# Surfacing norms to increase vaccine acceptance

**Alex Moehring**[a,b], **Avinash Collis**[b,c,1], **Kiran Garimella**[b,d,1], **M. Amin Rahimian**[b,e,1], **Sinan Aral**[a,b,d], **and Dean Eckles**[a,b,d,2]

[a]MIT Sloan School of Management; [b]MIT Initiative on the Digital Economy; [c]McCombs School of Business, The University of Texas at Austin; [d]Institute for Data, Systems, and Society, Massachusetts Institute of Technology; [e]Department of Industrial Engineering, University of Pittsburgh

Working paper
First version: 8 February 2021
This version: 23 February 2021

Despite the availability of multiple safe vaccines, vaccine hesitancy may present a challenge to successful control of the COVID-19 pandemic. As with many human behaviors, people's vaccine acceptance may be affected by their beliefs about whether others will accept a vaccine (i.e., descriptive norms). However, information about these descriptive norms may have different effects depending on people's baseline beliefs and the relative importance of conformity, social learning, and free-riding. Here, using a large, pre-registered, randomized experiment (N=349,664) embedded in an international survey, we show that accurate information about descriptive norms can substantially increase intentions to accept a vaccine for COVID-19. These positive effects (e.g., reducing by 5% the fraction of people who are "unsure" or more negative about accepting a vaccine) are largely consistent across the 23 included countries, but are concentrated among people who were otherwise uncertain about accepting a vaccine. Providing this normative information in vaccine communications partially corrects individuals' apparent underestimation of how many other people will accept a vaccine. These results suggest that public health communications should present information about the widespread and growing intentions to accept COVID-19 vaccines.

Keywords: COVID-19, descriptive norms, social influence, vaccine hesitancy

[1]A.C., K.G., and M.A.R. contributed equally to this work and are listed alphabetically.
[2]To whom correspondence should be addressed. E-mail: eckles@mit.edu

Nonpharmaceutical interventions in response to epidemics, such as the COVID-19 pandemic, often depend on the behavioral responses of the public for their effectiveness. Even with the availability of vaccines, success depends on people's choices to accept, or even seek out, the vaccine (1), since even low vaccine refusal rates can prevent achieving herd immunity (2, 3). Given the value of individual autonomy and the significant challenges of imposing vaccine mandates (4–6), it is important to understand how public health messaging can increase acceptance of safe and effective COVID-19 vaccines. Many messaging strategies address individual barriers to vaccination, such as complacency and inconvenience (7), as well as perceived risk of both vaccines and the disease (1). However, these strategies may have important limitations; for example, field studies show that corrective information about vaccine safety can effectively reduce misconceptions and false beliefs, though they are not as effective in changing vaccine-related intentions (8, 9). Messaging strategies that share recommendations from experts and emphasize reasons for accepting a vaccine have shown promising effects on increasing acceptance in the United States (10).

It may be important to look beyond individuals to consider how public health messaging can also leverage the significant roles of social networks (broadly defined) in shaping individual vaccination decisions (11–15). Rather than being a small factor, there is growing evidence that people's preventative health behaviors are dramatically influenced by many social and cultural factors, with implications for COVID-19 (16). In the United States, for example, analyses of mobility data during the COVID-19 pandemic revealed that people's mobility behaviors vary with their partisan affiliation (17) and media consumption (18, 19) and are affected by the behaviors of their social connections (20).

Acceptance of COVID-19 vaccines will likely involve substantial social influence, but it is not clear whether learning that others' are accepting a vaccine will increase or decrease acceptance. Positive peer effects can arise due to information diffusion (21, 22), conformity and injunctive norms (14), inferring vaccine safety and effectiveness from others' choices (23, 24), or pro-social motivations such as altruism (25, 26) and reciprocity (27). On the other hand, negative effects of others' acceptance can arise as a result of free-riding on vaccine-generated herd immunity, even if only partial or local (28, 29). The empirical evidence on when positive peer effects (24, 30, 31) or free-riding may dominate (28) is inconclusive. Furthermore, the effects of incorporating truthful information about others' into messaging strategies will depend on what that information is — how prevalent is vaccine acceptance in a given reference group? Thus, we need further empirical guidance about scalable and effective messaging strategies leveraging social influence. That is, some interpretations of the theoretical and empirical literature could motivate emphasizing high rates of vaccine acceptance in public health communications, little is known about how realistic interventions of this kind will affect intentions to accept new

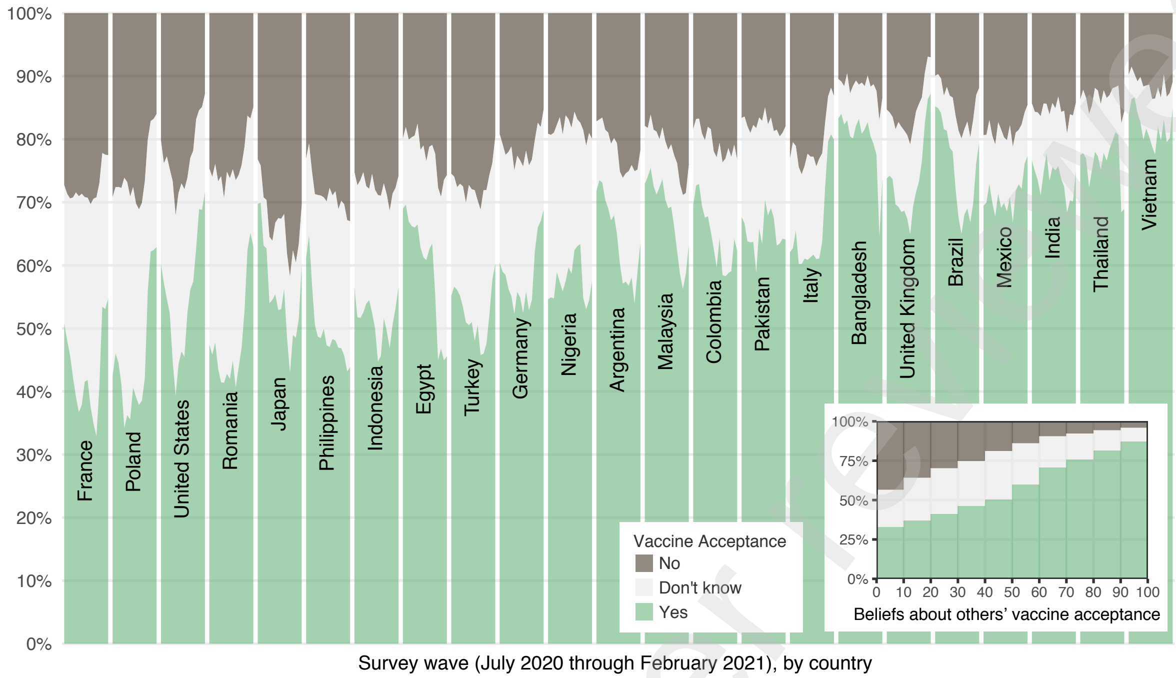


**Fig. 1.** There is substantial country-to-country variation in levels and trends in COVID-19 vaccine acceptance from July 2020 to February 2021, such that some countries have both recent increases and decreases in acceptance. Shown are the 23 countries with repeated data collection over time. "Yes" also includes respondents indicating they already received a vaccine. *(inset)* Pooling data from all 23 countries, people who believe a larger fraction of their community will accept a vaccine are on average more likely to say they will accept a vaccine; this is also true within each included country (Figure S13).

vaccines.

Here we provide evidence, from a large-scale randomized experiment embedded in an international survey, that information about descriptive norms — what other people do, believe, or say — can have substantial positive effects on intentions to accept new vaccines for COVID-19.

Through a collaboration with Facebook and Johns Hopkins University, and with input from experts at the World Health Organization and the Global Outbreak Alert and Response Network, we fielded a survey in 67 countries in their local languages, yielding over 1.7 million responses to date (32). This survey assessed people's knowledge about COVID-19, beliefs about and use of preventative behaviors, beliefs about others' behaviors and beliefs, and economic experiences and expectations. While it is often impossible to account for all factors that may jointly determine selection into the sample and survey responses, our collaboration with Facebook allows using state-of-the-art, privacy-preserving weighting for non-response using rich behavioral and demographic variables, as well as further weighting to target the adult population of each country (32, 33). All analyses presented here use these survey weights to ensure our results are as representative of these countries' adult populations as possible.

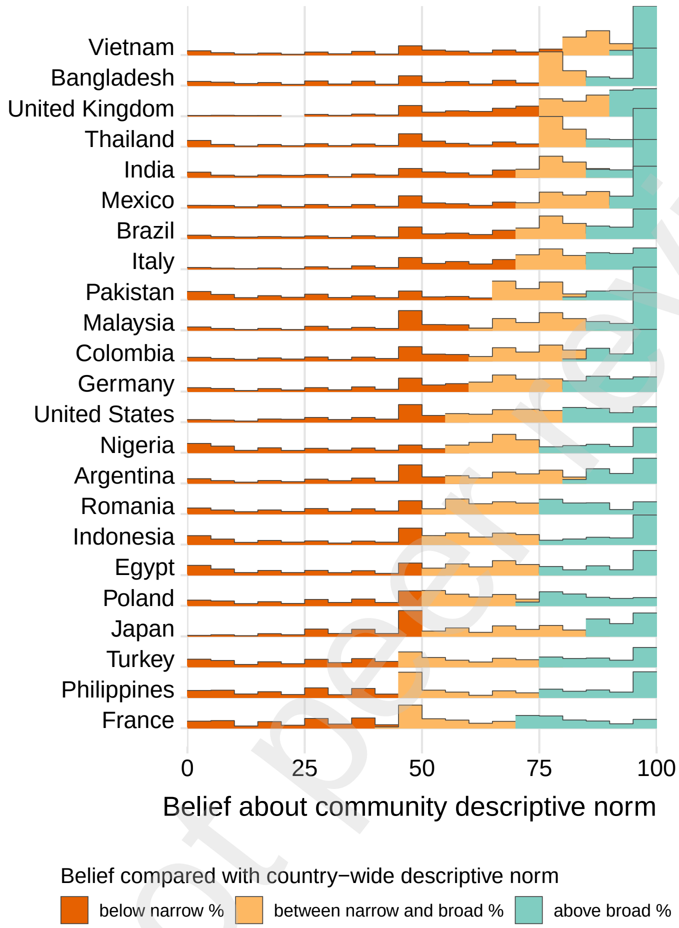


**Fig. 2.** Within-country distributions of beliefs about descriptive norms ("Out of 100 people in your community, how many do you think would take a COVID-19 vaccine if it were made available?") during the experimental period (October 2020 to February 2021). To enable comparison with actual country-wide potential vaccine acceptance, these histograms are colored by whether they are below (red) the narrow ("Yes" only) definition of vaccine acceptance, between (yellow) the narrow and broad ("Yes" and "Don't know") definitions, or above (teal) the broad definition.

This survey has documented substantial variation in stated intentions to take a vaccine for COVID-19 when one is available to the respondent, with, for example, some countries having much larger fractions of people saying they will take a vaccine than others (Figure 1); however, a plurality consistently say they will accept a vaccine and only a (often small) minority say they will refuse one. This is consistent with other smaller-scale national (10) and international (34) surveys. There is also substantial variation in what fraction of *other people* that respondents think will accept the vaccine, and these beliefs often substantially differ from country-wide levels of vaccine acceptance (Figure 2). This deviation can have multiple causes, including responding with round numbers; but we posit this is at least partially because some people have incorrect beliefs about descriptive norms. Underestimation of vaccine acceptance by others could be partially caused by processes — such as news coverage of the challenges posed by vaccine hesitancy or diffusion of anti-vaccine messages on social media — that make hesitancy more salient. Beliefs about descriptive norms are in turn positively correlated with vaccine acceptance (Figure 1 inset, Figure S13), likely reflecting many processes, like geographic and social clustering of vaccine hesitancy, but also including the causal effects of beliefs about others on intentions to accept a vaccine. Public health communications could present information about norms, perhaps correcting some people's overestimation of the prevalence of vaccine hesitancy. Unlike other ongoing, frequently observable preventative behaviors, like mask wearing, people may have little information about whether others intend to accept a vaccine — which suggests messages with this information could have particularly large effects.

## Randomized Experiment

To learn about the effects of providing normative information about new vaccines, beginning in October 2020, for the 23 countries with ongoing data collection in this study, we provided respondents with accurate information about how previous respondents in their country had responded to a survey question about vaccine acceptance, mask wearing, or physical distancing. We randomized at what point in the survey this information was provided, which behavior the information was about, and how we summarized previous respondents' answers — enabling us to estimate the effects of providing information about descriptive norms on people's stated intentions to accept a vaccine.

In the case of vaccine acceptance, we told some respondents, "Your responses to this survey are helping researchers in your region and around the world understand how people are responding to COVID-19. For example, we estimate from survey responses in the previous month that X% of people in your country say they will take a vaccine if one is made available", where X is the (weighted) percent of respondents saying "Yes" to a vaccine acceptance question. Other respondents received information on how many "say they *may* take a vaccine", which is the (weighted) percent who chose "Yes" or "Don't know" for that same question. Whether this information occurs before or after a more detailed vaccine acceptance question* and whether it uses the *broad* (combining "Yes" and "Don't know") or *narrow* ("Yes" only) definition of potential vaccine accepters is randomized — allowing us to estimate the causal effects of

*When the detailed vaccine acceptance question occurs after the normative information, it is always separated by at least one intervening screen with two questions, and it is often separated by several screens of questions (Figure S12a).

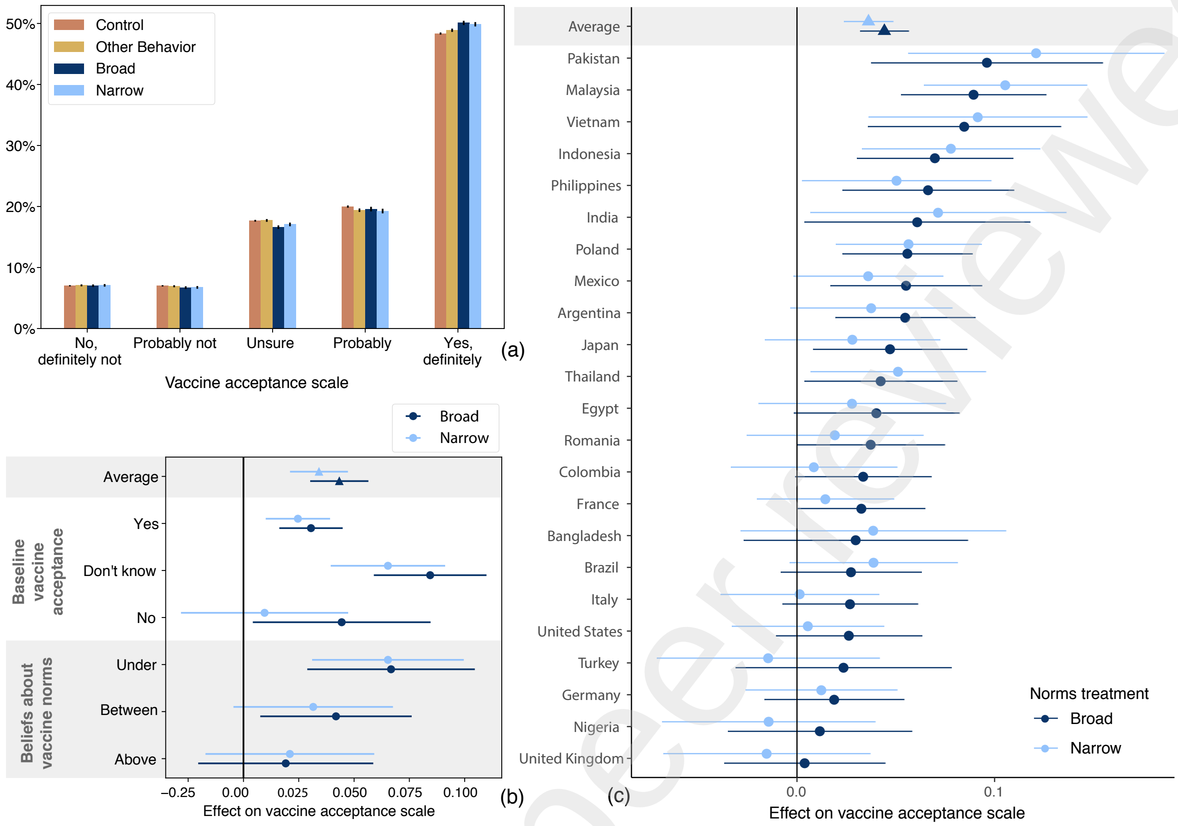


**Fig. 3.** (a) The normative information treatments shift people to higher levels of vaccine acceptance, whether compared with receiving no information (*control*) or information about other, non-vaccine-acceptance norms (*other behavior*). (b) These estimated effects are largest for respondents who are uncertain about accepting a vaccine at baseline and respondents with baseline beliefs about descriptive norms that are under (rather than above or between) both of the levels of normative information provided in the treatments. (c) While there is some country-level heterogeneity in these effects, point estimates of the effect of the broad normative information treatment are positive in all countries. Error bars are 95% confidence intervals.

this normative information. Here we focus on comparisons between providing the normative information about vaccines before or after measuring outcomes (e.g., vaccine acceptance); in the Supplementary Information (SI), we also report similar results when the control group consists of those who received information about other behaviors (i.e., about mask wearing and distancing), which can avoid concerns about differential attrition.

On average, presenting people with this normative information increases stated intentions to take a vaccine, with the broad and narrow treatments causing 0.04 and 0.03 increases on a five-point scale (95% confidence intervals: [0.03, 0.06] and [0.02, 0.05], respectively). The distribution of responses across treatments (Figure 3a) reveals that the effects of the broad (narrow) treatment are concentrated in inducing an additional 1.8% (1.2%) of people to say they will at least "probably" accept the vaccine, and moving 2.0% (1.9%) to "definitely" (Table

S8). This is a 5% relative reduction in the fraction of people choosing a response that is "unsure" or more negative. A post hoc analysis also concluded that these effects are largest among people who answer "Don't know" to the baseline vaccine acceptance question (Figure 3b, Table S11), consistent with the idea of targeting vaccine "fence-sitters" (35). These effects are relatively large and are of similar overall magnitude as global trends in vaccine acceptance over the course of the experiment (0.11 increase on the five-point scale) — a period that featured frequent and widely-distributed vaccine-related news.

The effects on vaccine acceptance can be at least partially explained by changes in respondents' beliefs about these descriptive norms. We can examine this because the survey also measured respondents' beliefs about vaccine acceptance in their communities (as displayed in Figure 2), and we randomized whether this was measured before or after providing the normative information. As expected, the normative information treatment increased the fraction of people that the respondents estimate will accept a vaccine (Figure S7). Among those respondents for whom we measured these normative beliefs prior to treatment, we can examine how treatment effects varied by this baseline belief. In particular, we classify respondents according to whether their baseline belief was *above* the broad ("may take") number, *under* the narrow ("will take") number, or *between* these two numbers.[†] Consistent with the hypothesis that this treatment works through revising beliefs about descriptive norms upwards, we find significant effects of the normative information treatment in the groups that may be underestimating vaccine acceptance — the *under* and *between* groups (Figure 3b), though the smaller sample sizes here (since these analyses are only possible for a random subset of respondents) only provide some evidence that the effect in the *under* group differs from that in the *above* group ($p = 0.09$ and $p = 0.09$ for broad and narrow treatments, respectively).[‡]

Having fielded this experiment in 23 countries, we can estimate and compare treatment effects internationally, which may be useful for both national and international communication efforts. Using a linear mixed-effects model, we estimate positive effects in the majority of countries (Figure 3c). While estimates for some countries are larger (e.g., Pakistan, Malaysia) and some are smaller (e.g., Nigeria, United Kingdom), most countries are statistically indistinguishable. Thus, we summarise the results as providing evidence that accurate normative information consistently increases intentions to accept COVID-19 vaccines.

[†] The question measuring beliefs about descriptive norms asks about "your community", while the information provided is for the country. Thus, for an individual respondent, these need not exactly match to be consistent.

[‡] We had also hypothesized that the broad and narrow treatments would differ from each other in their effects on respondents in the between group, but we found no such evidence, $p = 0.64$.

## Discussion

Framing vaccination as a social norm has been suggested as an effective approach to building COVID-19 vaccine confidence (15, 36, 37), but this recommendation has lacked direct evidence on a scalable messaging strategy, which this international randomized experiment now contributes. These results have implications for communication to the public through health messaging campaigns and the news media. For example, because very high levels of vaccine uptake are needed to reach herd immunity (3), it is reasonable for news media to cover the challenges presented by vaccine hesitancy; but our results suggest that it is valuable to contextualize such reporting by highlighting the widespread norm of accepting COVID-19 vaccines. Public health campaigns to increase acceptance of safe and effective vaccines can include information about descriptive norms. In an effort to influence the public, some public figures have already documented receiving a COVID-19 vaccine in videos on television and social media. The substantial positive effects of numeric summaries of everyday people's intentions documented here suggest that simple factual information about descriptive norms can similarly leverage social influence to increased vaccine acceptance. Some negative attitudes toward vaccination put disadvantaged communities at more risk and emphasizing country-wide vaccination norms may prove critical for removing susceptible pools and reducing the risk of endemic disease (3, 38).

The substantial effects of normative information about vaccine acceptance may reflect that people have little passive exposure to information about how many people in their communities and countries would accept a vaccine, or even have done so already. This result contrasts with other preventative behaviors (mask wearing and distancing), for which we observe smaller or no effects (see Supplementary Information Section S6), that are both ongoing (i.e., respondents have chosen whether to perform them before) and readily observable in public. However, it is possible that as people have more familiarity with social contacts choosing to accept a vaccine, this type of normative information will become less impactful, making the use of this communication strategy even more important in the early stages of a vaccine roll out. More generally, changes in stated intentions to accept a vaccine may likely not fully translate into actual take-up. Thus, we emphasize the need for a range of interventions that lower real and perceived barriers to vaccination, as well as leveraging descriptive norms and social contagion more generally, such as in spreading information about how to obtain a vaccine (21).

## Materials and Methods

**Experiment analysis.** The results presented in the main text and elaborated on in the supplementary materials each use a similar pre-registered methodology that we briefly describe here.

For the results in Figure 3a, we estimate the following linear regression for each behavior $k$

$$Y_{ik} = \delta_{0k} + \sum_{j \in J} \delta_{jk} D_{ik}^{j} + \gamma_k X_i + \sum_{j \in J} \eta_{jk} X_i D_{ik}^{j} + \varepsilon_{ik} \quad [1]$$

where $Y_{ik}$ is the outcome for individual $i$ and behavior $k \in K = \{\text{vaccine, distancing, masks}\}$, $D_{ik}^{j}$ is an indicator if individual $i$ received treatment $j \in J = \{\text{Broad, Narrow}\}$ for behavior $k$, and $X_i$ is a vector of centered covariates (39, 40). All statistical inference uses heteroskedasticity-consistent Huber–White "sandwich" estimates of the variance–covariance matrix.

For heterogeneous treatment effects (Figure 3b), we estimate a similar regression focusing on the vaccine behavior.

$$Y_i = \sum_{b \in B} 1[b_i = b] \left( \delta_0^b + \sum_{j \in J} \delta_j^b D_{ij}^b + \gamma_k X_i + \sum_{j \in J} \eta_j^b X_i D_{ij}^b \right) + \varepsilon_{ik} \quad [2]$$

**Mixed-effects model.** In the main text and Figure 3c, we report results from a linear mixed-effects model with coefficients that vary by country. This model is also described in our preregistered analysis plan. Note that the coefficients for the overall (across-country) treatments effects in this model differ slightly from the estimates from the model in equation 3; that is, the "Average" points in Figure 3b and 3c do not match exactly. As noted in our analysis plan, "sandwich" standard errors are not readily available here, so reported 95% confidence intervals are obtained by estimating the standard errors via a bootstrap.

**Data and materials and availability.** Documentation of the survey instrument and aggregated data from the survey are publicly available at https://covidsurvey.mit.edu. Researchers can request access to the microdata from Facebook and MIT at https://dataforgood.fb.com/docs/preventive-health-survey-request-for-data-access/. Preregistration details are available at https://osf.io/h2gwv/?view_only=f7d71d8684874b50be5981483613a80e. Analysis code for reproducing the results will be made public. The Committee on the Use of Humans as Experimental Subjects at MIT approved both the survey and embedded randomized experiment as exempt protocols.

**Acknowledgments.** The COVID-19 Global Beliefs, Behaviors, and Norms Survey was a collaborative effort involving contributions from individuals at multiple institutions, especially the Massachusetts Institute of Technology, Johns Hopkins University, the Global Outbreak Alert and Response Network, the World Health Organization, and Facebook, with key contributions from Stella Babalola, Nina Gobat, Esther Kim, Kelsey Mulcahy, Praveen Raja, Stephanie

Sasser, Dominick Shattuck, Jeni Stolow, Carlos Velasco, and Thomas Wynter. We thank Eytan Bakshy, Adam Berinsky, Daniel Bjorkegren, Alex Leavitt, Solomon Messing, and David Rand for helpful comments on this work. We thank the millions of respondents to this survey worldwide. This work was funded in part by a grant from Facebook to the MIT Initiative on the Digital Economy.

1. AA Malik, SM McFadden, J Elharake, SB Omer, Determinants of COVID-19 vaccine acceptance in the US. *EClinicalMedicine* **26**, 100495 (2020).
2. S Sanche, et al., High contagiousness and rapid spread of severe acute respiratory syndrome coronavirus 2. *Emerging Infectious Diseases* **26**, 1470–1477 (2020).
3. RM Anderson, C Vegvari, J Truscott, BS Collyer, Challenges in creating herd immunity to SARS-CoV-2 infection by mass vaccination. *The Lancet* **396**, 1614–1616 (2020).
4. C Signorelli, S Iannazzo, A Odone, The imperative of vaccination put into practice. *The Lancet Infectious Diseases* **18**, 26–27 (2018).
5. SB Omer, C Betsch, J Leask, Mandate vaccination with care. *Nature*, 469–472 (2019).
6. C Betsch, R Böhm, Detrimental effects of introducing partial compulsory vaccination: Experimental evidence. *The European Journal of Public Health* **26**, 378–381 (2016).
7. C Betsch, R Böhm, GB Chapman, Using behavioral insights to increase vaccination policy effectiveness. *Policy Insights from the Behavioral and Brain Sciences* **2**, 61–73 (2015).
8. B Nyhan, J Reifler, S Richey, GL Freed, Effective messages in vaccine promotion: a randomized trial. *Pediatrics* **133**, e835–e842 (2014).
9. B Nyhan, J Reifler, Does correcting myths about the flu vaccine work? An experimental evaluation of the effects of corrective information. *Vaccine* **33**, 459–464 (2015).
10. J Green, et al., Report 36: Evaluation of COVID-19 vaccine communication strategies. *The COVID States Project: A 50-State COVID-19 Survey Reports* (2021).
11. B Nyhan, J Reifler, S Richey, The role of social networks in influenza vaccine attitudes and intentions among college students in the southeastern united states. *Journal of Adolescent Health* **51**, 302–304 (2012).
12. EK Brunson, The impact of social networks on parents' vaccination decisions. *Pediatrics* **131**, e1397–e1404 (2013).
13. CT Bauch, DJ Earn, Vaccination and the theory of games. *Proceedings of the National Academy of Sciences* **101**, 13391–13394 (2004).
14. T Oraby, V Thampi, CT Bauch, The influence of social norms on the dynamics of vaccinating behaviour for paediatric infectious diseases. *Proceedings of the Royal Society B: Biological Sciences* **281**, 20133172 (2014).
15. NT Brewer, GB Chapman, AJ Rothman, J Leask, A Kempe, Increasing vaccination: Putting psychological science into action. *Psychological Science in the Public Interest* **18**, 149–207 (2017).
16. SS DeRoo, NJ Pudalov, LY Fu, Planning for a COVID-19 vaccination program. *JAMA* **323**, 2458–2459 (2020).
17. H Allcott, et al., Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics* **191**, 104254 (2020).
18. A Bridgman, et al., The causes and consequences of COVID-19 misperceptions: Understanding the role of news and social media. *Harvard Kennedy School Misinformation Review* **1** (2020).

19. A Simonov, SK Sacher, JPH Dubé, S Biswas, The persuasive effect of Fox News: Non-compliance with social distancing during the COVID-19 pandemic, (National Bureau of Economic Research), Technical report (2020).
20. D Holtz, et al., Interdependence and the cost of uncoordinated responses to COVID-19. *Proceedings of the National Academy of Sciences* **117**, 19837–19843 (2020).
21. A Banerjee, AG Chandrasekhar, E Duflo, MO Jackson, Using gossips to spread information: Theory and evidence from two randomized controlled trials. *The Review of Economic Studies* **86**, 2453–2490 (2019).
22. V Alatas, AG Chandrasekhar, M Mobius, BA Olken, C Paladines, When celebrities speak: A nationwide Twitter experiment promoting vaccination in Indonesia, (National Bureau of Economic Research), Technical Report 25589 (2019).
23. CT Bauch, S Bhattacharyya, Evolutionary game theory and social learning can determine how vaccine scares unfold. *PLoS Comput Biol* **8**, e1002452 (2012).
24. N Rao, MM Mobius, T Rosenblat, Social networks and vaccination decisions, (Federal Reserve Board of Boston), Technical report (2007).
25. JT Vietri, M Li, AP Galvani, GB Chapman, Vaccinating to help ourselves and others. *Medical Decision Making* **32**, 447–458 (2012).
26. E Shim, GB Chapman, JP Townsend, AP Galvani, The influence of altruism on influenza vaccination decisions. *Journal of The Royal Society Interface* **9**, 2234–2243 (2012).
27. L Korn, R Böhm, NW Meier, C Betsch, Vaccination as a social contract. *Proceedings of the National Academy of Sciences* **117**, 14890–14899 (2020).
28. Y Ibuka, M Li, J Vietri, GB Chapman, AP Galvani, Free-riding behavior in vaccination decisions: An experimental study. *PloS one* **9**, e87164 (2014).
29. R Böhm, C Betsch, L Korn, Selfish-rational non-vaccination: Experimental evidence from an interactive vaccination game. *Journal of Economic Behavior & Organization* **131**, 183–195 (2016).
30. JC Hershey, DA Asch, T Thumasathit, J Meszaros, VV Waters, The roles of altruism, free riding, and bandwagoning in vaccination decisions. *Organizational Behavior and Human Decision Processes* **59**, 177–187 (1994).
31. R Sato, Y Takasaki, Peer effects on vaccination behavior: Experimental evidence from rural nigeria. *Economic Development and Cultural Change* **68**, 93–129 (2019).
32. A Collis, et al., Global Survey on COVID-19 Beliefs, Behaviors, and Norms, (MIT Sloan School of Management), Technical report (2020).
33. N Barkay, et al., Weights and methodology brief for the COVID-19 symptom survey by University of Maryland and Carnegie Mellon University, in partnership with Facebook. *arXiv preprint arXiv:2009.14675* (2020).
34. JV Lazarus, et al., A global survey of potential acceptance of a COVID-19 vaccine. *Nature Medicine*, 1–4 (2020).
35. C Betsch, L Korn, C Holtmann, Don't try to convert the antivaccinators, instead target the fence-sitters. *Proceedings of the National Academy of Sciences* **112**, E6725–E6726 (2015).
36. AGCMH Wen-Ying Sylvia Chou, Caitlin E. Burgdorf, COVID-19 vaccination communication: Applying behavioral and social science to address vaccine hesitancy and foster vaccine confidence, (National Institutes of Health), Technical report (2020).
37. R Palm, T Bolsen, JT Kingsland, The effect of frames on COVID-19 vaccine hesitancy. *medRxiv* (2021).

38. E Paul, A Steptoe, D Fancourt, Attitudes towards vaccines and intention to vaccinate against COVID-19: Implications for public health communications. *The Lancet Regional Health-Europe*, 100012 (2020).
39. W Lin, Agnostic notes on regression adjustments to experimental data: Reexamining freedman's critique. *Annals of Applied Statistics* **7**, 295–318 (2013).
40. LW Miratrix, JS Sekhon, B Yu, Adjusting treatment effect estimates by post-stratification in randomized experiments. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **75**, 369–396 (2013).
41. J De Quidt, J Haushofer, C Roth, Measuring and bounding experimenter demand. *American Economic Review* **108**, 3266–3302 (2018).
42. J Mummolo, E Peterson, Demand effects in survey experiments: An empirical assessment. *American Political Science Review* **113**, 517–529 (2019).

# Supplementary Information for "Surfacing norms to increase vaccine acceptance"

## Contents

## S1. Experiment overview

During an update to the survey on October 28th, 2020, we introduced a prompt to all respondents that provided information about preventative behaviors in their country based on information from the survey. Although this information was provided to all respondents who completed the survey from an eligible country, the information was provided in a random order creating an experiment within the survey. For each eligible respondent, we showed the following message at a random position in the latter part of the survey:

> Your responses to this survey are helping researchers in your region and around the world understand how people are responding to COVID-19. For example, we estimate from survey responses in the previous month that [[country share]]% of people in your country say they [[broad or narrow]] [[preventative behavior]].

We filled in the blanks with one randomly chosen preventative behavior, a broad or narrow definition of the activity, and the true share of responses for the respondent's country. The three behaviors were vaccine acceptance, mask wearing, and social distancing. In the broad condition, we used a more inclusive definition of the preventative behavior and the narrow condition used a more restrictive definition. For example, for vaccine acceptance we either reported the share of people responding "Yes" or the share of people responding "Yes" or "Don't know" to the baseline vaccine acceptance question. The numbers shown are displayed in Figure S5.

We preregistered our analysis plan, which we also updated to reflect continued data collection and our choice to eliminate the distancing information treatment in later waves. While we describe some of the main choices here, all versions of our analysis plan can be viewed at https://osf.io/h2gwv/?view_only=f7d71d8684874b50be5981483613a80e. The analysis of the experiment in the main text that is not described in the analysis plan is labeled post hoc (in particular, heterogeneity by baseline vaccine acceptance). One set of more complex analyses speculatively described in the analysis plan (hypothesis 3, "may suggest using instrumental variables analyses") has not yet been pursued.

## S2. Data construction

Our dataset is constructed from the microdata described in Collis et al. (*S32*). We first code each outcome to a 5-point numerical scale. We then condition on being eligible for treatment and having a waves survey type (i.e. being in a country with continual data collection) to arrive at the full dataset of those eligible for treatment.[§] All randomization and balance checks

[§] Respondents in the snapshot survey may have received treatment if they self-reported being in a wave country. Their weights, however, will be wrong as their country will disagree with the inferred country so they are excluded.

described as "intent-to-treat" use this dataset. In our preregistered analysis plan, we described how the sample would be restricted to those who completed the survey and for whom we received a full survey completion weight from Facebook. This removes approximately 40% of respondents, resulting in 349,664,694 respondents. For the main analysis comparing users who received the vaccine information treatment to control users (e.g., in Figure 3b), there are 266,206 respondents.

As in our pre-analysis plan, the following variables are used in our analysis:

1. Outcomes

    (a) Over the next two weeks, how likely are you to wear a mask when in public? [Always, Almost always, When convenient, Rarely, Never]

    (b) Over the next two weeks, how likely are you to maintain a distance of at least 1 meter from others when in public? [Always, Almost always, When convenient, Rarely, Never]

    (c) If a vaccine against COVID-19 infection is available in the market, would you take it? [Yes, definitely, Probably, Unsure, Probably not, No, definitely not]

2. Mediators & Covariates

    (a) Baseline outcomes. These questions are similar to the outcome questions. Only the vaccine question always appears before the treatment in all cases; the others are in a randomized order. Thus, for use of the other covariates for increasing precision, mean imputation is required.

        - Masks. How often are you able to wear a mask or face covering when you are in public? How effective is wearing a face mask for preventing the spread of COVID-19?
        - Distancing: How often are you able to stay at least 1 meter away from people not in your household? How important do you think physical distancing is for slowing the spread of COVID-19?
        - Vaccine: If a vaccine for COVID-19 becomes available, would you choose to get vaccinated? This will be coded as binary indicators for the possible outcomes, grouping missing outcomes with "Don't know".

    (b) Beliefs about norms. These questions will be randomized to be shown before the treatment for some respondents and after treatment for other respondents.

This will allow us to study heterogeneity in baseline beliefs, as well as ensure our randomization does impact beliefs.

- Masks: Out of 100 people in your community, how many do you think do the following when they go out in public? Wear a mask or face covering.
- Distancing: Out of 100 people in your community, how many do you think do the following when they go out in public? Maintain a distance of at least 1 meter from others.
- Vaccine: Out of 100 people in your community, how many do you think would take a COVID-19 vaccine if it were made available?

When used in analysis, we require all covariates to be before both treatment and outcome. As the survey contains randomized order for these questions, this ensures that the distribution of question order is the same across treated and control groups and removes any imbalance created by differential attrition. Missing values are imputed at their (weighted) mean.

## S3. Randomization checks

Table S1 presents results of a test that the treatment and control shares were equal to 50% as expected. While the final dataset does have some evidence of imbalance that could be caused by differential attrition, the "robust" dataset (described in S6.2) is well balanced and the treatment is balanced across the three behaviors information could be provided about (Table S2). According to our pre-registered analysis plan, in the presence of evidence of differential attrition, we make use of additional analyses that use the information about other behaviors as an alternative control group throughout this supplement.

**Table S1. Randomization Tests**

| | p-val | Treated Share | Control Share |
|---|---|---|---|
| Full | 0.040 | 0.501 | 0.499 |
| Final | 0.028 | 0.498 | 0.502 |
| Robust | 0.222 | 0.499 | 0.501 |

The results of a test that the treated share and control shares equal 50%. The first row uses intent-to-treat on the full set of eligible respondents, the second row uses the final data set after conditioning on eligibility and completing the survey, and the third row uses the subset of responses in the final dataset that have at least one block between treatment and outcome.

In addition, baseline covariates measured before both treatment and the outcome are balanced across treatment and control groups (Table S3). The covariates are also balanced in the final analysis dataset (Table S4) and within treated users across the three possible treatment behaviors (Table S5).

**Table S2. Randomization Tests**

| | Vaccine | Masks | Distancing |
|---|---|---|---|
| Final | 0.510 | 0.509 | 0.441 |
| Robust | 0.523 | 0.308 | 0.519 |

The p-values of a test that each behavior was shown the expected number of times. This reports the results of a joint test that each period share was equal to the expected. For waves 9-12, each behavior was shown 1/3 of the time and for waves 12 on the vaccine treatments were shown to 2/3 of respondents and the mask treatments were shown to 1/3 of respondents. This table cannot include the full dataset intent-to-treat analysis because the behavior randomization occurred when the treatment was shown.

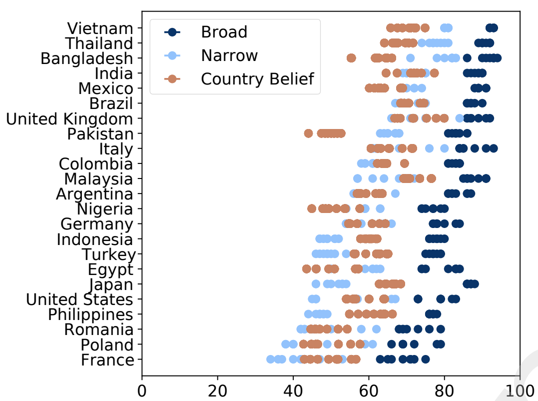


**(a)** Vaccine Treatments

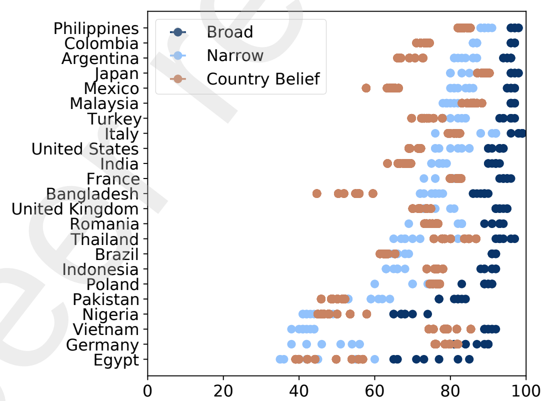


**(b)** Mask Treatments

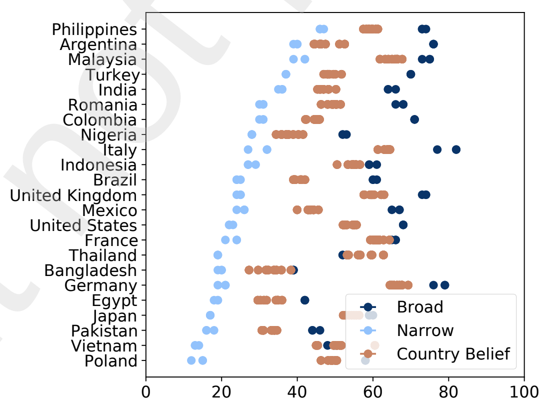


**(c)** Distancing Treatments

**Fig. S4.** Treatment Variation

For each behavior (Vaccine, Masks, Distancing), we plot the information provided to subjects based on the broad and narrow definitions of compliance. The treatments were updated every two weeks as new waves of data were included. The points labeled "country belief" display the weighted average belief in a country of how many people out of 100 practice (or will accept, for vaccines) each behavior.

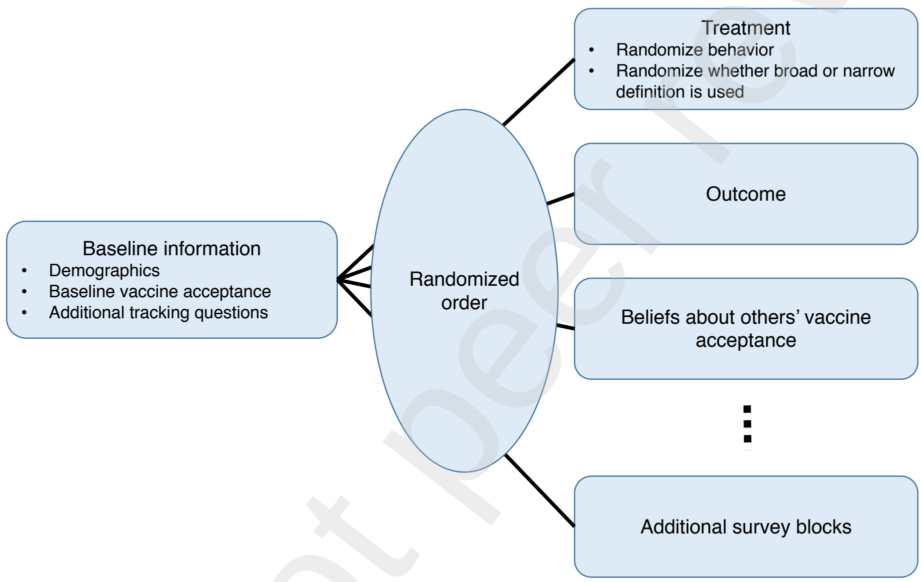


**Fig. S5.** Experiment Flow

Illustration of the flow of a respondent through the survey. First, they are presented with tracking and demographic questions. They then enter a randomized portion where blocks are in random order. This includes the treatment, outcome, and many of the baseline covariates included in regressions for precision. Recall all covariates used in analysis are only used if they are pre-treatment and outcome.

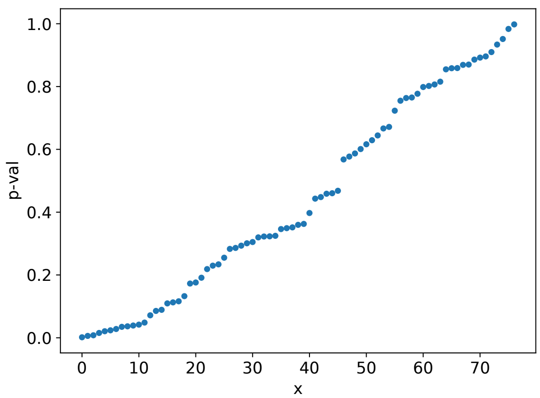


**(a)** Balance Tests: Intent-to-treat

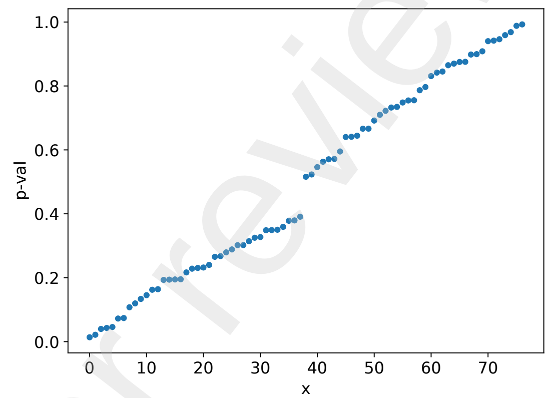


**(b)** Balance Tests: Final Sample

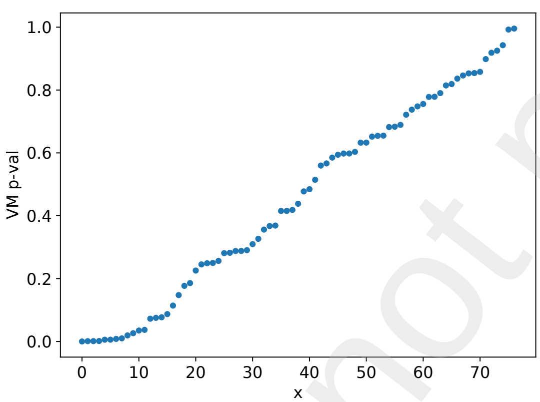


**(c)** Balance Tests: Vaccine vs Mask Treatments

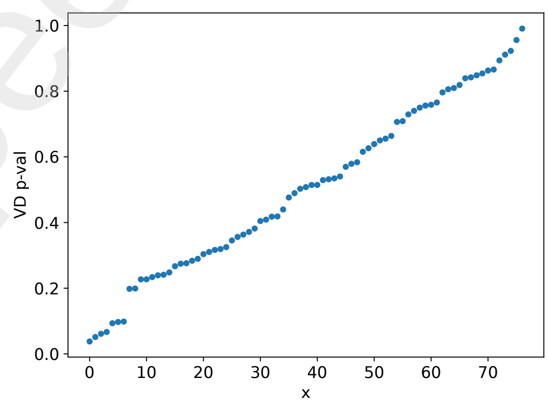


**(d)** Balance Tests: Vaccine vs Dist Treatments

**Fig. S6.** Balance Test p-Values

Ordered p-values for the balance tests described in Tables S3, S4, and S5 sorted in ascending order. All available pre-treatment covariates are included, which results in 76 tests. This includes roughly 40 covariates that are not presented in the tables for brevity. These are questions that permit multiple responses, including news media, sources, and trust, and a more detailed list of preventative measures taken.

**Table S3. Balance Tests: Intent-to-treat**

| | p-val | Control | | Treated | |
|---|---|---|---|---|---|
| age | 0.349 | 2.546 | (0.003) | 2.542 | (0.003) |
| gender | 0.219 | 1.441 | (0.001) | 1.439 | (0.001) |
| education | 0.577 | 2.784 | (0.002) | 2.783 | (0.002) |
| own health | 0.086 | 2.414 | (0.002) | 2.418 | (0.002) |
| vaccine accept | 0.983 | 1.467 | (0.001) | 1.467 | (0.001) |
| knowledge existing treatments | 0.468 | 0.282 | (0.001) | 0.282 | (0.001) |
| info exposure past week | 0.283 | 2.316 | (0.002) | 2.319 | (0.002) |
| info exposure more less wanted | 0.869 | 2.406 | (0.002) | 2.406 | (0.002) |
| know positive case | 0.320 | 1.259 | (0.002) | 1.256 | (0.002) |
| prevention mask | 0.629 | 3.604 | (0.003) | 3.606 | (0.003) |
| prevention distancing | 0.951 | 2.671 | (0.004) | 2.672 | (0.004) |
| prevention hand washing | 0.616 | 3.300 | (0.003) | 3.298 | (0.003) |
| effect mask | 0.293 | 2.978 | (0.003) | 2.973 | (0.003) |
| effect hand washing | 0.173 | 2.998 | (0.003) | 2.992 | (0.003) |
| country management | 0.234 | 1.808 | (0.004) | 1.800 | (0.004) |
| community management | 0.799 | 1.906 | (0.004) | 1.905 | (0.004) |
| community action importance | 0.601 | 3.362 | (0.003) | 3.360 | (0.003) |
| community action norms | 0.459 | 2.736 | (0.004) | 2.732 | (0.004) |
| distancing importance | 0.346 | 3.129 | (0.003) | 3.134 | (0.003) |
| norms dist | 0.042 | 48.775 | (0.107) | 49.079 | (0.107) |
| norms masks | 0.397 | 71.570 | (0.102) | 71.687 | (0.102) |
| norms vaccine | 0.667 | 60.728 | (0.104) | 60.663 | (0.103) |
| risk community | 0.286 | 2.557 | (0.006) | 2.549 | (0.006) |
| risk infection | 0.671 | 2.177 | (0.006) | 2.173 | (0.006) |
| control infection | 0.802 | 1.866 | (0.007) | 1.868 | (0.007) |
| infection severity | 0.024 | 1.279 | (0.004) | 1.267 | (0.004) |
| employed 2020 | 0.191 | 0.728 | (0.002) | 0.733 | (0.002) |

Pre-treatment covariate means for all respondents who were eligible for treatment in both the treatment and control groups along with the p-value for the test of the null that the means are equal. For each covariate, only responses where the covariate is not missing and occurs before both treatment and control are included. To account for changes to the sampling frequencies, these p-values are from the coefficient on the intent-to-treat term in a regression of the covariate on treatment, period, and centered interactions between treatment and period. As we do not have weights for all respondents, this is an unweighted regression.

**Table S4. Balance Tests: Final Dataset**

| | p-val | Control | | Treated | |
|---|---|---|---|---|---|
| age | 0.898 | 2.655 | (0.004) | 2.657 | (0.004) |
| gender | 0.797 | 1.441 | (0.001) | 1.440 | (0.001) |
| education | 0.516 | 2.830 | (0.002) | 2.829 | (0.002) |
| own health | 0.959 | 2.397 | (0.002) | 2.402 | (0.002) |
| vaccine accept | 0.231 | 1.487 | (0.002) | 1.485 | (0.002) |
| knowledge existing treatments | 0.946 | 0.277 | (0.001) | 0.275 | (0.001) |
| info exposure past week | 0.325 | 2.385 | (0.002) | 2.390 | (0.002) |
| info exposure more less wanted | 0.831 | 2.429 | (0.003) | 2.431 | (0.003) |
| know positive case | 0.014 | 1.308 | (0.002) | 1.301 | (0.002) |
| prevention mask | 0.266 | 3.637 | (0.003) | 3.641 | (0.003) |
| prevention distancing | 0.217 | 2.712 | (0.005) | 2.717 | (0.005) |
| prevention hand washing | 0.842 | 3.336 | (0.003) | 3.336 | (0.003) |
| effect mask | 0.134 | 2.990 | (0.004) | 2.981 | (0.004) |
| effect hand washing | 0.379 | 3.018 | (0.004) | 3.013 | (0.004) |
| country management | 0.572 | 1.772 | (0.005) | 1.758 | (0.005) |
| community management | 0.875 | 1.880 | (0.005) | 1.877 | (0.005) |
| community action importance | 0.988 | 3.381 | (0.003) | 3.379 | (0.004) |
| community action norms | 0.748 | 2.710 | (0.004) | 2.702 | (0.005) |
| distancing importance | 0.314 | 3.170 | (0.004) | 3.174 | (0.004) |
| norms dist | 0.022 | 49.266 | (0.127) | 49.716 | (0.127) |
| norms masks | 0.043 | 72.393 | (0.120) | 72.722 | (0.120) |
| norms vaccine | 0.641 | 61.258 | (0.121) | 61.124 | (0.121) |
| risk community | 0.359 | 2.580 | (0.007) | 2.565 | (0.007) |
| risk infection | 0.870 | 2.220 | (0.007) | 2.217 | (0.007) |
| control infection | 0.787 | 1.867 | (0.009) | 1.868 | (0.009) |
| infection severity | 0.040 | 1.274 | (0.004) | 1.261 | (0.004) |
| employed 2020 | 0.120 | 0.729 | (0.003) | 0.735 | (0.003) |

Pre-treatment covariate means for all respondents who were eligible for treatment, completed the entire survey, and received a full survey completion weight in both the treatment and control groups along with the p-value for the test of the null that the means are equal. For each covariate, only responses where the covariate is not missing and occurs before both treatment and control are included. To account for changes to the sampling frequencies, these p-values are from the coefficient on the treatment term in a regression of the covariate on treatment, period, and centered interactions between treatment and period. This is a weighted regression using full completion survey weights.

**Table S5. Balance Tests Between Treatments: Final Dataset**

| | VD p-val | VM p-val | Vaccine | Masks | Dist |
|---|---|---|---|---|---|
| age | 0.532 | 0.019 | 2.674 | 2.648 | 2.611 |
| gender | 0.239 | 0.415 | 1.439 | 1.440 | 1.445 |
| education | 0.863 | 0.438 | 2.826 | 2.830 | 2.839 |
| own health | 0.248 | 0.925 | 2.405 | 2.405 | 2.386 |
| vaccine accept | 0.476 | 0.001 | 1.498 | 1.479 | 1.442 |
| knowledge existing treatments | 0.304 | 0.819 | 0.213 | 0.273 | 0.555 |
| info exposure past week | 0.276 | 0.310 | 2.401 | 2.386 | 2.354 |
| info exposure more less wanted | 0.317 | 0.226 | 2.450 | 2.428 | 2.355 |
| know positive case | 0.842 | 0.560 | 1.316 | 1.302 | 1.233 |
| prevention mask | 0.866 | 0.369 | 3.650 | 3.639 | 3.609 |
| prevention distancing | 0.419 | 0.585 | 2.727 | 2.710 | 2.687 |
| prevention hand washing | 0.409 | 0.148 | 3.342 | 3.330 | 3.326 |
| effect mask | 0.356 | 0.515 | 2.990 | 2.985 | 2.933 |
| effect hand washing | 0.956 | 0.281 | 3.016 | 3.008 | 3.013 |
| country management | 0.382 | 0.186 | 1.761 | 1.751 | 1.765 |
| community management | 0.584 | 0.035 | 1.885 | 1.862 | 1.880 |
| community action importance | 0.740 | 0.598 | 3.383 | 3.378 | 3.367 |
| community action norms | 0.540 | 0.756 | 2.706 | 2.703 | 2.679 |
| distancing importance | 0.839 | 0.790 | 3.178 | 3.176 | 3.149 |
| norms dist | 0.529 | 0.779 | 49.783 | 49.759 | 49.237 |
| norms masks | 0.750 | 0.853 | 72.910 | 72.871 | 71.308 |
| norms vaccine | 0.534 | 0.652 | 61.120 | 61.230 | 60.816 |
| risk community | 0.854 | 0.654 | 2.577 | 2.565 | 2.515 |
| risk infection | 0.418 | 0.815 | 2.230 | 2.210 | 2.179 |
| control infection | 0.227 | 0.245 | 1.870 | 1.873 | 1.849 |
| infection severity | 0.290 | 0.682 | 1.261 | 1.259 | 1.264 |
| employed 2020 | 0.706 | 0.683 | 0.738 | 0.730 | 0.733 |

Pre-treatment covariate means for all respondents who were treated, completed the entire survey, and received a full survey completion weight along with the p-value for the test of the null that the means between treatment groups are equal. For each covariate, only responses where the covariate is not missing and occurs before both treatment and control are included. To account for changes to the sampling frequencies, these p-values are from the coefficient on the treatment behavior terms in a regression of the covariate on treatment behavior, period, and centered interactions between treatment behavior and period. This is a weighted regression using full completion survey weights.

## S4. Analysis methods

The results presented in the main text and elaborated on in sections S5, S6, and S6.1 each use a similar pre-registered methodology that we briefly describe here. For the results in sections S5 and S6, we estimate the following linear regression for each behavior $k$

$$Y_{ik} = \delta_{0k} + \sum_{j \in J} \delta_{jk} D_{ik}^{j} + \gamma_k X_i + \sum_{j \in J} \eta_{jk} X_i D_{ik}^{j} + \varepsilon_{ik} \tag{S3}$$

where $Y_{ik}$ is the outcome for individual $i$ and behavior $k \in K = \{\text{vaccine, distancing, masks}\}$, $D_{ik}^{j}$ is an indicator if individual $i$ received treatment $j \in J = \{\text{Broad, Narrow}\}$ for behavior $k$, and $X_i$ is a vector of centered covariates (*S39*, *S40*). In the figures and tables, we report the $\delta_{jk}$'s and suppress coefficients on covariates and interactions. All statistical inference uses heteroskedasticity-consistent Huber–White "sandwich" estimates of the variance–covariance matrix.

In section S6.1, we estimate a similar regression. As our analysis of heterogeneity focuses on the vaccine treatment, we will suppress the behavior index $k$.

$$Y_i = \sum_{b \in B} 1[b_i = b] \left( \delta_0^b + \sum_{j \in J} \delta_j^b D_{ij}^b + \gamma_k X_i + \sum_{j \in J} \eta_j^b X_i D_{ij}^b \right) + \varepsilon_{ik} \tag{S4}$$

**S4.1. Mixed-effects model.** In the main text and Figure 3c, we report results from a linear mixed-effects model with coefficients that vary by country. This model is also described in our preregistered analysis plan. Note that the coefficients for the overall (across-country) treatments effects in this model differ slightly from the estimates from the model; that is, the "Average" points in Figure 3b and 3c do not match exactly. As noted in our analysis plan, "sandwich" standard errors are not readily available here, so reported 95% confidence intervals are obtained by estimating the standard errors via a bootstrap.

## S5. Effects on beliefs about descriptive norms

Figure S7 present evidence that the treatments do update beliefs about the descriptive norms of survey respondents. The figures plot coefficients on treatment from a regression of survey norms on treatment status, including centered covariates and interactions as described in the pre-analysis plan. In this analysis, treated respondents are those who receive the treatment before the question eliciting beliefs about norms. This will not agree, in general, with the treatment status for the main analysis given the randomized question order in the survey. The covariates included in this analysis are pre-treatment and outcome relative to this treatment definition.

Figure S7a compares treatment and control respondents and figure S7b conditions on treated individuals and then uses individuals who received an information treatment for a different behavior as control. The coefficients plotted in figure S7b are smaller than in figure S7a, which indicates that normative information on other behaviors may induce an update in beliefs on the focal behavior.

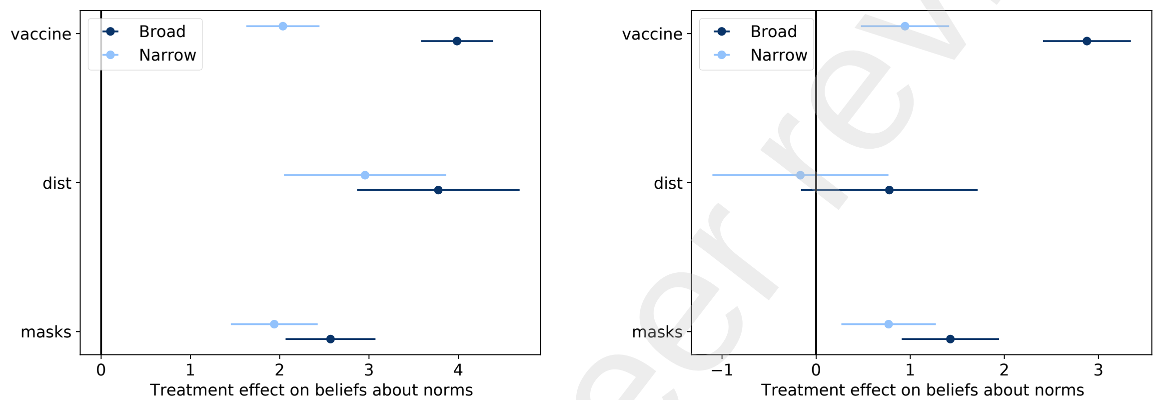


**(a)** Treatment vs. Control

**(b)** Treatment vs. Other Behavior Treatment

**Fig. S7.** Effects on beliefs about descriptive norms

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| Broad Treatment | 1.427*** | 0.778 | 2.878*** | 2.569*** | 3.776*** | 3.986*** |
| | (0.264) | (0.479) | (0.238) | (0.257) | (0.464) | (0.206) |
| Narrow Treatment | 0.771*** | -0.167 | 0.945*** | 1.939*** | 2.956*** | 2.036*** |
| | (0.256) | (0.477) | (0.239) | (0.248) | (0.463) | (0.209) |
| Control: Other Treatment | X | X | X | | | |
| Behavior | masks | dist | vaccine | masks | dist | vaccine |
| Number Controls | 106308 | 34002 | 69477 | 163222 | 52724 | 159903 |
| Number Treated | 53191 | 17354 | 87547 | 53191 | 17354 | 87547 |
| Observations | 159,499 | 51,356 | 157,024 | 216,413 | 70,078 | 247,450 |
| $R^2$ | 0.171 | 0.145 | 0.210 | 0.186 | 0.157 | 0.212 |
| Adjusted $R^2$ | 0.170 | 0.144 | 0.210 | 0.185 | 0.156 | 0.211 |
| Residual Std. Error | 24.987 | 27.046 | 24.229 | 25.402 | 27.085 | 24.656 |
| F Statistic | 105.393*** | 46.191*** | 163.326*** | 168.518*** | 64.202*** | 271.318*** |

*p<0.1; **p<0.05; ***p<0.01

**Table S6. Effects on beliefs about descriptive norms, for primary and alternative definitions of the control group**

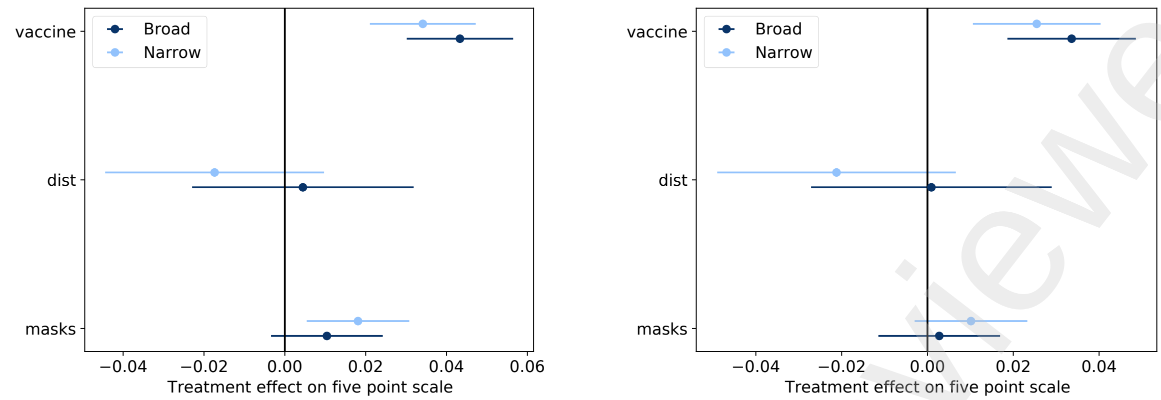


**(a)** Treatment vs. Control **(b)** Treatment vs. Other Behavior Treatment

**Fig. S8.** Treatment effects with primary and alternative definition of the control group

## S6. Effects on intentions

Figure S8 displays regression coefficients for the primary analysis, where the intention to partake in the outcome behavior is regressed on treatment, centered covariates, and their interactions. As discussed in the pre-analysis plan, we use both the randomized timing of information treatments and the randomized focal behavior of the intervention. Figure S8a uses respondents who receive the information after the outcome is measured as the control group and Figure S8b uses individuals who receive the information treatment for a different behavior as the control group. The results are largely consistent and suggest that the information treatment significantly increases reported vaccine acceptance, while effects for distancing and masks are smaller and not statistically distinguishable from zero.

Table S8 presents results from the same analysis after transforming the outcome variable into binary indicators. This allows us to understand across which thresholds the treatment has induced people to cross. The coefficients indicate that the treatment is inducing people to report they will at least probably take the vaccine and definitely take the vaccine. Similar regressions restricted to those who report they don't know if they will take the vaccine at baseline are presented in Table S9. Among this group, there is a larger effect and it is concentrated in moving people to say they will "probably" take the vaccine.

**S6.1. Heterogeneous treatment effects.** Figure S9 plots regression coefficients for estimates of heterogeneous treatment effects in equation 4 across different dimensions. We see the positive effects of our treatment concentrated in those with lower baseline beliefs about norms (Figure S9a) and in those who are unsure if they will accept a vaccine (Figure S9b). Estimates are also reported in Tables S10 and S11.

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| Broad Treatment | 0.003 | 0.001 | 0.034*** | 0.010 | 0.004 | 0.043*** |
| | (0.007) | (0.014) | (0.008) | (0.007) | (0.014) | (0.007) |
| Narrow Treatment | 0.010 | -0.021 | 0.025*** | 0.018*** | -0.017 | 0.034*** |
| | (0.007) | (0.014) | (0.008) | (0.006) | (0.014) | (0.007) |
| Control: Other Treatment | X | X | X | | | |
| Behavior | masks | dist | vaccine | masks | dist | vaccine |
| Number Controls | 115547 | 42237 | 78679 | 175125 | 64323 | 173252 |
| Number Treated | 58132 | 21296 | 92954 | 58132 | 21296 | 92954 |
| Observations | 173,679 | 63,533 | 171,633 | 233,257 | 85,619 | 266,206 |
| $R^2$ | 0.242 | 0.234 | 0.625 | 0.248 | 0.244 | 0.619 |
| Adjusted $R^2$ | 0.241 | 0.233 | 0.624 | 0.248 | 0.243 | 0.619 |
| Residual Std. Error | 0.693 | 0.855 | 0.795 | 0.695 | 0.856 | 0.797 |
| F Statistic | 118.650*** | 84.459*** | 872.011*** | 167.349*** | 113.450*** | 1332.485*** |

*p<0.1; **p<0.05; ***p<0.01

**Table S7. Treatment effects with primary and alternative definition of the control group**

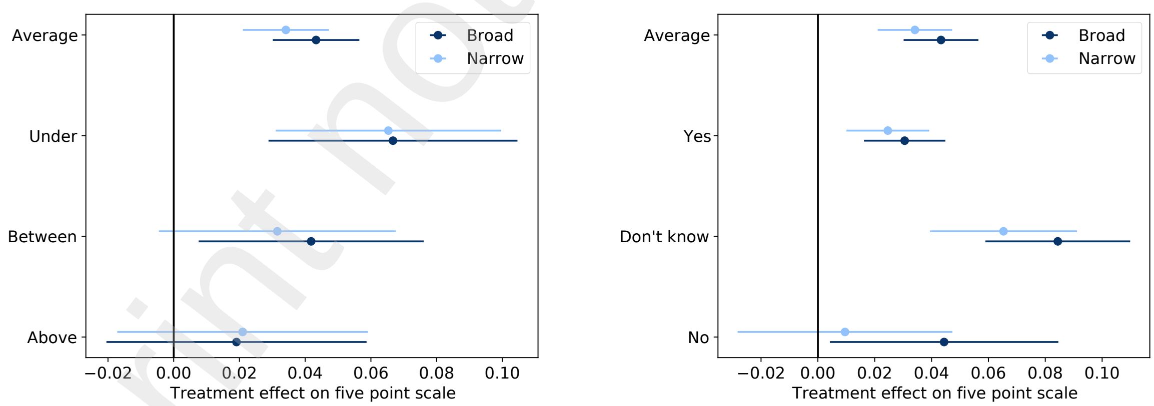


**(a)** Baseline Belief Partition

**(b)** Vaccine Acceptance

**Fig. S9**

| | > No, definitely not | > Probably not | > Unsure | > Probably |
|---|---|---|---|---|
| Intercept | 0.919*** | 0.845*** | 0.664*** | 0.472*** |
| | (0.001) | (0.001) | (0.001) | (0.001) |
| Narrow Treatment | -0.001 | 0.004** | 0.012*** | 0.019*** |
| | (0.002) | (0.002) | (0.003) | (0.003) |
| Broad Treatment | 0.001 | 0.005*** | 0.018*** | 0.020*** |
| | (0.002) | (0.002) | (0.003) | (0.003) |
| Observations | 266,206 | 266,206 | 266,206 | 266,206 |
| $R^2$ | 0.295 | 0.499 | 0.566 | 0.459 |
| Adjusted $R^2$ | 0.295 | 0.499 | 0.565 | 0.459 |
| Residual Std. Error | 0.230 | 0.255 | 0.310 | 0.367 |
| F Statistic | 126.209*** | 497.768*** | 1465.270*** | 1227.913*** |

*p<0.1; **p<0.05; ***p<0.01

Estimates of equation 3 with binary outcomes. The outcome variable for each column is an indicator equal to one if the respondent reported a value higher than the column name. For example, in the column "> Probably not" the outcome $Y_i$ equals one if the respondent answered "Unsure", "Probably", or "Yes, definitely".

**Table S8. Distributional treatment effects**

| | > No, definitely not | > Probably not | > Unsure | > Probably |
|---|---|---|---|---|
| Intercept | 0.970*** | 0.903*** | 0.294*** | 0.046*** |
| | (0.001) | (0.002) | (0.003) | (0.002) |
| Narrow Treatment | 0.004 | 0.005 | 0.034*** | 0.022*** |
| | (0.003) | (0.005) | (0.008) | (0.005) |
| Broad Treatment | 0.002 | 0.008* | 0.053*** | 0.020*** |
| | (0.003) | (0.005) | (0.008) | (0.004) |
| Observations | 54,015 | 54,015 | 54,015 | 54,015 |
| $R^2$ | 0.107 | 0.073 | 0.096 | 0.072 |
| Adjusted $R^2$ | 0.105 | 0.071 | 0.094 | 0.070 |
| Residual Std. Error | 0.160 | 0.286 | 0.445 | 0.218 |
| F Statistic | 5.544*** | 7.723*** | 22.786*** | 8.256*** |

*p<0.1; **p<0.05; ***p<0.01

Estimates of equation 3 with binary outcomes on sample of respondents who say they don't know if they will take a vaccine at baseline. The outcome variable for each column is an indicator equal to one if the respondent reported a value higher than the column name. For example, in the column "> Probably not" the outcome $Y_i$ equals one if the respondent answered "Unsure", "Probably", or "Yes, definitely".

**Table S9. Distributional treatment effects for "Don't know" respondents**

| | Average | Above | Between | Under |
|---|---|---|---|---|
| Broad Treatment | 0.043*** | 0.019 | 0.042** | 0.067*** |
| | (0.007) | (0.020) | (0.017) | (0.019) |
| Narrow Treatment | 0.034*** | 0.021 | 0.032* | 0.065*** |
| | (0.007) | (0.019) | (0.018) | (0.017) |
| Observations | 266,206 | 21,598 | 24,771 | 34,998 |
| $R^2$ | 0.619 | 0.410 | 0.642 | 0.653 |
| Adjusted $R^2$ | 0.619 | 0.406 | 0.640 | 0.651 |
| Residual Std. Error | 0.797 | 0.780 | 0.677 | 0.815 |
| F Statistic | 1332.485*** | 36.933*** | 159.012*** | 280.241*** |

*p<0.1; **p<0.05; ***p<0.01

The joint test that the broad and narrow coefficients are equal across groups has a p-value of 0.28, and the two-sided test that the broad (narrow) treatment effects in the Under and Above groups are equal has a p-value of 0.09 (0.09)

**Table S10. Heterogeneous Treatment Effects: Baseline Beliefs**

| | Average | No | Don't Know | Yes |
|---|---|---|---|---|
| Broad Treatment | 0.043*** | 0.044** | 0.084*** | 0.031*** |
| | (0.007) | (0.021) | (0.013) | (0.007) |
| Narrow Treatment | 0.034*** | 0.010 | 0.065*** | 0.025*** |
| | (0.007) | (0.019) | (0.013) | (0.007) |
| Observations | 266,206 | 40,104 | 54,015 | 169,400 |
| $R^2$ | 0.619 | 0.215 | 0.115 | 0.120 |
| Adjusted $R^2$ | 0.619 | 0.212 | 0.113 | 0.119 |
| Residual Std. Error | 0.797 | 0.979 | 0.733 | 0.700 |
| F Statistic | 1332.485*** | 43.359*** | 18.350*** | 34.791*** |

*p<0.1; **p<0.05; ***p<0.01

The joint test that the broad and narrow coefficients are equal across groups has a p-value of <0.001, and the two-sided test that the broad (narrow) treatment effects in the Yes and Don't know groups are equal has a p-value of <0.001 (<0.001). The two sided test that the broad (narrow) treatment effects in the Don't know and No groups are equal has a p-value of 0.01 (0.02).

**Table S11. Heterogeneous Treatment Effects: Baseline Vaccine Acceptance**

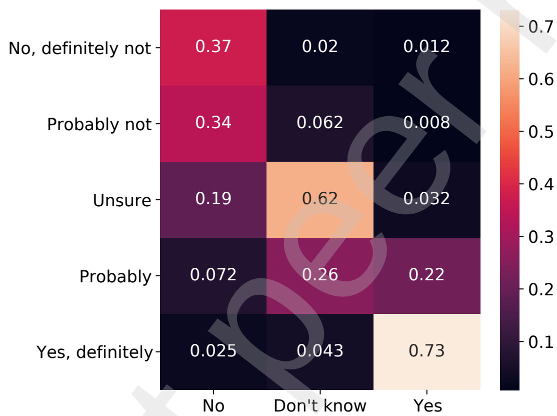


**Fig. S10.** Correlation of Baseline Vaccine Acceptance and Outcome (Detailed) Vaccine Acceptance

Heatmap showing relationship between baseline vaccine acceptance question (x-axis) and the outcome vaccine acceptance question (y-axis) for the control users. Each cell shows the probability of an outcome response conditional on the baseline response and each column sums to one.

**S6.2. Robustness checks.** One concern with survey experiments such as ours is the results could reflect researcher demand effects, where subjects respond how they think the researchers would want them to respond. While we cannot rule this out completely, we do not believe this is driving our results (cf. *S41*, *S42*). We may be less worried about researcher demand effects in this survey as it has a more general advertised purpose and it covers several topics, so normative information is not particularly prominent. Furthermore, unlike other sampling frames with many sophisticated study participants (e.g., Amazon Mechanical Turk), respondents are recruited from a broader population (Facebook users). Moreover, we may expect researcher demand effects to be smaller when the information treatment and the outcome are not immediately adjacent. In all cases, for the vaccine acceptance outcome, there is always at least one intervening screen of questions (the future mask-wearing and distancing intentions questions). Furthermore, they are often separated by more than this. We consider a subset of respondents where the treatment and the outcome are separated by at least one "block" of questions between them. Results of this analysis are presented in Figure S11 and Table S13. The treatment effect estimates on this smaller sample are less precise, but both positive. For vaccines, the p-values that the treatment effect is equal across this smaller sample and the broader sample are 0.03 and 0.05 for the broad and narrow treatments, respectively.

Figure S12 plots the distribution of the number of screens between treated and control. In Figure S12a, we plot the distribution for the entire sample and in Figure S12b we plot the distribution for the subset of those with at least one block between treatment and control. For this group there are at least three pages between the treatment and outcome

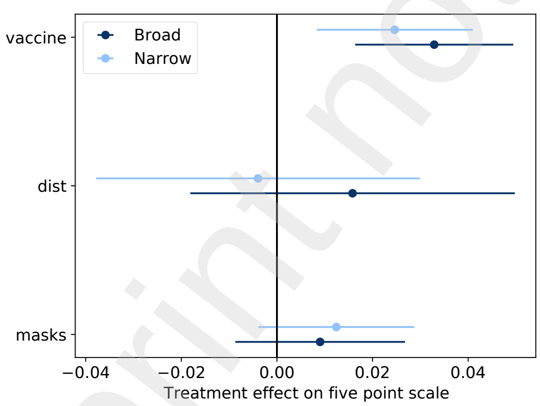


**(a)** Treatment vs. Control

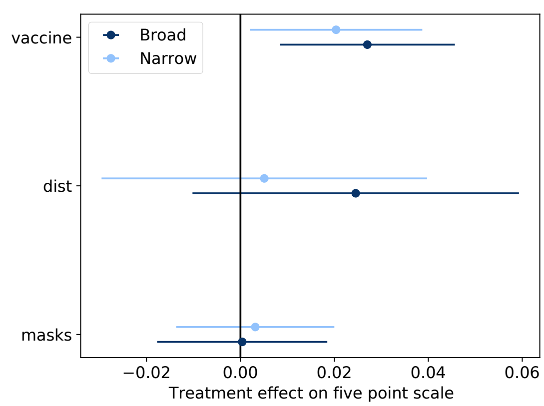


**(b)** Treatment vs. Other Behavior Treatment

**Fig. S11.** Robustness to Researcher Demand Effects

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| Broad Treatment | 0.000 | 0.025 | 0.027*** | 0.009 | 0.016 | 0.033*** |
| | (0.009) | (0.018) | (0.010) | (0.009) | (0.017) | (0.008) |
| Narrow Treatment | 0.003 | 0.005 | 0.020** | 0.012 | -0.004 | 0.025*** |
| | (0.009) | (0.018) | (0.009) | (0.008) | (0.017) | (0.008) |
| Control: Other Treatment | X | X | X | | | |
| Behavior | masks | dist | vaccine | masks | dist | vaccine |
| Number Controls | 76446 | 27441 | 51884 | 115840 | 41734 | 114567 |
| Number Treated | 38577 | 13784 | 61847 | 38577 | 13784 | 61847 |
| Observations | 115,023 | 41,225 | 113,731 | 154,417 | 55,518 | 176,414 |
| $R^2$ | 0.213 | 0.209 | 0.622 | 0.225 | 0.218 | 0.615 |
| Adjusted $R^2$ | 0.212 | 0.207 | 0.621 | 0.224 | 0.216 | 0.615 |
| Residual Std. Error | 0.708 | 0.865 | 0.801 | 0.710 | 0.870 | 0.804 |
| F Statistic | 72.340*** | 47.719*** | 578.290*** | 100.821*** | 60.786*** | 874.199*** |

*p<0.1; **p<0.05; ***p<0.01

Estimates of equation 3 on the restricted sample when outcome and treatment are separated by at least one additional block of questions.

**Table S12. Robustness to Greater Separation of Treatment and Outcome**

| | > No, definitely not | > Probably not | > Unsure | > Probably |
|---|---|---|---|---|
| Intercept | 0.918*** | 0.845*** | 0.664*** | 0.471*** |
| | (0.001) | (0.001) | (0.002) | (0.002) |
| Narrow Treatment | -0.001 | 0.003 | 0.007** | 0.015*** |
| | (0.002) | (0.003) | (0.003) | (0.004) |
| Broad Treatment | 0.000 | 0.003 | 0.014*** | 0.016*** |
| | (0.002) | (0.003) | (0.003) | (0.004) |
| Observations | 176,414 | 176,414 | 176,414 | 176,414 |
| $R^2$ | 0.291 | 0.497 | 0.565 | 0.456 |
| Adjusted $R^2$ | 0.291 | 0.497 | 0.565 | 0.456 |
| Residual Std. Error | 0.232 | 0.256 | 0.311 | 0.369 |
| F Statistic | 83.088*** | 324.582*** | 960.973*** | 807.161*** |

*p<0.1; **p<0.05; ***p<0.01

Estimates of equation 3 on the restricted sample when outcome and treatment are separated by at least one additional block of questions. The outcome variable in this analysis are binary indicators if the outcome was at least a certain response as in table S8.

**Table S13. Robustness to Greater Separation of Treatment and Outcome: Distributional Treatment Effects**

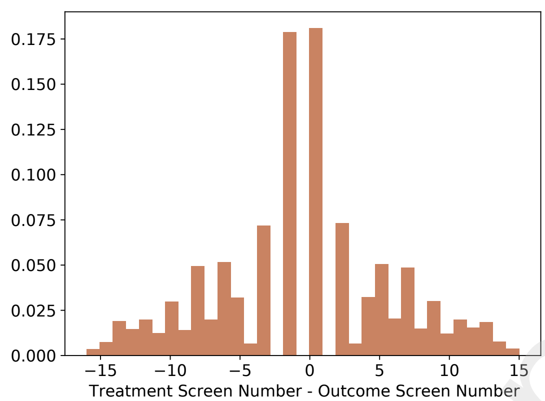


**(a)** Number of Screens Between Treatment and Outcome

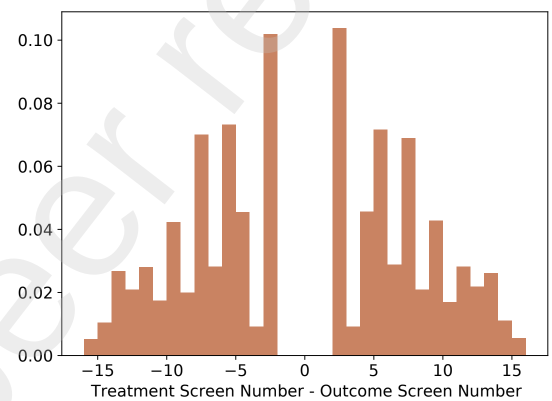


**(b)** Number of Screens Between Treatment and Outcome in Robustness Check Sample

**Fig. S12.** Separation of Treatment and Outcome

(a) Histogram of the number of screens between treatment and outcome. Negative numbers represent treated respondents and positive numbers are control respondents. The distribution is not smooth as the randomized order is at the block level, and blocks have varying number of screens (pages) within them. (b) The same histogram, but for the set of respondents with at least one block between treatment and outcome.

## S7. Norm–intention correlations

In the main text, Figure 1 (inset) shows the association between beliefs about descriptive norms and intentions to accept a COVID-19 vaccine. Figure S13 disaggregates this information by country. As in the main text, this is a purely observational association but is computed on the main experimental sample (i.e., starting in late October).

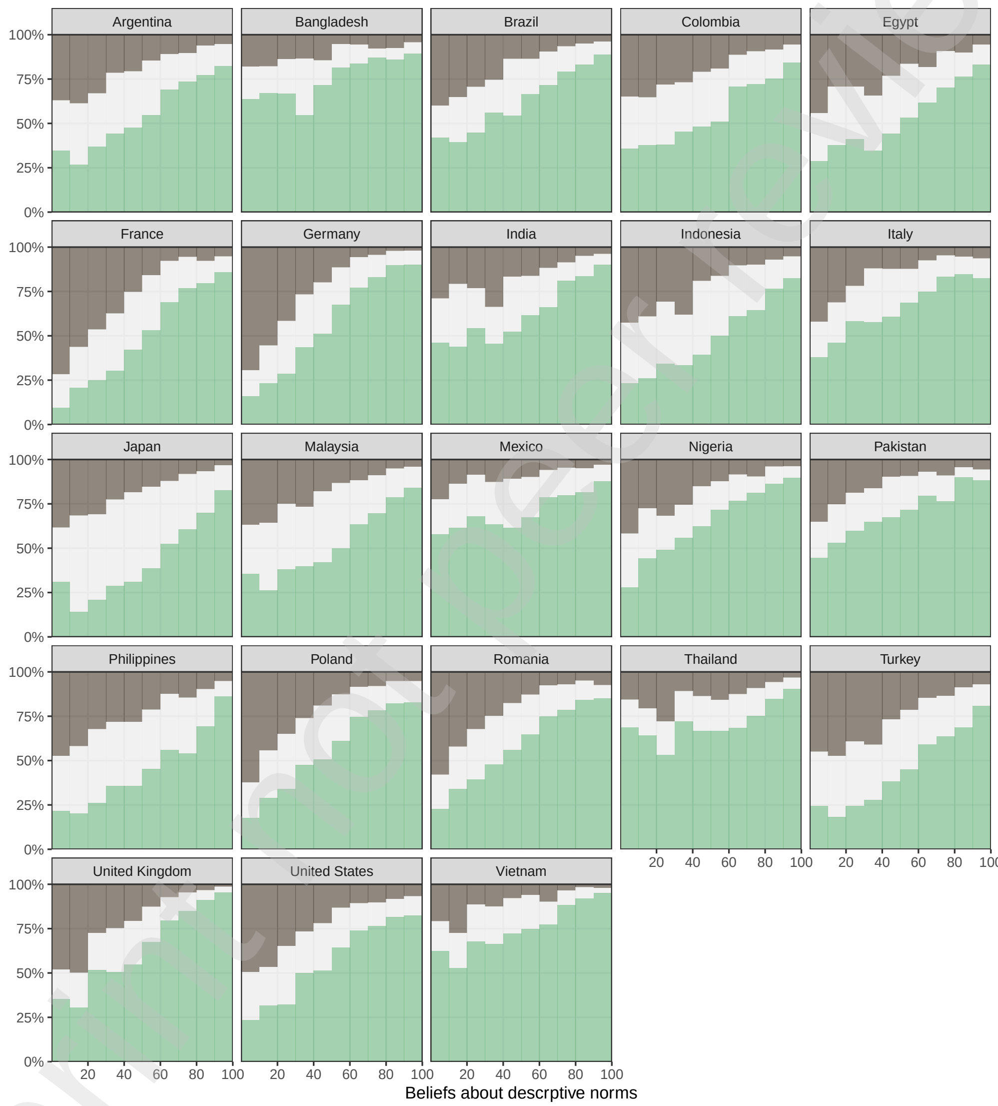


**Fig. S13.** People who believe a larger fraction of their community will accept a vaccine are on average more likely to say they will accept a vaccine, and this is true within the 23 included countries.

1. AA Malik, SM McFadden, J Elharake, SB Omer, Determinants of COVID-19 vaccine acceptance in the US. *EClinicalMedicine* **26**, 100495 (2020).
2. S Sanche, et al., High contagiousness and rapid spread of severe acute respiratory syndrome coronavirus 2. *Emerging Infectious Diseases* **26**, 1470–1477 (2020).
3. RM Anderson, C Vegvari, J Truscott, BS Collyer, Challenges in creating herd immunity to SARS-CoV-2 infection by mass vaccination. *The Lancet* **396**, 1614–1616 (2020).
4. C Signorelli, S Iannazzo, A Odone, The imperative of vaccination put into practice. *The Lancet Infectious Diseases* **18**, 26–27 (2018).
5. SB Omer, C Betsch, J Leask, Mandate vaccination with care. *Nature*, 469–472 (2019).
6. C Betsch, R Böhm, Detrimental effects of introducing partial compulsory vaccination: Experimental evidence. *The European Journal of Public Health* **26**, 378–381 (2016).
7. C Betsch, R Böhm, GB Chapman, Using behavioral insights to increase vaccination policy effectiveness. *Policy Insights from the Behavioral and Brain Sciences* **2**, 61–73 (2015).
8. B Nyhan, J Reifler, S Richey, GL Freed, Effective messages in vaccine promotion: a randomized trial. *Pediatrics* **133**, e835–e842 (2014).
9. B Nyhan, J Reifler, Does correcting myths about the flu vaccine work? An experimental evaluation of the effects of corrective information. *Vaccine* **33**, 459–464 (2015).
10. J Green, et al., Report 36: Evaluation of COVID-19 vaccine communication strategies. *The COVID States Project: A 50-State COVID-19 Survey Reports* (2021).
11. B Nyhan, J Reifler, S Richey, The role of social networks in influenza vaccine attitudes and intentions among college students in the southeastern united states. *Journal of Adolescent Health* **51**, 302–304 (2012).
12. EK Brunson, The impact of social networks on parents' vaccination decisions. *Pediatrics* **131**, e1397–e1404 (2013).
13. CT Bauch, DJ Earn, Vaccination and the theory of games. *Proceedings of the National Academy of Sciences* **101**, 13391–13394 (2004).
14. T Oraby, V Thampi, CT Bauch, The influence of social norms on the dynamics of vaccinating behaviour for paediatric infectious diseases. *Proceedings of the Royal Society B: Biological Sciences* **281**, 20133172 (2014).
15. NT Brewer, GB Chapman, AJ Rothman, J Leask, A Kempe, Increasing vaccination: Putting psychological science into action. *Psychological Science in the Public Interest* **18**, 149–207 (2017).
16. SS DeRoo, NJ Pudalov, LY Fu, Planning for a COVID-19 vaccination program. *JAMA* **323**, 2458–2459 (2020).
17. H Allcott, et al., Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics* **191**, 104254 (2020).
18. A Bridgman, et al., The causes and consequences of COVID-19 misperceptions: Understanding the role of news and social media. *Harvard Kennedy School Misinformation Review* **1** (2020).
19. A Simonov, SK Sacher, JPH Dubé, S Biswas, The persuasive effect of Fox News: Non-compliance with social distancing during the COVID-19 pandemic, (National Bureau of Economic Research), Technical report (2020).
20. D Holtz, et al., Interdependence and the cost of uncoordinated responses to COVID-19. *Proceedings of the National Academy of Sciences* **117**, 19837–19843 (2020).
21. A Banerjee, AG Chandrasekhar, E Duflo, MO Jackson, Using gossips to spread information: Theory and evidence from two randomized controlled trials. *The Review of Economic Studies* **86**, 2453–2490

(2019).
22. V Alatas, AG Chandrasekhar, M Mobius, BA Olken, C Paladines, When celebrities speak: A nationwide Twitter experiment promoting vaccination in Indonesia, (National Bureau of Economic Research), Technical Report 25589 (2019).
23. CT Bauch, S Bhattacharyya, Evolutionary game theory and social learning can determine how vaccine scares unfold. *PLoS Comput Biol* **8**, e1002452 (2012).
24. N Rao, MM Mobius, T Rosenblat, Social networks and vaccination decisions, (Federal Reserve Board of Boston), Technical report (2007).
25. JT Vietri, M Li, AP Galvani, GB Chapman, Vaccinating to help ourselves and others. *Medical Decision Making* **32**, 447–458 (2012).
26. E Shim, GB Chapman, JP Townsend, AP Galvani, The influence of altruism on influenza vaccination decisions. *Journal of The Royal Society Interface* **9**, 2234–2243 (2012).
27. L Korn, R Böhm, NW Meier, C Betsch, Vaccination as a social contract. *Proceedings of the National Academy of Sciences* **117**, 14890–14899 (2020).
28. Y Ibuka, M Li, J Vietri, GB Chapman, AP Galvani, Free-riding behavior in vaccination decisions: An experimental study. *PloS one* **9**, e87164 (2014).
29. R Böhm, C Betsch, L Korn, Selfish-rational non-vaccination: Experimental evidence from an interactive vaccination game. *Journal of Economic Behavior & Organization* **131**, 183–195 (2016).
30. JC Hershey, DA Asch, T Thumasathit, J Meszaros, VV Waters, The roles of altruism, free riding, and bandwagoning in vaccination decisions. *Organizational Behavior and Human Decision Processes* **59**, 177–187 (1994).
31. R Sato, Y Takasaki, Peer effects on vaccination behavior: Experimental evidence from rural nigeria. *Economic Development and Cultural Change* **68**, 93–129 (2019).
32. A Collis, et al., Global Survey on COVID-19 Beliefs, Behaviors, and Norms, (MIT Sloan School of Management), Technical report (2020).
33. N Barkay, et al., Weights and methodology brief for the COVID-19 symptom survey by University of Maryland and Carnegie Mellon University, in partnership with Facebook. *arXiv preprint arXiv:2009.14675* (2020).
34. JV Lazarus, et al., A global survey of potential acceptance of a COVID-19 vaccine. *Nature Medicine*, 1–4 (2020).
35. C Betsch, L Korn, C Holtmann, Don't try to convert the antivaccinators, instead target the fence-sitters. *Proceedings of the National Academy of Sciences* **112**, E6725–E6726 (2015).
36. AGCMH Wen-Ying Sylvia Chou, Caitlin E. Burgdorf, COVID-19 vaccination communication: Applying behavioral and social science to address vaccine hesitancy and foster vaccine confidence, (National Institutes of Health), Technical report (2020).
37. R Palm, T Bolsen, JT Kingsland, The effect of frames on COVID-19 vaccine hesitancy. *medRxiv* (2021).
38. E Paul, A Steptoe, D Fancourt, Attitudes towards vaccines and intention to vaccinate against COVID-19: Implications for public health communications. *The Lancet Regional Health-Europe*, 100012 (2020).
39. W Lin, Agnostic notes on regression adjustments to experimental data: Reexamining freedman's critique. *Annals of Applied Statistics* **7**, 295–318 (2013).
40. LW Miratrix, JS Sekhon, B Yu, Adjusting treatment effect estimates by post-stratification in randomized experiments. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **75**, 369–396 (2013).

41. J De Quidt, J Haushofer, C Roth, Measuring and bounding experimenter demand. *American Economic Review* **108**, 3266–3302 (2018).
42. J Mummolo, E Peterson, Demand effects in survey experiments: An empirical assessment. *American Political Science Review* **113**, 517–529 (2019).