellen Bates-Jefferys, Emily Crawford, et al. 2021. “Normalizing Community Mask-Wearing: A Cluster Randomized Trial in Bangladesh.” National Bureau of Economic Research Working Paper No. 28734. https://doi.org/10.3386/w28734.

Abu-Akel, Ahmad, Andreas Spitz, and Robert West. 2021. “The Effect of Spokesperson Attribution on Public Health Message Sharing during the COVID-19 Pandemic.” *PLOS ONE* 16 (2): e0245100. https://doi.org/10.1371/journal.pone.0245100.

Alatas, Vivi, Arun Chandrasekhar, Markus Mobius, Benjamin Olken, and Cindy Paladines. 2019. “When Celebrities Speak: A Nationwide Twitter Experiment Promoting Vaccination in Indonesia.” National Bureau of Economic Research Working Paper No. 25589. https://doi.org/10.3386/w25589.

Alsan, Marcella, and Sarah Eichmeyer. 2022. “Experimental Evidence on the Effectiveness of Non-Experts for Improving Vaccine Demand.” National Bureau of Economic Research Paper Working Paper No. 28593. https://doi.org/10.3386/w28593.

Alsan, Marcella, Fatima Cody Stanford, Abhijit Banerjee, Emily Breza, Arun G. Chandrasekhar, Sarah Eichmeyer, Paul Goldsmith-Pinkham, et al. 2021. “Comparison of Knowledge and Information-Seeking Behavior After General COVID-19 Public Health Messages and Messages Tailored for Black and Latinx Communities.” *Annals of Internal Medicine* 174 (4): 484–92. https://doi.org/10.7326/M20-6141.

Athey, Susan, Kristen Grabarz, Michael Luca, and Nils Wernerfelt. 2022. “The Effectiveness of Digital Interventions on COVID-19 Attitudes and

Beliefs." National Bureau of Economic Research Paper Working Paper No. 30273. https://doi.org/10.48550/ARXIV.2206.10214.

Banerjee, Abhijit, Marcella Alsan, Emily Breza, Arun G. Chandrasekhar, Abhijit Chowdhury, Esther Duflo, Paul Goldsmith-Pinkham, and Benjamin A. Olken. 2020. "Replication Data for: Can a Trusted Messenger Change Behavior When Information Is Plentiful? Evidence from the First Months of the COVID-19 Pandemic in West Bengal." Harvard Dataverse. https://doi.org/10.7910/DVN/AD79UA.

Banerjee, Abhijit, Emily Breza, Arun G Chandrasekhar, and Benjamin Golub. 2022. "When Less Is More: Experimental Evidence on Information Delivery During India's Demonetization." *Review of Economic Studies* forthcoming. https://doi.org/10.3386/w24679.

Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. "The Diffusion of Microfinance." *Science* 341 (6144): 1236498. https://doi.org/10.1126/science.1236498.

Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson. 2019. "Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials." *The Review of Economic Studies* 86 (6): 2453–90. https://doi.org/10.1093/restud/rdz008.

Beaman, Lori, Ariel BenYishay, Jeremy Magruder, and Ahmed Mushfiq Mobarak. 2021. "Can Network Theory-Based Targeting Increase Technology Adoption?" *American Economic Review* 111 (6): 1918–43. https://doi.org/10.1257/aer.20200295.

Bindlish, Vishva, and Robert E. Evenson. 1997. "The Impact Of T&V Extension in Africa: The Experience of Kenya and Burkina Faso." *World Bank Research Observer* 12 (2): 183–201. https://doi.org/10.1093/wbro/12.2.183.

Bird, Kelli A., Benjamin L. Castleman, Jeffrey T. Denning, Joshua Goodman, Cait Lamberton, and Kelly Ochs Rosinger. 2021. "Nudging at

Scale: Experimental Evidence from FAFSA Completion Campaigns." *Journal of Economic Behavior & Organization* 183 (March): 105–28. https://doi.org/10.1016/j.jebo.2020.12.022.

Bursztyn, Leonardo, Aakaash Rao, Christopher P Roth, and David H Yanagizawa-Drott. 2020. "Misinformation During a Pandemic." National Bureau of Economic Research Working Paper No. 27417. https://doi.org/10.3386/w27417.

Bush, Larry. 2020. "Biology of Infectious Disease." In *The Merck Manual of Diagnosis and Therapy*. Merck Sharp & Dohme Corp.

Chan, Jasper Fuk-Woo, Shuofeng Yuan, Kin-Hang Kok, Kelvin Kai-Wang To, Hin Chu, Jin Yang, Fanfan Xing, et al. 2020. "A Familial Cluster of Pneumonia Associated with the 2019 Novel Coronavirus Indicating Person-to-Person Transmission: A Study of a Family Cluster." *Lancet (London, England)* 395 (10223): 514–23. https://doi.org/10.1016/S0140-6736(20)30154-9.

Das, Jishnu, Benjamin Daniels, Monisha Ashok, Eun-Young Shim, and Karthik Muralidharan. 2020. "Two Indias: The Structure of Primary Health Care Markets in Rural Indian Villages with Implications for Policy." *Social Science & Medicine* 301 (May): 112799. https://doi.org/10.1016/j.socscimed.2020.112799.

Das, Jishnu, Alaka Holla, Aakash Mohpal, and Karthik Muralidharan. 2016. "Quality and Accountability in Health Care Delivery: Audit-Study Evidence from Primary Care in India." *American Economic Review* 106 (12): 3765–99. https://doi.org/10.1257/aer.20151138.

DellaVigna, Stefano, and Elizabeth Linos. 2022. "RCTs to Scale: Comprehensive Evidence from Two Nudge Units." *Econometrica* 90 (1): 81–116. https://doi.org/10.3982/ECTA18709.

Department of Telecommunications, India. 2019. “Telecom Statistics India 2019.” https://dot.gov.in/sites/default/files/Telecom%20Statistics%20India-2019.pdf?download=1.

Director of Census Operations, West Bengal. 2012. “Census of India 2011 - Administrative Atlas - West Bengal.” Administrative Atlas. Indian Administrative Service. https://censusindia.gov.in/nada/index.php/catalog/52.

Goldin, Jacob, Ithai Z Lurie, and Janet McCubbin. 2021. “Health Insurance and Mortality: Experimental Evidence from Taxpayer Outreach.” *The Quarterly Journal of Economics* 136 (1): 1–49. https://doi.org/10.1093/qje/qjaa029.

Guan, Wei-jie, Zheng-yi Ni, Yu Hu, Wen-hua Liang, Chun-quan Ou, Jian-xing He, Lei Liu, et al. 2020. “Clinical Characteristics of Coronavirus Disease 2019 in China.” *New England Journal of Medicine* 382 (18): 1708–20. https://doi.org/10.1056/NEJMoa2002032.

Kanich, Chris, Christian Kreibich, Kirill Levchenko, Brandon Enright, Geoffrey M. Voelker, Vern Paxson, and Stefan Savage. 2009. “Spamalytics: An Empirical Analysis of Spam Marketing Conversion.” *Commun. ACM* 52 (9): 99–107. https://doi.org/10.1145/1562164.1562190.

Lauer, Stephen A., Kyra H. Grantz, Qifang Bi, Forrest K. Jones, Qulu Zheng, Hannah R. Meredith, Andrew S. Azman, Nicholas G. Reich, and Justin Lessler. 2020. “The Incubation Period of Coronavirus Disease 2019 (COVID-19) From Publicly Reported Confirmed Cases: Estimation and Application.” *Annals of Internal Medicine* 172 (9): 577–82. https://doi.org/10.7326/M20-0504.

Li, Qun, Xuhua Guan, Peng Wu, Xiaoye Wang, Lei Zhou, Yeqing Tong, Ruiqi Ren, et al. 2020. “Early Transmission Dynamics in Wuhan, China, of Novel Coronavirus–Infected Pneumonia.” *New England Journal of Medicine* 382 (13): 1199–1207. https://doi.org/10.1056/NEJMoa2001316.

MacMonegle, Anna J., James Nonnemaker, Jennifer C. Duke, Matthew C. Farrelly, Xiaoquan Zhao, Janine C. Delahanty, Alexandria A. Smith, Pamela Rao, and Jane A. Allen. 2018. "Cost-Effectiveness Analysis of The Real Cost Campaign's Effect on Smoking Prevention." *American Journal of Preventive Medicine* 55 (3): 319–25. https://doi.org/10.1016/j.amepre.2018.05.006.

Meta. 2020. "Movement Maps." Meta. https://dataforgood.facebook.com/dfg/tools/movement-maps (accessed August 9, 2022).

Miller, Greg. 2020. "Researchers Are Tracking Another Pandemic, Too—of Coronavirus Misinformation." *Science,* March 24, 2020. https://www.science.org/content/article/researchers-are-tracking-another-epidemic-too-misinformation.

Ministry of Health and Family Welfare, India. 2020. "Micro Plan for Containing Local Transmission of Coronavirus Disease (COVID-19)." https://www.mohfw.gov.in/pdf/ModelMicroplanforcontainmentoflocaltransmissionofCOVID19.pdf.

Richardson, Matthew, Ewa Dominowska, and Robert Ragno. 2007. "Predicting Clicks: Estimating the Click-through Rate for New Ads." In *Proceedings of the 16th International Conference on World Wide Web*, 521–30. WWW '07. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/1242572.1242643.

Statista. 2022. "Mobile Monthly Data Usage per Smartphone 2021." Statista. February 2022. https://www.statista.com/statistics/489169/canada-united-states-average-data-usage-user-per-month/.

Sutton, Jeannette, C. Ben Gibson, Nolan Edward Phillips, Emma S. Spiro, Cedar League, Britta Johnson, Sean M. Fitzhugh, and Carter T. Butts. 2015. "A Cross-Hazard Analysis of Terse Message Retransmission on

Twitter." *Proceedings of the National Academy of Sciences* 112 (48): 14793–98. https://doi.org/10.1073/pnas.1508916112.

Torres, Carlos, Lucy Ogbu-Nwobodo, Marcella Alsan, Fatima Cody Stanford, Abhijit Banerjee, Emily Breza, Arun G. Chandrasekhar, et al. 2021. "Effect of Physician-Delivered COVID-19 Public Health Messages and Messages Acknowledging Racial Inequity on Black and White Adults' Knowledge, Beliefs, and Practices Related to COVID-19: A Randomized Clinical Trial." *JAMA Network Open* 4 (7): e2117115–e2117115. https://doi.org/10.1001/jamanetworkopen.2021.17115.

VanderWeele, Tyler J., and Nicholas A. Christakis. 2019. "Network Multipliers and Public Health." *International Journal of Epidemiology* 48 (4): 1032–37. https://doi.org/10.1093/ije/dyz010.

World Bank. 2022. "Literacy Rate, Adult Total (% of People Ages 15 and above) - India." The World Bank. June 2022. https://data.worldbank.org/indicator/SE.ADT.LITR.ZS?view=chart.

Zarocostas, John. 2020. "How to Fight an Infodemic." *The Lancet* 395 (10225): 676. https://doi.org/10.1016/S0140-6736(20)30461-X.

Tables and Figures

Table 1. Effect of Intervention on Reports Received by ASHAs Within 5 Days of Message

VARIABLES	(1) Number of Fever Cases	(2) Number of Respiratory Cases	(3) Any Symptoms Reported	(4) Total Number of Cases
TREATMENT	0.228	0.176	0.080	0.404
	(0.079)	(0.087)	(0.051)	(0.138)
	[0.004]	[0.045]	[0.117]	[0.004]
	{0.012}	{0.089}	{0.117}	{0.012}
Observations	675	675	675	675
District FE	✓	✓	✓	✓
Total rounds FE	✓	✓	✓	✓
Smartphone FE	✓	✓	✓	✓
Date FE	✓	✓	✓	✓
Control Mean	0.280	0.195	0.293	0.476

Notes: All columns include cases reported within a window of five days after treatment (May 7 - 10). We control for district, total rounds, smartphone, and date fixed effects as well as the total number of households each ASHA supervises. *Any Symptoms* is binary for whether either fever or respiratory cases were reported. *Total Number* is the sum of fever and respiratory cases. Standard errors are clustered at the PIN code level and reported in parentheses, *p*-values are reported in brackets and *q*-values in curly brackets.

Table 2. Effect of Intervention on Health Behaviors -- Phone Survey Respondents

Variables	(1) Did you travel outside your village?	(2) Number of interactions with people within 2 arms length	(3) Estimated % time washing hands upon returning home	(4) Did you use a mask?	(5) Number of conversations in person /online /mobile about COVID-19	(6) COVID-19 Knowledge Index
Panel A: Pooled						
Treatment	-0.074	-1.473	0.047	0.019	-2.099	0.097
	(0.033)	(1.164)	(0.023)	(0.009)	(1.303)	(0.123)
	[0.026]	[0.206]	[0.044]	[0.042]	[0.108]	[0.435]
	{0.089}	{0.248}	{0.089}	{0.089}	{0.162}	{0.435}
Observations	1,883	1,875	1,821	1,883	1,881	1,883
District FE	✓	✓	✓	✓	✓	✓
Survey Day FE	✓	✓	✓	✓	✓	✓
Jio FE	✓	✓	✓	✓	✓	✓
Control Mean	0.370	11.052	0.675	0.978	11.678	5.135
Panel B: Jio						
Treatment	-0.061	0.421	0.012	0.023	-0.539	0.227
	(0.043)	(1.527)	(0.030)	(0.014)	(1.604)	(0.140)
	[0.155]	[0.783]	[0.694]	[0.088]	[0.737]	[0.106]
	{0.311}	{0.783}	{0.783}	{0.311}	{0.783}	{0.311}
Observations	1,082	1,076	1,046	1,082	1,082	1,082
District FE	✓	✓	✓	✓	✓	✓
Survey Day FE	✓	✓	✓	✓	✓	✓
Control Mean	0.376	10.096	0.708	0.976	11.368	5.088
Panel C: Non-Jio						
Treatment	-0.094	-3.869	0.088	0.011	-3.794	-0.077
	(0.057)	(2.073)	(0.036)	(0.012)	(2.121)	(0.199)
	[0.100]	[0.063]	[0.015]	[0.349]	[0.074]	[0.699]
	{0.150}	{0.149}	{0.089}	{0.419}	{0.149}	{0.699}
Observations	801	799	775	801	799	801
District FE	✓	✓	✓	✓	✓	✓
Survey Day FE	✓	✓	✓	✓	✓	✓
Control Mean	0.362	12.190	0.636	0.981	12.048	5.190
Treat(Jio) = Treat(Non-Jio)	[0.645]	[0.112]	[0.100]	[0.492]	[0.221]	[0.174]
	{0.645}	{0.332}	{0.332}	{0.591}	{0.332}	{0.332}

Notes: All regressions control for district and survey date fixed effects. Panel A also includes a *Jio* access fixed effect. Respondent level controls also include age, gender, and smartphone access. The last two rows present *p*- and *q*-values for a test of equality between treatment effects in the Jio and non-Jio samples. Standard errors are clustered at the PIN code level and reported in parentheses. *p*-values are reported in square brackets; *q*-values are reported in curly brackets.

Figure 1. Effects of Intervention on Number of Reports Received by ASHAs Over Time

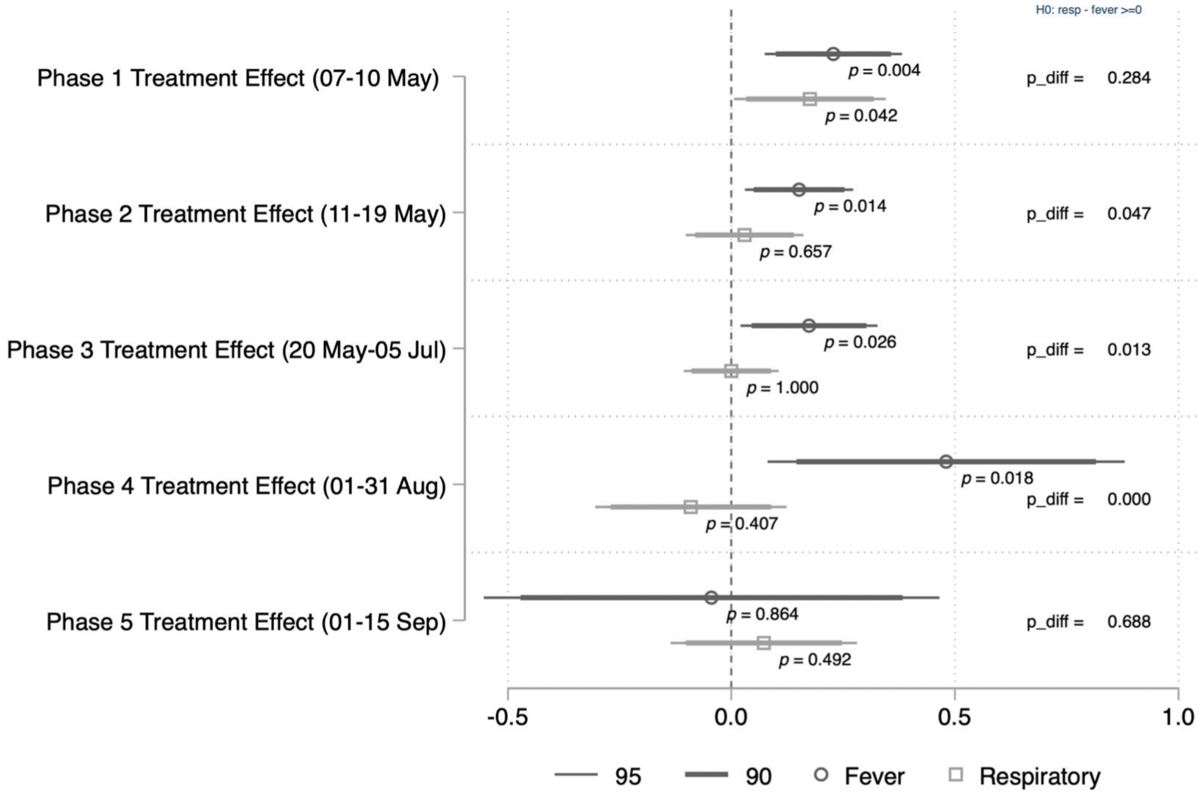

Notes: This figure plots treatment effect over time. District, date, smartphone, and survey round fixed effects are included as is a control for the number of households the ASHA supervises. Standard errors are clustered at the PIN code level. Thick and thin whiskers represent 90% and 95% confidence intervals, respectively. *p_diff* are p-values for the one-sided test that the treatment effect for Respiratory net of the one for Fever is strictly negative, as an alternative to the null that the treatment effect for Respiratory - treatment effect for Fever ≥ 0. See Appendix Table A4 for estimates and standard errors.

Figure 2. Effects on Population Over Time

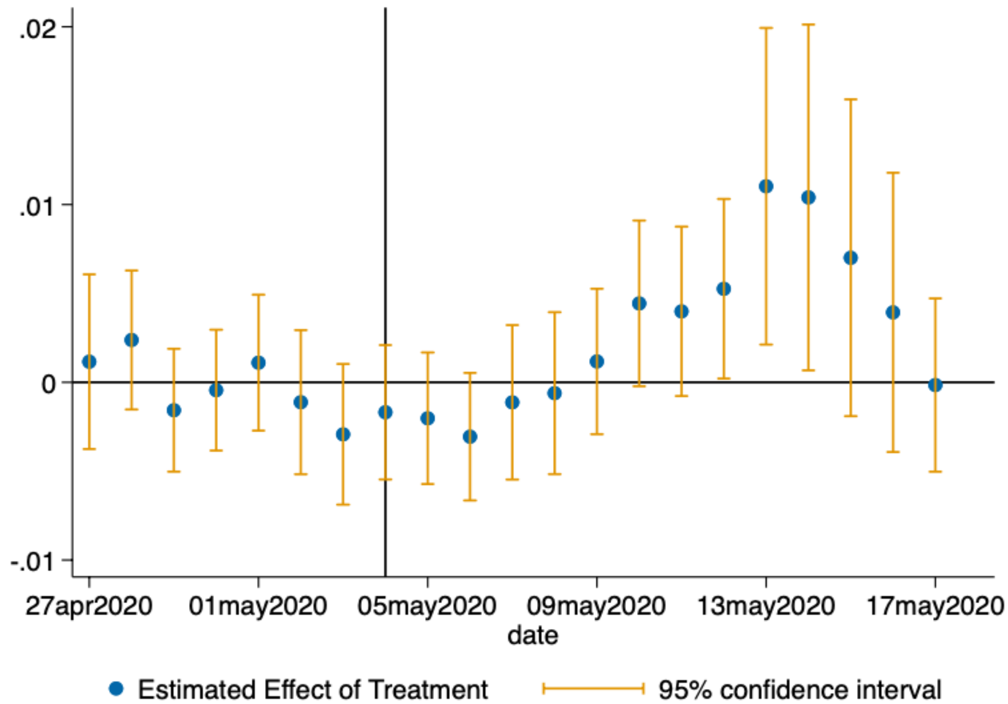

Notes: Estimated treatment effect on log of reported Facebook population within a given Bing tile, estimated day-by-day. District-by-date fixed effects, as well as the following controls interacted with date fixed effects are included: average treatment intensity in Bing tile *j* across 500 potential counterfactual randomized treatment assignments (Borusyak and Hull 2020), the dependent variable averaged over the same day of the week as day *t*, during Facebook's baseline period of the 90 days prior to March 20, 2020, and the lagged average outcome 30 to 15 days prior to the roll-out.

Figure 3. Effects of Intervention by Specific Message Content

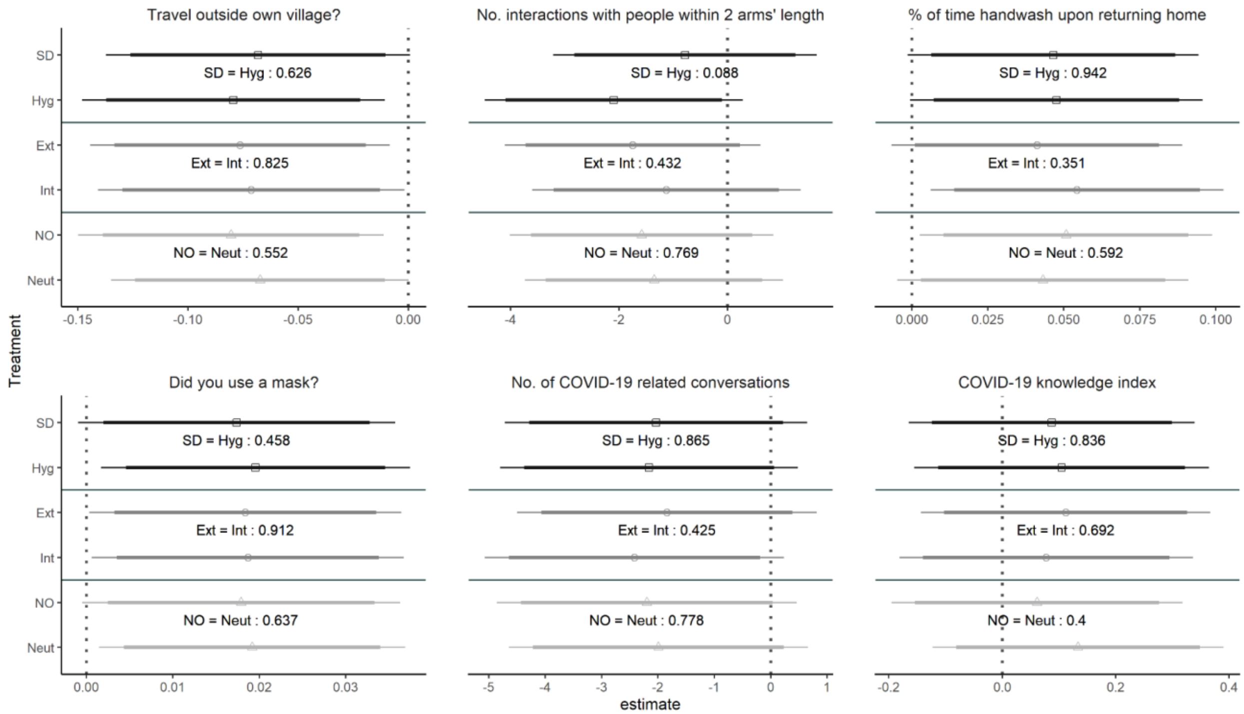

Notes: This figure presents estimated treatment effects by message content. Each panel depicts the treatment effect of a separate outcome. All regressions include controls for district and survey date, as well as for age, gender, smartphone access, and Jio access. The thick and thin whiskers represent 90% and 95% confidence intervals, respectively, and the horizontal lines separate estimates from different regressions. The numbers provided in each panel represent the *p*-value for a test of equality between the noted pairs of coefficients. Standard errors are clustered at the PIN code level. SD - "Social Distancing", Hyg - "Hygiene", Ext - "Externality + Internality", Int - "Internality Only", NO - "No Ostracism", and Neut - "Neutral". Sample size varies from 1,821 observations (estimated handwash) to 1,875 observations (number of interactions within two arms' length). See Appendix Tables A8-A11 for estimates and standard errors for the pooled, Jio, and non-Jio samples, as well as *F*-tests across treatment arms.