

Correcting Misperceptions about Support for Social Distancing to Combat COVID-19[†]

James Allen IV¹, Arlete Mahumane², James Riddell IV³, Tanya Rosenblat^{4,5}, Dean Yang^{5,6,7}, and Hang Yu^{8,9}

¹International Food Policy Research Institute (IFPRI)

²Beira Operational Research Center, National Institute of Health, Mozambique

³Division of Infectious Diseases, University of Michigan Medical School

⁴School of Information, University of Michigan

⁵Department of Economics, University of Michigan

⁶Ford School of Public Policy, University of Michigan

⁷Population Studies Center, University of Michigan

⁸National School of Development, Peking University

⁹Institute of South-South Cooperation and Development, Peking University

Version: July 28, 2023

Short Title: Correcting Social Distancing Misperceptions

Corresponding Author: Dean Yang

Address: 735 S. State St. 3315 Ann Arbor, MI 48109

Telephone: (734) 764-6158

Email: deanyang@umich.edu

JEL Classification: I12, D91, O12

Keywords: COVID-19, Social Distancing, Health Behavior, Health Policy, Behavioral Economics

[†]Contacts: James Allen IV: j.allen@cgiar.org; Dean Yang: deanyang@umich.edu. Faustino Lessitala provided top-notch leadership and field management. Patricia Freitag, Ryan McWay, and Maggie Barnard provided excellent research assistance. Julie Esch, Laura Kaminski, and Lauren Tingwall's grant management was world-class. We appreciate feedback from participants in Michigan's Health, History, Development, and Demography (H2D2) Seminar and the Conference on Experimental Insights from Behavioral Economics on COVID-19 (JHU-LSE). This work is supported by the Abdul Latif Jameel Poverty Action Lab (J-PAL) Innovation in Government Initiative through a grant from The Effective Altruism Global Health and Development Fund (grant number IGI-1366), the UK Foreign, Commonwealth & Development Office awarded through Innovations for Poverty Action (IPA) Peace & Recovery Program (grant number MIT0019-X9), the Michigan Institute for Teaching and Research in Economics (MITRE) Ulmer Fund (grant number G024289), and the National Institute on Aging of the National Institutes of Health (award number T32AG000221). Our protocols were reviewed and approved by Institutional Review Boards (IRBs) at the University of Michigan (Health Sciences and Social and Behavioral Sciences IRB, approval number HUM00113011) and the Mozambique Ministry of Health National Committee on Bioethics for Health (CNBS reference number 302/CNBS/20). The study was submitted to the AEA RCT Registry on May 26, 2020, registration ID number AEARCTR-0005862: 10.1257/rct.5862. Replication data, code, and codebooks are provided through Dataverse at <https://doi.org/10.7910/DVN/TMARZT>. The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of the aforementioned institutions.

Abstract

Can informing people of high community support for social distancing encourage them to do more of it? We randomly assigned a treatment correcting individuals' underestimates of community support for social distancing. In theory, informing people that more neighbors support social distancing than expected encourages *free-riding* and *lowers* the perceived benefits from social distancing. At the same time, the treatment induces people to revise their beliefs about the infectiousness of COVID-19 upwards; this *perceived infectiousness* effect as well as the *norm adherence* effect *increase* the perceived benefits from social distancing. We estimate impacts on social distancing, measured using a combination of self-reports and reports of others. While experts surveyed in advance expected the treatment to increase social distancing, we find that its average effect is close to zero and significantly lower than expert predictions. However, the treatment's effect is heterogeneous, as predicted by theory: it decreases social distancing where current COVID-19 cases are low (where free-riding dominates), but increases it where cases are high (where the perceived-infectiousness effect dominates). These findings highlight that correcting misperceptions may have heterogeneous effects depending on disease prevalence.

1 Introduction

Attitudes toward social distancing have changed rapidly during the pandemic (Janzwood, April 27, 2020). During such rapid change, people often underestimate support for social distancing in their communities. Early in the pandemic, 98% of our Mozambican sample thought that people should be social distancing, but estimated that only 69% of others in the community felt similarly. This gap motivates a public health policy: simply inform people of high rates of community support for social distancing. What impact would such messaging have on social distancing behavior?

In theory, the impact of such a "misperceptions correction" intervention on social distancing is ambiguous: on the one hand, informing people that more neighbors support social distancing than expected encourages *free-riding* and *lowers* the perceived benefits from social distancing. On the other hand, people should revise their belief about the seriousness of COVID-19 upwards in order to rationalize the observed number of infections in their neighborhood despite the higher than expected social distancing support. This *perceived infectiousness* effect *increases* the perceived benefits from social distancing and dominates free-riding in communities with high levels of infections.¹ Finally, the *norm adherence* effect should induce people to follow whatever local social norm is set by their neighbors - in our case this effect should always increase social distancing.

We implemented a randomized controlled trial testing the impact of informing people about high local support for social distancing. The treatment either updated beliefs upwards or confirmed beliefs about high rates of support for social distancing. Abiding by COVID-19 protocols, we conducted all treatments and surveys by phone among 2,117 Mozambican households.

Our outcome variable is the extent to which a household engages in social distancing. Measuring this behavior is challenging due to experimenter demand effects.² Yet most prior studies ask for self-reports about general social distancing compliance. When we do so, 95% claim to observe government social distancing recommendations. We therefore construct a novel measure of social distancing. First, we ask respondents to self-report several social distancing actions. Second, we ask *others* in the community to report on the respondent's social distancing. We are aware of no prior study that makes use of other-reports on a respondent's social distancing behavior. Incorporating self-reported actions and others' reports drops social distancing to a more discerning 8% (see Figure 1 and Section 3.3). Improved measurement leads the social distancing rate to fall by an order of magnitude.

1. Our model is related to the literature on decision-making under misspecified subjective models (Spiegler 2020). Agents hold incorrect assumptions on one model parameter (e.g., share of population social distancing), leading them to incorrect conclusions about other parameters (e.g., disease infectiousness).

2. Jakubowski et al. (2021) find that self-reported mask wearing is overstated relative to measures based on observations of others.

The average effect of the misperceptions correction treatment in the full sample is small and not statistically significantly different from zero. However, as theory predicts, there is substantial treatment effect heterogeneity: the treatment effect is statistically significantly more positive when local COVID-19 cases (per 100,000 population) are higher. In districts with few cases, the treatment effect is negative. In the district with the most COVID-19 cases, the treatment increases social distancing by 9.2 percentage points (statistically significant at the 5% level), a 70% increase over that district’s control-group mean.

This pattern is consistent with the theoretical prediction that as infection rates rise, the perceived-infectiousness effect should increasingly dominate the free-riding effect of the misperceptions correction treatment, leading the treatment effect to become more positive. We also test a further implication of the model: expectations of future infection rates should show similar treatment effect heterogeneity. Empirical analyses confirm this prediction, providing additional support for the theoretical model.

Alongside the social norm correction treatment, we also randomly assigned a “leader endorsement” treatment (an endorsement of social distancing by a community leader). The leader endorsement treatment has a very small effect on social distancing that is not statistically significantly different from zero. We also find no treatment effect heterogeneity for this treatment with respect to COVID-19 cases.

This paper contributes to understanding the impact of providing information about others’ beliefs and attitudes (Benabou and Tirole 2011; Bicchieri and Dimant 2019). In health settings, Yu (2020) and Yang et al. (2021) find (in an overlapping Mozambican sample) that correcting overestimates of stigmatizing attitudes promoted HIV testing, though Banerjee, La Ferrara, and Orozco-Olvera (2019) find that informing Nigerian young adults of peers’ attitudes on healthy sexual relationships did not change respondents’ own attitude.³ Regarding social distancing, Martinez et al. (2021) show that respondents are influenced by others’ social distancing actions in hypothetical vignettes; however, no prior study has tested the impact of providing information on community support of social distancing on respondent behavior.

Our emphasis on interactions between free-riding and perceived-infectiousness effects is novel, but each effect has been studied separately. Free-riding has been studied in the context of vaccination decisions (Hershey et al. 1994; Lau et al. 2019) and social distancing (Cato et al. 2020) and in similar Mozambican settings (Fafchamps, Vaz, and Vicente 2020). Perceived COVID-19 infection risk (e.g., due to vaccine anticipation, Andersson et al. (2021)) has been shown to lower social distancing intentions.

3. In other contexts, correcting misperceptions of community support or approval (i.e., the injunctive norm) has also been shown to change energy consumption (Schultz et al. 2007), female labor force participation (Bursztyn, González, and Yanagizawa-Drott 2020), donations to charities addressing climate change (Andre et al. 2021), and recycling program participation (Fuhrmann-Riebel et al. 2023).

2 Theory

Our model focuses on the interaction between the free-riding and perceived infectiousness effects for communities with low and high overall infection rates. We view norm-adherence as a uniform effect that should always increase social distancing.

We consider a community where people have random pairwise meetings. People believe that a share x of the population supports social distancing and that the probability of becoming infected from unprotected meetings is α . People treat x as given, but infer the infectiousness α from the current infection rate R in the community which they can observe (we describe this inference below). The true infectiousness of the disease is $\hat{\alpha}$.

Importantly, people in the community have *miscalibrated beliefs*: the true share of the population supporting social distancing is \hat{x} (we are interested in the case $\hat{x} > x$). People infer the true infectiousness $\hat{\alpha}$ of the disease only if they are correctly calibrated ($\hat{x} = x$).

Individual Effort A supporter engages in preventative effort e and assumes that other supporters choose effort e^* (in equilibrium we have $e = e^*$). Non-supporters choose effort $e = 0$.

When someone supporting social distancing meets another person, she escapes exposure with probability:

$$\begin{aligned}
 A(e, e_{other}) &= \sqrt{e + e_{other}} \\
 &= \begin{cases} \sqrt{e + e^*} & \text{if other person is supporter} \\ \sqrt{e} & \text{if other person is non-supporter} \end{cases} \quad (1)
 \end{aligned}$$

Hence, the marginal benefit of effort decreases both with own effort e as well as the other person's effort e^* .⁴

The expected probability of escaping exposure is therefore:

$$\bar{A}(e, e^*) = (1 - x)\sqrt{e} + x\sqrt{e + e^*} \quad (2)$$

An agent becomes exposed with probability $1 - \bar{A}(e, e_{other})$. If exposed she gets infected with probability α and suffers disutility $-C$ from infection.⁵ If she is not exposed then she does not get infected. Her baseline

4. We assume the other person's effort is unobservable. This is consistent with our finding that respondents underestimate the extent of social distancing.

5. For simplicity, we assume that infectiousness does not vary with the agent's type (supporter or non-supporter). Otherwise, we would need to keep track of two levels of infectiousness. The qualitative results would not change.

utility from no infection equals \bar{U} . The cost of preventative effort is e . Hence, her total utility equals:

$$\bar{U} - \alpha(1 - \bar{A}(e, e_{other}))C - e \quad (3)$$

The agent chooses e to maximize her utility, giving us the following first-order condition:

$$\frac{\alpha C}{2\sqrt{e}} \left[1 - x \left(1 - \frac{1}{\sqrt{1 + \frac{e^*}{e}}} \right) \right] = 1 \quad (4)$$

In equilibrium it has to be the case that the population effort e^* equals e . Hence, we can characterize equilibrium effort as:

$$e = \left(\frac{\alpha C}{2} \left[1 - x \left(1 - \frac{1}{\sqrt{2}} \right) \right] \right)^2 \quad (5)$$

This demonstrates the *free-riding effect*: increasing the share x of supporters *decreases* effort because the marginal utility from own effort decreases. Also, effort increases if the disease is more infectious (higher α) and if illness is costlier (higher C).

Infection Rate People observe the current infection rate in the community. Infections come from two sources: non-supporters become sick at rate $\alpha(1 - x\sqrt{e})$ while supporters become sick at rate $\alpha(1 - \bar{A}(e, e))$. Hence, people in the community assume that the current infection rate is generated by the following process:

$$\begin{aligned} R &= \alpha \left[\underbrace{(1-x)(1-x\sqrt{e})}_{\text{non-supporters}} + \underbrace{x(1-\sqrt{e}(1+(\sqrt{2}-1)x))}_{\text{supporters}} \right] \\ &= \alpha \left[1 - \sqrt{e} 2x \underbrace{\left(1 - x \left(1 - \frac{1}{\sqrt{2}} \right) \right)}_{=G(x)} \right] \end{aligned} \quad (6)$$

However, the true process determining current infections is actually:

$$R = \hat{\alpha} [1 - \sqrt{e}G(\hat{x})] \quad (7)$$

In other words, the true infection process is driven by the same social distancing effort of supporters but different infectiousness $\hat{\alpha}$ and different \hat{x} .

2.1 Basic Equilibrium

Supporters initially assume that the disease has low infectiousness and they adjust their estimate of α upwards until the current infection rate R stabilizes.

Proposition 1 *In equilibrium, effort level e , the current infection rate R , and the assumed infectiousness α satisfy Equations 5, 6 and 7. Moreover, $\hat{\alpha} > \alpha$ if $\hat{x} > x$.*

In equilibrium, both the assumed infection process (Equation 6) and the real infection rate (Equation 7) must produce observed infection rate R . For the second part, note that $G(x)$ is increasing in $x \in [0, 1]$: hence, $\hat{x} > x$ implies $\hat{\alpha} > \alpha$ to generate the same infection rate R .

2.2 Treatment Effect

We now consider the effect of our treatment informing people that the population share supporting social distancing is really $\hat{x} > x$.

Proposition 1 implies that if supporters are informed that the true population share supporting social distancing is $\hat{x} > x$, they must infer higher disease infectiousness than they initially assumed (because their estimated disease infectiousness immediately jumps from α to true $\hat{\alpha}$). This is the *perceived-infectiousness effect*.

Supporters of social distancing will adjust their effort level to a new level \hat{e} , but there are two countervailing effects:

1. Holding assumed infectiousness α constant, the free-riding effect *decreases* effort.
2. The perceived-infectiousness effect *increases* effort, because the agent now believes the disease is more infectious than initially thought (perceived α increases), increasing the gain from social distancing.

Intuitively, the perceived-infectiousness effect varies monotonically with R : when infections are low, supporters' effort is low, and both supporters and non-supporters get infected at similar rates. Hence, agents revise the estimate of infectiousness α only slightly upwards in response to the treatment. On the other hand, when infections are high, supporters' effort is high and the upward revision will be larger.

The following theorem makes this intuition precise. Instead of doing comparative statics on R (which is determined in equilibrium) we state the comparative statics results in terms of the infectiousness $\hat{\alpha}$ (for given x and \hat{x}). Note that R increases with $\hat{\alpha}$.

Theorem 1 *Assume an agent is informed that a share $\tilde{x} > x$ of the population supports social distancing. Then there is a threshold $\hat{\alpha}^*$ such that for any $\hat{\alpha} < \hat{\alpha}^*$ the free-riding effect dominates and equilibrium effort decreases, and for $\hat{\alpha} > \hat{\alpha}^*$ the perceived-infectiousness effect dominates and the equilibrium effort increases.*

See Appendix A for the proof.

The interplay between free-riding and perceived-infectiousness effects also yields analogous predictions about a central belief about COVID-19: the future infection rate. In the endline survey, we ask respondents to estimate this. The expected future rate differs from the current infection rate R , because this study occurs at a point when infection rates are clearly evolving. The misperceptions correction treatment changes respondent beliefs about social distancing support and about infectiousness, and therefore should change expected future infection rates. Recall that non-supporters are always infected with higher probability than supporters. The higher the infectiousness parameter $\hat{\alpha}$, the higher should be future infection rates for both groups. When $\hat{\alpha}$ is currently small, the perceived-infectiousness effect is small. Simultaneously, the treatment corrects beliefs about the share of social-distancing supporters upwards, which should *reduce* estimates of future infection rates because supporters have lower infection rates. Thus, the expected future infection rate *decreases* when $\hat{\alpha}$ is currently small. In contrast, when $\hat{\alpha}$ is currently large, the treatment leads to a large increase in perceived infectiousness, implying that the disease will infect higher shares of both supporters and non-supporters. This will tend to *increase* expected future infection rates.

To summarize, the misperceptions correction treatment effect on the expected future infection rate should show heterogeneity similar to that described in Theorem 1. The treatment effect on the expected future infection rate is strictly negative if the current local infection rate (R) (which moves monotonically with $\hat{\alpha}$) is small enough. The treatment effect on the expected future infection rate increases with the current infection rate, and can become positive if current infection rates are sufficiently high.

In our empirical analyses, we test these predictions regarding heterogeneity in the misperceptions correction treatment effect.

3 Sample and Data

3.1 Data

We implemented three rounds of surveys by phone in July–November 2020: a pre-baseline, baseline and endline survey (see Figure A.2 for a study timeline). Respondents were drawn across 76 communities in Central Mozambique from a sample of a prior study (Yang et al. 2021) that focused on HIV-vulnerable households—a policy-relevant sample especially vulnerable to COVID-19.⁶ To avoid risk of spreading COVID-19 via in-person interaction with study participants, we also limited the sample to those households with phones.

6. AEA RCT Registry for Yang et al. (2021): <https://doi.org/10.1257/rct.3990-5.1>. In that prior study, we run a randomized evaluation of a bundled community-level HIV/AIDS program whose main component was home visits by case care workers to promote HIV testing to HIV-vulnerable households, such as those with HIV-positive or other chronically ill members, orphaned children, or a grandparent as the household head. In this study, we use community-stratified randomization and regress with community fixed effects to rule out the influence of this prior intervention on our results.

Thus both HIV-vulnerability and phone ownership are two relevant factors to bear in mind when considering the external validity of the results. We surveyed one adult per household. Appendix B provides details on the COVID-19 context, study communities and the study timeline.

Between a pre-baseline survey and baseline survey, we randomly assigned households to treatments and registered a pre-analysis plan (PAP). The baseline survey was immediately followed by over-the-phone treatment implementation. There was a minimum of 3.0 weeks and average of 6.3 weeks between baseline and endline surveys for all respondents. Baseline and endline surveys occurred when COVID-19 cases were rising rapidly.

The endline sample size is 2,117 respondents, following a sample size of 2,226 at baseline. The retention rate between baseline and endline is 95.1% overall, at least 94.4% in each of the seven districts surveyed, and balanced across treatment conditions. We also surveyed 145 community opinion leaders over the 76 study communities—at least one, an average of 2.11, and at most 4 per community—for inputs to the primary outcome and treatments as described below.

3.2 Measuring Misperceptions

We measure both true and perceived support for social distancing as follows. First, to measure actual community support for social distancing, we asked respondents *"Do you support the practice of social distancing to prevent the spread of coronavirus? (Yes, No, Don't know, Refuse to Answer)"*, which captures the respondent's first-order belief of the injunctive norm for social distancing. We then calculated the fraction of "Yes" responses across the sample and within each community.⁷ Directly after, to measure perceived community support, we asked respondents *"For every 10 households in your community, how many do you think support the practice of social distancing to prevent the spread of coronavirus? (integer 0-10)"*, capturing the respondent's second-order beliefs of the injunctive norm for social distancing within their community. The difference between the true and perceived community support for social distancing is the respondent's misperception of the social norm.

Three possible concerns with our measure of perceived support for social distancing include the role of uncertainty, the restricted scale, and bias from experimenter demand effects. First, a possible concern is that unawareness and uncertainty around new social norms and others' beliefs—plausible at the start of the pandemic—may lead respondents away from the extreme points of the answer scale. However, in Appendix C, we present a cumulative distribution of our perceived community support measure across survey rounds that shows that respondents readily utilized the extreme ends of the scale, with 8% and 35% of the sample

7. The fraction was calculated by dividing the number "Yes" responses by the number of all responses (i.e., Yes, No, Don't know, Refuse to Answer).

at pre-baseline reporting perceived community support of 0% and 100%, respectively, and 51% of the sample at baseline reporting 100%. Second, despite more common use of a 0-100 scale when measuring perceived norms (e.g., Andre et al. 2021; Fuhrmann-Riebel et al. 2023), we simplified our scale to an 11-point 0-10 scale due to past difficulties eliciting "percentage" measures in this context, repeated feedback from our field team and local partners that a 0-100 scale was too complex, and the inability to use a "slider" mechanism over the phone. Given the high concentration of perceived community support at 100% at baseline, the restricted scale may attenuate the treatment effect of the misperceptions correction on perceived community support given that there is "little room to improve" for many respondents in the sample. Third, experimenter demand effects may have led respondents to report higher shares of perceived support for social distancing in order to make their communities look favorable. Such action would lead to an upward-biased estimate of true perceptions of community support and, in turn, an underestimate of the misperception of the social norm, which would also lead to an attenuation of the treatment effect for the misperceptions correction intervention.⁸ We ask the reader to bear in mind these possible limitations when interpreting the results.

3.3 Primary Outcome

The primary outcome is an indicator that the respondent practiced social distancing, as pre-specified in our PAP. It is constructed from self-reports of social distancing as well as others' reports of the respondent's social distancing. The outcome is equal to one if the respondent is practicing social distancing according to both self-reports and other-reports, and zero otherwise.

Respondents are social distancing according to their self-report if both of the following are true: 1) they answer "yes" to "In the past 14 days, have you observed the government's recommendations on social distancing?", and 2) they report doing more than the sample median number of "social distancing actions" in the past seven days. A list of eight social distancing actions and their corresponding summary statistics are presented in Appendix D. At pre-baseline and baseline, respondents were asked about a randomly selected four social distancing actions and, with a sample median of three for both surveys, had to report doing all four actions to be considered social distancing. At endline, respondents were asked about all eight social distancing actions and, with a sample median of six, had to report doing seven or eight actions to be considered social distancing.⁹

8. See Section 4.1 for a description of the misperceptions correction treatment. If upward-biased estimates of perceived support remain less than or equal to true community support, then the misperceptions correction is implemented and may still boost respondents' true perception of community support in a way not captured by our measure; however, if the bias leads to overestimating true community support, then respondents will become ineligible for the misperceptions correction treatment thereby attenuating the treatment effect (but not biasing upwards).

9. While this threshold was pre-specified, results are robust to alternate definitions of this component (see Appendix G.3), such as a threshold of six, or dropping social distancing actions #4 and #6 for which respondents might misinterpret and answer "no" if not showing symptoms.

To collect others' reports on a respondent's social distancing, study participants were asked about their social interactions with ten other community study participants. These ten others were identified from social network data and geographic proximity. Additionally, community leaders were also asked about social interactions with all study participants in their respective community.¹⁰ At baseline, the average respondent household was known by 0.98 community leaders and 3.21 neighboring survey respondents. Other-reports were collected at baseline and endline.

In collecting other-reports, we asked others whether they had seen anyone from the respondent household in the last 14 days.¹¹ If so, we then asked: 1) Did he/she come closer than 1.5 meters to you or others not of his/her household at any point in the last 14 days?; 2) Did he/she shake hands, try to shake hands, or touch you or others not of his/her household in the last 14 days?; and 3) In general, did he/she appear to be observing the government's recommendations on social distancing (avoid large gatherings and keep at least 1.5 meters distance from people not of his/her household)? Respondents are considered to be social distancing according to others if all others responded "no", "no", and "yes" (respectively) to these three questions, reported having not seen the respondent in the past 14 days, or reported not knowing the respondent.¹²

Figure 1 displays how these questions lead to the social distancing outcome. First, 95% of respondents say "yes" to the self-report on general social distancing. When considering self-reports above the sample median number of social distancing actions, the social distancing rate falls to 36%. Finally, incorporating others' reports reduces the rate further to 8%. Limited overlap between self-reports and others' reports of social distancing suggests that each is providing different sets of information. We suspect that self-reports likely over-report social distancing due to experimenter demand bias, whereas others' reports are likely less biased by experimenter demand and rather over-report due to recall bias or lack of observation (as respondents not known or not seen in the past 14 days were not assumed to violate social distancing behavior).¹³ Together, we believe the combined measure is a novel improvement from simple self-reports, though we leave comparison of both measurement methods to observed behavior as an avenue for future work. Incorporating additional information into the social distancing measure—using self-reports of specific social distancing behaviors as well as other-reports—leads to substantially lower social distancing rates.

10. The average community leader was asked about 33.90 households (std. dev.=22.10, minimum=2, second-highest=99, maximum=228—a special case where one individual was the traditional leader across multiple communities). To mitigate survey fatigue, leaders were told upfront of the number and offered a stepwise incentive that increased for each additional set of 25 study households.

11. As is common in this context, households were identified by the name of the household head and a list of other known household members.

12. At baseline, 90.55% of respondent households were known by some other respondent or community leader.

13. For example, complete lack of observation by others was true for 9% of the sample (see footnote 12).

4 Research Design

4.1 Treatments

We implemented a randomized controlled trial estimating impacts on social distancing of two treatments: 1) misperceptions correction, and 2) leader endorsement.¹⁴ Before the baseline survey, we randomly assigned 30% of households completing the pre-baseline survey each to one of two treatments and the remaining 40% to a control group. Sample sizes by treatment condition were as follows: misperceptions correction (N=628, 29.7% of sample), leader endorsement (N=637, 30.1%), and control group (N=852, 40.3%). Treatment scripts are located in Appendix E.

For the misperceptions correction treatment, we used the following data: 1) respondents' own support for social distancing from the pre-baseline survey, from which we estimated the true share of community support for social distancing (as the fraction of respondents expressing support within the community), and 2) respondents' perceived share of community support for social distancing at baseline (reported as an integer out of 10). Immediately after completing the baseline survey, treated individuals underestimating the share were told the true share supporting social distancing, rounded to an integer out of 10.¹⁵ Treated individuals correctly estimating the share were also told that they were correct. In practice, 92.4% of treated respondents received this treatment, 53.2% of whom underestimated community support for social distancing and 46.8% of whom correctly estimated it. The small minority overestimating the share were not provided additional information.¹⁶

For the leader endorsement treatment, we identified and surveyed community opinion leaders prior to the baseline survey and requested their permission to tell others in their community that they "support social distancing, are practicing social distancing, and encourage others to do the same". Then, in this treatment, we reported this endorsement to respondents, mentioning the community leader(s) by name.¹⁷

Attrition between baseline and endline is low (4.9%). In Appendix F, we show that attrition and key baseline variables are balanced across treatment conditions. Further, at endline, 97.9% recall receiving the baseline survey and, of those, 99.4% report trusting the COVID-19 information we provided.¹⁸

14. These two treatments were registered in a pre-analysis plan uploaded to the AEA RCT Registry (registration ID number AEARCTR-0005862: <https://doi.org/10.1257/rct.5862-3.0>) prior to the start of the intervention at baseline. Previously, our AEA RCT pre-trial profile had also included a third social distancing treatment arm proposing to provide individuals with information on the private and public value of social distancing; however, we cut this treatment to improve power for the remaining treatments prior to registering the pre-analysis plan.

15. In 63 out of 76 communities (82.9%) the number we convey to respondents is 10 out of 10, and in 13 communities (17.1%) the number is 9 out of 10.

16. While respondents were not incentivized to truthfully guess community support (for scalability), true beliefs can still be updated for all except those who overestimated true community support with an upward biased guess; however, the latter case should only attenuate our treatment effect and not bias it upward.

17. Communities had at least one and an average of 2.09 endorsements from community leaders (std. dev.=0.94, maximum=4).

18. Trust may have arisen from multiple in-person household surveys since 2017 (see Yang et al. 2021).

4.2 Regressions

A pre-specified ordinary-least-squares regression equation provide treatment effect estimates:¹⁹

$$Y_{ijd} = \beta_0 + \beta_1 T1_{ijd} + \beta_2 T2_{ijd} + \eta B_{ijd} + \delta_{ijd}^{others} + \delta_{ijd}^{leaders} + \gamma_{jd} + \varepsilon_{ijd} \quad (8)$$

where Y_{ijd} is the social distancing indicator for respondent i in community j and district d ; $T1_{ijd}$ and $T2_{ijd}$ are indicator variables for the misperceptions correction and leader endorsement treatment groups, respectively; B_{ijd} is the baseline value of the dependent variable; δ_{ijd}^{others} is a vector of dummy variables for the number of other respondents who report knowing the respondent's household from 0 to 8; $\delta_{ijd}^{leaders}$ is a vector of indicators for the number of community leaders who report knowing the respondent's household from 0 to 4;²⁰ γ_{jd} are community fixed effects; and ε_{ijd} is a mean-zero error term. We report robust standard errors.²¹

Coefficients β_1 and β_2 represent the intent-to-treat impacts of the misperceptions correction and leader endorsement treatments (respectively) on social distancing.

We modify Equation 8 to estimate heterogeneity in treatment effects with respect to local COVID-19 case loads:

$$Y_{ijd} = \beta_0 + \beta_1 T1_{ijd} + \beta_2 T2_{ijd} + \beta_3 (T1_{ijd} * Covid_d) + \beta_4 (T2_{ijd} * Covid_d) + \eta B_{ijd} + \delta_{ijd}^{others} + \delta_{ijd}^{leaders} + \gamma_{jd} + \varepsilon_{ijd} \quad (9)$$

Equation 9 adds interactions between treatment indicators and the cumulative number of district-level COVID-19 cases per 100,000 population at the start of the endline survey.²² Coefficients β_1 and β_2 in Equation 9 now represent the impacts of the treatments in districts where COVID-19 cases are zero (slightly out of sample); β_3 and β_4 represent the change in the respective treatment effect for a one-unit increase in district-level COVID-19 cases per 100,000 population.

4.3 Hypotheses

We pre-specified the hypothesis that each treatment (β_1 and β_2 in Equation 8) would have positive effects. Subject-matter experts (surveyed without knowing results) concurred with this expectation.²³ The mean expert predictions were that the misperceptions correction and leader endorsement treatments would increase

19. Appendix G.1 shows that all conclusions are robust to logit and probit specifications.

20. As pre-specified, we cap δ_{ijd}^{others} at the first integer that covers over 90% of the sample, and $\delta_{ijd}^{leaders}$ at the maximum number of leaders found in any community.

21. Appendix G.2 shows that clustering standard errors by the 76 communities or 7 districts has minimal impact on standard errors and does not affect whether any coefficients are statistically significant at conventional levels.

22. The main effect of $Covid_d$ is absorbed by γ_{jd} .

23. Predictions by 71 individuals provided at <https://socialscienceprediction.org/> (survey closing date January 2, 2021).

social distancing by 5.23 and 5.56 percentage points, respectively.

We also test the hypotheses that the impact of the misperceptions correction treatment on social distancing and on the expected future infection rate will be greater in areas with a higher current COVID-19 infection rate (β_3 in Equation 9 will be positive). We did not pre-specify these hypotheses, but advance them on the basis of our theoretical model.

5 Results

5.1 Pre-Treatment Descriptives

Table 1 presents pre-treatment summary statistics of social distancing support, perceptions and behavior in the first six months of the COVID-19 pandemic. First, we document a large and statistically significant gap between actual and perceived support for social distancing: at both pre-baseline and baseline, over 97% of respondents support social distancing; however, respondents underestimate the community share expressing such support, on average estimating 69% in a pre-baseline survey and 80% at baseline. Second, we observe a large and statistically significant 11 percentage point increase in the perceived share of community support between pre-treatment survey rounds, consistent with the idea that misperceptions for new public health behaviors are most prevalent at the start of the public health crisis and then diminish over time as social networks share information. Third, despite increases in reported and perceived support for social distancing, we see small decreases in self-reported social distancing behavior; in the theoretical model, this behavior is predicted where the current local infection rate is low, as was indeed the case for all study communities prior to the endline survey.²⁴

5.2 Average Treatment Effects

In Table 2 Column (1), we present regression estimates for our primary outcome.²⁵ Both treatment coefficients are small in magnitude and neither is statistically significantly different from zero. These findings diverge from expert predictions of treatment effects. We strongly reject the null that our T1 and T2 treatment effect estimates are equal to the positive mean expert predictions (p -value <0.001 in each case).

However, we find the misperceptions correction has a positive effect on measures of perceived community support for social distancing. Analyses presented in Appendix C (not pre-specified) shows that the treatment effect is concentrated on the lower end of the distribution, having a significant positive effect on a respondent

24. See Figure A.2 to see relatively low levels of new COVID-19 cases in Mozambique during the pre-baseline and baseline relative to the endline survey.

25. The complete set pre-specified analyses are presented in Appendix H.

perceiving that at least 50% of households in their community support social distancing.

5.3 Treatment Effect Heterogeneity

In Table 2 Column (2), we present regression estimates of treatment effect heterogeneity (Equation 9) with respect to the local infection rate, measured as COVID-19 cases per 100,000 population in the respondent’s district.

The misperceptions correction treatment effect is heterogeneous with respect to local COVID-19 cases. The coefficient on the interaction term with $T1_{ijd}$ is positive and statistically significant at the 1% level. The coefficient on the $T1_{ijd}$ main effect is the predicted effect of misperceptions correction in a district with zero cases (slightly out of sample), and suggests that the misperceptions correction would reduce social distancing by 3.4 percentage points in such a location (statistically significant at the 5% level).

Figure 2 displays this treatment effect heterogeneity. We plot district-specific treatment effects (estimating Equation 8 separately in each of seven districts) on the y-axis (with 95% confidence intervals) against district case counts on the x-axis. In the six districts with the lowest case counts, coefficients are negative. By contrast, in Chimoio, the district with the most cases (39.08/100,000) that also accounts for one-quarter of the sample, we estimate a large positive effect: 9.2 percentage points—a 70% increase over that district’s control group (statistically significant at the 5% level).

This heterogeneous treatment effect holds up to various robustness checks (presented in Appendix G). First, we run logit and probit specifications of the primary results. Second, we cluster standard errors by community and district. Third, we vary the threshold by which self-reported "social distancing actions" were incorporated in the social distancing indicator. Fourth, we test four alternative measures of the local COVID-19 infection rate, including the simple case count (not per capita) and high-case-count indicators, to show that the treatment effect heterogeneity is not unique to our preferred measure. Fifth, we exclude the top-COVID-19 and largest-sampled district of Chimoio to verify that it alone is not driving our results. In all cases, we find that our primary results are very similar.

By contrast, the leader endorsement treatment effect is not heterogeneous with respect to local case loads. The coefficient on the corresponding interaction term in Column (2) is small in magnitude and not statistically significantly different from zero. Some reasons why this treatment may not be effective, even when COVID-19 cases are high, include limited familiarity of leaders among all community members or limited confidence that the leader’s endorsement reflected true beliefs rather than political “lip service”. Coupled with findings from Banerjee et al. (2019) on gossips spreading information, the result suggests that network-central individuals may be effective at transmitting information but not necessarily because their

opinions have a dominating influence on community members' beliefs.

The interplay between the free-riding and perceived-infectiousness effects is the distinctive feature of our theoretical model. When the perceived-infectiousness effect is large enough, it overcomes the countervailing free-riding effect, and the misperceptions correction treatment leads to more social distancing. An additional implication of the theory is that the treatment should have similar heterogeneous effects on the expected future infection rate.

We conduct this additional test of the theory, examining treatment effects on the expected future infection rate.²⁶ In Columns (3) and (4) of Table 2, the outcome is the share of the community the respondent thinks will get sick from COVID-19 (responses were integers out of 10; we divide by 10 to yield a 0-1 scale). In Column (3), we estimate average treatment effects. Each coefficient is small in magnitude and not statistically significantly different from zero.

In Column (4), we estimate heterogeneity in treatment effects with respect to local cases, and find the same pattern as in Column (2). The misperceptions correction decreases the expected future infection rate in districts with no cases, and this impact becomes more positive as current cases rise (the $T1_{ijd}$ main effect and interaction term coefficients are both statistically significant at the 5% level).

These treatment effect heterogeneity findings in Table 2 Columns (2) and (4) jointly support the theoretical model. When current infection rates are low, the misperceptions correction treatment does not change perceived infectiousness much, but leads to realizations that social distancing support is higher than previously thought. People therefore reduce estimates of the future infection rate, and also reduce their own social distancing (choosing to free-ride). By contrast, when current infection rates are high, the treatment causes larger increases in perceived infectiousness. Notwithstanding an increase in the share of social distancing supporters, people increase their estimate of the future infection rate and increase their social distancing.

6 Conclusion

Support for social distancing increased rapidly during the COVID-19 pandemic. If people are unaware of the extent to which others' beliefs on social distancing have changed, would revealing true high rates of such support lead to more social distancing? In theory, the impact of providing such information is ambiguous: it could reduce social distancing if free-riding effects dominate, but could have a positive effect on social distancing if perceived-infectiousness effects dominate. Perceived-infectiousness effects are more likely to dominate when the current local infection rate is higher.

26. The question is "For every 10 people in your community, how many do you think would get sick from coronavirus?" Sample sizes in these regressions are smaller. We implemented this question midway through the endline survey, after finding preliminary evidence suggesting the need to explore mechanisms behind treatment effect heterogeneity.

We implemented a randomized controlled trial testing the impact of a “misperceptions correction” treatment revealing high community support for social distancing. The treatment effect on social distancing exhibits the spatial heterogeneity predicted by theory: negative in areas with low infection rates (reflecting the dominance of free-riding effects), and more positive in areas with higher rates (as perceived-infectiousness effects become increasingly prominent). In the area with the most cases, amounting to one-quarter of our sample, the treatment effect is positive and large in magnitude. The treatment effect on the expected future infection rate shows similar heterogeneity, confirming an additional theoretical prediction.

Our results suggest that when local infection rates are high, health policies shifting perceptions of community support for social distancing upwards could help promote social distancing behavior. Future research is needed to confirm the external validity of these findings and determine how the results translate to other contexts. For example, in cities, looser social networks among neighbors might lead to larger misperceptions of community support while population-dense housing might further activate the perceived-infectiousness effect; alternatively, in communities with lower baseline support for social distancing, a misperceptions correction treatment may be less motivating but may also potentially "gain more ground" among those with the lowest support who also underestimate the social norm. These findings may also help predict the impacts of analogous public health messaging revealing community support for preventive measures against other infectious diseases.

References

- Andersson, Ola, Pol Campos-Mercade, Armando Meier, and Erik Wengstrom. 2021. "Anticipation of COVID-19 Vaccines Reduces Social Distancing." *SSRN Working Paper*, no. 3765329, <https://doi.org/10.2139/ssrn.3765329>.
- Andre, Peter, Teodora Boneva, Felix Chopra, and Armin Falk. 2021. "Fighting climate change: The role of norms, preferences, and moral values," <https://doi.org/10.2139/ssrn.3885418>. <https://ssrn.com/abstract=3885418>.
- Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson. 2019. "Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials." *The Review of Economic Studies* 86 (6): 2453–2490. <https://doi.org/10.1093/restud/rdz008>.
- Banerjee, Abhijit, Eliana La Ferrara, and Victor H Orozco-Olvera. 2019. "The Entertaining Way to Behavioral Change: Fighting HIV with MTV." *NBER Working Paper*, no. 26096, <https://doi.org/10.3386/w26096>. <https://www.nber.org/papers/w26096>.
- Benabou, Roland, and Jean Tirole. 2011. "Identity, Morals, and Taboos: Beliefs as Assets." *Quarterly Journal of Economics* 126 (2): 805–855. <https://doi.org/10.1093/qje/qjr002>. <https://academic.oup.com/qje/article/126/2/805/1867845>.
- Bicchieri, Cristina, and Eugen Dimant. 2019. "Nudging with Care: The Risks and Benefits of Social Information." *Public choice*, 1–22.
- Bursztyn, Leonardo, Alessandra L González, and David Yanagizawa-Drott. 2020. "Misperceived social norms: Women working outside the home in Saudi Arabia." *American economic review* 110 (10): 2997–3029.
- Cato, Susumu, Takashi Iida, Kenji Ishida, Asei Ito, Kenneth Mori McElwain, and Masahiro Shoji. 2020. "Social Distancing as a Public Good Under the COVID-19 Pandemic." *Public Health* 188:51–53. <https://doi.org/10.1016/j.puhe.2020.08.005>.
- Fafchamps, Marcel, Ana Vaz, and Pedro C Vicente. 2020. "Voting and peer effects: Experimental evidence from Mozambique." *Economic Development and Cultural Change* 68 (2): 567–605.
- Fuhrmann-Riebel, Hanna, Ben D'Exelle, Kristian López Vargas, Sebastian Tonke, and Arjan Verschoor. 2023. "Correcting Misperceptions About Trends and Norms to Address Weak Collective Action." *Available at SSRN 4311583*, <https://doi.org/10.2139/ssrn.4311583>. <https://ssrn.com/abstract=4311583>.

- Hershey, John C, David A Asch, Thi Thumasathit, Jascqueline Meszaros, and Victor W Waters. 1994. “The Roles of Altruism, Free Riding, and Bandwagoning in Vaccination Decisions.” *Organizational Behavior and Human Decision Processes* 59 (2): 177–187. <https://doi.org/10.1006/obhd.1994.1055>.
- Jakubowski, Aleksandra, Dennis Egger, Carolyne Nekesa, Layna Lowe, Michael Walker, and Edward Miguel. 2021. “Self-Reported Mask Wearing Greatly Exceeds Directly Observed Use: Urgent Need for Policy Intervention in Kenya.” *medRxiv Preprint*, <https://doi.org/10.1101/2021.01.27.21250487>. <https://www.medrxiv.org/content/10.1101/2021.01.27.21250487v1>.
- Janzwood, Scott. April 27, 2020. *The Social Distancing Norm Cascade: The Role of Belief Systems in Accelerating Normative Change during the COVID-19 Pandemic*. Canada: Cascade Institute. <https://cascadeinstitute.org/wp-content/uploads/2020/05/Janzwood-Social-Distancing-Norm-Cascade-0514.pdf>.
- Lau, Krystal, Marisa Miraldo, Matteo M Galizzi, and Katharina Hauck. 2019. “Social Norms and Free-riding in Influenza Vaccine Decisions in the UK: An Online Experiment.” *The Lancet* 394 (2): 65. [https://doi.org/10.1016/S0140-6736\(19\)32862-4](https://doi.org/10.1016/S0140-6736(19)32862-4).
- Martinez, Deborah, Cristina Parilli, Carlos Scartascini, and Alberto Simpser. 2021. “Let’s (not) get together! The role of social norms on social distancing during COVID-19.” *PLoS ONE* 16 (3): 1–14. <https://doi.org/10.1371/journal.pone.0247454>.
- Schultz, Wesley P, Jessica M Nolan, Robert B Cialdini, Noah J Goldstein, and Vladas Griskevicius. 2007. “The Constructive, Destructive, and Reconstructive Power of Social Norms.” *Psychological Science* 18 (5): 429–434. <https://doi.org/10.1111/j.1467-9280.2007.01917.x>. <https://pubmed.ncbi.nlm.nih.gov/17576283/>.
- Spiegler, Ran. 2020. “Behavioral Implications of Causal Misperceptions.” *Annual Review of Economics* 12 (1): 81–106. <https://doi.org/10.1146/annurev-economics-072219-111921>. eprint: <https://doi.org/10.1146/annurev-economics-072219-111921>. <https://doi.org/10.1146/annurev-economics-072219-111921>.
- Yang, Dean, James Allen IV, Arlete Mahumane, James Riddell IV, and Hang Yu. 2021. “Knowledge, Stigma and HIV Testing: An Analysis of a Widespread HIV/AIDS Program.” *NBER Working Paper*, no. 28716.
- Yu, Hang. 2020. *Three Essays in Development Economics. Dissertation*. Ann Arbor, MI: University of Michigan.

Tables and Figures

Table 1: **Summary Statistics of Pre-Treatment Social Distancing Measures**

VARIABLES	Pre-Baseline			Baseline			T-test
	N	Mean	SD	N	Mean	SD	p-value
(1) Respondent supports social distancing (SD)	2,117	0.976 ^a	0.153	2,117	0.989 ^b	0.104	0.001
(2) Perceived share of community supporting SD	2,109	0.689 ^a	0.313	2,114	0.800 ^b	0.262	0.000
(3) Primary SD indicator: if (4) & (7)				2,117	0.078	0.269	
(4) → Self-report SD indicator: if (5) & (6)	2,117	0.383	0.486	2,117	0.355	0.479	0.045
(5) → Self-report: Followed govt rules last 14 days	2,117	0.952	0.214	2,117	0.949	0.219	0.692
(6) → Self-report: SD behaviors above median	2,117	0.396	0.489	2,117	0.361	0.481	0.012
(7) → Others SD indicator: if (8) & (9)				2,117	0.232	0.422	
(8) → Other households' report of SD				2,117	0.378	0.485	
(9) → Leaders' report of SD				2,117	0.519	0.500	

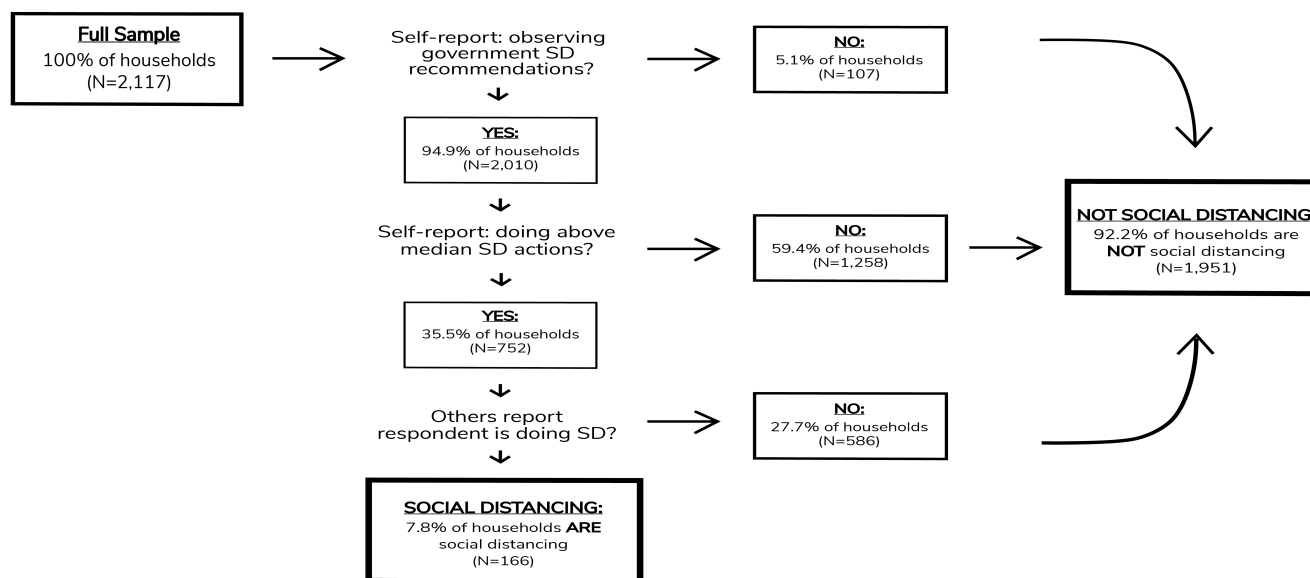
Notes: Pre-baseline data collected from July 10 to August 16, 2020, and baseline data collected from August 26 to October 4, 2020. Variables are as follows. Variables are as follows. Row 1: indicator equal to one if respondent answers “yes” to supporting “the practice of social distancing to prevent the spread of coronavirus” and zero otherwise. Row 2: perceived share of households (asked as “for every 10 households”) in community that support social distancing (SD). Row 3: indicator for SD equal to one if respondent is SD according to self (Row 4) and others’ reports (Row 7), and zero otherwise. Row 4: indicator for SD according to self if respondent answered “yes” to observing the government’s recommendations on SD in the last 14 days (Row 5) and report doing more than the sample median number of SD behaviors (Row 6), and zero otherwise. Row 7: indicator for SD according to others if all other respondents (Row 8) and community leaders (Row 9) reported not knowing the respondent household, not seeing the respondent household in the past 14 days, or—if seen—that the respondent household 1) did NOT come closer than 1.5 meters to others outside their household; 2) did NOT shake hands, try to shake hands, or touch others outside their household; and 3) appeared to be observing the government’s recommendations on SD, and zero otherwise. All variables have a minimum of 0 and a maximum of 1. Last column displays the p-value of a paired t-test on the difference between pre-baseline and baseline measure (where pre-baseline data are available). Superscripts ^a and ^b indicate paired t-tests comparing reported and perceived support for social distancing at pre-baseline and baseline, respectively, which are significantly different (p-value=0.000).

Table 2: **Treatment Effects on Social Distancing and Expected COVID-19 Illnesses**

VARIABLES	(1) Primary SD Indicator	(2) Primary SD Indicator	(3) Perceived share of people in community that will get sick from COVID-19	(4) Perceived share of people in community that will get sick from COVID-19
T1: Misperceptions Correction	0.0042 (0.0140)	-0.0466** (0.0191)	0.0418 (0.0322)	-0.1936** (0.0944)
T2: Leader Endorsement	-0.0054 (0.0137)	-0.0258 (0.0198)	-0.0209 (0.0308)	-0.0598 (0.0944)
T1 × District COVID-19 Cases		0.0030*** (0.0011)		0.0073** (0.0029)
T2 × District COVID-19 Cases		0.0012 (0.0010)		0.0013 (0.0029)
Observations	2,117	2,117	812	812
R-squared	0.158	0.163	0.146	0.152
Control Mean DV	0.0857	0.0857	0.3590	0.3590
Control SD DV	0.2801	0.2801	0.3685	0.3685


Notes: Dependent variable in Columns 1-2 defined in Table 1. Dependent variable in Columns 3-4 is the expected future infection rate: “For every 10 people in your community, how many do you think would get sick from coronavirus?” (converted to share from 0 to 1). “T1: Misperceptions Correction” is equal to one if respondent was randomly assigned to the misperceptions correction treatment, and zero otherwise. “T2: Leader Endorsement” is equal to one if respondent was randomly assigned to the leader endorsement treatment, and zero otherwise. “T1 x District COVID-19 Cases” and “T2 x District COVID-19 Cases” are the respective treatment indicators interacted with district-level cumulative COVID-19 cases per 100,000 population at the start of the endline survey (see Appendix B.3, Table A.1, Column 2). All regressions control for a baseline measure of the dependent variable, a vector of indicators for number of community leaders knowing the respondent at baseline (0 through 4), and a vector of indicators for number of other respondents knowing the respondent at baseline (0 through 8). All regressions also include community fixed effects. Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 1: The Social Distancing Measure



Notes: Figure shows the breakdown of the social distancing measure at baseline. As pre-specified, respondents considered social distancing (SD) if: 1) self-report “yes” to “In the past 14 days, have you observed the government’s recommendations on social distancing?”, 2) self-report doing more than the sample median number of “social distancing actions” in the past seven days, and 3) are considered to be SD according to leaders and other respondents in the community. Percentages reported are all shares of full sample (N=2,117). See Table 1 and Section 3.3 of the main text for social distancing question definitions.

Figure 2: District-Level Misperceptions Correction Treatment Effects by COVID-19 Cases



Notes: Misperceptions correction treatment effects (triangles) estimated separately for each of seven districts (with 95% confidence intervals). District-level treatment effects plotted on vertical axis against district-level cumulative COVID-19 case loads at start of endline survey (per 100,000 population) on horizontal axis.

Online Appendix

This is the online appendix to: Allen IV, James, Arlete Mahumane, James Riddell IV, Tanya Rosenblat, Dean Yang, and Hang Yu. "Correcting Perceived Social Distancing Norms to Combat COVID-19." *Economic Development and Cultural Change* (2023).

A Proofs

A.1 Proof of Theorem 1

The agent will adjust her effort level in response to the treatment to $\sqrt{\hat{e}} = \frac{\hat{\alpha}C}{2}H(\hat{x})$ where $H(x) = 1 - x(1 - \frac{1}{\sqrt{2}})$. Hence, the prior and posterior effort levels satisfy:

$$\frac{\sqrt{\hat{e}}}{\sqrt{e}} = \frac{\hat{\alpha} H(\hat{x})}{\alpha H(x)} \quad (\text{A.1})$$

We take the ratios of Equations 6 and 7:

$$\frac{\hat{\alpha}}{\alpha} = \frac{1 - \sqrt{e}G(x)}{1 - \sqrt{\hat{e}}G(\hat{x})} \quad (\text{A.2})$$

We therefore obtain:

$$\frac{\sqrt{\hat{e}}}{\sqrt{e}} = \frac{H(\hat{x})(1 - \sqrt{e}G(x))}{H(x)(1 - \sqrt{\hat{e}}G(\hat{x}))} \quad (\text{A.3})$$

Effort increases iff $\frac{\sqrt{\hat{e}}}{\sqrt{e}} > 1$:

$$\begin{aligned} \frac{H(\hat{x})(1 - \sqrt{e}G(x))}{H(x)(1 - \sqrt{\hat{e}}G(\hat{x}))} &> 1 \\ \sqrt{e}[H(x)G(\hat{x}) - H(\hat{x})G(x)] &> H(x) - H(\hat{x}) \end{aligned}$$

Now note that $G(x) = 2xH(x)$ such that:

$$\begin{aligned} \sqrt{e}2H(x)H(\hat{x})(\hat{x} - x) &> \left(1 - \frac{1}{\sqrt{2}}\right)(\hat{x} - x) \\ \sqrt{e} &> \frac{1 - \frac{1}{\sqrt{2}}}{H(x)H(\hat{x})} \end{aligned} \quad (\text{A.4})$$


This shows that the perceived-infectiousness effect dominates if the initial effort level e is high enough. Effort is determined by Equation 5 and increases with α (which increases with $\hat{\alpha}$). Therefore, for sufficiently large $\hat{\alpha}$ the perceived-infectiousness effect dominates.

B Study Context

B.1 Study Area

Study participants come from 76 communities in central Mozambique. The study communities are in seven districts of three provinces: Dondo and Nhamatanda in Sofala province; Gondola, Chimoio and Manica in Manica province; and Namacurra and Nicoadala in Zambezia province. These 76 communities are mapped in Figure A.1. Compared to other communities in Mozambique, the study areas are relatively accessible to main transport corridors (highways and ports), and are thus important geographic conduits for infectious disease.

Figure A.1: Study Area




B.2 Study Timeline

We collected survey data in three rounds between July 10 and November 18, 2020. Figure A.2 depicts the study timeline below a rolling average of new Mozambican COVID-19 cases. We piloted surveys in Round 1 (pre-baseline). Immediately before the Round 2 survey, we randomly assigned households to treatments and submitted our pre-analysis plan to the AEA RCT Registry. The Round 2 survey served as a baseline and was immediately followed (on the same phone call) by our treatment interventions. Round 3 was our endline

survey. Surveys collected data on COVID-19 knowledge, beliefs, and behaviors. While data collection for Round 3 began only one day after completion of Round 2, there was a minimum of 3.0 weeks and average of 6.3 weeks between Rounds 2 and 3 surveys for any given respondent. While the Round 1 survey occurred when new COVID-19 cases remained relatively steady, both the Round 2 and Round 3 surveys occurred during a period of substantial growth in new COVID-19 cases.

Figure A.2: Study Timeline



Notes: Round 1 is pre-baseline survey to collect social distancing support data, Round 2 is baseline survey, and Round 3 is endline survey. There is at least a three week gap between baseline and endline survey for any given study participant. Pre-analysis plan uploaded and treatments randomly assigned immediately prior to start of Round 2 baseline survey, on Aug. 25, 2021. Treatments implemented immediately following baseline survey on same phone call. Baseline measures reported in Table 1 come from Round 2 surveys and endline measures come from Round 3 surveys.

B.3 COVID-19 Context

The Mozambican government declared a State of Emergency due to the COVID-19 pandemic on March 31, 2020, recommending social distancing (at least 1.5 meters) and requiring it at public and private institutions and gatherings. The government also suspended schools, required masks at funerals and markets, banned gatherings of 20 or more, and closed bars, cinemas and gymnasiums (Republic of Mozambique, 4/1/2020). The government stopped short of implementing a full economic “lockdown” due to its economic costs (Jones, Egger, and Santos 2020). On August 5, 2020, the government renewed the State of Emergency, called for improved mask-wearing, and announced a schedule for loosening restrictions (Nyusi, 8/5/2020). In September, the government loosened some restrictions including resuming religious gatherings at 50% capacity (U.S Embassy in Mozambique). Throughout this period, the government’s social distancing recommendation remained constant.

COVID-19 cases by district at the start of the Round 3 (endline) survey are estimated as follows. Data on

district-level population come from Mozambique’s 2017 Census (National Institute of Statistics (INE) 2017). District COVID-19 case counts come from the government’s COVID-19 Mozambique dashboard (Ministry of Health 2020) and correspondence with provincial health offices. Each district’s case count is from the start date of the endline survey in the district (ranging from October 5 to November 1, 2020). We also show the number of respondents in our study sample in each district.

Table A.1: **COVID-19 Cases by District**

	(1)	(2)	(3)	(4)
DISTRICT	Cumulative COVID-19 Cases	Cases per 100,000 people	Population	Number of Study Respondents
Sofala Province				
Dondo	8	4.137	193,382	323
Nhamatanda	12	4.300	279,081	214
Manica Province				
Gondola	3	3.553	84,429	224
Chimoio	142	39.082	363,336	524
Manica	20	9.292	215,239	290
Zambezia Province				
Namacurra	4	1.652	242,126	244
Nicoadala	52	28.779	180,686	298

Notes: COVID-19 cases by district at the start of the Round 3 (endline) survey. Column 1: District COVID-19 case counts come from the government’s COVID-19 Mozambique dashboard (Ministry of Health 2020) and correspondence with provincial health offices, measured at the start date of the endline survey in the district (ranging from October 5 to November 1, 2020). Column 2: Calculated from Columns 1 and 3. Column 3: District-level population come from Mozambique’s 2017 Census (National Institute of Statistics (INE) 2017). Column 4: Number of respondents in our study sample in each district.

C Effect on Perceived Community Support

Table A.2 presents the cumulative distribution of this perceived community support measure in the final samples at pre-baseline, baseline and endline, and subdivided by treatment arm at endline. Even at pre-baseline, the distribution is skewed upwards with over 80% of the sample reporting that the majority (50% or greater) of households in their community support social distancing, though a sizable 8% use the extreme lower end of the scale to report that none (0%) of the households in their community support social distancing. At baseline, the distribution is further skewed upwards with over 90% of the sample reporting that the majority (50% or greater) of households in their community support social distancing and over half of the sample reporting that 100% of households do the same.

Table A.2: **Sample Distribution (Cumulative %) by Perceived Community Support**


Perceived Share	Pre-Baseline	Baseline	Endline			
	Total	Total	Total	Control	T1	T2
0%	8.0	2.7	2.8	3.5	2.2	2.5
10%	9.0	3.1	3.1	3.6	2.4	3.0
20%	10.7	4.4	4.4	5.9	2.9	3.9
30%	14.0	6.5	6.5	7.4	5.1	6.6
40%	19.1	9.6	8.8	9.5	6.9	9.9
50%	34.3	21.1	19.0	19.3	18.3	19.2
60%	41.8	27.1	23.9	24.2	23.0	24.5
70%	49.7	33.4	30.3	29.6	30.5	31.1
80%	59.7	43.4	40.8	39.8	39.1	44.0
90%	65.2	48.9	46.8	45.8	45.3	49.6
100%	100.0	100.0	100.0	100.0	100.0	100.0

Notes: Perceived share of households in the community who support social distancing is estimated by dividing responses to the question “For every 10 households in your community, how many support social distancing?” by 10, and hence has 11 categories from 0%, 10%... 90%, 100%. Cells report cumulative percentages from 0% up to the row in question. Pre-baseline “Total” refers to all responses from the final sample in Round 1 (N=2,109), and Baseline “Total” refers to all responses from the final sample in Round 2 (N=2,114). At endline, “Total” refers to Round 3 responses from the final sample (N=2,116), “Control” from the control group, “T1” from the misperceptions correction treatment group, and “T2” from the leader endorsement treatment group.

We find that the misperceptions correction treatment did increase respondents’ perceived community support, particularly for those at the lower end of the distribution. Figure A.3 shows the cumulative distribution function for the perceived community support measure at endline. Relative to the control group, those receiving the misperceptions correction treatment were less likely to report that fewer than 50% of households in their community supported social distancing, instead reporting higher perceptions of commu-

nity support. Further, Table A.3 presents three regressions estimating the treatment effects on the perceived community support. In Column (1), the dependent variable is the perceived share of households in the community who support social distancing. The coefficient is positive and marginally statistically significant (p -value=0.12). Regressions in Columns (2) and (3) find that the misperceptions correction treatment has a positive effect on an indicator for the respondent believing the majority (50% or more) of households in their community support social distancing, and an indicator that the respondent’s perceived community support increased between baseline and endline (both coefficients are statistically significantly different from zero at the 5% level).

Figure A.3: **Cumulative Distribution of Perceived Community Support by Treatment**



Notes: Perceived share of households in the community who support social distancing is estimated by dividing responses to the question “For every 10 households in your community, how many support social distancing?” by 10, and hence has 11 categories from 0%, 10%... 90%, 100%. Figure depicts the cumulative distribution function of this variable for the “Control” group and “Misperceptions Correction” treatment arm. The leader endorsement treatment is excluded for clarity.

Table A.3: **Treatment Effects on Perceived Community Support (PCS)**

VARIABLES	(1)	(2)	(3)
	Continuous	Indicator if	Indicator if
	PCS	PCS \geq 50%	PCS increased
T1: Misperceptions Correction	0.0196 (0.0128)	0.0291** (0.0138)	0.0507** (0.0241)
T2: Leader Endorsement	0.0041 (0.0128)	0.0036 (0.0149)	0.0358 (0.0236)
Observations	2,116	2,116	2,113
R-squared	0.164	0.118	0.043
Control Mean DV	0.8115	0.9049	0.2550
Control SD DV	0.2681	0.2935	0.4361

Notes: Dependent variables are defined as follows. Column 1 is the perceived share of households in community that support social distancing, which takes on the values shown in Table A.2. Column 2 is an indicator equal to one if respondent reports that majority (50% or more) of households in community support social distancing, and zero otherwise. Column 3 is an indicator equal to one if the respondent's perceived community support increased between the baseline (pre-treatment) and endline (post-treatment) surveys. "T1: Misperceptions Correction" and "T2: Leader Endorsement" and controls are as defined in Table 2, except column 3 does not include a baseline value of the outcome as a control as it was used to calculate the outcome. All regressions also include community fixed effects. Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D Social Distancing Index

The list of actions included in the Social Distancing Index and their corresponding summary statistics are presented below.

Social Distancing Actions: Is this something your household has been doing for the last seven days? (Answers indicating social distancing in parentheses.¹)

1. Shop in crowded areas like informal markets (No)
2. Gather with several friends (No)
3. Help the elderly avoid close contact with other people, including children (Yes)
4. If show symptoms of coronavirus, immediately inform my household and avoid people (Yes)
5. Drink alcohol in bars (No)
6. Wear a face mask if showing symptoms of coronavirus (Yes)
7. Instead of meeting in person, call on the phone or send text message (Yes)
8. Allow children to build immunity by playing with children from other households (No)

Below are the summary statistics for the questions that comprise the self-reported social distancing index at baseline and endline. Respondents were asked “Is this something your household has been doing for the last seven days?” about a randomly determined four social distancing actions at baseline and all eight social distancing actions at endline. Responses were coded as indicators equal to one if indicative of social distancing (answers that indicate social distancing shown in parentheses), and zero otherwise.

1. For items 4 and 6 that are conditional on showing symptoms, survey staff instructed respondents to answer “Yes” (doing social distancing) if not showing symptoms.

Table A.4: **Summary Statistics for Components of Social Distancing Index**

VARIABLES	Baseline			Endline			T-test
	N	Mean	SD	N	Mean	SD	p-value
Shop in crowded areas like informal markets (No)	1,032	0.642	0.480	2,115	0.678	0.467	0.5011
Gather with several friends (No)	1,047	0.349	0.477	2,113	0.414	0.493	0.0357
Help the elderly avoid close contact with other people, including children (Yes)	1,094	0.877	0.329	2,114	0.923	0.266	0.0000
If show symptoms of coronavirus, immediately inform my household and avoid people (Yes)	1,050	0.836	0.370	2,113	0.860	0.347	0.0314
Drink alcohol in bars (No)	1,082	0.226	0.419	2,113	0.272	0.445	0.0152
Wear a face mask if showing symptoms of coronavirus (Yes)	1,034	0.902	0.297	2,114	0.885	0.319	0.3993
Instead of meeting in person, call on the phone or send text message (Yes)	1,039	0.935	0.247	2,112	0.930	0.255	0.5922
Allow children to build immunity by playing with children from other households (No)	1,070	0.439	0.497	2,113	0.456	0.498	0.0814

Notes: Variables are coded as indicators equal to one if indicative of social distancing (answers that indicate social distancing shown in parentheses), and zero otherwise. Respondents were asked “Is this something your household has been doing for the last seven days?” about a randomly determined four social distancing actions at baseline and all eight social distancing actions at endline. The baseline sample was asked a subset of these questions which explains the smaller number of observations at baseline. Last column displays the p-value of a paired t-test on the difference between baseline and endline measure (where baseline data are available).

E Treatment Details and Scripts

Both the misperceptions correction and leader endorsement treatments were implemented directly following the baseline survey, on the same phone call. If a respondent was randomly assigned to a treatment, the corresponding intervention text would appear on the enumerator’s tablet. Enumerators read a script aloud exactly as shown below. Following the treatment, respondents were asked if they would like the information repeated. Of the N=628 receiving the misperceptions correction and N=637 receiving the leader endorsement, only 8.6% and 9.4% asked for the script to be repeated, respectively.

Script for T1: Misperceptions Correction – “Now I want to give you some information about social distancing. In this survey, you indicated that you think *<insert respondent’s answer here>* of every 10 households in your community support the practice of social distancing.”

- *If response UNDERESTIMATES community support for social distancing:* “However, more households support social distancing than you think! Based on the results of our first COVID-19 survey, approximately *<insert actual community support for social distancing here>* of every 10 households in your community support social distancing to prevent the spread of the coronavirus.”
- *If response CORRECTLY ESTIMATES community support for social distancing:* “You are correct! Based on the results of our first COVID-19 survey, approximately *<insert actual community support for social distancing here>* of every 10 household in your community support social distancing to prevent the spread of the coronavirus.”
- *If response OVERESTIMATES community support for social distancing: (no information given)*

Script for T2: Leader Endorsement – “Our research team recently called and talked to your *<list leaders’ titles and names here>*. They said that they support social distancing, are practicing social distancing themselves, and encourage others to do the same.”

F Attrition and Balance

Appendix Table A.5 presents regressions examining whether attrition and baseline variables are balanced with respect to treatment assignment.² Attrition between Round 2 (baseline) and Round 3 (endline) is only 4.9% and is less than 5.6% in each of the seven districts surveyed. Balance in attrition is confirmed in Column (1), which starts with the Round 2 (baseline) sample and regresses treatments on an indicator equal to one if the respondent was not reached for the Round 3 (endline) survey. Balance in baseline social distancing outcomes is confirmed in Columns (2)-(4), which examines the Round 2 social distancing outcomes. Balance in baseline household characteristics is confirmed in Columns (6)-(8), which examines the final Round 3 sample and regresses treatments on Round 1 measures of household income, an index of food insecurity, and an indicator for presence of an older adult over 60 years. In not a single regression in the table is a coefficient on a treatment indicator statistically significant at conventional levels.

2. Figure A.2 shows the study timeline for the three survey rounds collected. Round 1 is a pre-baseline measure, Round 2 measures baseline values and Round 3 measures endline outcomes.

Table A.5: **Treatment Effect on Attrition and Balance**

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Attrition	Primary SD Indicator	Others' Report of SD	Self-Report of SD	Perceived Social Norm	Hh Income	Food Insecurity	Older Adult in Hh
T1: Misperceptions Correction	-0.0127 (0.0111)	-0.0176 (0.0134)	-0.0005 (0.0203)	-0.0096 (0.0247)	-0.0101 (0.0138)	-159.46 (181.66)	0.0011 (0.0191)	-0.0029 (0.0250)
T2: Leader Endorsement	-0.0015 (0.0113)	-0.0032 (0.0143)	0.0090 (0.0206)	0.0042 (0.0249)	-0.0201 (0.0137)	-39.95 (181.76)	-0.0240 (0.0193)	0.0240 (0.0252)
Observations	2,226	2,117	2,117	2,117	2,114	1,873	2,117	2,096
R-squared	0.030	0.096	0.199	0.076	0.047	0.043	0.090	0.058
Control Mean DV	0.0533	0.0833	0.2289	0.3556	0.8095	1176	0.8415	0.3424
Control SD DV	0.2248	0.2765	0.4204	0.4790	0.2618	4029	0.3654	0.4748

Notes: Dependent variables are as follows. Column 1: indicator if respondent attrited from the sample between baseline and endline. Columns 2-4: baseline SD outcomes defined in Table 1. Column 5: baseline perceived share of community supporting SD, defined further in Table 1. Column 6: at pre-baseline, self-reported total income for the previous week (in Mozambican meticaís). Column 7: at pre-baseline, indicator if, in the last 7 days, household has 1) lacked food; 2) reduced number of meals/portions; or was unable to buy their usual amount of food due to 3) market shortages, 4) high prices, 5) reduced income. Column 8: at pre-baseline, indicator if adult age 60 or older is present in the household. Controls are as defined in Table 2. All regressions also include community fixed effects. Robust standard errors in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

G Robustness of Treatment Effects Estimates

In this appendix, we show that our primary results are robust to 1) using a logit or probit specification, 2) clustering standard errors by community or district, 3) two alternative measures of the social distancing indicator as the primary outcome, 4) four alternative measures of COVID-19 case intensity used to test the interaction, and 5) excluding the district with the highest number of COVID-19 cases.

G.1 Logit and Probit Specifications

The primary social distancing indicator is a binary variable that is analyzed using an ordinary least-squares (OLS) regression, as pre-specified. As a robustness check, we adapt Equation 8 to be run using logit and probit regressions.

Table A.6 presents results from the logistic regression on the primary outcomes, while Table A.7 presents corresponding probit regression results. Regression coefficients are presented as marginal effects. Results in both tables are consistent with the results from OLS linear probability models presented in Table 2.

Table A.6: **Treatment Effects Estimated Using Logistic Regression**

VARIABLES	(1) Primary SD Indicator	(2) Primary SD Indicator	(3) Perceived share of households in community that will get sick from COVID-19	(4) Perceived share of households in community that will get sick from COVID-19
T1: Misperceptions Correction	0.0100 (0.0221)	-0.0756** (0.0376)	0.0270 (0.0395)	-0.4034*** (0.1376)
T2: Leader Endorsement	-0.0069 (0.0222)	-0.0398 (0.0349)	-0.0274 (0.0394)	-0.3005** (0.1354)
T1 × District COVID-19 Cases		0.0038*** (0.0013)		0.0132*** (0.0040)
T2 × District COVID-19 Cases		0.0016 (0.0013)		0.0084** (0.0040)
Observations	1,285	1,285	806	806
Control Mean DV	0.1415	0.1415	0.3563	0.3563
Control SD DV	0.3488	0.3488	0.3680	0.3680

Notes: Dependent variables are defined in Tables 1 and 2. Coefficients presented are marginal effects from logit regression. Social distancing treatments and controls are as defined in Table 2. All regressions also include community fixed effects. Robust standard errors in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table A.7: **Treatment Effects Estimated Using Probit Regression**

VARIABLES	(1) Primary SD Indicator	(2) Primary SD Indicator	(3) Perceived share of households in community that will get sick from COVID-19	(4) Perceived share of households in community that will get sick from COVID-19
T1: Misperceptions Correction	0.0089 (0.0212)	-0.0709** (0.0347)	0.0288 (0.0390)	-0.4015*** (0.1316)
T2: Leader Endorsement	-0.0084 (0.0214)	-0.0356 (0.0330)	-0.0298 (0.0392)	-0.3057** (0.1346)
T1 × District COVID-19 Cases		0.0037*** (0.0013)		0.0132*** (0.0038)
T2 × District COVID-19 Cases		0.0014 (0.0012)		0.0085** (0.0040)
Observations	1,285	1,285	806	806
Control Mean DV	0.1415	0.1415	0.3563	0.3563
Control SD DV	0.3488	0.3488	0.3680	0.3680

Notes: Dependent variables are defined in Tables 1 and 2. Coefficients presented are marginal effects from probit regression. Social distancing treatments and controls are as defined in Table 2. All regressions also include community fixed effects. Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

G.2 Clustering Standard Errors

In our primary analysis, we report robust standard errors, as pre-specified. As a robustness check, Table A.8 shows our main regressions clustering standard errors by the study’s 76 communities or 7 districts. The results show that clustering has minimal impact on standard errors and does not affect whether any coefficients are statistically significant at conventional levels.

Table A.8: **Treatment Effects Estimated with Clustered Standard Errors**

VARIABLES	Primary SD Indicator			
	(1) Clustered at Community	(2) Clustered at District	(3) Clustered at Community	(4) Clustered at District
T1: Misperceptions Correction	0.0042 (0.0133)	0.0042 (0.0269)	-0.0466*** (0.0150)	-0.0466*** (0.0124)
T2: Leader Endorsement	-0.0054 (0.0132)	-0.0054 (0.0133)	-0.0258 (0.0169)	-0.0258 (0.0163)
T1 × District COVID-19 Cases			0.0030*** (0.0009)	0.0030*** (0.0005)
T2 × District COVID-19 Cases			0.0012 (0.0010)	0.0012** (0.0005)
Observations	2,117	2,117	2,117	2,117
R-squared	0.158	0.158	0.163	0.163
Control Mean DV	0.0857	0.0857	0.0857	0.0857
Control SD DV	0.2801	0.2801	0.2801	0.2801

Notes: Standard errors (in parentheses) are clustered at the level of 76 communities (Columns 1 and 3) or 7 districts (Columns 2 and 4). Dependent variable defined in Table 1, and variables and suppressed controls (including community fixed effects) are defined in Table 2. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

G.3 Social Distancing Indicator Alternatives

For the social distancing indicator in the primary outcome, one condition is that respondents must report doing more than the sample median number of “social distancing actions” in the past seven days, as pre-specified, which worked out to being at least seven out of eight actions (as the sample median number was six). One concern might be that this relatively arbitrary threshold of the sample median may be driving the primary results.

Table A.9 shows that our primary results are robust to alternative definitions of the social distancing indicator based on the a respondent’s self-reported number of social distancing actions. First, the dependent variable in Columns (1)-(2) is social distancing measure in which threshold number of self-reported actions to be considered social distancing is six out of eight (i.e., at or above the sample median). Second, the dependent variable in Columns (3)-(4) is social distancing measure which excludes social distancing actions #4 and #6 from Section D as these are conditional on experiencing symptoms and thus might be inadvertently misreported, thereby the threshold number of self-reported actions is changed to five out of six actions.

Under both alternative social distancing indicators, the main treatment effects in Columns (1) and (3) remain statistically insignificant while the coefficients relating to the misperceptions correction treatment in Columns (2) and (4) demonstrate similar treatment effect heterogeneity with respect to the local infection rate. Interestingly, Column (4) also shows treatment effect heterogeneity for the leader endorsement.

Table A.9: **Treatment Effects with Alternative Social Distancing Measures**

VARIABLES	(1)	(2)	(3)	(4)
	Alternative SD Indicator 1	Alternative SD Indicator 1	Alternative SD Indicator 2	Alternative SD Indicator 2
T1: Misperceptions Correction	-0.0060 (0.0159)	-0.0459** (0.0215)	-0.0013 (0.0144)	-0.0452** (0.0198)
T2: Leader Endorsement	-0.0073 (0.0159)	-0.0313 (0.0226)	-0.0018 (0.0144)	-0.0372* (0.0204)
T1 \times District COVID-19 Cases		0.0024** (0.0012)		0.0026** (0.0011)
T2 \times District COVID-19 Cases		0.0014 (0.0012)		0.0021** (0.0011)
Observations	2,117	2,117	2,117	2,117
R-squared	0.197	0.199	0.163	0.167
Control Mean DV	0.1244	0.1244	0.0939	0.0939
Control SD DV	0.3302	0.3302	0.2919	0.2919

Notes: Dependent variable in Columns 1-2 is social distancing measure in which threshold number of self-reported actions to be considered social distancing is 6 out of 8 (instead of 7 out of 8). Dependent variable in Columns 3-4 is social distancing measure which excludes social distancing actions 4 and 6 that are conditional on experiencing symptoms (threshold number of self-reported actions is changed to 5 out of 6 actions). Variables and suppressed controls (including community fixed effects) are defined in Table 2. Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

More out of curiosity than robustness, Table A.10 also presents separate regression estimates on the two components of the primary social distancing indicator: the self-report of social distancing in Columns (1)-(2) and the other's report of social distancing in Columns (3)-(4). The main treatment effects in Columns (1) and (3) remain statistically insignificant. Meanwhile, coefficients relating to the misperceptions correction treatment in Columns (2) and (4) reveal that heterogeneity with respect to the local infection rate is largely driven by its effect on self-reported social distancing, with point estimates on other's reports in Column (4) being in the same direction but statistically insignificant. Interestingly, Column (2) also shows treatment effect heterogeneity for the leader endorsement on the self-reported social distancing outcome.

Table A.10: **Treatment Effects on Self and Others' Reports of Social Distancing**

VARIABLES	(1) Self-Report of SD	(2) Self-Report of SD	(3) Others' Report of SD	(4) Others' Report of SD
T1: Misperceptions Correction	0.0134 (0.0238)	-0.0541 (0.0349)	0.0010 (0.0181)	-0.0271 (0.0259)
T2: Leader Endorsement	-0.0189 (0.0234)	-0.0707** (0.0344)	0.0145 (0.0183)	0.0134 (0.0261)
T1 \times District COVID-19 Cases		0.0040*** (0.0016)		0.0017 (0.0013)
T2 \times District COVID-19 Cases		0.0031** (0.0015)		0.0001 (0.0013)
Observations	2,117	2,117	2,117	2,117
R-squared	0.211	0.214	0.333	0.334
Control Mean DV	0.4061	0.4061	0.2113	0.2113
Control SD DV	0.4914	0.4914	0.4084	0.4084

Notes: Dependent variables in Columns 1-2 and Columns 3-4 are defined in Rows 4 and 7 of Table 1, respectively. Variables and suppressed controls (including community fixed effects) are defined in Table 2. All regressions also include community fixed effects. Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

G.4 Local COVID-19 Infection Rate Alternatives

In the main paper, we measure the local COVID-19 infection rate as the district-level cumulative COVID-19 cases per 100,000 population at the start of the endline survey. One concern might be that we only observe treatment effect heterogeneity using this one measure of the local infection rate but not other potentially justifiable measures.

Table A.11 instead shows that the primary results are robust to various other measures of the local infection rate: 1) "District COVID-19 Case Count" is the total count of district-level cumulative COVID-19 cases at the start of the endline survey (i.e., Table A.1, Column 1); 2) " $\mathbb{1}(\text{Cases Above District Median})$ " is an indicator for if the respondent's district has above median COVID-19 cases relative to the sample's seven districts, thus applying to the top three districts; 3) " $\mathbb{1}(\text{Cases Above Sample Median})$ " is an indicator for if the respondent's district has above median COVID-19 cases relative to all sample respondents, applying to the top two districts (due to the large sample in the top-COVID-19 district); and 4) " $\mathbb{1}(\text{Cases Above National Average})$ " is an indicator if the respondent's district-level cumulative COVID-19 cases per 100,000 population is estimated at above Mozambique's national average at the start of the endline survey, which applies to only the top-COVID-19 district in the sample.³

3. The national average of COVID-19 cases per capita at the start of the endline survey is estimated by dividing the 9,296 cumulative COVID-19 cases in Mozambique on October 5, 2020 (Johns Hopkins University, 2022) by the World Bank 2021 population estimate of 32.08 million (World Bank, 2023). Thus, we estimate 28.98 cases per 100,000 as the national average at this time.

In Columns (1)-(4), we observe that the misperception correction intervention has statistically significant treatment effect heterogeneity with respect to the local infection rate, as before, except in Column (4) where the standalone treatment effect is marginally significant (p-value=0.101). We conclude then that the finding is not an exception driven by our specific measure of the local COVID-19 infection rate.

Table A.11: **Treatment Effects with Alternative Local COVID-19 Infection Rates**

VARIABLES	(1)	(2)	(3)	(4)
	Primary SD Indicator	Primary SD Indicator	Primary SD Indicator	Primary SD Indicator
T1: Misperceptions Correction	-0.0375** (0.0164)	-0.0342* (0.0188)	-0.0312* (0.0159)	-0.0218 (0.0133)
T2: Leader Endorsement	-0.0208 (0.0172)	-0.0212 (0.0199)	-0.0203 (0.0168)	-0.0140 (0.0141)
T1 × District COVID-19 Case Count	0.0009*** (0.0003)			
T2 × District COVID-19 Case Count	0.0003 (0.0003)			
T1 × 1(Cases Above District Median)		0.0727*** (0.0281)		
T2 × 1(Cases Above District Median)		0.0292 (0.0274)		
T1 × 1(Cases Above Sample Median)			0.0909*** (0.0305)	
T2 × 1(Cases Above Sample Median)			0.0378 (0.0288)	
T1 × 1(Cases Above National Average)				0.1034** (0.0410)
T2 × 1(Cases Above National Average)				0.0358 (0.0375)
Observations	2,117	2,117	2,117	2,117
R-squared	0.163	0.161	0.163	0.163
Control Mean DV	0.0857	0.0857	0.0857	0.0857
Control SD DV	0.2801	0.2801	0.2801	0.2801

Notes: Dependent variable defined in Table 1, and treatment indicators and suppressed controls (including community fixed effects) are defined in Table 2. Remaining displayed variables interact treatment indicators with ways to specify the local COVID-19 infection rate other than our preferred measure: district-level cumulative COVID-19 cases per 100,000 population. "District COVID-19 Case Count" is the total count of district-level cumulative COVID-19 cases at the start of the endline survey (i.e., Table A.1, Column 1). "1(Cases Above District Median)" is equal to one if the respondent's district has above median COVID-19 cases relative to the sample's seven districts, and zero otherwise. "1(Cases Above Sample Median)" is equal to one if the respondent's district has above median COVID-19 cases relative to all sample respondents, and zero otherwise. "1(Cases Above National Average)" is equal to one if the respondent's district-level cumulative COVID-19 cases per 100,000 population is estimated at above Mozambique's national average at the start of the endline survey, and zero otherwise. Robust standard errors in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

G.5 Excluding Chimoio District

A central finding of the paper is the heterogeneity in the treatment effect of the misperceptions correction treatment with respect to local COVID-19 cases per 100,000 population (Table 2 Column 2). A question

that arises is whether this heterogeneity is entirely driven by the Chimoio district, which has the highest COVID-19 case load in the sample by a fair margin (see Figure 2 and Appendix B.3). We therefore test the robustness of our findings to excluding from the sample the 524 respondents in Chimoio district (one-quarter of the sample), thereby only exploiting the more limited variation in district-level case loads across the remaining six districts.

Table A.12 below presents coefficient estimates from this restricted sample. First of all, Column (1) reveals that the coefficient on the misperceptions correction treatment is negative and statistically significant at the 10% level. Because this sample drops the district with the highest case loads, this result is consistent with theoretical predictions and previous findings that at lower case loads, the misperceptions correction treatment effect is more likely to be negative.

In Column (2), where we test for heterogeneity in the treatment effect, results are quite similar to the findings in Column (2) of Table 2 in the main text. The T1 main effect and interaction term coefficients are of similar magnitudes to those in Column (2) of Table 2, and maintain statistical significance at conventional levels (the T1 interaction term coefficient is now significant at the 10% instead of 5% level).

In sum, our central findings regarding heterogeneity in the treatment effect of the misperceptions correction treatment are robust to excluding from the sample respondents from the district (Chimoio) with the highest COVID-19 case loads.

Table A.12: Treatment Effects Excluding Chimoio District

VARIABLES	(1) Primary SD Indicator	(2) Primary SD Indicator
T1: Misperceptions Correction	-0.0237* (0.0131)	-0.0410** (0.0194)
T2: Leader Endorsement	-0.0150 (0.0141)	-0.0263 (0.0208)
T1 × District COVID-19 Cases		0.0019* (0.0010)
T2 × District COVID-19 Cases		0.0012 (0.0010)
Observations	1,593	1,593
R-squared	0.141	0.142
Control Mean DV	0.0710	0.0710
Control SD DV	0.2570	0.2570

Notes: Regressions exclude 524 respondents from Chimoio district. Dependent variable defined in Table 1, and variables and suppressed controls (including community fixed effects) are defined in Table 2. Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

H Populated Pre-analysis Plan

On August 25, 2020, prior to baseline data collection, we uploaded our pre-analysis plan (PAP) “Accelerating Changes in Norms about Social Distancing to Combat COVID-19” to the American Economic Association’s RCT Registry, registration ID number AEARCTR-0005862: <https://doi.org/10.1257/rct.5862-3.0>.

In our PAP, we specify the following regression for our primary analysis, which is the same as Equation 8 in the main text:

$$Y_{ijd} = \beta_0 + \beta_1 T1_{ijd} + \beta_2 T2_{ijd} + \eta B_{ijd} + \delta_{ijd}^{others} + \delta_{ijd}^{leaders} + \gamma_{jd} + \varepsilon_{ijd} \quad (\text{H.1})$$

where Y_{ijd} is the social distancing indicator for household i in community j and district d ; $T1_{ijd}$ and $T2_{ijd}$ are indicator variables for the misperceptions correction and leader endorsement treatment groups, respectively; B_{ijd} is the baseline value of the dependent variable; γ_{jd} are community fixed effects; and ε_{ijd} is a mean-zero error term. We report robust standard errors. The regression also controls for the number of other survey respondents and community leaders who report knowing the survey respondent at baseline (in Round 2). Specifically, δ_i^{others} is a vector of dummy variables for the distinct number of other surveyed study respondents who report knowing the household (0, 1, 2, . . . , 7, 8 or more; where 8 is the first integer where over 90% of the sample is represented by previous non-negative integers), and $\delta_i^{leaders}$ is a vector of dummy variables for the distinct number of community leaders who report knowing the household (0, 1, 2, 3, 4; where 4 is maximum number of leaders found within one of the 76 sample communities). Including this control variable helps reduce residual variance in the dependent variable, because respondents who are known by more others in the community will also have more reports of social interactions with others. These results are presented in the main paper in Table 2 Column (1) and are also replicated in Column (1) of Table A.13.

Additionally, we pre-specified the following secondary analyses. First, we analyze impacts of the social distancing treatments on the separate components of the social distancing index—the self and others’ report. These results are presented in Table A.13 Columns (2) and (3), respectively. Treatment effects on these outcomes are very similar to those in Column (1) in that they are small in magnitude and statistically insignificant. Second, we also pool SD1 and SD2 together to examine the effect of some endorsement of social distancing (whether by other community members or by community leaders) on the primary social distancing outcome. The coefficient in Table A.13 Column (4) is also small in magnitude and not statistically significantly different from zero at conventional levels.

Table A.13: **Additional Pre-specified Analyses**

VARIABLES	(1) Primary SD Indicator	(2) Self-Report of SD	(3) Others' Report of SD	(4) Primary SD Indicator
T1: Misperceptions Correction	0.0042 (0.0140)	0.0134 (0.0238)	0.0010 (0.0181)	
T2: Leader Endorsement	-0.0054 (0.0137)	-0.0189 (0.0234)	0.0145 (0.0183)	
Pooled SD Treatments				-0.0006 (0.0116)
Observations	2,117	2,117	2,117	2,117
R-squared	0.158	0.211	0.333	0.158
Control Mean DV	0.0857	0.4061	0.2113	0.0857
Control SD DV	0.2801	0.4914	0.4084	0.2801

Notes: Dependent variables are defined in Table 1. “T1: Misperceptions Correction” is an indicator equal to one if respondent was randomly assigned to the misperceptions correction treatment, and zero otherwise. “T2: Leader Endorsement” is an indicator equal to one if respondent was randomly assigned to the leader endorsement treatment, and zero otherwise. “Pooled SD Treatments” is an indicator equal to one if respondent was randomly assigned to the misperceptions correction treatment or leader endorsement treatment, and zero otherwise. Controls are as defined in Table 2. All regressions also include community fixed effects. Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

We also randomly assigned a family of treatments to improve COVID-19 knowledge in the same study population.⁴ Randomization of the misperceptions correction and leader endorsement treatments were stratified within 76 communities and within the separate knowledge treatment conditions (i.e., the knowledge and social distancing treatments were cross-randomized). As pre-specified, we run a regression on the primary social distancing outcome with indicators for social distancing treatments, the cross-randomized knowledge treatments and their interaction terms. Results are presented in Table A.14, and show no large or statistically significant interaction effects between the social distancing and knowledge treatments.

4. The pre-analysis plan (PAP) for the knowledge study can be found here: <https://fordschool.umich.edu/mozambique-research/combating-COVID-19>.

Table A.14: **Interactions between Social Distancing and Knowledge Treatments**

(1)	
VARIABLES	Primary SD Indicator
T1: Misperceptions Correction	-0.0237 (0.0214)
T2: Leader Endorsement	-0.0210 (0.0222)
K1: Incentive	-0.0218 (0.0241)
K2: Feedback	-0.0025 (0.0251)
K3: Incentive & Feedback	-0.0144 (0.0238)
T1 × K1	0.0545 (0.0390)
T2 × K1	0.0249 (0.0372)
T1 × K2	0.0467 (0.0397)
T2 × K2	0.0139 (0.0385)
T1 × K3	0.0404 (0.0382)
T2 × K3	0.0374 (0.0372)
Observations	2,117
R-squared	0.160
Control Mean DV	0.0857
Control SD DV	0.2801

Notes: Dependent variable is defined in Table 1. Social distancing treatments are defined in Table 2. “K1 Incentive”, “K2 Feedback”, and “K3 Incentive & Feedback” are indicators equal to one if respondent was randomly assigned to one of these knowledge treatments, and zero otherwise. Remaining regressors represent interactions between social distancing treatments and the knowledge treatments. Controls are as defined in Table 2. Regression also includes community fixed effects. Robust standard errors in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

References for the Online Appendix

- Jones, Sam, Eva-Maria Egger, and Ricardo Santos. 2020. "Is Mozambique Prepared for a Lockdown During the COVID-19 Pandemic?" *UNU-WIDER Blog*, <https://www.wider.unu.edu/publication/mozambique-prepared-lockdown-during-covid-19-pandemic>.
- Ministry of Health. 2020. *COVID-19 Epidemiological Situation, Republic of Mozambique*. Maputo, Mozambique: Government of Mozambique.
- National Institute of Statistics (INE). 2017. *Mozambique Population and Housing Census 2017*. Maputo, Mozambique: Government of Mozambique.
- Nyusi, Felipe Jacinto. August 5, 2020. *Communication to the Nation of His Excellency Philip Jacinto Nyusi, President of Republic of Mozambique, on the New State of Emergency, within the Scope of the Coronavirus Pandemic COVID-19*. Maputo, Mozambique: Maputo Mozambique.
- Republic of Mozambique. April 2, 2020. "*Bulletin of the Republic*", *I Series, No. 64*. Maputo, Mozambique.
- U.S Embassy in Mozambique. 2020. "COVID-19 Information," <http://mz.usembassy.gov/covid-19-information/>.