Journal of Health Economics 96 (2024) 102883

Contents lists available at ScienceDirect

Journal of Health Economics

journal homepage: www.elsevier.com/locate/jhe

Religious proximity and misinformation: Experimental evidence from a mobile phone-based campaign in India[☆]

Alex Armand[a,*], Britta Augsburg[b], Antonella Bancalari[c], Kalyan Kumar Kameshwara[d]

[a] Nova School of Business and Economics, Universidade Nova de Lisboa, NOVAFRICA and CEPR, Rua da Holanda 1, 2775-405 Carcavelos, Portugal
[b] Institute for Fiscal Studies, 7 Ridgmount St, WC1E 7AE, London, United Kingdom
[c] Institute for Fiscal Studies and IZA, WC1E 7AE, 7 Ridgmount St, London, United Kingdom
[d] Westminster Business School, University of Westminster, 35 Marylebone Rd, NW1 5LS, London, United Kingdom

ARTICLE INFO

JEL classification:
C93
D91
I12
I15
O12

Keywords:
Health campaign
Information
Religion
India
COVID-19
Field experiment

ABSTRACT

We investigate how religion concordance influences the effectiveness of preventive health campaigns. Conducted during the early stages of the COVID-19 pandemic in two major Indian cities marked by Hindu–Muslim tensions, we randomly assigned a representative sample of slum residents to receive either a physician-delivered information campaign promoting health-related preventive practices, or uninformative control messages on their mobile phones. Messages, introduced by a local citizen (the sender), were cross-randomized to start with a greeting signaling either a Hindu or a Muslim identity, manipulating religion concordance between sender and receiver. We found that doctor messages increased compliance with recommended practices and beliefs in their efficacy. Our findings suggest that the campaign's impact is primarily driven by shared religion between sender and receiver, leading to increased message engagement and compliance with recommended practices. Additionally, we observe that religion concordance helps protect against misinformation.

1. Introduction

Interacting with familiar and predictable individuals facilitates communication and enables behavioral change in various spheres, including nation-building processes (Bazzi et al., 2019; Mousa, 2020; Lowe, 2021), financial decision-making (Fisman et al., 2017,

☆ We would like to thank Oriana Bandiera, Michael Callen and Maithreesh Ghatak, as well as seminar participants at the LSE-STICERD Development Economics seminar series and at IFS for helpful comments. This project would not have been possible without the insightful support of our partners in the field, Tatheer Fatima, Yashashvi Singh, Bhartendu Trivedi, and the field team of Morsel Research and Development. We are grateful to Julia Loh for outstanding research assistance. We are also indebted to study participants for their willingness to contribute to the study. This work was supported by the International Growth Centre [grant number IND-20077]; and the LSE Special Grant for Coronavirus Research [grant number 106404]. Kameshwara acknowledges support from the ESRC's SWDTP Placement Scheme. Armand acknowledges funding from the Fundação para a Ciência e a Tecnologia [grant numbers UIDB/00124/2020, UIDP/00124/2020 and Social Sciences DataLab – PINFRA/22209/2016], POR Lisboa and POR Norte [grant numbers Social Sciences DataLab – PINFRA/22209/2016]. Augsburg and Bancalari acknowledge support from the ESRC's Centre for the Microeconomic Analysis of Public Policy at IFS [grant number ES/T014334/1]. Ethics approval was secured from the London School of Economics and Political Science [REC reference number 1132]. The pre-analysis plan was registered on the AEA RCT registry (Armand et al. 2020)

* Corresponding author.
E-mail addresses: alex.armand@novasbe.pt (A. Armand), britta_a@ifs.org.uk (B. Augsburg), antonella.bancalari@ifs.org.uk (A. Bancalari), k.kameshwara@westminster.ac.uk (K.K. Kameshwara).

https://doi.org/10.1016/j.jhealeco.2024.102883
Received 6 December 2023; Received in revised form 27 March 2024; Accepted 17 April 2024
Available online 9 May 2024
0167-6296/

2020), and experimental games (Habyarimana et al., 2007; Bicchieri et al., 2022). The propensity to adapt behavior based on shared characteristics and identities is notable in health-related interactions such as those between patients and doctors (Greenwood et al., 2018; Alsan et al., 2019; Greenwood et al., 2020; Hill et al., 2020). Leveraging race, gender or class concordance has been found crucial when promoting preventive healthcare (Alsan and Wanamaker, 2018; Torres et al., 2021; Alsan et al., 2021; Alsan and Eichmeyer, 2021). Religion remains understudied in this context, despite its significant historical influence, its heightened importance in times of unpredictable events (Bentzen, 2021), and its centrality to public health in low-income settings (Iyer, 2016; Benjamin et al., 2016; Banerjee et al., 2022; Taragin-Zeller et al., 2023). In particular, little is known about the role of shared religious identity in the diffusion of health information and the spread of misinformation about preventive health.[1]

This paper examines the effectiveness of a physician-delivered information campaign that promotes health-related preventive practices. We investigate how introducing religion concordance between the sender and the recipient enhances the campaign's effectiveness. We do so among residents of densely populated informal settlements, often referred to as 'slum dwellers', a largely understudied population (Lilford et al., 2017). We document that promoting preventive behavior can increase compliance with recommended practices and beliefs about their efficacy. Our findings indicate that the campaign's impact is primarily driven by a shared religion between the sender and the receiver. In this case, recipients listen to a greater portion of the message and are more compliant with recommended practices. Furthermore, we find that religion concordance helps to protect against misinformation.

We implement a field experiment in the Indian state of Uttar Pradesh (UP) in the context of a global outbreak of an infectious disease: the COVID-19 pandemic. At the onset of the pandemic, we designed a mobile-phone-based information campaign to raise citizens' awareness about evidence-based practices to mitigate the spread of the virus, and to counteract the sudden rise in misinformation surrounding the pandemic (World Health Organization, 2020).[2] To this purpose, between October 2020 and January 2021, we sent two pre-recorded voice messages to a representative sample of slum residents, in the two major cities of the state. The campaign held particular importance in this context, not only due to the overcrowded living conditions that made physical distancing challenging, but also due to the low-income and marginalized nature of the setting, which limited access to healthcare and adequate hygienic conditions.

Each voice message consists of two components: an introduction by a local citizen, the *sender*, followed by the *content* of the message. Using cross-randomization, we vary both components. First, to obtain exogenous variation in religion concordance between the sender and receiver, we randomly vary the greeting used by the sender at the beginning of the message to signal either a Muslim or Hindu identity. Religion is highly salient in our setting, particularly at the time of the experiment. In India, Hindu–Muslim tensions have been present since the pre-partition era, and are particularly relevant for UP, home to the largest Muslim population in India (Jha, 2013; Mitra and Ray, 2014). In line with religion being salient in the presence of unpredictable events (Sinding Bentzen, 2019; Atkin et al., 2021), the onset of the pandemic saw a sudden increase in these inter-religious tensions: misleading claims about the role of Muslim citizens in the spread of the virus were the primary driver of fake news on social media and spurred further violence (e.g. Yasir, 2020).

Second, to obtain exogenous variation in the content of the message, we randomize whether the receiver is sent messages about preventive practices or uninformative content. In the former, which we label as *doctor* messages, the content is provided by doctors of locally renowned hospitals, provides reminders about evidence-based policy recommendations, and debunks common misconceptions about the virus. The religious identity of doctors is not revealed. In the latter, which we label as *control* messages, the content consists of Bollywood gossip unrelated to the pandemic. Thanks to cross-randomization, both the doctor and control messages are either religion-concordant or religion-discordant.[3]

We gathered information about participants' behavior related to preventive practices, particularly the extent to which respondents wear a face mask when going out, the frequency of hand-washing, and the extents to which they stay in the slum, do not receive visitors from outside the slum, and do not meet anybody from outside the slum. We aggregate these individual reports into an index of compliance with recommended practices. Additionally, we collected data on beliefs over the efficacy of both recommended and non-evidence-based practices, and about participants' response to misinformation about the pandemic, during a baseline and two follow-up surveys. We base our main analysis on intention to treat (ITT) effects, which capture the effect of sending the messages. Using administrative data on the take-up of the interventions, we complement ITT estimates with local average treatment estimates (LATE) of the effect among compliers.

The design of the experiment allows us first to study the overall effect of promoting preventive practices and then to estimate the effect of introducing shared religion between the sender and receiver, a novel set up in the literature. Providing informative content via mobile phones is effective at promoting welfare-improving behavior. Compared with control messages, doctor messages significantly increase compliance with recommended practices and update recipients' beliefs about the efficacy of these practices positively. However, despite being debunked in the message, doctor messages have no significant effect on the degree to which respondents believe that non-evidence-based practices such as relying on vegetarianism or on a stronger immune system can protect from infection, indicating the persistence of these beliefs to new information.

[1] Research on the mechanisms of (mis)information more generally remains limited and predominantly focused on higher-income countries (DellaVigna and Kaplan, 2007; Allcott and Gentzkow, 2017; Lazer et al., 2018; Bursztyn et al., 2023).

[2] In India, the spread of misinformation about COVID-19 was so severe that it compelled the Prime Minister Narendra Modi to address the nation urging everyone to rely only on credible medical advice and demanding social media companies curb misinformation on their platforms (Al Jazeera, 2020; Mahapatra and Plagemann, 2019). Internet penetration rates went from 4% in 2007 to 50% in 2020, raising social media platforms as a primary source of news and as a key means of communication for all political party actors (Statista, 2021; Akbar et al., 2020).

[3] The experimental design also cross-randomized whether the receiver was incentivized with lower or higher monetary incentives to listen to the message. Refer to Section 4.

To assess the added benefit of shared religion, we focus on the sample that was sent the doctor message and we exploit the cross randomization in the religion concordance between the sender and the receiver of the information. First, we find that religion concordance leads participants to listen to a larger portion of the doctor message, an increase of 13.3% compared with religion-discordant messages. Second, the effect of doctor messages on compliance with recommended practices is primarily driven by religion-concordant messages. Third, religion concordance in the doctor messages effectively reduced beliefs over the efficacy of non-evidence-based practices, particularly those with a religious connotation.

The last two results are specific to the combination of informative content provided by the doctor and religion concordance. Studying the differential effects of religion concordance in the control messages, which serves as a placebo test, indicates no effect in any of the outcomes studied. In addition, the effects are specific to misinformation. In fact, none of the interventions influences agreement with non-factual opinions about the spreading of COVID-19, by definition more persistent and harder for information campaigns to influence than pure misinformation (e.g. Walter and Salovich, 2021). Finally, we provide evidence that spillover effects were not present in the interventions, suggesting that mobile-phone campaigns are effective at targeting individuals rather than communities.

To understand the drivers behind these impacts, we first analyze respondents' fact-checking behavior, an important determinant of factual knowledge (Barrera et al., 2020). The findings reveal that doctor messages significantly reduce the likelihood of verifying the truthfulness of information. This reduction is likely because individuals, having heard the messages from doctors, feel more confident in dismissing misinformation. We further use a novel survey instrument to measure whether respondents agree with misinformation shared by other citizens and show that doctor messages reduce agreement with misinformation shared by citizens *outside* the religious group of the respondent (*out-group* citizens), while keeping unchanged their level of agreement with citizens of the same religion (*in-group* citizens). Religion concordance in the doctor messages is effective at detaching in-group norm compliance in the response to misinformation. When the sender and the receiver have the same religion, doctor messages reduce agreement with misinformation shared by in-group citizens by 4.6% compared with religion-discordant messages. This finding aligns with existing research and for high-income countries, which emphasizes that the perceived credibility of information is influenced by the social distance between the communicator and the recipient (Tabellini, 2008; Alsan et al., 2019).

Our results suggest that the information campaign somewhat reduces the effort to verify information's truthfulness while creating a protective layer against misinformation. However, this layer is crucially affected by salience within a group, suggesting a high level of in-group norm compliance in our setting (e.g. Akerlof and Kranton, 2000). However, this compliance can be reduced through a carefully designed information campaign that takes into account social proximity with the objective of leveraging social norms, challenging the assumption that in- and out-groups agree with prevailing norms.

To address concerns related to the treatment group exerting more social desirability bias in the self-reported outcomes, we collect measures of the (Crowne and Marlowe, 1960) social desirability scale at baseline. Although individuals with a strong tendency toward social desirability may show more endorsement for recommended practices or widespread beliefs, we demonstrate that this pattern is not more pronounced in the treatment group tan in the control group. In addition, we show that, at baseline, social desirability does not influence reporting differently depending on the religion, the gender, and the caste of the respondents.

Our findings offer novel insights into the design of information campaigns, an instrument that has been extensively used to communicate risk and best practices for health behavior (Dupas, 2011). We complement available evidence on the effectiveness of communication technology to raise health awareness in the US (Alsan et al., 2020; Breza et al., 2021; Torres et al., 2021), in the Indian state of West Bengal (Banerjee et al., 2020), and in rural India and Bangladesh (Siddique et al., 2022). We further the understanding of these interventions by providing novel evidence on how the effectiveness of information campaigns on preventive behavior is crucially influenced by shared identity. Our design is unique in the literature because it allows identification of the effect of the initial signal of shared religion (i.e. the first word of the message), while keeping the content of the message indistinguishable in terms of religious identity. Previous literature focuses instead on *micro-targeting* (i.e. the shaping of both the sender and the information content to the individual characteristics of the receiver). This approach has been used to influence interactions with patients (Yom-Tov et al., 2018; Alsan and Eichmeyer, 2021).

By linking compliant behavior with beliefs and response to misinformation, we provide novel evidence not only on the drivers of information, but also on the mechanisms of misinformation, whose persistence remains a puzzling result in the literature (Van der Linden et al., 2017; Zhuravskaya et al., 2020). In particular, despite the recognition that understanding how beliefs are affected by information is crucial, few studies explicitly elicit the effect of information on beliefs over practices and on (dis)agreement with misinformation (Kremer et al., 2019).

Finally, highlighting the role of religion also complements available evidence on the role of identity in decision-making. The literature shows how identity affects cooperation, political mobilization trust, and violence (Philpott, 2007; Bhalotra et al., 2014; Lowe, 2021; Alsan and Wanamaker, 2018), but there is limited evidence on information-sharing. We reinforce the role of religious identity among interacting citizens, a growing field of study in both economics and political science (Iyer, 2016). The specific focus on the use of religion for spreading information through mobile phones furthers our understanding of how these technologies stimulate social mobilization (e.g. Enikolopov et al., 2020; Manacorda and Tesei, 2020).

2. Conceptual framework

Following the frameworks of Pauly and Blavin (2008) and Baicker et al. (2015), we assume that agents have inaccurate beliefs about or salience of the value of preventive health practices in a global outbreak of an infectious disease, the COVID-19 pandemic. Wrong beliefs about the returns of preventive practices can lead to under-adoption, i.e., lower take-up than the socially optimal level.[4] If these are binding constraints to preventive care, an information campaign could promote adoption by correcting beliefs about the returns of these practices or by raising their salience (Haaland et al., 2023).

We study two hypotheses related to this mechanism. The first hypothesis is that messages from doctors are effective at promoting the adoption of preventive practices. This hypothesis depends primarily on three factors. First is whether doctors are considered a credible and trusted source of information (O'Keefe, 2016; Khan et al., 2021). This is crucial as in our information campaign 95% of the targeted population report doctors as the most trusted source of COVID-19 information. Second is the degree of malleability of the beliefs that are causing under-adoption. Information campaigns are more effective at influencing beliefs based on misconceptions or incomplete understanding than at changing views that are less grounded on facts or knowledge (Walter and Salovich, 2021). This factor demands distinguishing between these two dimensions when analyzing the campaign's impacts. Third is whether messages influence an individual's attitude towards checking the truthfulness of new information, which demands studying how the targeted population reacts when facing misinformation. For instance, a campaign may increase fact-checking if individuals become more aware of the degree of misinformation flowing in their social network, or decrease it if the ability to recognize false or inaccurate information is improved.

The second hypothesis is that messages from doctors are more effective when the sender and the receiver of the message are socially close and that such closeness becomes salient. If beliefs or salience are binding constraints to preventive care, then social proximity could enhance the effectiveness of the campaign by increasing the degree of credibility of information, particularly in the face of high levels of parasite stress (Fincher and Thornhill, 2012) and when the target group is more marginalized and less educated, and thus more socially distant from doctors (Lazer et al., 2018; Bavel et al., 2020). The enhancing effect of social proximity can also operate by raising the salience of group identity, with important consequences for norm compliance (e.g. Akerlof and Kranton, 2000; Chen and Li, 2000), but also for the updating of beliefs. For instance, social proximity could correct beliefs that have a close connection to the in-group or out-group identities. In our setting, beliefs over the effectiveness of vegetarianism in protecting against COVID-19 have a strong salience in Hindu communities, but not in Muslim communities. More generally, religious beliefs and practices tend to increase in times of crisis, and the COVID-19 pandemic was no exception (Bentzen, 2021). Adherence to recommendations was shown to be higher among more religious individuals in the context of the US (DeFranza et al., 2021).

3. Context

Our research setting is slums in the two largest urban agglomerations in the Indian state of Uttar Pradesh (UP), the cities of Lucknow and Kanpur. Appendix Figure A1 shows their geographic location. The setting is highly relevant for contagious diseases as, similar to many expanding cities in low- and middle-income countries, Lucknow and Kanpur are characterized by a relatively high prevalence of informal settlements and the prospect of rapid population growth.[5] While UP has a higher poverty rate than the average for India (29.43% versus 21.92%; Reserve Bank of India, 2019), its slum population is highly comparable to the average slum population in the country (Armand et al., 2023).[6]

We draw a random sample from the slum population of the two cities, as described in more detail in Section 4. Appendix Table C1 presents descriptive statistics of the sample. Almost 80% of respondents are male, mostly being the household head, with an average age of 40 years. In terms of income, 73% live in a dwelling not shared with other families, 61% have access to a private latrine, and 38% have a ration card (i.e. an official document giving access to the subsidized purchase of essential commodities). The social composition of the targeted area is heterogeneous, with an average share of Muslim residents in a slum of 21%, 22% of slums having no Muslim residents, and no slum having no Hindu residents. The distribution of different religions and castes in these populations is shown in Appendix Figure A1. The sample also presents high levels of religiosity as 64% of respondents strongly agree or agree with the statements "My religious faith/philosophy of life has a pronounced impact on my daily life" and "When I take important decisions, my religious faith/philosophy of life plays a considerable role". This proportion is higher for Muslim respondents (77% compared with 54%), and falls over time as restrictions are eased.[7] The decrease in religiosity is consistent with observations that religiosity is higher in times of crisis (Bentzen, 2021). In line with high levels of religiosity, the average trust in COVID-19 information shared by religious leaders is 0.53 out of 1. However, this level of trust is lower than the average trust in information shared by the government (0.73) and by doctors (0.85).

During the period of the study and similar to other Indian states, UP was hit hard by the pandemic. The number of COVID-19 cases and related deaths rose rapidly (Appendix Figure A2). At the time of the baseline survey, 12% of respondents reported that at least one member was experiencing COVID-19 symptoms. To address the emergency, the Government of India introduced guidelines for social distancing and wearing of face masks, which remained in place throughout the study period (see Fig. 1). The salience of these guidelines was particularly high in UP due to the features of its population. Out of 29 states, UP is the largest (home to 200 million people), the fourth most-densely populated, and the sixth in terms of share of population living in slums, totaling more than 6 million people (Government of India, 2011).

[4] Under-adoption in informal settlements, or 'slums', can also be driven by limited access to clean water and safe sanitation, and overcrowding (Patel, 2020; Wasdani and Prasad, 2020; Armand et al., 2023).

[5] In 2015, Lucknow and Kanpur were the 129th and 141st cities worldwide in terms of population (United Nations, 2019), with expected growth in the period 2015–35 of 59% and 37%, respectively. Across agglomerations of similar size, this growth prospect is comparable to cities such as Accra (Ghana) or Amman (Jordan).

[6] The shares of adult males (0.53 in UP versus 0.52 in India), adult females (0.47 versus 0.48), and children (0.14 versus 0.12), as well as the sex ratio (1.12 versus 1.08) and the share belonging to Scheduled Castes (0.22 versus 0.20), are indicative of close similarities between these two populations. In terms of literacy rates, the average slum in UP outperforms the average for the whole of India (0.78 versus 0.69).

[7] Religiosity declines from 72% in the first follow-up to 58% in the second follow-up. The measure of religiosity is not available at baseline.

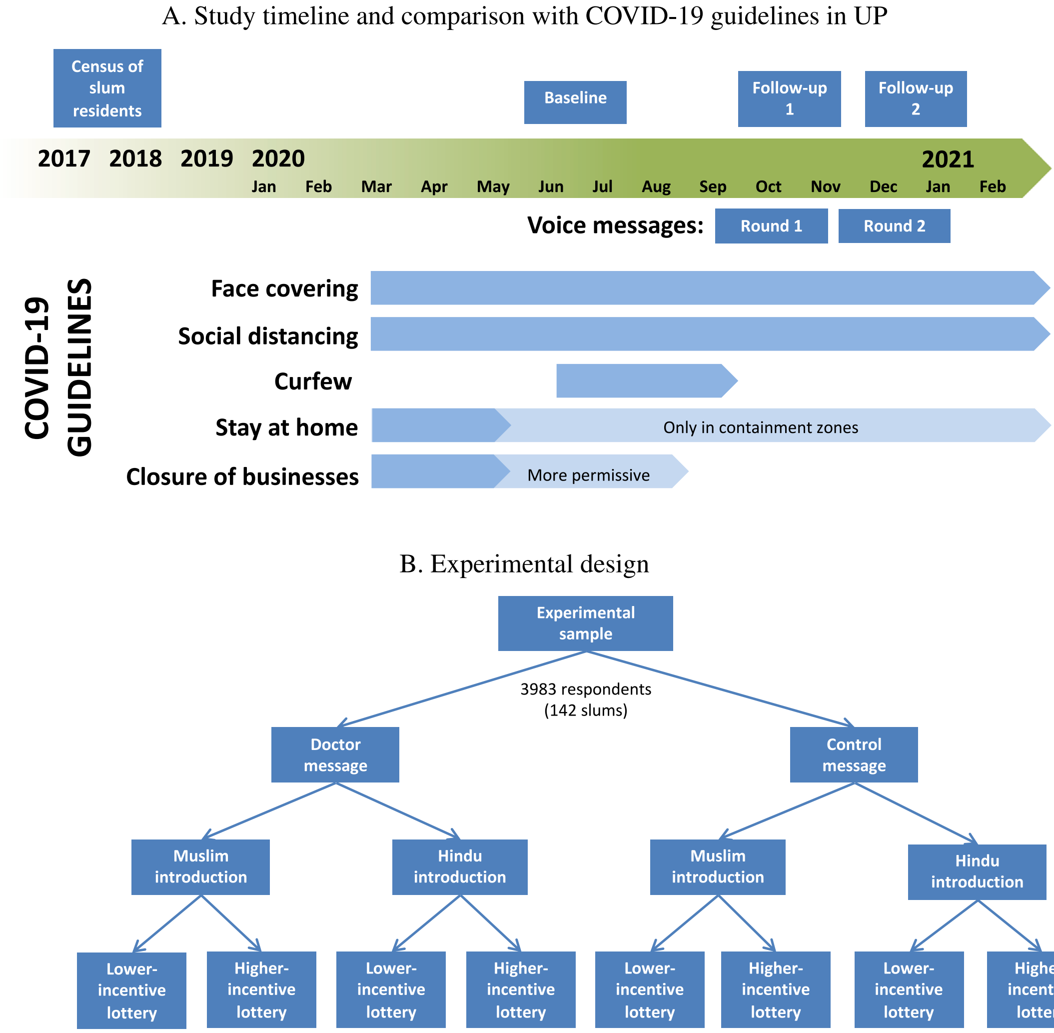

Allocated	490	505	503	498	501	503	479	504
Interviewed at follow-up 1	474	474	493	457	477	478	483	465
Interviewed at follow-up 2	496	468	485	504	473	493	503	477

Fig. 1. Study timeline and experimental design.
Notes. Panel A shows a summary of the study timeline with a comparison with COVID-19 guidelines in UP, while panel B shows a summary of the experimental design. Guidelines are compiled from official sources (Awasthi, 2020; Government of India, 2021). Lucknow and Kanpur were included in the red zone in May 2020. Red zones are the areas with high coronavirus cases and high doubling rate in the previous 21 days. The first phase of the closure of businesses included all businesses apart from essential shops and services, while the second more permissive phase allowed the re-opening of the following activities: shopping malls, religious places, hotels, and restaurants in June 2020 (unlock phases 1 and 2); gyms and yoga centers in August 2020 (unlock phase 3); entertainment, sport, political, academic, and social functions and gatherings with a limited number of participants in September 2020 (unlock phases 4, 5, and 6). Curfews were first characterized by night curfews from 9pm to 5am in June and July 2020, and then weekend curfews until September 2020. Local authorities had the power to impose curfews based on local conditions.

The onset of the pandemic was accompanied by the spread of misinformation regarding the causes and prevention of COVID-19. The diffusion of fake news was facilitated by the relatively low literacy levels and the dramatic increase in internet penetration rates experienced by India, which went from 4% in 2007 to 50% in 2020 and raised social media platforms as a primary source of news and as a key mean of communication for the Government of India and other political party actors (Mahapatra and Plagemann, 2019; Statista, 2021). The wave of misinformation became so severe that PM Narendra Modi addressed the nation urging everyone to rely only on credible medical advice and demanding social media companies to curb misinformation on their platforms (Akbar et al., 2020; Al Jazeera, 2020).

The primary drivers of the increase in fake news on social media were misleading claims about the role of Muslim citizens in the spread of the virus. As evidenced by trend analysis of social media interactions in Facebook-related media (Appendix Figure A2), the targeting of the Muslim population spiked during the onset of the COVID-19 pandemic, with them being blamed for the spread of the virus. This trend is primarily driven by UP, where these tensions fueled pre-existent tensions that spurred violence against the Muslim population (Banaji and Bhat, 2020; Menon, 2020) and have had an impact on public health, affecting its provision (such as hospitals in the state reportedly segregating Hindu and Muslim COVID-19 patients; Withnall, 2020) and hindering it (Sarkar, 2020).

Religious tensions and misinformation centering on religion are not specific to the pandemic. First, Hindu–Muslim conflict in India goes back to the pre-partition era and has flared up regularly since (e.g.) Mitra and Ray, 2014). UP stands out as one of the states where Hindu–Muslim tensions have been particularly long and severe (see, for instance, Narayan, 2014). Second, misconceptions centering on religion have important links with political mobilization in India, as politics and religious (Hindu) nationalism are deeply connected (Philpott, 2007; Laborde, 2021), and misinformation campaigns led by political actors are often targeted at religious minorities (Poonam and Bansal, 2019; Al-Zaman, 2021).

4. Intervention and experimental design

The intervention is designed to test the hypotheses discussed in Section 2. It took place during the initial phase of the COVID-19 pandemic and consists of sharing voice messages via calls targeted at individual citizens using mobile-phone technology.[8] Fig. 1 shows the study timeline, comparing it with COVID-19 regulations in UP in the corresponding period, and summarizes the experimental design. Each message has two components: the *introduction* delivered by a local citizen (the *sender*) and the *content* of the message. The full scripts of the messages are reported in Appendix Section A.2.

To introduce variation in social proximity associated with the message, we exploit religious diversity in UP. In the slum setting, the representation of religious groups is comparable to that of the whole state, with 79% of the sample being represented by Hindu citizens and 21% by Muslim citizens. Members of these religious groups tend to use distinct greetings. We exploit this characteristic by introducing two variations in the introduction of the message. The sender signals either a Hindu identity by using the greeting "namaste" at the start of the message or a Muslim identity by using the greeting "salam alaykum". The remaining part of the introduction is kept constant, including the language spoken. We refer to *religion concordance* of the message when the initial greeting of the sender is signaling the same religion as the receiver of the message an to *religion discordance* when it is signaling a different religion.[9]

To separately introduce variation in the message content, we varied the content following the introduction to be either informative (with the objective of raising preventive health awareness) or uninformative. In the informative version, labeled as the *doctor* messages, the content is presented by doctors from locally renowned medical institutions, debunks common misconceptions about ways to prevent COVID-19, and provides reminders about the confirmed ways to protect against infection. Qualified medical practitioners were chosen for the informative content to guarantee that information was shared by trusted sources (see Section 2). We sent two rounds of messages. Each message reminded the receiver about the World Health Organization (WHO) recommended practices to avoid contagion and, in addition, the first message highlighted that eating a vegetarian diet does not protect against COVID-19 (sent in October–November 2020) and the second message debunked the fake news that the immune system of Indians is resilient to COVID-19 (sent in December 2020–January 2021).[10] At baseline, relying on vegetarianism and on the Indian immune system were the two most prevalent non-evidence-based preventive practices to avoid contagion from COVID-19 (Appendix Figure A3). All participants allocated to the *doctor* messages received messages from the same set of three doctors. We did not randomize the religious identity of doctors in order to disentangle the effects of identity from other doctor-specific characteristics (e.g. doctor from religion A also being more charismatic than doctor from religion B). Instead, we used messages from religion-neutral doctors (i.e. doctors did not reveal their religious identity, neither through salutation nor by their name).[11]

In the uninformative version, labeled as the *control* messages, the recording begins with the same introduction from the local citizen as in the doctor message, but the message content is unsubstantiated, religiously neutral gossip about Bollywood stars. The choice of content was based on suggestions from our experienced data collection partner, local to the study site. Sending a control message, rather than no message, allows us to disentangle the effects of the intervention from the effect of receiving a message.

The length of the recordings was 1.58 min (or 95 s) for the first round of the doctor message and 1.55 min (or 93 s) for the second round of the doctor message. Though ideally the length of the control message had been the same, it ended up shorter in our

[8] Alternative remote approaches include live phone calls (Sadish et al., 2021), communication via instant messaging platforms (Bowles et al., 2020), and pedagogical interventions (Badrinathan, 2021).

[9] Religious identity is actively expressed in everyday life in India through dietary restrictions (beef for Hindus, pork for Muslims), language preferences (Arabic/Urdu for Muslims, Sanskrit/Hindi for Hindus), attire, rituals, customs, and religious holidays, among other attributes, as recently highlighted in a representative survey of 30,000 adults (Sahgal et al., 2021).

[10] The content for these messages was built by first asking several doctors from renowned local institutions to reply unscripted to the questions "iIs it true that eating a vegetarian diet protects against COVID-19?" and "Is it true that the immune system of Indians is resilient to COVID-19?". Responses were collated ensuring that every message consisted of a first part debunking the misconception and a second part on policy recommendations.

[11] It is possible that respondents infer that the doctors providing the answers belong to the same religious group as the sender. This dimension is not observed in our data. We take a conservative approach and interpret the results as the religion concordance being between the participant and the sender only, as intended in the intervention design.

design, at 1.28 min (or 77 s) in both rounds.[12] Sharing the audio messages via phone calls allowed us to know which participants answered the call and to measure the duration of the audio message that was played (see Section 7.1).

To reduce the risk of low uptake of the information campaign, all messages were incentivized to increase attention paid to the message by giving participants the chance to enter a lottery if they replied correctly to a follow-up question about the message. Respondents were randomly allocated to either a lower-incentive or a higher-incentive lottery. The research design is therefore a $2 \times 2 \times 2$ randomized controlled trial using household-level randomization after stratifying by religion of the household head and city of residence. We adopted the following procedure: first, we randomly allocated targeted households to receive either doctor or control messages; second, we cross randomized households in both the doctor and control message groups to receive a message introduced by a Hindu or a Muslim greeting, thus creating exogenous variation in religion concordance; third, we cross-randomized households in both the doctor and control message groups into a lower-incentive lottery with a value of Rs. 2500 (US$32) or a higher-incentive lottery with a value of Rs. 5000 (US$64).

Appendix Table A2 shows descriptive statistics of the take-up of voice messages, both on the extensive margin (i.e. whether a person picked up the phone) and the intensive margin (i.e. conditional on picking up the phone, what share of the message is listened to). The table also shows conditional correlations between these variables and individual characteristics. On average, 36.2% of participants picked up the phone when sent the first message and 38.4% picked up the phone when sent the second message. Conditional on picking up, participants listened to 60.9% of the first message and 50.0% of the second message. Given the soft nature of the intervention, our take-up is relatively high compared with other information experiments and mass information campaigns (e.g. Azrieli et al., 2018). For instance, in the context of unincentivized video messages sent to Indian citizens by SMS urging them to comply with COVID-19 policies, Banerjee et al. (2020) achieved a viewing rate of just 1.1%.

To test whether demographic characteristics are predictors of take-up, we perform F-tests for the joint equality to zero of the coefficients on the characteristics included in the regressions explaining whether the respondent picked up the call and the share of the message that is listened to. In the full sample, we reject this hypothesis only for the share of the message that is listened to and exclusively for the second round of messages. This suggests that participants might have responded in terms of take-up of the message, but only in the second round.

In line with the pre-analysis plan (Armand et al., 2020) and to obtain a standard level of statistical power, in Section 7 we discuss treatment effects up to the second level of randomization. We focus on the effect of the content of the message and its combination with either the sender's religion or the level of monetary incentives. For the latter, because the lottery amounts are both sizable, and therefore differential impacts are marginal, we present the results in Appendix Section D.5 and discuss them in Section 7 when relevant. Section 7.5 discusses potential threats of spillover effects deriving from the experimental design, and how we exploit household-level randomization to test for spillover effects.

5. Data

We draw on two data sources, summarized in this section: a panel survey of slum residents from a sampling frame carrying unique information for more than 30,000 households living in the slums of the study area before the beginning of the pandemic,[13] and administrative data on the implementation of interventions. Appendix B offers detailed description of each variable, including the type (self-reported, elicited, or from administrative records) and the round (baseline or follow-up), and elaborates on the ethical considerations related to data collection activities.

Primary panel data. We collected primary data among slum residents on households' experiences during the COVID-19 pandemic, such as their knowledge on how to prevent the virus, compliance with policies, their sources of information, and trust and beliefs. We collected a baseline survey in June–July 2020, reaching 3983 households. Two waves of follow-up panel data were collected in October–November 2020 and December 2020–January 2021 (3.5 and 5.5 months after the baseline survey), reaching 3801 households during the first follow-up and 3899 during the second follow-up survey. To keep the time gap between the intervention and follow-up data collection similar across individuals, we split the sample into four batches determined by the operational capacity of the field team. In each batch, we interviewed households two weeks after sending the voice messages by conducting phone conversations. The sampled households that were not reachable at the time of the survey were replaced with replacement households randomly selected from the sampling frame described above.

Combining both follow-up surveys, we re-interviewed 87% of residents at least once, with a low implied attrition rate (13%) compared with phone surveys conducted in similar settings. Response rates are typically around 50% in non-crisis contexts, while

12 In addition to sharing voice messages, the original intervention also included sending the video underlying the voice messages through a WhatsApp chatbot, i.e. a software purposely programmed for the intervention that runs on the encrypted WhatsApp platform and in which users can communicate with the software through the chat interface. In addition to the variation induced by the initial greeting, videos also varied the name (as printed in the video) and the clothes of the sender to signal either a Muslim or Hindu identity. Yet, videos were only visualized by a very small share of participants due to the WhatsApp policy requiring each chat to start with a generic greeting "Hi" and to share the rest of the chatbot message and the video message only if the respondent replied to the initial greeting. Previous studies using WhatsApp make use of subscribers, thus by-passing this precondition (e.g. Bowles et al., 2020). We sent the video message to all phone numbers in the sample, 38.9% received the chatbot message saying "Hi" (i.e. this share had a smartphone, WhatsApp installed on their phone, and a data package activated or an internet connection), and just 2.5% replied to the initial greeting and received the rest of the chat message and the video. This percentage did not vary by treatment arm. We cannot verify the share that downloaded and watched the video, but, in line with the literature, we can reasonably assume it to be much smaller than 2.5%. Such low uptake is a common risk in experiments (e.g. Azrieli et al., 2018). Including controls for the receipt of the video message on WhatsApp or excluding these participants from the sample does not alter any of the results.

13 Refer to Solís Arce et al. (2021) and Armand et al. (2023) for further details about this population and the census procedures.

they are expected to be lower during crisis contexts. For instance, a study during the Ebola crisis was able to re-interview only 38% (Himelein et al., 2020). Attrition is orthogonal to treatment allocation, while being female and not sharing a dwelling significantly reduces attrition (Appendix C).

The primary outcome is compliance with recommended practices to avoid spreading COVID-19, as highlighted in the doctor messages. We collected information about behavior related to hygiene and physical distance split in two modules. To guarantee both a high quality of information and a concise interview, each module was administered to a random subset of households only. We build a compliance index for all respondents that responded to one of the modules. The index captures the extent to which respondents wear a face mask when going out, the frequency of hand-washing, and the extents to which they stay in the slum, do not receive visitors from outside the slum, and do not meet anybody from outside the slum. Individual questions are detailed in Appendix B. To build the index, we aggregate individual variables using an index of z-scores following (Kling et al., 2007), by first normalizing individual variables in standard deviations from the control group, and then averaging available information.

We supplement this index with information on beliefs over the efficacy of different ways to prevent infection from COVID-19. We asked respondents about their level of agreement with various recommended preventive practices (i.e. those present in policy recommendations) and non-evidence-based preventive practices (i.e. those not present in policy recommendations), all of them discussed in the doctor messages (see Section 4). The evidence-based practices were wearing a face mask, hand-washing, and keeping physical distance. The non-evidence-based practices were the two most-common views collected at baseline on how to protect from the virus, which are also the ones that the doctor messages debunked: relying on vegetarianism or on the Indian immune system.[14] The beliefs over the efficacy of recommended practices are strongly positively correlated with the compliance index, and the beliefs over the efficacy of non-evidence-based practices are negatively correlated with the compliance index, validating the index (Appendix Table D5).

Finally, we measure how participants respond to misinformation about COVID-19. We gather information on *fact-checking*, a proxy for evidence-based behavior related to misinformation. Additionally, we introduce a novel survey instrument to elicit how participants respond when facing misinformation shared by other citizens. In line with the literature (e.g. Scheufele and Krause, 2019), we define misinformation as incorrect views based on faulty knowledge or understanding. We present respondents with two statements attributed to a third person living in UP, whom we refer to as the *interlocutor*, and we then elicit their level of agreement with each statement. Statements are presented in a random order during the interview to avoid question-order bias.[15] The content of the statements was chosen to reflect common claims by the media, including some with significant religious salience. The first statement, “If you are vegetarian, you do not need to worry about the coronavirus”, carries specific religious salience since, in the context of India, vegetarianism is widely associated with the dominant ideology of Hinduism. The second statement, “If you are a good person, you do not need to worry about the coronavirus”, carries general religious salience, with the idea that religion helps in becoming a good person.[16]

As agreement with misinformation is often associated with motivated thinking (i.e. the set of emotional biases leading individuals to agree with views based on desirability rather than evidence), agreement with these statements may vary based on the interlocutor’s identity. This aspect is crucial in our context, where religious tensions can blur the lines between misinformation agreement and group identity, often linked to religion (Tankard and Paluck, 2016; Nyhan, 2021). To investigate this, we choose the name of the interlocutor to signal different religious identities using five options: 1 male Muslim name, 1 female Muslim name, 1 male Hindu name, 1 female Hindu name, or a generic “people”. Names were selected using information on the most common names by religion from the census of slum residents (see Section 4). For each respondent, statements are randomly allocated to one of these 5 options. Because the list of statements is constant in the survey, but interlocutors vary in each interview, we can measure agreement with statements shared by citizens that are either *in-group* or *out-group* interlocutors, depending on whether the respondent shares the religion signaled by the interlocutor. When two interlocutors fall in the same religious identity category, we average agreement with their individual statements.

Fig. 2 provides descriptive statistics on these variables for the control group. Panel A focuses on the index of compliance with recommended practices and on respondents’ levels of agreement with evidence- and non-evidence-based preventive practices over the course of study. Panel B focuses on fact-checking and respondents’ levels of agreement with misinformation shared by in-group and out-group citizens.[17] A few observations are worth highlighting. First, likely because some of the restrictions were removed in the follow-up (e.g. the self-employed were allowed to work, and offices, supermarkets and entertainment industries reopened), the average level of compliance with preventive practices reduces over time. At the same time, the level of agreement with evidence-based ways to protect from the virus (second and third figures in Panel A) remains significantly higher than agreement with non-evidence-based practices (last two figures in Panel A). Moreover, the level of misinformation is noteworthy, as people on

[14] Baseline information for these variables is not available because the baseline questionnaire elicited practices through an open-ended question, rather than in levels of agreement with their efficacy.

[15] The exact script of the question reads as follows: “We have surveyed a few people from UP and we would like to hear if you agree with their opinion. Note that responses to the statements are a matter of opinion. There is no scientific evidence about their truthfulness. On a scale of 1 to 5, where 1 means you strongly disagree and 5 you strongly agree, how much do you agree or disagree with the following statements. [Interlocutor] says that [statement]”.

[16] We elicit agreement with three further statements, which contain views about COVID-19 that are not necessarily based on facts or knowledge. We label these statements as *opinions*. Because opinions are harder to influenced with information campaigns and fact-checking (Walter and Salovich, 2021), we use them as placebo statements. Impacts on these variables are discussed in Section 7.3, while Appendix Section D.6 presents descriptive statistics.

[17] Appendix Figure A4 shows respondents’ levels of agreement with each statement, distinguishing by whether the interlocutor is in- or out-group.

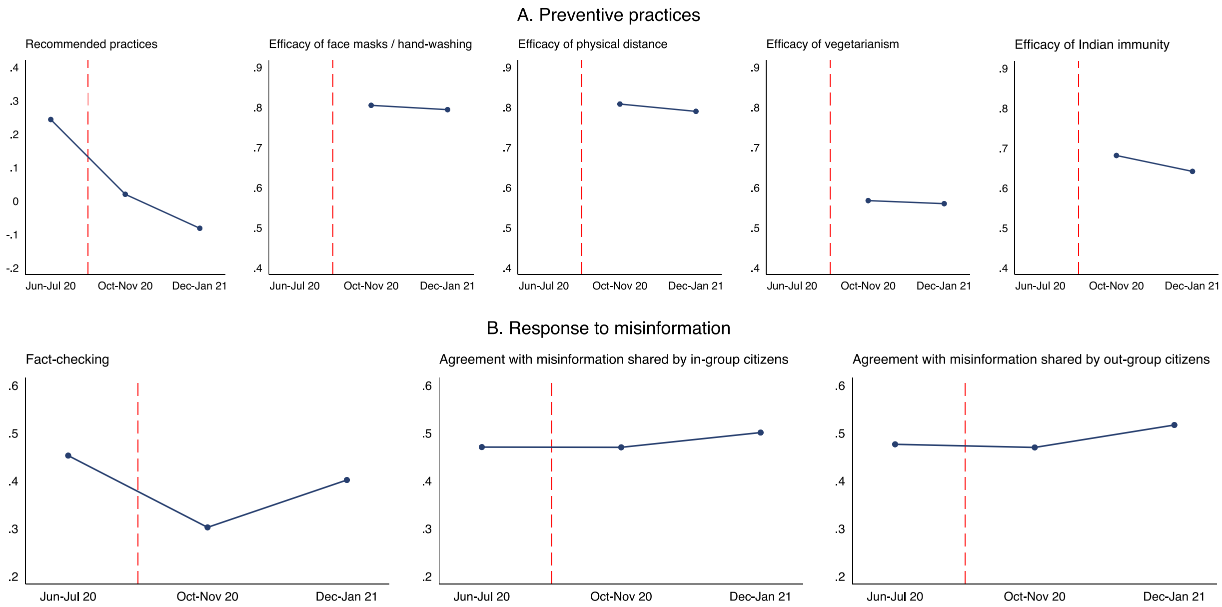

Fig. 2. Preventive practices and response to misinformation (control group).
Notes. Each figure shows the average of outcome variables measured at different points in time. The vertical line separates the baseline measurement from the follow-up measurements. In both panels, the sample is restricted to the control group that did not receive the doctor message. Further details about these variables are provided in Section 5. Variables are defined in Appendix B.

average neither agree nor disagree with misconceptions shared by both in-group and out-group citizens, and it increases slightly over time (last two figures in Panel B).[18]

While compliance and beliefs are based on self-reported data, it is important to highlight that this information was collected two weeks after exposure to the interventions. This extended period reduces concerns regarding experimenter demand effects (i.e. changes in behavior by experimental subjects due to cues about what constitutes appropriate behavior), as well as spurious priming effects (i.e. effects that dissipate within hours after the intervention and are only driven by the salience of the message, not by a change in knowledge, attitudes or behavior). To rigorously address concerns about experimenter demand effects, we collected baseline data on social desirability using the Marlowe–Crowne Social Desirability Scale, a survey module developed by social psychologists to measure a person's propensity to give socially desirable answers (Crowne and Marlowe, 1960). Due to time constraints in phone-based surveys, we employ a shortened version of the module, the Socially Desirable Response Set Five-Item Survey (SDRS-5; Hays et al., 1989). Shorter versions of the module have also been validated by Fischer and Fick (1993). The module prompts respondents with statements regarding traits that appear exceptionally positive or idealized, such as always being polite and a good listener, or never being jealous or resentful. Respondents who often agree with these statements receive higher scores, indicating a propensity to provide socially desirable answers.

In line with Hays et al. (1989), we collapse responses in an SDRS-5 score ranging from 0 to 1, with higher scores indicating more social desirability in responses. Scores are highly balanced across treatment arms (Appendix Table C1). In addition, we show evidence that, at baseline, social desirability does not influence reporting differently depending on the religion, the gender, and the caste of the respondents (Appendix Section D.10), excluding the possibility that these characteristics elicit different answers (e.g. Fowler and Mangione, 1990).[19]

Administrative data. The voice messages were sent to the whole sample in two rounds using an automated system. For each round, the system provides information about the delivery of voice messages, and about the duration and share of the voice message that each user played. Differential effects of the interventions on the take-up of messages are discussed in Section 7.1.

[18] We find differences by religion in beliefs for non-evidence-based practices as well as for misinformation (Appendix D.1). On average, Hindu respondents are significantly more likely to agree with non-evidence-based ways and with misinformation shared by other citizens, a difference that is mainly driven by beliefs about vegetarianism, the predominant diet among the Hindu population. Finally, agreement with misinformation tend to be relatively constant over time and similar across different types of interlocutors.

[19] Information about who is present at the moment of the interview is not available. Therefore, we cannot exclude the possibility that bystanders could have influenced responses (e.g. Tourangeau and Yan, 2007). We proxy for the presence of bystanders using a dummy for whether the interview is at a weekend. Results confirm the absence of social desirability bias along this dimension (Appendix Section D.10). We discuss heterogeneous treatment effects by social desirability in Section 7.

6. Empirical approach

The primary objective is to test different hypotheses on how the interventions translate into behavioral impacts, as discussed in Section 2. The first hypothesis is that the doctor message, which carries informational content related to COVID-19, impacts health-related behavior (compared with the control message, which has no content related to preventive practices). The second hypothesis is that the doctor message with religion concordance between the sender and the receiver generates different impacts from the doctor message in which the religion of the sender is different from one of the receiver's. In the experimental design, there exists another hypothesis in which the control message with religion concordance generates differential impacts compared with a control message in which the religion of the sender is different from the one of the receiver. However, because the control message has no content related to preventive practices and it is not expected to impact health-related behavior, we expect no differential impact. We in fact treat this comparison as a placebo comparison and discuss it in Appendix Section D.2.

For the first hypothesis, we estimate the impact of the doctor message using the following specification:

$$Y_{ijt} = \beta_D \, doctor_i + \alpha \, \mathbf{X}_{ij} + \delta_t + \epsilon_{ijt} \tag{1}$$

where Y_{ijt} are outcomes for interest of respondent i in slum j at time t. The variable $doctor_i$ is an indicator variable equal to 1 if receiver i is in the doctor message treatment group, and 0 otherwise. $\mathbf{X}_{ij}$ is a set of control variables, and δ_t are period-of-survey indicator variables. In the main analysis, $\mathbf{X}_{ij}$ includes only the indicator variables for randomization strata.[20] Adding more control variables selected with the post-double selection LASSO (PDSL) procedure (Tibshirani, 1996; Belloni et al., 2013) or controlling for the baseline value of the outcome variable (ANCOVA specification) does not affect the results; if anything, precision improves (Appendix Section D.7). The error term ϵ_{ijt} is assumed to be clustered at the slum level, but results are robust to alternative assumptions about standard errors, such as clustering at the individual level.

For the second hypothesis, we estimate the role of religion concordance with the sender of the doctor message by restricting the sample to the doctor message group, therefore focusing on a group that received the same informational content, and estimating the following specification:

$$Y_{ijt} = \beta_C \, concordance_i + \alpha \, \mathbf{X}_{ij} + \delta_t + \epsilon_{ijt} \tag{2}$$

where $concordance_i$ is an indicator variable equal to 1 if receiver i was sent a message in which the sender and the receiver have the same religion, and 0 otherwise. The parameter β_C captures the differential effect of receiving a religion-concordant doctor message compared with a religion-discordant doctor message.[21] It is therefore testing whether religion concordance, compared with discordance, creates differences in the effects of the doctor messages estimated in Eq. (1). We note that this approach complements the pre-specified one, which proposed an interacted model, imposing that the main effect of religion concordance (i.e. the effect of sending a message with concordance independently from the content) is the same in the doctor message and in the control group. Results using this approach are in line with the ones presented in the main text, but less precise for some outcomes (Appendix Section D.4). Because in the final design of the experiment, the content in control messages is very different and significantly shorter than the one in the doctor messages, and assuming homogeneity of the main effect of religion concordance reduces precision, our preferred strategy remains that of presenting the results using Eqs. (1) and (2) separately, assuming that the main effect of religion concordance is heterogeneous in the doctor and the control messages. In line, religion concordance has differential effects on the take-up of interventions depending on the content of the message (see Section 7.1).

We estimate both Eqs. (1) and (2) by pooling data from the two follow-up surveys together, thus estimating the average impact in the follow-up period (i.e. assuming β_D and β_C are constant over time). When outcome variables are measured in close temporal proximity, this approach allows averaging out the noise in the outcome variables and increases power (McKenzie, 2012). Appendix Section D.4 shows results for each follow-up survey separately. Appendix Section D.1 shows how estimates vary in sub-samples defined by pre-specified variables (religion of the respondent and percentage of residents in the slum who are Muslim), and by other relevant dimensions (caste, strength of religious identity, trust in the government), which we discuss in the next section.

Because not everybody listens to the message that is sent (Section 7.1 provides details about treatment compliance), as is standard in mass information campaigns, we supplement the main estimates with instrumental variable (IV) estimates that consider the actual exposure to the interventions. Using administrative data, we compute $share_{ijt}$, i.e. the (endogenous) share of each message that is effectively listened to on the phone by respondent i in slum j at time t.[22]

We then estimate versions of Eqs. (1) and (2) in which the treatment indicators are multiplied by $share_{ijt}$. To estimate the effect of listening to a doctor message, we define actual exposure as $shareD_{ijt} = share_{ijt} \cdot doctor_i$, and instrument it with the treatment indicator $doctor_i$. We estimate the following equations using two-stage least squares (2SLS):

$$\begin{aligned} Y_{ijt} &= \beta_D^{IV} \, shareD_{ijt} + \alpha \, \mathbf{X}_{ij} + \delta_t + \epsilon_{ijt} \\ \widehat{shareD}_{ijt} &= \gamma_D \, doctor_i + \lambda \, \mathbf{X}_{ij} + \delta_t + v_{ijt} \end{aligned} \tag{3}$$

[20] We include the indicator variable for the city of residence, and an indicator for whether the household is of Muslim religion as defined in the census of households (see Section 5). These indicators were used for stratified randomization (see Section 4).

[21] Appendix D.5 provides estimates of the effect of a Hindu versus a Muslim greeting, independently from the religion of the recipient. We observe no effect for these comparisons.

[22] Appendix Section D.9 provides a similar analysis using as the measure of actual exposure to the interventions an indicator variable as to whether the respondent listened to any positive share of the message.

Table 1
Treatment effects on the take-up of messages.

	Picked up (1)	% listened to (2)	Duration (minutes) (3)
A. Full sample			
Doctor message	−0.016	−0.246	0.301
	(0.013)	(0.014)	(0.027)
	[0.22 , 0.22]	[0.00 , 0.00]	[0.00 , 0.00]
Mean (control message)	0.381	0.674	0.551
Observations	7700	2873	2873
B. Sample restricted to doctor message group			
Religion concordance	−0.029	0.053	0.122
	(0.016)	(0.021)	(0.047)
	[0.09 , 0.09]	[0.01 , 0.02]	[0.01 , 0.02]
Mean (religion discordance)	0.377	0.398	0.790
Observations	3851	1406	1406

Notes. Estimates based on OLS regressions using Eq. (1) in Panel A and Eq. (2) restricting the sample to participants allocated to the doctor message in Panel B (see Section 6). Standard errors clustered at the slum level are reported in parentheses. P-values are presented in brackets, the first from individual testing and the second adjusting for testing that each treatment is jointly different from zero for all outcomes presented in each panel. The dependent variables are: in column (1) *Picked up*, an indicator variable equal to 1 if the respondent picked up the call in any of the two rounds of interventions, and 0 otherwise; in column (2) *% listened to*, the share of the message that is listened to, conditional on having picked up; in column (3) *Duration (minutes)*, the duration of the call, conditional on having picked up. Note that the doctor messages have different duration from the control messages (see Section 4). All specifications include indicator variables for data collection rounds, and strata indicators (city and religion of respondent).

To estimate the effect of listening to a religion-concordant doctor message, we instead restrict the sample to the doctor message group and define actual exposure as $shareC_{ijt} = share_{ijt} \cdot concordance_i$. We instrument this variable with the treatment indicator $concordance_i$ and estimate the following equations using 2SLS:

$$\begin{aligned} Y_{ijt} &= \beta_C^{IV}\, shareC_{ijt} + \alpha \mathbf{X}_{ij} + \delta_t + \epsilon_{ijt} \\ \widehat{shareC}_{ijt} &= \gamma_C\, concordance_i + \lambda \mathbf{X}_{ij} + \delta_t + \upsilon_{ijt} \end{aligned} \tag{4}$$

The parameters β_D and β_C in Eq. (1) and Eq. (2) can be interpreted as ITT effects (i.e. they capture the effect of sending a message, independently from whether a person listens to it). Conversely, β_D^{IV} and β_C^{IV} inform about the magnitude of the effects in the presence of full compliance. In light of the likely heterogeneity in the (potential) impacts of the intervention, these estimates can be interpreted as the local average treatment effects (LATE) for participants that comply with the interventions (e.g. Imbens and Angrist, 1994).

We find balance in terms of observable characteristics across groups, both in the allocation to the doctor and control messages, and across Muslim and Hindu senders within the doctor message group. Appendix Table C1 shows mean differences at baseline between the different treatment arms for respondent characteristics.

For statistical inference, we supplement in each table standard inference for the ITT estimates of Eqs. (1) and (2) with multiple hypothesis testing adjusting p-values for the significance of each individual coefficient in the table using the (List et al., 2019) bootstrap-based procedure. To this end, we categorize hypotheses by grouping variables into three groups and present the results in Section 7. First, in Section 7.1, we test whether the interventions impacted take-up of the messages. Second, in Section 7.2, we test whether the interventions changed compliance with recommended practices and belief over the efficacy of preventive behavior. Third, in Section 7.3, we focus on whether interventions influenced the response of study participants to misinformation. Fourth, in Section 7.4 we look at whether effects of religion concordance vary by whether a respondent is Muslim or Hindu. Finally, in Section 7.5, we verify whether estimates are influenced by potential threats from spillover effects.

7. Results

7.1. Take-up of the campaign

Table 1 shows estimates of the effect of the doctor message and of religion concordance in the doctor message on the probability of having picked up the call and on the share and on the duration (in minutes) of the message that is listened to. These variables are computed from administrative data (see Section 5). Heterogeneity of treatment effects on the take-up of the interventions by the round of messages and by religion – a pre-specified heterogeneity dimension – are reported in Appendix Section D.1.

We begin by focusing on the effect of sending doctor messages versus control messages. In Panel A, we estimate Eq. (1) using the full sample of respondents. On average, 38.1% of respondents in the control group picked up the call at least once. Conditional on having picked up the call, they listened to the message for 0.55 min (33 s) or 67.4% of the message. Sending a doctor message did not shift the probability of picking up the call, but did significantly decrease the share of the message that is listened to by 24.6

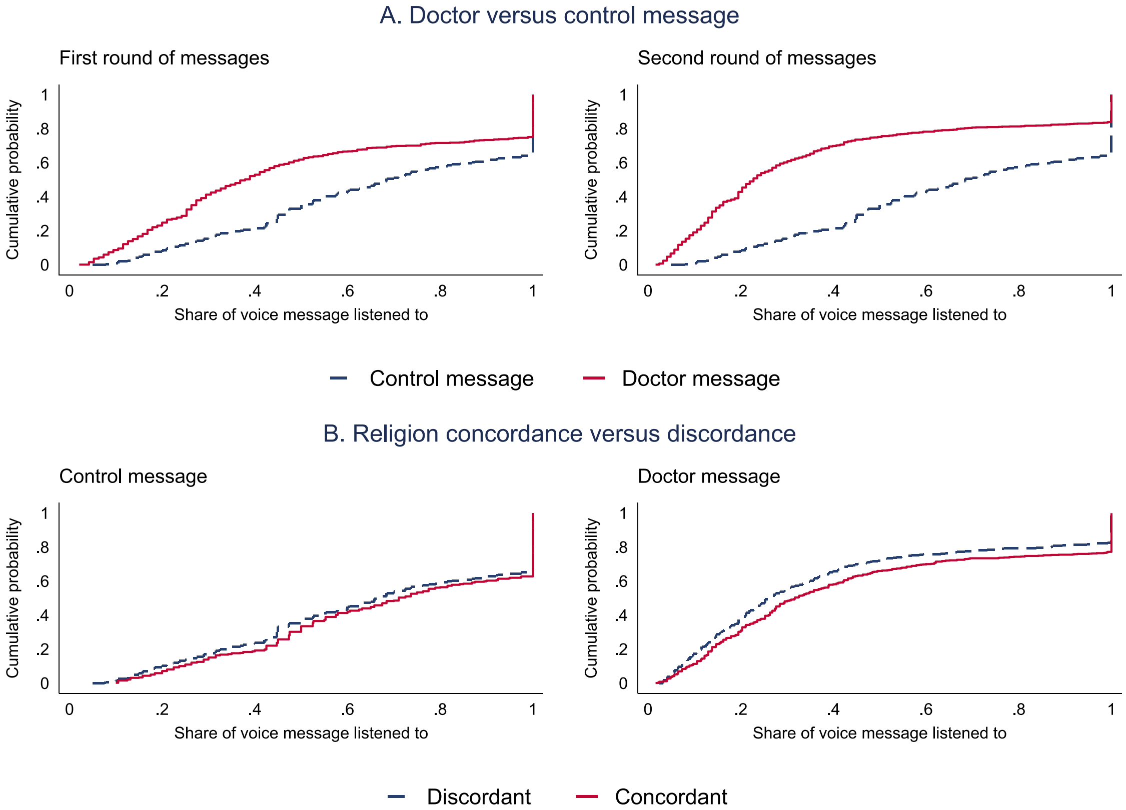

Fig. 3. Share of voice messages listened to by study participants.
Notes. The figures show the share of the messages listened to by study participants, conditional on having picked up the call. Information is based on administrative data from the intervention. Panel A includes the full sample separated by round of intervention. Panel B restricts the sample to the control message group in the left figure and to the doctor message group in the right figure. Treatment effects on the take-up of messages are reported in Table 1. The duration of the call can be longer than the duration of the recorded message if the receiver spends time replying to the question at the end of the message. The p-values of Kolmogorov–Smirnov tests for the equality of distributions in each panel are smaller than 0.001 in both figures of Panel A, 0.002 in the left figure of Panel B, and 0.035 in the right figure of Panel B. If we exclude participants who listened to the full message, p-values are smaller than 0.001 in both figures of Panel A, 0.033 in the left figure of Panel B, and 0.020 in the right figure of Panel B.

percentage points. While the doctor's message keeps the respondent on the phone for an additional 0.30 min (18 s) on average, this extended duration does not result in a higher proportion of the message being listened to. This seemingly counter-intuitive result can be explained by the fact that the doctor's message is longer than the control message. Panel A in Fig. 3 highlights these differences separately for the first and second round of messages. Kolmogorov–Smirnov tests of the equality of the distributions of the share of the message that is listened to in the control and doctor message groups is rejected at the 1% confidence level for both the first and second round of messages.

In Panel B of Table 1, we focus on the introduction of religious proximity with the sender in the doctor message and estimate treatment effects using Eq. (2). Since we do not include the control group in this estimation, the length of the message is the same across groups. On average, 37.6% of respondents that received a doctor message with an introduction from a different religion picked up the call and, conditional on having picked up the call, they listened to 39.8% of the message, corresponding to 0.79 min (47 s). Religion concordance changes exposure to the doctor message significantly. For one, we find that the share of respondents that picked up the call is reduced by 2.8 percentage points (a decrease of 7.4% over the mean for messages with religion discordance). While people would not know about the source of the call ex ante, and thus one would not expect any difference across treatment groups, we show in Appendix Section D.1 that this reduced probability of picking up the phone when there is religion concordance is driven by the second call. This suggests that some respondents may have recognized the number and opted not to answer again. It is plausible that religious proximity heightened the call's salience, potentially prompting individuals to save the number for future recognition. Alternatively, as we later demonstrate, the effects of religion concordance on various behavioral responses could influence the decision to answer subsequent calls.

Importantly, conditional on having picked up the call, religion concordance leads to a significantly larger exposure to the doctor message. The share of the message that is listened to increases by 5.3 percentage points, corresponding to an additional 0.12 min (7 s). These effects, corresponding to an increase of 13.3% and 15.0% over the means for messages with religion discordance, are specific to the doctor message. In fact, in Appendix Section D.2, we show that religion concordance in the control messages

Table 2
Preventive practices: doctor versus control message.

	Compliance	Beliefs over the efficacy of...			
		Recommended practices		Non-evidence-based practices	
	Recommended practices (1)	Face masks/ hand-washing (2)	Physical distancing (3)	Vegetarianism (4)	Indian immunity (5)
A. OLS					
Doctor message	0.051	0.006	0.005	0.003	−0.007
	(0.022)	(0.003)	(0.004)	(0.006)	(0.005)
	[0.02 , 0.08]	[0.01 , 0.06]	[0.19 , 0.36]	[0.60 , 0.62]	[0.13 , 0.36]
Mean (control message)	−0.032	0.799	0.799	0.563	0.661
Observations	5125	7700	7698	7692	7697
B. IV					
% listened · doctor message	0.326	0.041	0.031	0.020	−0.046
	(0.138)	(0.016)	(0.024)	(0.038)	(0.030)
	[0.02]	[0.01]	[0.19]	[0.60]	[0.13]
Mean (not listened)	−0.032	0.799	0.799	0.563	0.661
Effect size (avg. exposure)	0.137	0.017	0.013	0.008	−0.020
Observations	5125	7700	7698	7692	7697

Notes. Estimates based on OLS regressions using equation Eq. (1) in Panel A and on 2SLS regressions in Panel B (see Section 6). Standard errors clustered at the slum level are reported in parentheses. *P*-values are presented in brackets. In Panel A, the first value is from individual testing and the second is adjusted for testing that each treatment is jointly different from zero for all outcomes presented in the panel. The dependent variables are: in column (1) *Recommended practices*, an index capturing adherence to WHO's recommendations for protecting against infection, built using the procedure of Kling et al. (2007) described in Section 5; column (2) *Face masks and hand-washing*, the average level of agreement with wearing face masks in crowded places and washing hands with soap more frequently and for longer; column (3) *Physical distancing*, which concerns keeping physical distance from other people; column (4) *Vegetarianism*, the level of agreement with relying on eating a vegetarian diet; column (5) *Indian immunity*, the level of agreement with relying on the Indian immune system. The level of agreement in columns (2)–(5) is measured using a re-scaled Likert scale where 0 refers to strongly disagree and 1 refers to strongly agree. *Effect size (avg. exposure)* rescales the IV estimate to the (estimating sample) average share of the doctor message that is listened to, conditional on picking up the call. All specifications include indicator variables for data collection rounds, and strata indicators (city and religion of respondent).

had no effect on the probability of picking up the call, nor on the share of the message that is listened to. We conclude that it is the combination of religion concordance with relevant informational content that is driving respondents to listen for longer to the information campaign.

We observe that it is a full shift in the distribution of listening time that is driving these results. This highlights the importance of not only the very first seconds of the call, when the sender introduces the message, but also the content that follows the introduction. Panel B in Fig. 3 presents the distribution of the share of each message that is listened to by study participants in the presence of religion concordance or religion discordance for both the control group (left figure) and the doctor message group (right figure). Kolmogorov–Smirnov tests of the equality of the distributions in the presence of religion concordance or discordance is rejected at the 1% confidence level in the control group, and at the 5% confidence level in the doctor message group. If we exclude respondents that listened to the full message, we can still reject equality at the 5% confidence level in both figures of Panel B.

While we do not observe any difference in the distribution for the control message, we observe a difference for the doctor message group.

7.2. Compliance with preventive practices and beliefs about their efficacy

We now turn to compliance with and views about preventive practices. We first focus on the effect of sending doctor messages versus control messages (Table 2) before turning to the impacts of sending a doctor message that is religion-concordant (Table 3). In each Table, Panel A presents ITT estimates and Panel B shows IV estimates of the effect of doctor messages on compliance with preventive practices and on beliefs about their efficacy in fighting COVID-19. In column (1), we focus on compliance with recommended practices using the index that aggregates different indicators of preventive behavior (see Section 5). In columns (2)–(5) we focus on respondents' beliefs over the efficacy of different preventive practices, in columns (2)–(3) on recommended practices and in columns (4)–(5) on non-evidence-based practices, such as relying on vegetarianism or on Indian immunity to the virus.

Sending the doctor messages increases significantly the compliance with recommended practices by 0.05 standard deviations relative to the control group percentage. This effect is driven by increases in both hand-washing and physical distancing (Appendix Section D.3), indicating that doctor messages were effective at promoting recommended practices to avoid contagion.[23]

[23] On average, in the follow-up surveys, respondents in the control group reported that 70% wore a face mask when leaving the house, 73% washed hands frequently, 8% did not leave the slum during the week previous to the interview, 24% did not receive a visit from outside the slum during the week previous to the interview, and 8% did not meet anybody from outside the slum the day before the interview (Appendix Table D4).

Table 3
Preventive practices: the effect of religion concordance in the doctor message.

	Compliance	Beliefs over the efficacy of...			
		Recommended practices		Non-evidence-based practices	
	Recommended practices (1)	Face masks/ hand-washing (2)	Physical distancing (3)	Vegetarianism (4)	Indian immunity (5)
A. OLS					
Religion concordance	0.101	−0.004	−0.006	−0.017	−0.001
	(0.033)	(0.004)	(0.004)	(0.008)	(0.007)
	[0.00 , 0.01]	[0.33 , 0.53]	[0.20 , 0.46]	[0.04 , 0.12]	[0.84 , 0.86]
Mean (religion discordance)	−0.032	0.807	0.806	0.571	0.654
Observations	2544	3851	3849	3846	3849
B. IV					
% listened · religion concordance	0.619	−0.023	−0.035	−0.108	−0.008
	(0.207)	(0.023)	(0.027)	(0.051)	(0.042)
	[0.00]	[0.33]	[0.20]	[0.03]	[0.84]
Mean (religion discordance)	−0.032	0.807	0.806	0.571	0.654
Effect size (avg. exposure)	0.281	−0.010	−0.016	−0.049	−0.004
Observations	2544	3851	3849	3846	3849

Notes. The sample is restricted to respondents in the doctor message group. Estimates based on OLS regressions using Eq. (1) in Panel A and on 2SLS regressions in Panel B (see Section 6). Standard errors clustered at the slum level are reported in parentheses. *P*-values are presented in brackets. In Panel A, the first value is from individual testing and the second is adjusted for testing that each treatment is jointly different from zero for all outcomes presented in the panel. The dependent variables are: in column (1) *Recommended practices*, an index capturing adherence to WHO's recommendations for protecting against infection, built using the procedure of Kling et al. (2007) described in Section 5; column (2) *Face masks and hand-washing*, the average agreement with wearing face masks in crowded places and washing hands with soap more frequently and for longer; column (3) *Physical distancing*, which concerns keeping physical distance with other people; column (4) *Vegetarianism*, the level of agreement with relying on eating a vegetarian diet; column (5) *Indian immunity*, the level of agreement with relying on the Indian immune system. The level of agreement in columns (2)–(5) is measured using a re-scaled Likert scale where 0 refers to strongly disagree and 1 refers to strongly agree. *Effect size (avg. exposure)* rescales the IV estimate to the (estimating sample) average share of the religion-concordant doctor message that is listened to, conditional on picking up the call. All specifications include indicator variables for data collection rounds, and strata indicators (city and religion of respondent).

The increase in compliance with recommended practices is accompanied by changes in beliefs over the efficacy of evidence-based practices. We find a significant increase in agreement with using face masks and practicing hand-washing to protect against the virus of 0.6 percentage points (0.75% over the control mean), while agreement with social distancing also increased, though not significantly. This result may be influenced by the constraints of living in overcrowded spaces, as is the case in the slums where the study was conducted. We do not observe any effect on beliefs over the efficacy of non-evidence based practices at conventional significance levels. Inference for these effects is robust to multiple hypothesis testing at standard confidence levels.[24] Estimates increase significantly in magnitude when considering IV estimates (Panel B). Listening to the full doctor message increases compliance with recommended practices by 0.33 standard deviations and in the level of agreement with using face masks and practicing hand-washing by 4.1 percentage points (corresponding to a 5.1% increase relative to the control mean). When re-scaling the IV coefficient to the (estimating sample) average share of each message that is listened to, conditional on picking up the call (row 'Effect size (avg. exposure)' in the table), the estimated effect sizes are 0.14 standard deviations for compliance and 1.8 percentage points for belief in the efficacy of wearing face-masks and hand-washing. Overall, while the effects on compliance are large in magnitude, the effects on beliefs are either absent or relatively small.

Table 3 focuses on the effect of sending a doctor message that is religion-concordant, compared with one that is religion-discordant. Religion concordance in the doctor message increases the compliance with recommended practices by 0.10 standard deviations. Again, inference for this effect is robust to multiple hypothesis testing. Because this effect is almost double the estimate of the effect of the doctor message, it indicates that the efficacy of doctor messages in promoting compliance is almost wholly driven by messages in which the sender and the receiver have the same religion. This result is confirmed by estimating the effects with an interaction model (Appendix Section D.4). This finding is possibly due to the fact that these receivers listen to a larger proportion of the message (Section 7.1) and/or attach stronger importance to the message. IV estimates indicate a large magnitude of the effect when the whole message is listened to by the receiver, leading to an increase in compliance with recommended practices of 0.63 standard deviations (or 0.29 standard deviations when re-scaled).

As compared with religion discordance, religion concordance in the doctor messages does not alter beliefs over the efficacy of recommended practices, but it does reduce agreement with non-evidence-based practices to some extent. We observe a reduction of 1.7 percentage points in the agreement with vegetarianism being a way to prevent contagion (10.8 percentage points with the IV estimate, and 4.9 percentage points when re-scaled), an effect that is significant only at the 13% level when corrected for multiple

[24] While the effect on beliefs over the efficacy of face masks and hand-washing is stronger in the first follow-up round, compliance with recommended practices is significantly affected in both rounds (Appendix D7).

Table 4
Response to misinformation: doctor versus control message.

	Fact-checking	Agreement with misinformation shared by ...	
		In-group citizens	Out-group citizens
	(1)	(2)	(3)
A. OLS			
Doctor message	−0.022	0.002	−0.015
	(0.010)	(0.007)	(0.006)
	[0.03 , 0.06]	[0.80 , 0.80]	[0.01 , 0.03]
Mean (control message)	0.352	0.485	0.494
Observations	7700	5180	6709
B. IV			
% listened · doctor message	−0.145	0.012	−0.093
	(0.066)	(0.047)	(0.037)
	[0.03]	[0.80]	[0.01]
Mean (not listened)	0.352	0.485	0.494
Effect size (avg. exposure)	−0.061	0.005	−0.040
Observations	7700	5180	6709

Notes. Estimates based on OLS regressions using Eq. (1) in Panel A and on 2SLS regressions in Panel B (see Section 6). Standard errors clustered at the slum level are reported in parentheses. *P*-values are presented in brackets. In Panel A, the first value is from individual testing and the second is adjusted for testing that each treatment is jointly different from zero for all outcomes presented in the panel. The dependent variables are: in column (1) *Fact-checking*, an indicator variable equal to 1 if the respondent always or very frequently check the truthfulness of information shared or discussed with family and friends, and 0 otherwise; in columns (2)–(3) *Agreement with misinformation shared by [...]*, the average level of agreement with statements including incorrect views based on faulty knowledge or understanding, where 0 refers to strongly disagree and 1 refers to strongly agree. In column (2), the outcome variables include only statements from an interlocutor with the same religion as the respondent. In column (3), the outcome variables include only statements from an interlocutor with a religion different from the respondent's or from the generic term "people". Individual statements and categorization are described in Appendix Section A.1. *Effect size (avg. exposure)* rescales the IV estimate to the (estimating sample) average share of the doctor message that is listened to, conditional on picking up the call. All specifications include indicator variables for data collection rounds, and strata indicators (city and religion of respondent).

hypothesis testing. The magnitude of this effect is larger than the effect of doctor messages alone, as it corresponds to a reduction in beliefs that vegetarianism is effective protection of 3.0% over the mean for religion-discordant doctor messages.

These effects are not driven by changes in perceptions of the risk of contagion with COVID-19, which is unaffected by the interventions (Appendix Section D.6). In addition, the placebo test confirms that the effects of religion concordance on beliefs are specific to the doctor message; i.e. similarly to the case of the take-up of messages, we observe no differential effect of religion concordance in the control group (Appendix Section D.2).

Although the campaign influences behavior by shaping beliefs about the efficacy of recommended practices, we also find that it is largely ineffective in altering beliefs regarding non-evidence-based preventive practices. Beliefs in these unproven practices continue to persist among the study population.

We rule out that these effects are driven by more social desirability bias in the treatment group. Although individuals with a strong tendency toward social desirability may show more endorsement for recommended practices or widespread beliefs, we demonstrate that this pattern is not more pronounced in the treatment group compared to the control group. In Appendix Figure D5 we show that the treatment effects on self-reported compliance and beliefs are of similar magnitude for respondents with a low versus high propensity for social desirability bias. This test serves as a crucial validation of our findings, as it enables us to assess bias across all outcomes.

7.3. Response to misinformation

Our results indicate that the campaign's informative content more effectively shifted compliance behaviors when the sender and receiver shared a religious proximity. In this section, we focus on whether the campaign was also effective at protecting against misinformation.

We begin by studying whether sending doctor messages is more effective at achieving this than sharing gossip in the control message. Table 4 presents the results. In column (1), fact-checking is measured as an indicator variable equal to 1 if the respondent always or very frequently checks the truthfulness of the information he/she shares or discusses, and 0 otherwise. In columns (2)–(3), we focus on the level of agreement with misinformation shared by in-group citizens and by out-group citizens. The elicitation procedure for these outcomes is described in Section 5. Panel A presents ITT estimates, while Panel B provides IV estimates of LATE effects.

Sending messages from a trusted source, in this case doctors, crowds out fact-checking.[25] The inclination of respondents to verify information shared by and discussed with family and friends decreases significantly by 2.2 percentage points (6.3% over the control mean), and this decrease remains significant after adjusting p-values for multiple hypothesis testing. This effect, which is driven by the second round of data collection (Appendix Table D8), translates into a LATE estimate of 14.5 percentage points across survey rounds when the respondent listens to the full doctor message (or 41.2% over the control mean) or 6.1 percentage points when considering the average time listened to the message. Perhaps this crowding-out happens because individuals, having heard the messages from doctors, feel more confident in dismissing misinformation. Reductions in fact checking following the doctor message are slightly higher, but not robust to multiple hypothesis testing, when respondents were incentivized with the higher lottery amount (Appendix Table D9), potentially driven by participants paying closer attention to the campaign.

In terms of agreement with misinformation shared by other citizens, we observe that doctor messages do not impact this dimension when shared by in-group citizens, but they do reduce agreement when misinformation is shared by out-group citizens. Doctor messages lead to a significant reduction in agreement of 1.6 percentage points (an effect of 3.2% over the control mean) when the statement is made by a citizen of a different religion. This effect is robust to multiple hypothesis testing and corresponds to a LATE estimate of 10.0 percentage point reduction when the doctor message is listened to fully, corresponding to a reduction of 20.2% over the control mean, or a reduction of 4.3 percentage points when considering the average listening time. Thanks to the design of the survey instrument, the content of the statements used to measure how respondents react to misinformation is orthogonal to the citizen sharing it being in- or out-group (i.e. statements are constant, while the citizen varies exogenously). These results highlight how measuring impacts on the response to misinformation requires consideration of social norms and group identity. The information campaign promoted by the doctor messages is effective at protecting against misinformation carrying little group identity (i.e. shared by citizens of a different religious group). However, doctor messages alone are ineffective at protecting against misinformation carrying stronger group identity (i.e. shared by citizens of the same religious group as the respondent).

We now turn attention to whether the doctor message has religion concordance or discordance impacts the response to misinformation. Table 5 provides estimates of the effects by restricting the sample to recipients of the doctor message. Religion concordance does not introduce, on average, any significant differential effect for fact-checking or for agreement with misinformation shared by out-group citizens. However, religion concordance protects against misinformation shared by citizens with the same identity. While the doctor message decreases agreement with misinformation shared by out-group citizens, it is only in the presence of religion concordance that a doctor message also influences agreement with misinformation reported by in-group citizens. When the doctor message is introduced by a religion-concordant greeting, agreement with this type of misinformation is reduced by 2.3 percentage points compared with a doctor message introduced by a religion-discordant greeting (corresponding to a decrease of 4.6% over the mean for religion-discordant messages). This effect is highly significant and robust to multiple hypothesis testing. The magnitude of the LATE estimate is a reduction of 14.9 percentage points in agreement after listening fully to a doctor message introduced by a religion-concordant greeting (corresponding to a decrease by 30.0% over the mean for religion-discordant messages) or 6.7 percentage points when considering the average listening time. In contrast, religion concordance does not further shift disagreement with misconceptions reported by out-group citizens compared with the main effect provided by the informative content.

Similar to the effects presented in Sections 7.1 and 7.2, these effects are specific to the combination of religion concordance with a doctor message, as receiving the religion-concordant greeting with the control message does not affect agreement with any of the variables presented in Table 5 (Appendix Section D.2).

To verify whether these effects are specific to misinformation, we present a placebo test by estimating treatment effects on agreement with a different type of statement shared by citizens. We focus on opinions related to COVID-19 rather that misinformation because these are harder to influence via fact-checking (e.g. Walter and Salovich, 2021). Appendix Table D10 shows that doctor messages, with or without religion concordance, have no impact on agreement with opinions, independently from whether these are reported by an in-group interlocutor or an out-group interlocutor. This finding reinforces that the pattern of effects observed is specific to misinformation about COVID-19. It also suggests that the limited effectiveness of the information campaign in influencing beliefs over the efficacy of non-evidence-based practices might be related to non-factual opinions, which are more persistent and harder to influence by information campaigns.

In summary, protection against misinformation can be more effectively achieved through informative content shared by sources that are trusted. However, in order to fully safeguard against misinformation and break the connection between beliefs and group identity, we need to factor in religious proximity in information campaigns. Only in the presence of religion concordance is the agreement with misinformation shared by both in- and out-group citizens reduced by the campaign's informative content. In the absence of concordance, the reduction in agreement with misconceptions occurs solely when the misinformation originates from out-group citizens.

We rule out that these effects are driven by more social desirability bias in the treatment group. In Appendix Figure D5 we show that the treatment effects on agreement with misinformation (both shared by in-group and out-group citizens) are of similar magnitude for respondents with a low versus high propensity for social desirability bias.

[25] The interventions have no effect on the level of trust. We do not find any effect on reported levels of trust in information shared by different groups, including doctors and health experts and other citizens of UP (Appendix Section D.6).

Table 5
Response to misinformation: the effect of religion concordance in the doctor message.

	Fact-checking	Agreement with misinformation shared by ...	
		In-group citizens	Out-group citizens
	(1)	(2)	(3)
A. OLS			
Religion concordance	0.006	−0.026	0.007
	(0.015)	(0.009)	(0.008)
	[0.69 , 0.69]	[0.00 , 0.02]	[0.40 , 0.65]
Mean (religion discordance)	0.326	0.498	0.476
Observations	3851	2588	3341
B. IV			
% listened · religion concordance	0.037	−0.169	0.043
	(0.091)	(0.059)	(0.051)
	[0.69]	[0.00]	[0.40]
Mean (religion discordance)	0.326	0.498	0.476
Effect size (avg. exposure)	0.017	−0.076	0.020
Observations	3851	2588	3341

Notes. The sample is restricted to respondents in the doctor message group. Estimates based on OLS regressions using Eq. (1) in Panel A and on 2SLS regressions in Panel B (see Section 6). Standard errors clustered at the slum level are reported in parentheses. *P*-values are presented in brackets. In Panel A, the first value is from individual testing and the second is adjusted for testing that each treatment is jointly different from zero for all outcomes presented in the panel. The dependent variables are: in column (1) *Fact-checking*, an indicator variable equal to 1 if the respondent always or very frequently check the truthfulness of information shared or discussed with family and friends, and 0 otherwise; in columns (2)–(3) *Agreement with misinformation shared by [...]*, the average level of agreement with statements including incorrect views based on faulty knowledge or understanding, where 0 refers to strongly disagree and 1 refers to strongly agree. In column (2), the outcome variables include only statements from an interlocutor with the same religion as the respondent. In column (3), the outcome variables include only statements from an interlocutor with a religion different from the respondent's or from the generic term "people". Individual statements and categorization are described in Appendix Section A.1. *Effect size (avg. exposure)* rescales the IV estimate to the (estimating sample) average share of the religion-concordant doctor message that is listened to, conditional on picking up the call. All specifications include indicator variables for data collection rounds, and strata indicators (city and religion of respondent).

Table 6
Religion concordance and religious affiliation.

	Compliance	Beliefs over the efficacy of ...				Fact-checking	Agreement with Misinformation shared by ...	
		Recommended practices		Non-evidence-based practices				
	Recommended practices	Face masks /hand-washing	Physical distancing	Vegetarianism	Indian immunity		In-group citizens	Out-group citizens
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Religion concordance	0.104	0.001	−0.003	−0.018	0.005	0.003	−0.023	0.007
	(0.036)	(0.004)	(0.005)	(0.009)	(0.007)	(0.016)	(0.010)	(0.009)
	[0.00 , 0.02]	[0.87 , 0.87]	[0.55 , 0.77]	[0.05 , 0.19]	[0.53 , 0.88]	[0.84 , 0.85]	[0.03 , 0.07]	[0.48 , 0.73]
x Muslim respondent	−0.012	−0.021	−0.013	0.000	−0.029	0.010	−0.020	0.000
	(0.079)	(0.010)	(0.013)	(0.019)	(0.015)	(0.035)	(0.022)	(0.020)
	[0.88 , 0.99]	[0.04 , 0.20]	[0.31 , 0.68]	[0.99 , 0.99]	[0.05 , 0.19]	[0.76 , 0.94]	[0.37 , 0.75]	[0.98 , 0.98]
Muslim respondent	0.052	0.016	−0.004	−0.074	−0.009	−0.134	−0.055	−0.030
	(0.112)	(0.015)	(0.017)	(0.023)	(0.021)	(0.043)	(0.030)	(0.021)
	[0.65 , 0.94]	[0.27 , 0.72]	[0.80 , 0.81]	[0.00 , 0.02]	[0.68 , 0.89]	[0.00 , 0.00]	[0.07 , 0.14]	[0.16 , 0.18]
Mean (religion discordance)	−0.032	0.807	0.806	0.571	0.654	0.326	0.498	0.476
Observations	2544	3851	3849	3846	3849	3851	2588	3341

Notes. Estimates based on OLS regressions using Eq. (2) restricted to the doctor message group, adding an interaction term between the religion concordance indicator and an indicator variable for whether the respondent is Muslim. Standard errors clustered at the slum level are reported in parentheses. *P*-values are presented in brackets. The first value is from individual testing and the second is adjusted for testing that each variable is jointly different from zero for all outcomes grouped according to Tables 3 and 5. Dependent variables in columns (1)–(5) are defined in Table 3, while dependent variables in columns (6)–(8) are defined in Table 5. All specifications include indicator variables for data collection rounds, and strata indicators (city and religion of respondent).

7.4. Religion concordance and religious affiliation

This section complements the results on the effect of religion concordance in the doctor message discussed in Sections 7.2 and 7.3 by looking at whether the effect of concordance varies by whether a respondent is Muslim or Hindu.

Table 6 replicates the estimation in Panel A of Tables 3 and 5, but introducing in Eq. (2) an interaction term between the religion concordance indicator and an indicator variable for whether the respondent is Muslim, to capture heterogeneity in the effect.

We highlight significant differences between Muslim and Hindu respondents in beliefs over the efficacy of vegetarianism and in agreement with misinformation. On average, Muslim respondents tend to have significantly less agreement with these dimensions. Compliance with and beliefs over the efficacy of recommended practices are instead comparable across Muslim and Hindu respondents. This suggests the existence of significant differences in the beliefs of Muslim and Hindu respondents, but primarily for topics with religious salience.

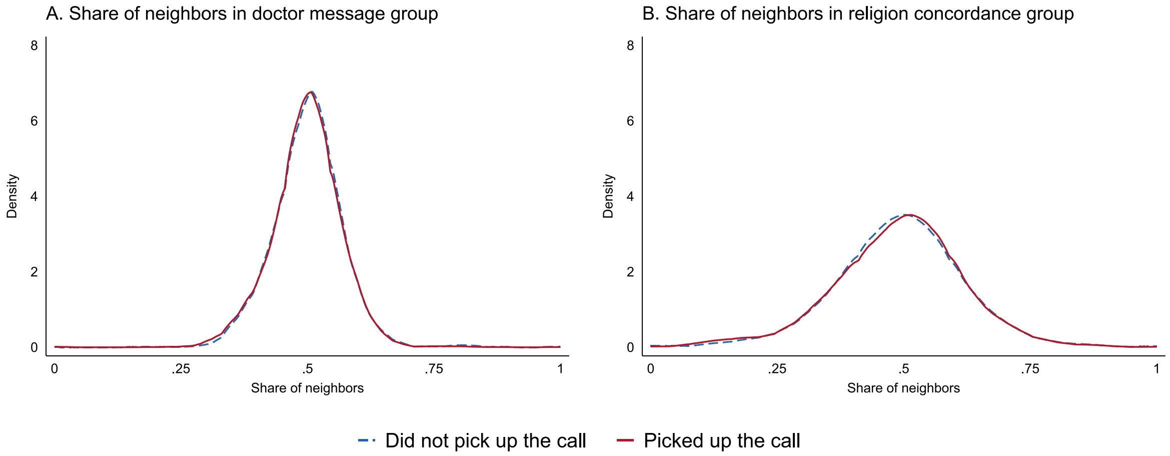

Fig. 4. Treatment allocation among neighbors, by respondent's group.
Notes. The figures show the distribution of the share of households living in the same slum as the respondent that are allocated to the doctor message group (Panel A) or to the religion concordance group (Panel B), depending on whether the respondent picked up or did not pick up the intervention call. In Panel B, the sample is restricted to the doctor message group. Distributions are estimated non-parametrically using kernel density estimation, assuming an Epanechnikov kernel function with a bandwidth of 0.02 in Panel A and 0.04 in Panel B. The *p*-values of Kolmogorov–Smirnov tests for equality of distributions are 0.18 in Panel A and 0.75 in Panel B.

Looking at heterogeneity in the effect of religion concordance in the doctor message, the results indicate that the effect is relatively homogeneous across Muslim and Hindu respondents. In fact, the effect of religion concordance is statistically different across religions only for beliefs over the efficacy of face masks and hand-washing. In this case, the effect for Muslim respondents is 2.0 percentage points lower than for Hindu respondents.

Further analysis of heterogeneity in the effect of both the doctor message and religion concordance in the doctor message is presented in Appendix D.1.

7.5. *Information spillovers*

Randomization into the experimental arms is conducted at the household level because the intervention is directed one-to-one through mobile phones and we wanted to prevent informational spillovers within households. The interpretation of the estimates of treatment effects discussed in Sections 7.1–7.4 would be affected by the presence of information spillovers (e.g. Vazquez-Bare, 2022). Spillover effects are mitigated by the voice messages being automatic calls that cannot be forwarded or shared, but there remains the possibility of word-of-mouth information sharing, particularly within communities.

While this study was not specifically designed to capture spillover effects in information campaigns, we test for their presence by leveraging variation in intervention exposure across slums, induced by the household-level randomization not being stratified at slum level. The availability of precise geo-location of slum borders, as well as where each household resides, allows us to measure the share of households living in the same slum as the respondent that is allocated to the doctor message group and, conditional on being allocated to the doctor message group, the share that also receives a religion-concordant message.[26] By design, the probability of neighbors being in each of these groups is on average 0.5. However, household-level randomization allows for random variation in this probability across respondents. Fig. 4 shows the distribution of these variables by whether the respondent took-up the intervention. The distributions confirm not only the random pattern of treatment allocation among neighbors, but also the similarity of this pattern across respondents that did and did not pick up the intervention call. Kolmogorov–Smirnov tests fail to reject the equality of the distributions along this dimension.

Exploiting this variation, we estimate both Eqs. (1) and (2) controlling for this measure of 'neighbor treatment'. Rejecting the null hypothesis of a zero coefficient for this measure indicates the presence of information spillovers. Table 7 presents the results. The estimates of treatment effects discussed in Sections 7.1–7.3 are unaffected by controlling for the treatment allocation among neighbors. In addition, the effect of treatment allocation among neighbors is not statistically significantly different from zero for most of the outcomes, indicating limited importance of community-level information sharing. These results highlight that, despite interventions having the potential to spread information across individuals, community-level spillovers do not play a central role, and we can interpret our main results as consistent estimates of causal effects of the intervention campaign.

[26] Results using the treatment allocation of the nearest neighbor are in line (Appendix Section D.8).

Table 7
Spillover effects of the interventions.

	Compliance	Beliefs over the efficacy of ...				Fact-checking	Agreement with Misinformation shared by ...	
		Recommended practices		Non-evidence-based practices				
	Recommended practices	Face masks /hand-washing	Physical distancing	Vegetarianism	Indian immunity		In-group citizens	Out-group citizens
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Full sample								
Doctor message	0.057	0.008	0.006	0.003	−0.006	−0.022	0.003	−0.012
	(0.022)	(0.003)	(0.004)	(0.007)	(0.005)	(0.011)	(0.008)	(0.006)
	[0.01 , 0.05]	[0.01 , 0.03]	[0.17 , 0.45]	[0.68 , 0.69]	[0.28 , 0.49]	[0.05 , 0.14]	[0.68 , 0.67]	[0.07 , 0.13]
Doctor message (% neighbors)	0.268	0.068	0.048	−0.015	0.074	0.014	0.062	0.126
	(0.206)	(0.043)	(0.053)	(0.083)	(0.070)	(0.146)	(0.088)	(0.087)
	[0.19 , 0.56]	[0.11 , 0.41]	[0.36 , 0.61]	[0.86 , 0.86]	[0.29 , 0.61]	[0.93 , 0.93]	[0.48 , 0.75]	[0.15 , 0.38]
Mean (control message)	−0.032	0.799	0.799	0.563	0.661	0.352	0.485	0.494
Observations	5125	7700	7698	7692	7697	7700	5180	6709
B. Sample restricted to doctor message group								
Religion concordance	0.109	−0.004	−0.005	−0.018	−0.002	0.005	−0.028	0.005
	(0.034)	(0.004)	(0.005)	(0.008)	(0.007)	(0.015)	(0.009)	(0.008)
	[0.00 , 0.01]	[0.35 , 0.59]	[0.24 , 0.54]	[0.03 , 0.11]	[0.83 , 0.84]	[0.74 , 0.75]	[0.00 , 0.01]	[0.53 , 0.78]
Religion concordance (% neighbors)	0.270	−0.004	0.004	−0.037	−0.010	−0.037	−0.082	−0.063
	(0.161)	(0.029)	(0.035)	(0.048)	(0.051)	(0.081)	(0.052)	(0.053)
	[0.09 , 0.36]	[0.88 , 0.98]	[0.91 , 0.92]	[0.43 , 0.86]	[0.85 , 1.00]	[0.65 , 0.65]	[0.12 , 0.32]	[0.23 , 0.42]
Mean (religion discordance)	−0.032	0.807	0.806	0.571	0.654	0.326	0.498	0.476
Observations	2544	3851	3849	3846	3849	3851	2588	3341

Notes. Estimates based on OLS regressions using Eq. (1) in Panel A and Eq. (2) restricted to the doctor message group in Panel B (see Section 6). *% neighbors* is the share of households living in the same slum as the respondent that are allocated to the corresponding group. Standard errors clustered at the slum level are reported in parentheses. *P*-values are presented in brackets. The first value is from individual testing and the second is adjusted for testing that each variable is jointly different from zero for all outcomes grouped according to Tables 2 and 4. Dependent variables in columns (1)–(5) are defined in Table 2, while dependent variables in columns (6)–(8) are defined in Table 4. All specifications include indicator variables for data collection rounds, and strata indicators (city and religion of respondent).

8. Conclusions

We demonstrate that a physician-delivered information campaign promoting health-related preventive practices among slum dwellers in India is effective at improving compliance with recommended practices and beliefs about their efficacy. Importantly, we show that the campaign's efficacy is primarily driven by religious proximity between the sender and the receiver of information, and that this religion concordance helps to protect individuals against misinformation.

These findings open new avenues for future research to explore both the effectiveness of information campaigns and the role of social proximity in decision-making. In particular, while the novelty of our study is to focus on religion, future research could delve into the relative effectiveness of different dimension of social proximity as well as their interaction, and could be tested as a tool to counter medical mistrust, which can be particularly strong within specific sub-populations (Alsan and Wanamaker, 2018; Jaiswal and Halkitis, 2019). Understanding how various social factors influence information dissemination can more comprehensively guide the design of information campaigns. It is also important to understand whether tailoring messages and leveraging social proximity to delivery them could lead to unintended consequences in the longer term, such as segregation between communities, particularly if the dimensions chosen are the basis of tensions.

Understanding how social proximity interacts with information campaigns and health-related behaviors offers opportunities for targeted policy interventions. Policymakers can leverage these insights to create more effective and culturally attuned campaigns, thereby enhancing public health outcomes across diverse communities. In particular, policymakers should consider incorporating religious proximity into the design of information campaigns, ensuring that messages resonate based on the audience's identity. At the same time, in light of the ongoing challenges posed by misinformation, policy interventions should be aimed not only at disseminating accurate information but also at effectively countering misinformation.

While our evidence suggests that such a light-touch intervention has limited positive externalities, it remains a very low-cost intervention. The cost of setting up the intervention (including recording and editing the message, interactive voice response set-up fees, and monthly rental) was less than US$600, the smaller incentive cost US$32 for 1000 respondents, and sending two messages cost US$0.028 per respondent at the time of the experiment, totaling US$0.21 per participant. Additional costs to consider are staff to set up the message sending.

This study underscores the potential of mobile-based campaigns as effective tools in low-income areas, offering scalable and low-cost methods for widespread information dissemination.

CRediT authorship contribution statement

Alex Armand: Writing – review & editing, Writing – original draft, Validation, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Britta Augsburg:** Writing – review & editing, Writing – original draft, Validation, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Data curation, Conceptualization. **Antonella Bancalari:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Kalyan Kumar Kameshwara:** Validation, Software, Project administration, Investigation, Formal analysis.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jhealeco.2024.102883.

References

Akbar, S.Z., Kukreti, D., Sagarika, S., Pal, J., 2020. Increase in communal tone in COVID-19 fake news in India: Study. In: The Quint.

Akerlof, G.A., Kranton, R.E., 2000. Economics and identity. Q. J. Econ. 115 (3), 715–753.

Al Jazeera, 2020. Misinformation, fake news spark India coronavirus fears. Al Jazeera news, URL https://www.aljazeera.com/news/2020/03/misinformation-fake-news-spark-india-coronavirus-fears-200309051731540.html. (Accessed 13 July 2021).

Al-Zaman, S., 2021. COVID-19-related social media fake news in India. J. Media 2, 100–114.

Allcott, H., Gentzkow, M., 2017. Social media and fake news in the 2016 election. J. Econ. Perspect. 31 (2), 211–236.

Alsan, M., Eichmeyer, S., 2021. Experimental evidence on the effectiveness of non-experts for improving vaccine demand. National Bureau of Economic Research, Working Paper 28593.

Alsan, M., Garrick, O., Graziani, G., 2019. Does diversity matter for health? Experimental evidence from oakland. Amer. Econ. Rev. 109 (12), 4071–4111.

Alsan, M., Stanford, F.C., Banerjee, A., Breza, E., Chandrasekhar, A.G., Eichmeyer, S., 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: a randomized controlled trial. Ann. Int. Med. 174 (4), 484–492.

Alsan, M., Stantcheva, S., Yang, D., Cutler, D., 2020. Disparities in coronavirus 2019 reported incidence, knowledge, and behavior among US adults. JAMA Netw. Open 3 (6), e2012403.

Alsan, M., Wanamaker, M., 2018. Tuskegee and the health of black men. Q. J. Econ. 133 (1), 407–455.

Armand, A., Augsburg, B., Bancalari, A., 2020. COVID-19, fake news and religious tensions: experimental evidence from India. AEA RCT Registry, http://dx.doi.org/10.1257/rct.6564-1.0, October 08.

Armand, A., Augsburg, B., Bancalari, A., Ghatak, M., 2023. Public Service Delivery, Exclusion and Externalities: Theory and Experimental Evidence from India. Centre for Economic Policy Research Discussion Paper DP18636.

Atkin, D., Colson-Sihra, E., Shayo, M., 2021. How do we choose our identity? A revealed preference approach using food consumption. J. Polit. Econ. 129 (4), 1193–1251.

Awasthi, P., 2020. UP seals off hotspots in 15 districts, including Lucknow, Noida to contain spread. The Week April 08.

Azrieli, Y., Chambers, C.P., Healy, P.J., 2018. Incentives in experiments: A theoretical analysis. J. Polit. Econ. 126 (4), 1472–1503.

Badrinathan, S., 2021. Educative interventions to combat misinformation: Evidence from a field experiment in India. Am. Political Sci. Rev. 115 (4), 1325–1341.

Baicker, K., Mullainathan, S., Schwartzstein, J., 2015. Behavioral hazard in health insurance. Q. J. Econ. 130 (4), 1623–162.

Banaji, S., Bhat, R., 2020. WhatsApp vigilantes: An exploration of citizen reception and circulation of WhatsApp misinformation linked to mob violence in India. LSE Media Commun. 2, 1–14.

Banerjee, A., Alsan, M., Breza, E., Chandrasekhar, A.G., Chowdhury, A., Duflo, E., Goldsmith-Pinkham, P., Olken, B.A., 2020. Messages on COVID-19 Prevention in India Increased Symptoms Reporting and Adherence to Preventive Behaviors Among 25 Million Recipients with Similar Effects on Non-recipient Members of Their Communities. Working Paper Series 27496.

Banerjee, P., Seth, R., Dhaliwal, B.K., Sullivan, A., Yawar, Q., Thankachen, B., Closser, S., Shet, A., 2022. Vaccine acceptance in rural India: Engaging faith leaders as vaccine ambassadors. Front. Public Health (10), 979424.

Barrera, O., Guriev, S., Henry, E., Zhuravskaya, E., 2020. Facts, alternative facts, and fact checking in times of post-truth politics. J. Public Econ. 182, 104123.

Bavel, J.J.V., Baicker, K., Boggio, P.S., Capraro, V., Cichocka, A., Cikara, M., Crockett, M.J., Crum, A.J., Douglas, K.M., Druckman, J.N., Drury, J., Dube, O., Ellemers, N., Finkel, E.J., Fowler, J.H., Gelfand, M., Han, S., Haslam, S.A., Jetten, J., Kitayama, S., Mobbs, D., Napper, L.E., Packer, D.J., Pennycook, G., Peters, E., Petty, R.E., Rand, D.G., Reicher, S.D., Schnall, S., Shariff, A., Skitka, L.J., Smith, S.S., Sunstein, C.R., Tabri, N., Tucker, J.A., van der Linden, S., van Lange, P., Weeden, K.A., Wohl, M.J., Zaki, J., Zion, S.R., Willer, R., 2020. Using social and behavioural science to support COVID-19 pandemic response. Nat. Hum. Behav. 4 (5), 460–471.

Bazzi, S., Gaduh, A., Rothenberg, A.D., Wong, M., 2019. Unity in diversity? How intergroup contact can foster nation building. Amer. Econ. Rev. 109 (11), 3978–4025.

Belloni, A., Chernozhukov, V., Hansen, C., 2013. Inference on treatment effects after selection among high-dimensional controls. Rev. Econ. Stud. 81 (2), 608–650.

Benjamin, D.J., Choi, J.J., Fisher, G., 2016. Religious identity and economic behavior. Rev. Econ. Stat. 98 (4), 617–637.

Bentzen, J.S., 2021. In crisis, we pray: Religiosity and the COVID-19 pandemic. J. Econ. Behav. Organ. 192, 541–583.

Bhalotra, S., Clots-Figueras, I., Cassan, G., Iyer, L., 2014. Religion, politician identity and development outcomes: Evidence from India. J. Econ. Behav. Organ. 104, 4–17.

Bicchieri, C., Dimant, E., Gaechter, S., Nosenzo, D., 2022. Social proximity and the erosion of norm compliance. Games Econom. Behav. 132, 59–72.

Bowles, J., Larreguy, H., Liu, S., 2020. Countering misinformation via WhatsApp: Preliminary evidence from the COVID-19 pandemic in Zimbabwe. PLoS One 15 (10), e0240005.

Breza, E., Stanford, F.C., Alsan, M., Alsan, B., Banerjee, A., Chandrasekhar, A.G., Eichmeyer, S., Glushko, T., Goldsmith-Pinkham, P., Holland, K., et al., 2021. Effects of a large-scale social media advertising campaign on holiday travel and COVID-19 infections: a cluster randomized controlled trial. Nat. Med. 1–7.

Bursztyn, L., Rao, A., Roth, C., Yanagizawa-Drott, D., 2023. Opinions as facts. Rev. Econ. Stud. 90 (4), 1832–1864.

Chen, Y., Li, S.X., 2000. Group identity and social preferences. Am. Econ. Rev. 99 (1), 431–457.

Crowne, D.P., Marlowe, D., 1960. A new scale of social desirability independent of psychopathology. J. Consult. Psychol. 4 (24), 349–354.

DeFranza, D., Lindow, M., Harrison, K., Mishra, A., Mishra, H., 2021. Religion and reactance to COVID-19 mitigation guidelines. Am. Psychol. 76 (5), 744–754.

DellaVigna, S., Kaplan, E., 2007. The fox news effect: Media bias and voting. Q. J. Econ. 122 (3), 1187–1234.

Dupas, P., 2011. Health behavior in developing countries. Ann. Rev. Econ. 3 (1), 425–449.

Enikolopov, R., Makarin, A., Petrova, M., 2020. Social media and protest participation: Evidence from Russia. Econometrica 88 (4), 1479–1514.

Fincher, C.L., Thornhill, R., 2012. Parasite-stress promotes in-group assortative sociality: The cases of strong family ties and heightened religiosity. Behav. Brain Sci. 35 (2), 61–79.

Fischer, D.G., Fick, C., 1993. Measuring social desirability: Short forms of the marlowe-crowne social desirability scale. Edu. Psychol. Measur. 53 (2), 417–424.

Fisman, R., Paravisini, D., Vig, V., 2017. Cultural proximity and loan outcomes. Amer. Econ. Rev. 107 (2), 457–492.

Fisman, R., Sarkar, A., Skrastins, J., Vig, V., 2020. Experience of communal conflicts and intergroup lending. J. Polit. Econ. 128 (9), 3346–3375.

Fowler, F.J., Mangione, T.W., 1990. Standardized survey interviewing: Minimizing interviewer-related error. Sage, Newbury Park.

Government of India, 2011. Census of India 2011.

Government of India, 2021. Orders prescribing lockdown for containment of COVID-19 epidemic in the country and guidelines for phased re-opening (unlocks). In: Retrieved from MHA Press Releases. Last accessed 2022.

Greenwood, B.N., Carnahan, S., Huang, L., 2018. Patient–physician gender concordance and increased mortality among female heart attack patients. Proc. Natl. Acad. Sci. USA 115 (34), 8569–8574.

Greenwood, B.N., Hardeman, R.R., Huang, L., Sojourner, A., 2020. Physician-patient racial concordance and disparities in birthing mortality for newborns. Proc. Natl. Acad. Sci. USA 117 (35), 21194–21200.

Haaland, I., Roth, C., Wohlfart, J., 2023. Designing information provision experiments. J. Econ. Lit. 61 (1), 3–40.

Habyarimana, J., Humphreys, M., Posner, D.N., Weinstein, J.M., 2007. Why does ethnic diversity undermine public goods provision? Am. Political Sci. Rev. 101 (4), 709–725.

Hays, R.D., Hayashi, T., Stewart, A.L., 1989. A five-item measure of socially desirable response set. Edu. Pholog. Measur. 49 (3), 629–636.

Hill, A., Jones, D., Woodworth, L., 2020. Physician-patient race-match and patient outcomes. (ISSN: 1556-5068) SSRN Working Paper.

Himelein, K., Eckman, S., Lau, C., McKenzie, D., 2020. Mobile phone surveys for understanding COVID-19 impacts: Part ii response, quality, and questions. In: Development Impact Blog. World Bank, https://blogs.worldbank.org/impactevaluations/mobile-phone-surveys-understanding-covid-19-impacts-part-i-sampling-and-mode. (Accessed 13 July 2021).

Imbens, G.W., Angrist, J.D., 1994. Identification and estimation of local average treatment effects. Econometrica 62 (2), 467–475.

Iyer, S., 2016. The new economics of religion. J. Econ. Lit. 54 (2), 395–411.

Jaiswal, J., Halkitis, P.N., 2019. Towards a more inclusive and dynamic understanding of medical mistrust informed by science. J. Behav. Med. 45 (2), 79–85.

Jha, S., 2013. Trade, institutions, and ethnic tolerance: Evidence from South Asia. Am. Political Sci. Rev. 107 (4), 806–832.

Khan, A., Nasim, S., Shaukat, M.I., Stegmann, A., 2021. Building trust in the state with information: Evidence from urban punjab. J. Public Econ. 202, 104494.

Kling, J.R., Liebman, J.B., Katz, L.F., 2007. Experimental analysis of neighborhood effects. Econometrica 75 (1), 83–119.

Kremer, M., Rao, G., Schilbach, F., 2019. Behavioral development economics. In: Handbook of Behavioral Economics: Applications and Foundations 1, vol. 2, Elsevier, pp. 345–458.

Laborde, C., 2021. Minimal secularism: Lessons for, and from, India. Am. Political Sci. Rev. 115 (1), 1–13.

Lazer, D.M., Baum, M.A., Benkler, Y., Berinsky, A.J., Greenhill, K.M., Menczer, F., Metzger, M.J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S.A., Sunstein, C.R., Thorson, E.A., Watts, D.J., Zittrain, J.L., 2018. The science of fake news: Addressing fake news requires a multidisciplinary effort. Science 359 (6380), 1094–1096.

Lilford, R., Oyebode, O., Satterthwaite, D., Melendez-Torres, G., Chen, Y.-F., Mberu, B., et al., 2017. Improving the health and welfare of people who live in slums. Lancet 389 (10068), 559–570.

List, J.A., Shaikh, A.M., Xu, Y., 2019. Multiple hypothesis testing in experimental economics. Exper. Econ. 22 (4), 773–793.

Lowe, M., 2021. Types of contact: A field experiment on collaborative and adversarial caste integration. Amer. Econ. Rev. 111 (6), 1807–1844.

Mahapatra, S., Plagemann, J., 2019. Polarisation and politicisation: The social media strategies of Indian political parties, 3.

Manacorda, M., Tesei, A., 2020. Liberation technology: Mobile phones and political mobilization in africa. Econometrica 88 (2), 533–567.

McKenzie, D., 2012. Beyond baseline and follow-up: The case for more T in experiments. J. Dev. Econ. 99 (2), 210–221.

Menon, S., 2020. Coronavirus: The human cost of fake news in India. In: BBC News. BBC Reality Check, URL https://www.bbc.co.uk/news/world-asia-india-53165436. (Accessed 14 March 24).

Mitra, A., Ray, D., 2014. Implications of an economic theory of conflict: Hindu-muslim violence in India. J. Polit. Econ. 122 (4), 719–765.

Mousa, S., 2020. Building social cohesion between christians and muslims through soccer in post-ISIS Iraq. Science 369 (6505), 866–870.

Narayan, B., 2014. Communal riots in uttar pradesh. Econ. Political Weekly 49 (37), 29–32.

Nyhan, B., 2021. Why the backfire effect does not explain the durability of political misperceptions. Proc. Natl. Acad. Sci. 118 (15), e1912440117.

O'Keefe, D.J., 2016. Persuasion: Theory and Research, third ed. SAGE Publications, Inc, Thousand Oaks, CA.

Patel, A., 2020. Preventing COVID-19 amid public health and urban planning failures in slums of Indian cities. World Med. Health Policy 12 (3), 266–273.

Pauly, M.V., Blavin, F.E., 2008. Moral hazard in insurance, value-based cost sharing, and the benefits of blissful ignorance. J. Health Econ. 27 (6), 1407–1417.

Philpott, D., 2007. Explaining the political ambivalence of religion. Am. Political Sci. Rev. 101 (3), 505–525.

Poonam, S., Bansal, S., 2019. Misinformation is endangering India's election. Atlantic April 01.

Reserve Bank of India, 2019. State-wise poverty rate. National Sample Survey Organisation and NITI Aayog, Delhi, India.

Sadish, D., Adhvaryu, A., Nyshadham, A., 2021. (Mis)information and anxiety: Evidence from a randomized Covid-19 information campaign. J. Dev. Econ. 152, 102699.

Sahgal, N., Evans, J., Salazar, A., Starr, K., Corichi, M., 2021. Religion in India: Tolerance and Segregation. Pew Research Center, Washington D.C..

Sarkar, S., 2020. Religious discrimination is hindering the covid-19 response. Br. Med. J. 369, m2280.

Scheufele, D.A., Krause, N.M., 2019. Science audiences, misinformation, and fake news. Proc. Natl. Acad. Sci. USA 116 (16), 7662–7669.

Siddique, A., Rahman, T., Pakrashi, D., Islam, A., Ahmed, F., 2022. Raising health awareness in rural communities: A randomized experiment in Bangladesh and India. Rev. Econ. Stat. 1–45.

Sinding Bentzen, J., 2019. Acts of god? Religiosity and natural disasters across subnational world districts. Econ. J. 129 (622), 2295–2321.

Solís Arce, J.S., Warren, S.S., Meriggi, N.F., Scacco, A., McMurry, N., Voors, M., Syunyaev, G., Malik, A.A., Aboutajdine, S., Adeojo, O., et al., 2021. COVID-19 vaccine acceptance and hesitancy in low- and middle-income countries. Nat. Med. 27 (8), 1385–1394.

Statista, 2021. Main sources of news across India in 2019, by age.

Tabellini, G., 2008. The scope of cooperation: Values and incentives. Q. J. Econ. 123 (3), 905–950.

Tankard, M.E., Paluck, E.L., 2016. Norm perception as a vehicle for social change. Soc. Issues Policy Rev. 10 (1), 181–211.

Taragin-Zeller, L., Berenblum, T., Brasil, E., Rozenblum, Y., Baram-Tsabari, A., 2023. Religious diversity and public health: Lessons from COVID-19. PLoS One 18 (8), e0290107.

Tibshirani, R., 1996. Regression shrinkage and selection via the LASSO. J. R. Stat. Soc. Ser. B Stat. Methodol. 58 (1), 267–288.

Torres, C., Ogbu-Nwobodo, L., Alsan, M., Stanford, F.C., Banerjee, A., Breza, E., Chandrasekhar, A.G., Eichmeyer, S., Karnani, M., Loisel, T., Goldsmith-Pinkham, P., Olken, B.A., Vautrey, P.L., Warner, E., Duflo, E., 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 Netw. Open 4 (7), e2117115.

Tourangeau, R., Yan, T., 2007. Sensitive questions in surveys. Psychol. Bull. 133, 859–883.

United Nations, 2019. World urbanization prospects: The 2018 revision. In: World Urbanization Prospects: The 2018 Revision. Department of Economic and Social Affairs, Population Division, New York, US.

Van der Linden, S., Leiserowitz, A., Rosenthal, S., Maibach, E., 2017. Inoculating the public against misinformation about climate change. Global Challenges 1 (2), 1600008.

Vazquez-Bare, G., 2022. Identification and estimation of spillover effects in randomized experiments. J. Econometrics 237 (1), 105237.

Walter, N., Salovich, N.A., 2021. Unchecked vs. uncheckable: How opinion-based claims can impede corrections of misinformation. Mass Commun. Soc. 24 (4), 500–526.

Wasdani, K.P., Prasad, A., 2020. The impossibility of social distancing among the urban poor: the case of an Indian slum in the times of COVID-19. Local Environ. 25 (5), 414–418.

Withnall, A., 2020. India coronavirus hospital 'segregating hindu and muslim patients'. In: INDEPENDENT. Independent, April 16.

World Health Organization, 2020. COVID-19 Response. Seventy-Third World Health Assembly (A73.CONF./1 Rev. 1), World Health Organization.

Yasir, S., 2020. India is scapegoating muslims for the spread of the coronavirus. Accessed 13 July 2021.

Yom-Tov, E., Shembekar, J., Barclay, S., Muennig, P., 2018. The effectiveness of public health advertisements to promote health: a randomized-controlled trial on 794,000 participants. NPJ Dig. Med. 1 (24).

Zhuravskaya, E., Petrova, M., Enikolopov, R., 2020. Political effects of the internet and social media. Annu. Rev. Econ. 12, 415–438.