Encouraging the resumption of economic activity after COVID-19: Evidence from a large scale-field experiment in China

Citation: Juan Palacios, Yichun Fan, Erez Yoeli, Jianghao Wang, Yuchen Chai, Weizeng Sun, David G. Rand, Siqi Zheng, Encouraging the resumption of economic activity after COVID-19: Evidence from a large scale-field experiment in China, Proceedings of the National Academy of Sciences Feb 2022, 119 (5) e2100719119

As Published: 10.1073/pnas.2100719119

Publisher: Proceedings of the National Academy of Sciences

Persistent URL: https://hdl.handle.net/1721.1/139852

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Encouraging the resumption of economic activity after COVID-19: Evidence from a large scale-field experiment in China

Juan Palacios[a], Yichun Fan[a], Erez Yoeli[b], Jianghao Wang[a,c], Yuchen Chai[a], Weizeng Sun[d], David G. Rand[b,e], and Siqi Zheng[a,1]

[a]Sustainable Urbanization Lab, Center for Real Estate, Department of Urban Studies and Planning, Massachusetts Institute of Technology, Cambridge, MA 02139; [b]Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02142; [c]Institute of Geographic Sciences & Natural Resources Research, Chinese Academy of Sciences, 100101 Beijing, China; [d]School of Economics, Central University of Finance and Economics, 100081 Beijing, China; and [e]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139

Edited by Elke Weber, Departments of Psychology & Public Affairs, Andlinger Center for Energy and Environment, Princeton University, Princeton, NJ; received January 13, 2021; accepted November 16, 2021

As the COVID-19 pandemic comes to an end, governments find themselves facing a new challenge: motivating citizens to resume economic activity. What is an effective way to do so? We investigate this question using a field experiment in the city of Zhengzhou, China, immediately following the end of the city's COVID-19 lockdown. We assessed the effect of a descriptive norms intervention providing information about the proportion of participants' neighbors who have resumed economic activity. We find that informing individuals about their neighbors' plans to visit restaurants increases the fraction of participants visiting restaurants by 12 percentage points (37%), among those participants who underestimated the proportion of neighbors who resumed economic activity. Those who overestimated did not respond by reducing restaurant attendance (the intervention yielded no "boomerang" effect); thus, our descriptive norms intervention yielded a net positive effect. We explore the moderating role of risk preferences and the effect of the intervention on subjects' perceived risk of going to restaurants, as well as the contrast with an intervention for parks, which were already perceived as safe. All of these analyses suggest our intervention worked by reducing the perceived risk of going to restaurants.

COVID-19 | descriptive norms | field experiment | voluntary economic resumption | policy

The respiratory syndrome-coronavirus 2 (SARS-CoV-2) and the policies implemented to reduce its transmission (1, 2) have sent the global economy into its deepest recession since the Great Depression (3). In the second quarter of 2020, the US economy alone exhibited its greatest gross domestic product reduction in modern history due to the COVID-19 pandemic (4). Although mandatory government shutdowns are commonly thought to be the cause of reduced activity, there is mounting evidence that most of the reduction in economic activity was voluntary, spurred by fears of the virus (5, 6). At the time of writing, there are still many countries with high rates of COVID-19, but, in some countries, the virus has been contained by low rates of community transmission (e.g., Korea, Japan, and China) and/or vaccination (e.g., the European Union, United Kingdom, the United States, and Israel) (7–9). Governments are thus devoting increasing attention to reopening their services and the economy more broadly. In the United States, for instance, President Joe Biden has given prime billing to reopening schools and economic recovery (10). However, economic recovery will remain stymied so long as the public remains skeptical that it is safe to engage in normal economic activity (11)—a skepticism that is exacerbated by concerns over the safety of vaccines (12), the appearance of new variants of the virus (e.g., the "Delta" variant), and a general distrust of government (13). As former chair of the Council of Economic Advisors Austan Goolsbee recently said, after unemployment numbers surged alongside virus cases in December 2020, "Jobs day lesson 1 million: The virus is the boss" (14).

Can government officials and others effectively communicate when it is safe for individuals to resume normal economic activity? There has been almost no research on this question to date (15) and, to the best of our knowledge, no work that measures the effect of messaging on actual behavior rather than just self-report intentions.

In this study, we explore the effect of "descriptive norms" on individuals' willingness to resume economic activity in the aftermath of a COVID-19 lockdown. Descriptive norm interventions provide information about how common a behavior is, with the goal of encouraging people to behave in the way others are behaving (e.g., to engage in the behavior if the norm states that most other people do, and to avoid the behavior if the norm states that few people are engaging in it). Descriptive norms are a commonly employed "nudge" and have been successfully used to influence a variety of behaviors, including reducing alcohol consumption (16), improving tax compliance (17), increasing the reuse of hotel towels (18), conserving energy (19), reducing smoking, and engaging in other health-related behaviors (20). Descriptive norms are thought to be effective messages because

Significance

When the COVID-19 pandemic ends and it becomes safe to resume economic behavior, we will need to find effective ways of communicating that it is truly safe to do so. In this study, we tested a simple "nudge" that informed people of the proportion of their neighbors who were planning to visit a restaurant over the weekend, so that those who felt that restaurants were very unsafe would realize that many others felt comfortable visiting them. Our nudge successfully motivated such hesitant individuals to increase restaurant visits by 37%, and additional analyses indicate they indeed felt that visiting restaurants was safer.

Author contributions: J.P., Y.F., E.Y., and S.Z. designed research; J.P., Y.F., E.Y., J.W., Y.C., and W.S. performed research; J.P. and Y.F. analyzed data; J.P., E.Y., D.G.R., and S.Z. wrote the paper; J.W. and Y.C. designed the smartphone app for study; and Y.F., Y.C., and W.S. designed the sample recruitment.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: sqzheng@mit.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2100719119/-/DCSupplemental.

Published January 26, 2022.

ECONOMIC SCIENCES

others' behavior can inform whether it is beneficial to engage in a behavior—whether the behavior is, e.g., safe, helpful, or expected (21). Further, conforming to others' behavior is a fundamental feature of how people learn and an important driver of human cultural evolution (22). Descriptive norms are also appealing from a pragmatic perspective because they are inexpensive to implement: It is virtually costless to include descriptive norms messages with other outgoing messages and usually relatively low cost to collect the requisite data on how common a behavior is.

Descriptive norms, however, do not always generate large behavioral effects (23–26) and can sometimes even backfire (27), a phenomenon sometimes called "the boomerang effect." This is particularly a concern for people with relatively high prior beliefs on how common the behavior is. In the context of resuming economic activity following a pandemic, there is another particularly powerful reason to fear that descriptive norms will backfire: The more people go to public spaces like restaurants, the higher the objective risk of being infected if one joins them. If descriptive norm information causes people to believe that more people are going out, this could quite rationally lead them to be discouraged from visiting such spaces due to risk avoidance (28, 29). Thus, it is unclear ex ante whether descriptive norms will help or hurt the resumption of economic activity.

Experimental Design

We test the impact of descriptive norms in a field experiment in the Chinese city of Zhengzhou (population 10.1 million), just after the end of a government-mandated COVID-19 shutdown (Fig. 1*A*). We recruited 622 subjects from the city's residents who work in the central business district of the city. Recruitment was done via text messages, phone calls, and advertisements on social media. Recruitment lasted a total of 5 wk (see *Methods* and *SI Appendix*, section H.4 for a detailed description of the recruitment procedure and statistics of sample attrition and new participants recruited by time). Each participant received a recruitment text message with instructions to install the smartphone app, which we designed especially for the study. The app is able to collect the global positioning system (GPS) coordinates of participants during the study period, send tailored questionnaires, and pay "red-pocket" money, a commonly used form of payments in China. Upon recruitment, subjects were randomized into either a treatment group or a control with equal probability via block randomization (*Methods* for details).

Fig. 1*B* displays an overview of the experimental design. For all subjects—those in both the control and the treatment group—on the first full week following their recruitment, we performed a survey on Wednesday, in which we collected whether they plan to go to a restaurant in the coming weekend and their belief on the percentage of their neighbors who plan to do so. Belief elicitation was incentive compatible: Subjects received a small (2 Chinese Yuan, equivalent to roughly 0.30 USD; see *SI Appendix*, Table H.3 for a description of monetary compensations to participants) bonus if their guess was within 2 percentage points (pp) of the truth.

Then, on Friday of that week, we re-elicited subjects' beliefs about the percentage of their neighbors who will actually go to restaurants over the weekend. Finally, on Sunday evening, we asked the respondents to report their restaurant visitation over

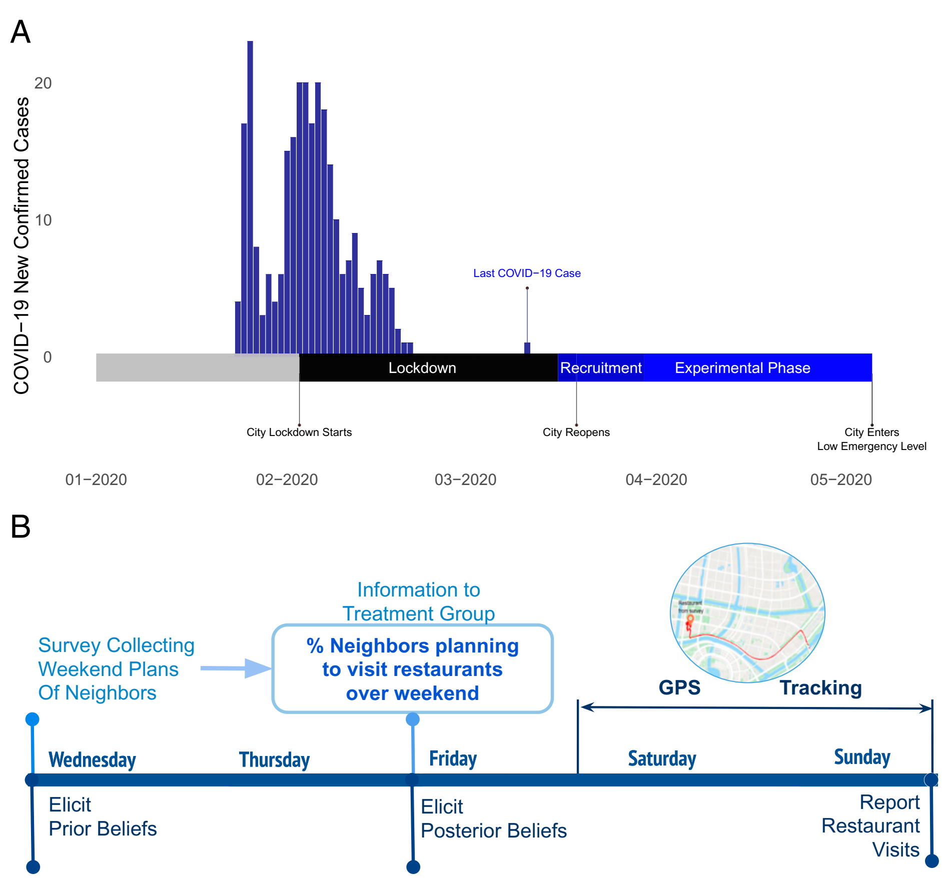

Fig. 1. (*A*) The timeline of the study. (*B*) The outline of questionnaires, messages, and GPS tracking time in an experimental week. The full description of the data collection and construction of the descriptive norm intervention is presented in *Methods*. The exact wording of the descriptive norm is "For each 100 people filling the questionnaire living in your urban district, [[Information Descriptive Norm]] plan to a restaurant this weekend." The full display of the intervention is described in *SI Appendix*, section H.2.

the weekend. We also collect GPS data from their smartphones to cross-validate the stated behaviors. Further details of the experimental design can be found in *Methods*.

All materials (i.e., survey questions, rewards, messages, etc.) were identical across the treatment and control groups, except for one difference: The treatment group received the descriptive norms intervention in advance of the Friday survey re-eliciting their beliefs (whereas the control group proceeded directly to re-eliciting their beliefs, without receiving the descriptive norms information). The descriptive norm intervention was composed of a single sentence that informed subjects of the true proportion of their neighbors who were planning to go to restaurants over the weekend. The true proportion was computed by synthesizing the results from the Wednesday survey, together with results from an extra survey we did on a random sample of another 500 citizens living in the same neighborhoods (i.e., urban district) as our participants. Besides that sentence composed of the descriptive norm intervention, the content and appearance of all materials were identical across both the treatment and control groups (see *SI Appendix*, section H.2 for the wording and display of intervention).

Besides the main outcomes, we collected a rich set of participant characteristics, preferences, beliefs, and attitudes to understand which factors moderate the effectiveness of descriptive norms in our setting. Specifically, we collected 13 moderators at the recruitment date, including demographics (age, gender, having kids or not, education, income), economic preferences (general and health-specific risk preferences, time preference, altruism), knowledge about coronavirus, pre-COVID activity frequency, and community and media trust (see *SI Appendix*, Table H.1 for how each dimension was measured).

Prior to the intervention, our subjects reported a meaningfully curtailed baseline level of restaurant visits: 49.1% of subjects in the control group reported attending restaurants every week on average before the pandemic, but this number was just 31.18% at the time our intervention began, a 17.92 pp (or 57.47%) decline (*SI Appendix*, Table E.1). Our primary hypotheses were that subjects in the treatment group who underestimated others' behaviors (i.e., whose beliefs that others planned to visit restaurants were below the true percentage) would respond to the descriptive norms intervention by increasing these beliefs and modifying their behavior. We preregistered these hypotheses as well as our 13 moderators (randomized controlled trial [RCT] ID: AEARCTR-0005644) and tested these hypotheses by comparing treated individuals with people in the control group with similar prior beliefs.

Results

Our main result is that the descriptive norms intervention effectively corrected subjects' beliefs about others' restaurant visits and that these corrections translated into large increases in their own restaurant visits. Specifically, for those subjects with initial beliefs below the true value, the treatment shifted beliefs upward by 4.64 pp (95% CI: 0.97 to 8.31 pp, $P = 0.013$), with greater increases associated with more downward-biased initial beliefs (Fig. 2*A*). These subjects were subsequently 12.5 pp (95% CI: 1.54 to 23.3 pp, $P = 0.025$) more likely to report visiting a restaurant over the weekend (Fig. 2*B*) relative to the control group, a 37.0% increase. Meanwhile, subjects whose beliefs were above the truth also corrected their beliefs, decreasing them in response to the intervention by 10.9 pp (95% CI: -15.46 to -6.52 pp, $P < 0.001$). However, these subjects did not alter their visits to restaurants almost at all ($\beta = -1.25$, 95% CI: -13.70 to 11.18 pp, $P = 0.842$). Thus, a boomerang effect was not present in our intervention—subjects with higher-than-truth prior beliefs were not deterred from going to restaurants, and the intervention ultimately resulted in a net positive effect on restaurant visits. In *SI Appendix*, sections A.1 and A.2, we confirm the robustness of these results using a variety of parametric and nonparametric specifications.

Next, we perform three secondary analyses that, in combination, suggest that our intervention worked by increasing subjects' perception that restaurants are safe. The first of these analyses focuses on the moderating role of risk preferences. In this analysis, we looked for evidence of moderation by the characteristics, preferences, beliefs, and attitudes that we collected in our surveys on the recruitment date and found that subjects' risk tolerance may moderate the effect of the intervention. Fig. 3*A* shows that risk-averse individuals did not adjust their behaviors when they learned that more people than they expected were going to restaurants (general risk, $\beta = 0.18$ pp; CI, -15.03 to 15.39 pp; $P = 0.981$) (health risk, $\beta = 0.42$ pp; CI, -14.93 to 15.76 pp; $P = 0.958$), whereas risk-tolerant people increased restaurant visits substantially (general risk, $\beta = 21.0$ pp; CI, 4.05 to 37.95 pp; $P = 0.017$) (health risk, $\beta = 21.5$ pp; CI, 5.47 to 37.53 pp; $P < 0.01$). The estimated interaction effect between the intervention dummy and the general risk measure is 24.7 pp (95% CI: 2.28 to 47.2 pp, P = 0.031); for the health risk measure it is 22.2 pp (95% CI: 0.2 to 44.2 pp, P = 0.048). As can be seen in *SI Appendix*, Fig. A.7, subjects' belief adjustment did not depend on their risk tolerance, but whether this adjustment translated to a change in behavior very much did: Only subjects who were relatively risk tolerant were more likely to visit restaurants. *SI Appendix*, section A.4 describes the estimation for the full set of preregistered moderators—besides risk aversion, we find only evidence for education as a significant moderator in our sample (interaction effect for above college education, 30.6 pp; 95% CI, 5.97 to 55.22 pp; P = 0.015). Once again, we undertook a number of tests to assess the robustness of the results (*SI Appendix*, section A.4), and the significant effect of the treatment among risk-seeking individuals with prior beliefs below the truth was robust to correcting SEs for multiple-hypothesis testing (results included in *SI Appendix*, Table A.3).

Second, we consider the effect of the intervention on subjects' stated perceptions of the risk associated with visiting a restaurant. On each Friday, immediately after our intervention, we asked individuals the question, "How would you rate the infectious risk of dining out in restaurants nearby?" Subjects recorded their replies with a seven-point Likert scale. We found that subjects who received the descriptive norms intervention reported somewhat lower perceived risk than those in the control, although this result is only marginally significant (standardized coefficient = -0.160; 95% CI, -0.35, 0.026; P = 0.093) (*SI Appendix*, Table D.1).

Third, and finally, we contrast the impact of our intervention with an additional treatment that repeated the descriptive norms intervention in the same sample of participants, but for parks instead of restaurants. Visiting parks is far less risky than visiting restaurants (30). Parks are open-air spaces where social interactions can take place at a distance, resulting in much less risk of transmission than in restaurants. Subjects' own perceptions of the relative risks conform with this: 82.8% of participants considered visiting local restaurants riskier than visiting parks in the neighborhood (*SI Appendix*, Fig. E.1). Accordingly, park visitation returned to normal very quickly in the aftermath of the pandemic: Within the control group, 33.3% of subjects reported going to parks once a week on average before the pandemic; this number was 37.6% at the time our intervention began, a 4.6 pp (or 13.9%) increase (*SI Appendix*, Table E.1). All this suggests that the descriptive norm intervention would be unlikely to change subjects' beliefs about risk and behaviors, since they already viewed parks as relatively safe and this was reflected in their behavior. For these reasons, we think the intervention serves as a useful comparison with our main intervention, akin to a "placebo" intervention (we acknowledge that we became aware of park visitation having returned to normal only after data collection began, and thus the use of parks as a placebo was post hoc and not preregistered).

ECONOMIC SCIENCES

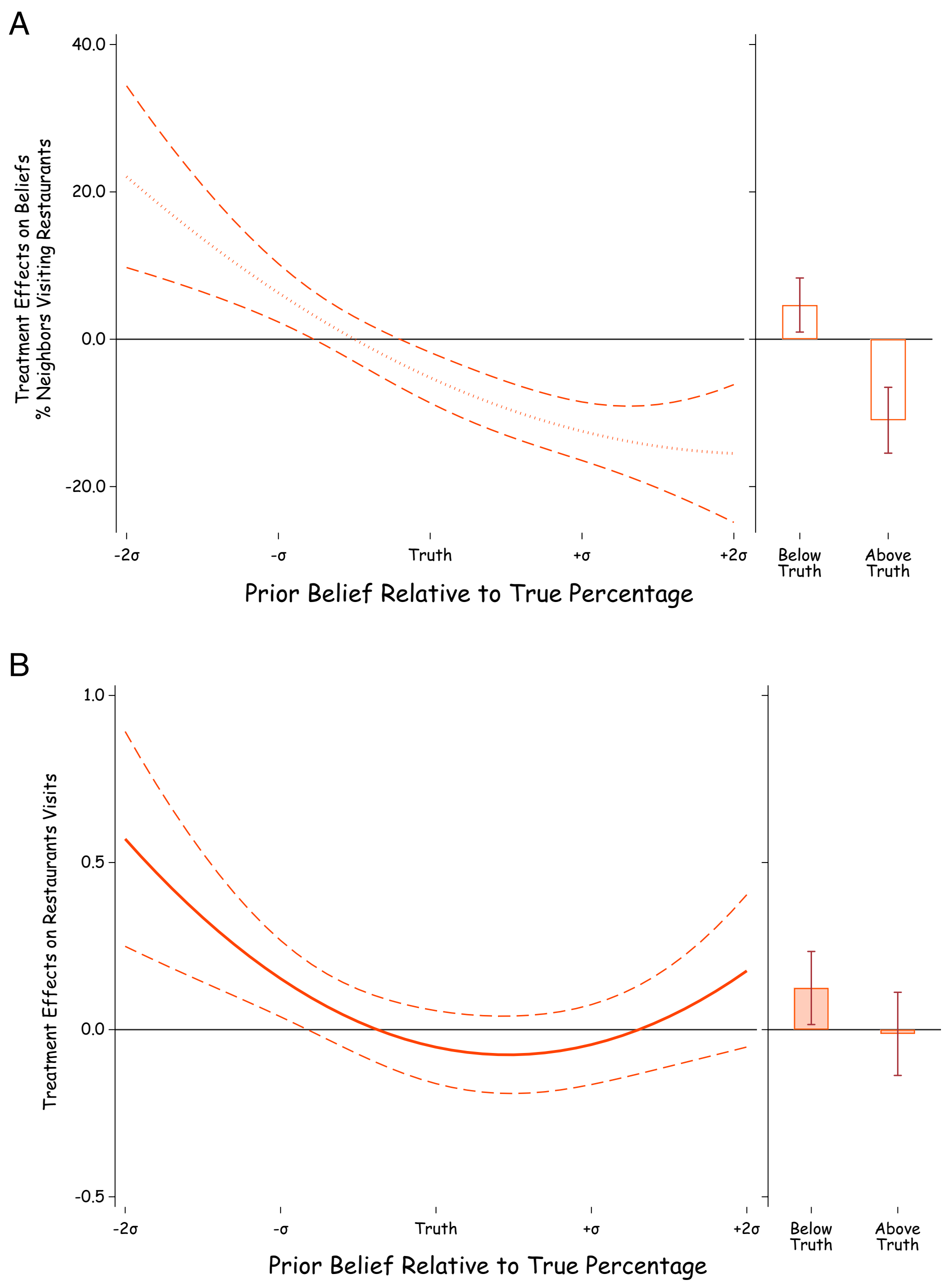

Fig. 2. (*A*) Vertical axes display the change in subjects' beliefs from Wednesday (prior) to Friday (posterior) regarding the percentage of neighbors who are actually going to go to restaurants over the weekend. The horizontal axes display the bias in beliefs prior to our intervention, computed by subtracting subjects' prior beliefs regarding the percentage of neighbors that are planning to go to restaurants from the true percentage of neighbors planning to go to restaurants. The true percentages, used to build the descriptive norm information, are computed from the proportion of individuals that reported in our Wednesday questionnaire to have plans to go to restaurants over the weekend. The sample includes only responses of individuals that were treated for the first time and control subjects that were in the sample for 1 wk. The dotted line represents the average differences between treatment and control groups calculated in a regression analysis interacting the treatment dummy with a polynomial form of the bias in prior beliefs. *SI Appendix*, section A.3 describes the polynomial regression used to construct the plot. Dashed lines represent the 95% confidence interval. The bar graphs included in *A*, *Right* indicate the average change in posterior beliefs for individuals with prior beliefs below and above the truth, separately; the error bars describe the 95% confidence intervals. Nonparametric comparisons between treatment and control groups are displayed in *SI Appendix*, Fig. A.2. *B* describes the change in subjects' visits to restaurants, after our first intervention, as a function of their bias in initial beliefs relative to the intervention information (i.e., prior belief minus true percentage). The solid line represents the average differences between treatment and control groups, calculated in a regression analysis interacting the treatment dummy with a polynomial form of the bias in prior beliefs. *SI Appendix*, section A.3 describes the polynomial regression used to construct the plot. Dashed line represents the 95% confidence interval. The bar graphs included in *B*, *Right* indicate the average change in real visitation rate for individuals with prior beliefs below and above the truth, separately; the error bars describe the 95% confidence intervals. The estimated treatment effects on behavior based on Logit and Probit regressions show the similar levels of statistical significance, as reflected by the *P* values associated with the coefficients of interest (*SI Appendix*, Table A.1). Nonparametric comparisons between treatment and control groups are displayed in *SI Appendix*, Fig. A.4.

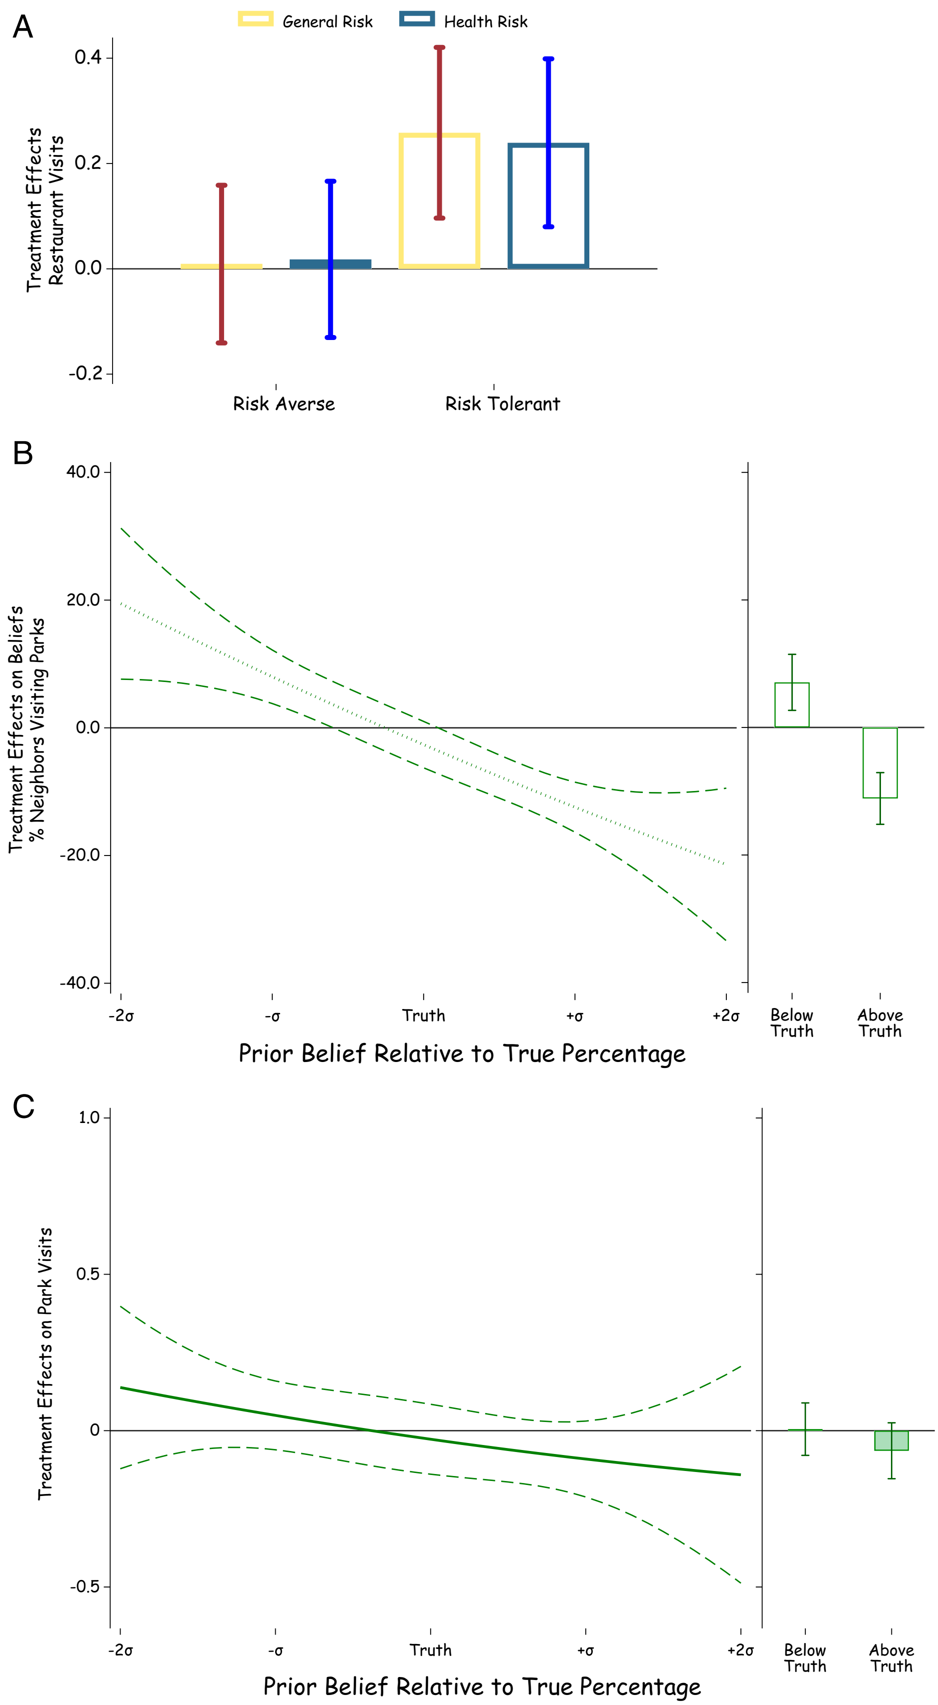

Fig. 3. *A* displays the treatment effects for individuals with low (below median) and high (above median) risk tolerance separately. We display the results for both the general and health-specific risk aversion. Bars describe point estimates and error bars describe the 95% confidence intervals. Sample is restricted to individuals with prior beliefs below the truth. There are significant differences of the treatment effect across the risk-averse and risk-tolerance subgroups based on interaction terms (Eq. **2**), also when using a continuous rather than a group indicator of risk (*SI Appendix*, Fig. A.7), and after adjusting for multiple-hypothesis testing following the techniques proposed by ref. 40 (*SI Appendix*, Table A.3). This moderator analysis has been preregistered at the American Economic Association Registry for randomized control trials (RCT ID: AEARCTR-0005644). *B* describes the treatment effect of descriptive norms on beliefs regarding the percentage of neighbors visiting a park visitation green reported in our Friday survey. *SI Appendix*, section A.3 describes the polynomial regression used to construct the plot. The dotted line represents the average differences between treatment and control groups calculated in a regression analysis interacting the treatment dummy with a polynomial form of the bias in prior beliefs. Dashed lines represent the 95% confidence interval. The bar graphs included in *B*, *Right* indicate the average change in posterior beliefs for individuals with prior beliefs below and above the truth, separately; the error bars describe the 95% confidence intervals. Nonparametric comparisons between treatment and control groups are displayed in *SI Appendix*, Fig. A.2. *C* displays the differences in treatment effects across the distance between individuals' prior beliefs and the true percentage of neighbors planning to go to parks (used to construct the descriptive norm intervention). The sample includes only responses of individuals that were treated for the first time and control subjects that were in the sample for 1 wk. The solid line represents the average differences between treatment and control groups calculated in a regression analysis interacting the treatment dummy with a polynomial form of the bias in prior beliefs. *SI Appendix*, section A.3 describes the polynomial regression used to construct the plot. Dashed lines represent the 95% confidence interval. The bar graphs included in *C*, *Right* indicate the average change in park visits for individuals with prior beliefs below and above the truth, separately; the error bars describe the 95% confidence intervals. Regression tables with different model specifications are displayed in *SI Appendix*, Table A.2, and nonparametric comparisons between treatment and control groups are included in *SI Appendix*, Fig. A.4.

Indeed, we found that the effect of the descriptive norm intervention for parks was, at most, muted. For those subjects with prior beliefs below the truth, the treatment also successfully shifted beliefs upward by 7.33 pp (95% CI: 2.92 to 11.73 pp, $P < 0.001$) (Fig. 3*B*). Yet unlike a restaurant, knowing more people going to parks hardly translated into a change in behavior: We observe only a 3.39 pp increase in park visits, and this increase is not statistically significant (95% CI: −6.7 to 13.51 pp, $P = 0.511$) (Fig. 3*C*). In *SI Appendix*, section A.2, we show that these results are robust to employing additional regression specifications. In *SI Appendix*, Table A.2, we estimate the difference between the effects of the intervention for restaurants and for parks, which we find to be relatively large in magnitude (−0.09 pp, or roughly three-quarters of the main treatment effect on restaurants), but not statistically significant ($P = 0.23$), likely due to lack of statistical power for detection interactions (which require more power to detect than main effects).

Although risk appears to be the most important channel by which our descriptive norms intervention operated, there are other potential channels. For instance, our intervention might change subjects' perceptions of others' normative expectations—that is, whether others think it is desirable, right, fair, etc., to go to a restaurant. This channel is often very powerful (31), and while we cannot rule it out, we did not find evidence for it in our sample. In particular, the treatment effect is not significantly larger among more prosocial individuals, who might be expected to be particularly responsive to information about normative expectations (32) (*SI Appendix*, section A.4). Nor did the treatment significantly change participants' perceptions of the behavior desired by the experimenters (*SI Appendix*, section A.5).

We note that our field experiment included an additional, unrelated intervention that informed subjects of nearby restaurants that participated in a certification program that affirmed these restaurants employed best practices for reducing virus transmission. The intervention was adapted from a program that was developed by Meituan (a popular third-party platform for restaurant reviews similar to Yelp or Tripadvisor) and was widely available in that city at the time of experiment (for implementation details, see *Methods* and *SI Appendix*, section B). We included this treatment to help local policymakers evaluate the effectiveness of this widespread certification practice; it was included in the same experiment as our social norms treatment for pragmatic purposes. Although it does not shed light on the role of social norms, for completeness we report that the certification intervention had no significant effect on restaurant visits (β = 6.00 pp; 95% CI, −2.62 to 13.64 pp; $P = 0.184$; *SI Appendix*, Table B.2).

To mitigate concerns about self-reporting biases, we used the GPS coordinates of participants who agreed to share their data with the research team to evaluate their correspondence with self-reported visits ($N = 175$). GPS validation reaffirms the quality of the self-reported outcomes, showing a high correspondence between the reported restaurant visited and the GPS coordinates of participants. The complete description of the validation exercise is in *SI Appendix*, section G. Finally, using data from our exit survey, we provide evidence showing experimenter demand effects are not biasing our results (*SI Appendix*, section A.5).

Discussion

Overall, our experiment suggests that descriptive norms can serve as an effective tool for motivating a return to normal economic activity by reducing people's fear that such activity is unsafe. By conducting a randomized controlled trial, and combining self-reported data with GPS validations, we provide high-quality evidence of the causal effect of social norms on the resumption of economic activity. We hope our experiment informs politicians' and policymakers' efforts to motivate a return to normal activity in the aftermath of COVID-19.

Our field experiment also provides a particularly detailed look into how descriptive norms worked in a particularly challenging setting. Because individual actions carry substantial health risks both for participants and for those people around them, and because objective infection risk increases with the number of others going out, this is a context where descriptive norms could reasonably be expected not to work as intended. Interestingly, however, learning that others planned to resume activity led participants to think that this behavior was actually safer, rather than more dangerous. The intervention had a positive net effect, motivating behavior change among those who underestimated how many others had resumed activity, but without meaningfully changing the behavior of those who overestimated this proportion. This speaks to the power of social norms to change beliefs even in uncertain times, even in the face of countervailing forces.

Our study has a number of limitations. One particularly substantial limitation is that it was run in just one setting at a very particular moment in time. Caution needs to be taken when generalizing our results to other cultures and time periods. People in different cultures may have greater distrust of authorities attempting to implement such an intervention (13), be less likely to be influenced by others' behavior (33–35), or have more accurate beliefs about others' activity than did the individuals in our Chinese sample. Moreover, the increasing availability of vaccines may fundamentally alter the risk-related inferences that individuals draw from the behavior of others, particularly if others' vaccination status is either known or unknown. Such factors could change the impact of a descriptive norms intervention like ours.

Methods

Study Context. Our study was conducted in Zhengzhou, the capital city of Henan province in central China, which is 500 km from the city of Wuhan where the first outbreak of COVID-19 was documented. The city had a total of 158 infected confirmed COVID cases, with the last case reported on March 11. The city government imposed strict lockdown measures on citizens from January 26 to March 19, 2020. During the city lockdown, nonessential businesses were closed, and citizens were not allowed to leave their homes, except for essential purchases or health reasons. In addition, participants' residential communities imposed a series of measures on residences to avoid the spread of the virus (e.g., limiting access to the property).

Starting from March 19, restaurants and parks in the city could start to reopen (Fig. 1*A* displays the timeline of the study, together with the evolution of the number of daily COVID-19 cases). However, customers were obliged to take extensive precautionary measures, imposed by the city government or their residential community, when leaving their homes. Customers going to restaurants were required to wear masks and show their health certificates, a color scheme embedded in citizens' smartphones to signal infection levels based on their presence in areas of high infection risk (e.g., Wuhan). Individuals that had fever, cough, and other COVID-related symptoms were not allowed to enter the dining area. Starting from May 6, 2020 (end date of our experiment), the measures to control the coronavirus spread in the city were further relaxed: The city started loosening the remaining restrictions and requirements for citizens in public spaces (e.g., lifting requirements to wear masks in public spaces). We implemented our experimental intervention for the first time on March 30, 10 d after the local government began to loosen the lockdown measures and dining-out services were allowed to reopen in the city. At that time, around 50% of restaurants were opened, based on data from a sample of 5,000 restaurants collected from the restaurant portal Dianping.com (the Chinese version of Yelp.com). Over the 5 wk of the experiment, the percentage of restaurants reopening increased to nearly 70%. Our last round of intervention took place on May 1, the last weekend before the loosening of COVID precautionary policies in the city.

Recruitment. Our main recruitment started on March 16 and finished on March 22 (see Fig. 1*A* for the study timeline). In addition, we refreshed the sample every week during the experiment to keep the sample size stable (see *SI Appendix*, Table H.1 for an overview of the number of participants over the experimental weeks). In total, we recruited 622 valid residents working around the central business district of the city, via text messages and phone calls, to participate in the study. Each participant received a recruitment text message with instructions to install the smartphone app designed for the study, and a description of the compensation for being part of the experiment (see *SI Appendix*, Table H.3 for a description of all monetary compensations).

When individuals agreed to participate and signed the informed consent, a group of local assistants explained the study purpose in more detail, answered general questions about the experiment, supported individuals with technical difficulties in installing the smartphone app, and asked the participants to complete a short screening questionnaire. The smartphone app developed for the study is able to collect the precise GPS coordinates of participants during the study period, send tailored questionnaires, and pay red -pocket money. We used the GPS data collected from the app to validate the self-reported weekend plans of individuals (results from validation exercise are reported in *SI Appendix*, section G).

Participants completed two pretreatment surveys measuring detailed individual information, including their sociodemographic characteristics, perception and knowledge regarding the COVID-19 pandemic, pre-COVID habits for going to parks and restaurants, perceived social norms, and community trust (36). In addition, we include the survey instruments from the Global Preference Survey (37), a globally representative dataset with experimentally validated measures of risk attitude, altruism, and time preferences (questionnaires are attached to Questionnaire File) (*SI Appendix*, section H). These measures serve as important controls and factors to be used for heterogeneity analysis, as prespecified in our preregistered analysis plan (RCT ID: AEARCTR-0005644). At the end of the experiment, participants received, again, the same questions about attitudes, norms, and economic preferences, to test for any changes during the experiment.

Descriptive Norm Intervention. In this study, we manipulate the beliefs of our participants regarding the proportion of neighbors that are visiting restaurants over the weekend. We follow the design introduced by ref. 38, to provide a random subset of individuals in our sample with truthful information intended to shift beliefs regarding others' restaurant visitation. The treatment consists of providing truthful personalized information regarding the proportion of neighbors planning to visit restaurants, plausibly affecting beliefs regarding the number of neighbors actually going to restaurants over the weekend.

Following the timeline of Fig. 1*B*, we implemented the descriptive norm treatment as follows:

- From Tuesday to Thursday, we took the following steps.
  1) We elicited the individuals' plans to visit restaurants in the upcoming weekend to build our information treatment. We asked them to report whether they intended to go to a restaurant during the upcoming weekend.
  2) We elicited subjects' prior beliefs regarding the percentage of their neighbors who were, at that point, planning to go to restaurants over the weekend.
  3) We also elicited subjects' beliefs regarding the percentage of neighbors that will actually go to restaurants.

All participants in our experiment were incentivized to make the right guess by receiving monetary compensations for each accurate guess (within two percentage points of the real proportion of neighbors; *SI Appendix*, Table H.3 describes all monetary compensations in the study). The monetary compensation of two Chinese Yuan for right guesses was given at the end of the experiment to avoid biasing participants during the experiment. Prior beliefs about planned and actual behavior of neighbors are highly correlated (*SI Appendix*, Fig. A.3).

To enhance the representativeness of the treatment information, we complemented the information from our experimental sample with a survey of a random sample of 500 citizens living in the same neighborhoods as our participants, who also reported their plans for the weekend (i.e., we have about 1,000 individuals, in total, per week, answering regarding their plans). The percentage of neighbors planning to go to restaurants over the weekend for each neighborhood (i.e., urban district) is used to construct our information treatment.

The sample's average prior belief regarding the percentage of neighbors going out was quite close to the truth but with substantial variations across individuals. The standard deviation for restaurant beliefs was 23 percentage points; full distribution of priors and posteriors is displayed in *SI Appendix*, Fig. A.1.

- On Friday morning, just before the weekend started, we took the following steps.
  1) Everyone in our sample was reminded of their prior beliefs regarding the number of people that will actually go to restaurants over the weekend, as reported in the midweek survey.
  2) People in the treatment group were informed about the actual proportion of neighbors who planned to go to restaurants the following weekend.
  3) We reelicited beliefs regarding the proportion of neighbors who will actually go to restaurants over the weekend. Again, all participants in our experiment were incentivized to make the right guess (*SI Appendix*, Table H.3 describes all monetary compensations in the study).

Responses from our midweek survey indicated that an average of 38% (51%) of all subjects planned to go to restaurants (parks) over the weekend across the five urban districts (see *SI Appendix*, Table H.2 for the exact percentages computed for each urban district and week separately for restaurants and parks).

- Over the weekend, we took the following steps.
  1) The GPS coordinates of participants were monitored from Friday afternoon to Sunday evening for those subjects that granted the app designed for the study access to their GPS data.
  2) On Sunday night, all participants reported whether they had been to any restaurants or parks during the weekend and, if so, the time of their visit and the name of the place they went to. We cross-validated the self-reported weekend behaviors with the GPS data retrieved from the smartphones of participants via our study app (see *SI Appendix*, section G for a detailed description of the validation exercise).

For parks, our placebo treatment, we repeated the entire procedure described above, with identical questions, and incentives.

Randomization Strategy. We used a block randomization strategy, where individuals are randomly assigned to each experimental arm, stratified by the geographical area of their home address (i.e., five urban districts in the city) and pre-COVID frequency of restaurant visits (i.e., a dummy variable indicating going to restaurants at least once a week or less).

The experiment involved two interventions. Our primary intervention was the descriptive norm intervention. A secondary intervention was the restaurant certificate intervention (see *Results* and *SI Appendix*, section B for details). We randomized individuals into these interventions independently, as follows. For each intervention, within each of 10 strata, individuals were randomly assigned to the intervention with 50% likelihood—that is, they were equally likely to receive the intervention or to be in the control. Consequently, for each individual within a stratum, there were the following four possible assignments, with equal likelihood: 1) control arm for both interventions, 2) descriptive norm intervention only, 3) restaurant certificate intervention only, and 4) both the descriptive norm and restaurant certificate interventions. The balance tests between control and treatment subsamples show the randomization was well executed (see *SI Appendix*, Table H.1 for balance tests).

Main Econometric Analysis. We implement a series of regression analyses to examine changes in beliefs and behavior associated with our treatment interventions. We run the following regression separately for restaurant and park activities:

$$Y_i = \alpha T_i + \beta T_i \times AboveTruth_i + \delta AboveTruth_i + X_i\Gamma + E_{it}, \qquad [1]$$

where Y_i includes the two main outcomes of interest $\{PosteriorBelief_i, Behavior_i\}$. The variable $PosteriorBelief_i$ is individual *i*'s belief about the proportion of neighbors who will actually go to the restaurants, elicited on Friday post treatment survey. $Behavior_i$ is a dummy variable indicating whether individual *i* visits a restaurant on the weekend. T_i is a dummy variable taking a value of one for people who receive the descriptive norm treatment information and zero otherwise; $AboveTruth_i$ is a dummy variable indicating whether individual *i*'s pretreatment belief on Wednesday regarding the proportion of neighbors planning to go to restaurants is higher than the actual percentage; X_i describes the set of control variables. To calculate the main treatment effects on behavior ($Behavior_i$), the list control variables include the strata variables (neighborhood fixed effects and a dummy variable indicating whether the individual went to restaurants at least once a week before the COVID pandemic), whether the individual reported having plans to go to a restaurant in our Wednesday survey, the individual's education level, and calendar week fixed effects. The coefficients of interest, $\alpha(\alpha+\beta)$, capture the impact of the treatment on an individual's $PosteriorBelief_{it}$ and $Behavior_{it}$ for individuals with beliefs below (above) the truth. Our main estimation of treatment effects focuses on the participant's first week after he/she entered our experiment, to avoid the confounding impacts of the previous weeks' information treatment. For the two additional preregistered information treatments (i.e., restaurant certificate and expanded prosocial norm), similar regression models have been run. Results are reported in *SI Appendix*, sections B and C.

Moderator Analysis. To understand how the treatment effects were moderated by individual risk preferences, we decomposed individuals into two subsamples: risk averse and risk tolerant, with the median level as the splitting point. We then estimated the treatment effect of the descriptive norm by each subgroup respectively using Eq. **1**.

As a robustness check, we also conducted interaction analysis with both discrete (dummy variable indicating the individual is above median risk tolerance) and continuous definitions of risk preferences, using the following modified version of Eq. **1**:

$$Y_i = \alpha T_i + \beta T_i \times R_i + \gamma R_i + X_i \Gamma + E_i. \quad [2]$$

Here R_i denotes the value of the risk preference moderator for individual i. Our parameter of interest, β, denotes whether people who are more risk tolerant responded significant differently to our treatment. A more detailed description of our moderation analysis, together with extra robustness checks, is included in *SI Appendix*, section A.4. Results of moderator analysis for the full set of variables can be found in *SI Appendix*, Figs. A.9 and A.10.

Ethics Statement. This project has human subjects approval from the Committee on the Use of Humans as Experimental Subjects at Massachusetts Institute of Technology (Protocol 2002000100).

Data Availability. Anonymized survey data have been deposited in Harvard Dataverse (DOI: 10.7910/DVN/NCKGHA) (39).

ACKNOWLEDGMENTS. This work was supported by the "MIT China Future City Zhengzhou City Living Lab Program."

1. S. Hsiang *et al.*, The effect of large-scale anti-contagion policies on the coronavirus (COVID-19) pandemic. *Nature* **584**, 262–267 (2020).
2. S. Flaxman *et al.*; Imperial College COVID-19 Response Team, Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* **584**, 257–261 (2020).
3. World Bank, *Global Economic Prospects, June 2020* (World Bank, Washington, DC, 2020).
4. Bureau of Economic Analysis, "Gross domestic product, 2nd quarter 2020 (advance estimate) and annual update" (Bureau of Economic Analysis, 2020).
5. A. Goolsbee, C. Syverson, Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. *J. Public Econ.* **193**, 104311 (2021).
6. R. Chetty, J. N. Friedman, N. Hendren, M. Stepner, the Opportunity Insights Team, Real-time economics: A new platform to track the impacts of COVID-19 on people, businesses, and communities using private sector data. https://www.hks.harvard.edu/centers/mrcbg/programs/growthpolicy/real-time-economics-new-platform-trackimpacts-covid-19-people. Accessed 14 January 2022.
7. K. Dooling, The advisory committee on immunization practices' updated interim recommendation for allocation of COVID-19 vaccine—United States, December 2020. *Morb. Mortal. Wkly. Rep.* **69**, 1657–1660 (2020).
8. The Times of Israel, Israel aims to vaccinate 25% of population in a month; 250,000 had shots so far. *The Times of Israel*, 26 December 2020. https://www.timesofisrael.com/over-quarter-million-israelis-said-vaccinated-another-3994-virus-cases-found/. Accessed 5 January 2021.
9. M. Eddy, M. Santora, E.U. Starts Effort to Vaccinate 450 Million. *The New York Times*, 27 December 2020. https://www.nytimes.com/live/2020/12/27/world/covid-19-coronavirus-updates. Accessed 5 January 2021.
10. US Embassy, "Biden-Harris transition. The Biden-Harris plan to beat COVID-19" (2020). https://it.usembassy.gov/the-biden-harris-plan-to-beat-covid-19/. Accessed 22 January 2021.
11. T. Fetzer, L. Hensel, J. Hermle, C. Roth, Coronavirus perceptions and economic anxiety. *Rev. Econ. Stat.* **103**, 968–978 (2021).
12. J. V. Lazarus *et al.*, A global survey of potential acceptance of a COVID-19 vaccine. *Nat. Med.* **27**, 225–228 (2020).
13. L. Hensel *et al.*, Global behaviors, perceptions, and the emergence of social norms at the onset of the COVID-19 pandemic. *J. Econ. Behav. Organ*. **193**, 473–496 (2022).
14. A. Goolsbee, "Jobs day lesson 1 million: The virus is the boss." Twitter.com (2021). https://twitter.com/austan_goolsbee/status/1347548098978062336. Accessed 14 January 2022.
15. N. H. L. Leung *et al.*, Respiratory virus shedding in exhaled breath and efficacy of face masks. *Nat. Med.* **26**, 676–680 (2020).
16. H. W. Perkins, A. D. Berkowitz, Perceiving the community norms of alcohol use among students: Some research implications for campus alcohol education programming. *Int. J. Addict.* **21**, 961–976 (1986).
17. M. Hallsworth, J. A. List, R. D. Metcalfe, I. Vlaev, The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *J. Public Econ.* **148**, 14–31 (2017).
18. N. J. Goldstein, R. B. Cialdini, V. Griskevicius, A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *J. Consum. Res.* **35**, 472–482 (2008).
19. K. S. Wolske, K. T. Gillingham, P. W. Schultz, Peer influence on household energy behaviours. *Nat. Energy* **5**, 202–212 (2020).
20. R. C. Dempsey, J. McAlaney, B. M. Bewick, A critical appraisal of the social norms approach as an interventional strategy for health-related behavior and attitude change. *Front. Psychol.* **9**, 2180 (2018).
21. J. J. V. Bavel *et al.*, Using social and behavioural science to support COVID-19 pandemic response. *Nat. Hum. Behav.* **4**, 460–471 (2020).
22. M. Muthukrishna, T. J. H. Morgan, J. Henrich, The when and who of social learning and conformist transmission. *Evol. Hum. Behav.* **37**, 10–20 (2016).
23. J. Beshears, J. J. Choi, D. Laibson, B. C. Madrian, K. L. Milkman, The effect of providing peer information on retirement savings decisions. *J. Finance* **70**, 1161–1201 (2015).
24. C. A. Gravert, L. Olsson Collentine, When nudges aren't enough: Incentives and habit formation in public transport usage. SSRN [Preprint] (2019). http://dx.doi.org/10.2139/ssrn.3500699 (Accessed 14 January 2022).
25. A. S. Kristal, A. V. Whillans, What we can learn from five naturalistic field experiments that failed to shift commuter behaviour. *Nat. Hum. Behav.* **4**, 169–176 (2020).
26. M. A. Andor, A. Gerster, J. Peters, C. M. Schmidt, Social norms and energy conservation beyond the US. *J. Environ. Econ. Manage.* **103**, 102351 (2020).
27. P. W. Schultz, J. M. Nolan, R. B. Cialdini, N. J. Goldstein, V. Griskevicius, The constructive, destructive, and reconstructive power of social norms. *Psychol. Sci.* **18**, 429–434 (2007).
28. E. P. Fenichel, Economic considerations for social distancing and behavioral based policies during an epidemic. *J. Health Econ.* **32**, 440–451 (2013).
29. E. P. Fenichel *et al.*, Adaptive human behavior in epidemiological models. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 6306–6311 (2011).
30. S. G. Benzell, A. Collis, C. Nicolaides, Rationing social contact during the COVID-19 pandemic: Transmission risk and social benefits of US locations. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 14642–14644 (2020).
31. G. Kraft-Todd, E. Yoeli, S. Bhanot, D. Rand, Promoting cooperation in the field. *Curr. Opin. Behav. Sci.* **3**, 96–101 (2015).
32. D. G. Rand, E. Yoeli, M. Hoffman, Harnessing reciprocity to promote cooperation and the provisioning of public goods. *Policy Insights Behav. Brain Sci.* **1**, 263–269 (2014).
33. H. Kim, H. R. Markus, Deviance or uniqueness, harmony or conformity? A cultural analysis. *J. Pers. Soc. Psychol.* **77**, 785 (1999).
34. R. E. Nisbett, T. Masuda, Culture and point of view. *Proc. Natl. Acad. Sci. U.S.A.* **100**, 11163–11170 (2003).
35. M. J. Gelfand *et al.*, *The Relationship between Cultural Tightness–Looseness and COVID-19 Cases and Deaths: A Global Analysis* (The Lancet Planetary Health, 2021).
36. J. M. Jachimowicz, S. Chafik, S. Munrat, J. C. Prabhu, E. U. Weber, Community trust reduces myopic decisions of low-income individuals. *Proc. Nat. Acad. Sci. U.S.A.* **114**, 5401–5406 (2017).
37. A. Falk *et al.*, Global evidence on economic preferences. *Q. J. Econ.* **133**, 1645–1692 (2018).
38. D. Cantoni, D. Y. Yang, N. Yuchtman, Y. J. Zhang, Protests as strategic games: Experimental evidence from Hong Kong's antiauthoritarian movement. *Q. J. Econ.* **134**, 1021–1077 (2019).
39. J. Palacios, Data for "Encouraging the resumption of economic activity after COVID-19: Evidence from a large scale-field experiment in China." Harvard Dataverse. https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/NCKGHA. Deposited 5 January 2022.
40. J. A. List, A. M. Shaikh, Y. Xu, Multiple hypothesis testing in experimental economics. *Exp. Econ.* **22**, 773–793 (2019).