

God is in the Rain: The Impact of Rainfall-Induced Early Social Distancing on COVID-19 Outbreaks

Rolly Kapoor*
Haedong Aiden Rho*
Kinpritma Sangha†
Bhavyaa Sharma*
Ajay Shenoy‡
Guanghong Xu*

July 8, 2020

First Version: May 18, 2020

Abstract

We test whether earlier social distancing affects the progression of a local COVID-19 outbreak. We exploit county-level rainfall on the last weekend before statewide lockdown. After controlling for historical rainfall, temperature, and state fixed-effects, current rainfall is a plausibly exogenous instrument for social distancing. Early distancing causes a large reduction in cases and deaths that persists for weeks. The effect is driven by a reduction in the chance of a very large outbreak. Our estimates provide an empirical target for epidemiological models. Their size suggests early distancing may have sizable returns, and that random events early in an outbreak can have persistent effects on its course.

Keywords: COVID-19, coronavirus, social distancing, rainfall

JEL Codes: I12, I18, H75

*University of California, Santa Cruz

†Anlitiks, Inc.

‡University of California, Santa Cruz; Corresponding Author: email at azshenoy@ucsc.edu. Postal Address: Rm. E2-455, University of California, M/S Economics Department, 1156 High Street, Santa Cruz CA, 95064. We are grateful to SafeGraph for providing access to their data. We also appreciate helpful comments from Laura Giuliano, Justin Wolfers, and innumerable individuals on Twitter.

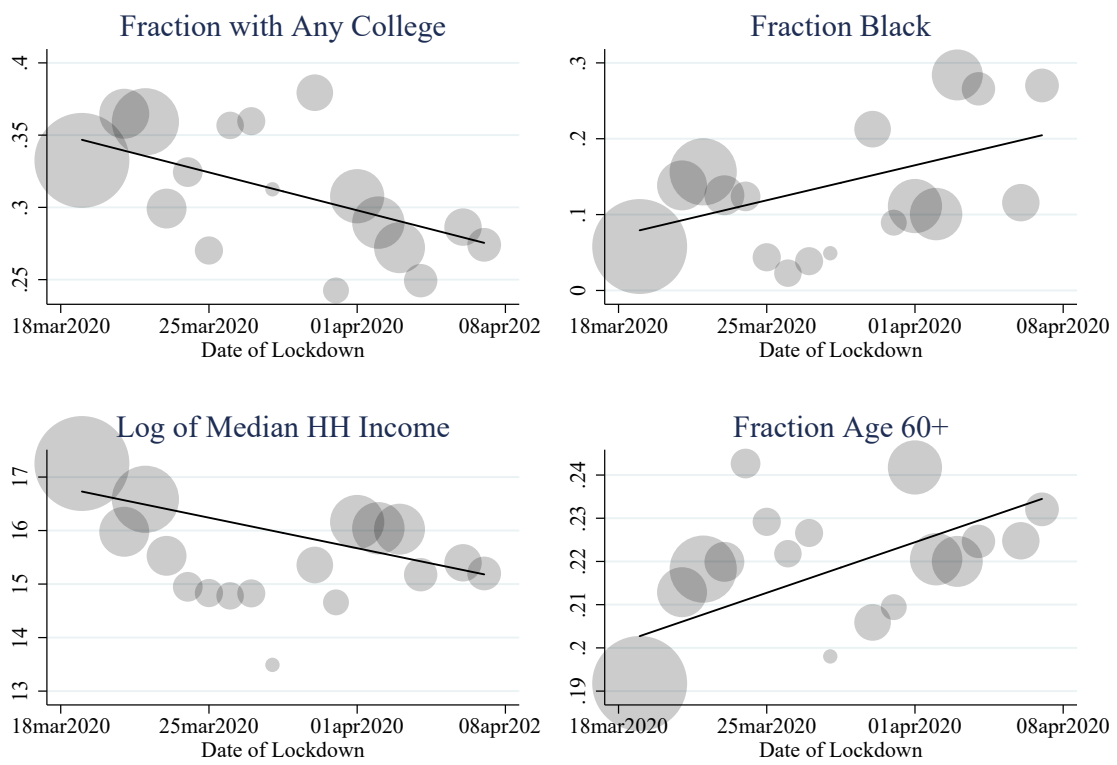
1 Introduction

As COVID-19 outbreaks have spiraled in some regions while sparing others, it remains an open question whether earlier social distancing played a key role. California's statewide lockdown began only 3 days earlier than New York's, yet California has suffered far fewer deaths. Six Bay Area counties within California (as well as Santa Cruz County) began their lockdown a few days earlier than the rest of the state. They have had a more favorable trajectory than other parts of the state despite suffering much earlier exposure to the disease.

But naïve comparisons between states risk conflating the impact of earlier distancing with differences in state characteristics. Figure 1 shows that states that issued earlier lockdowns have higher median incomes and more college degree holders, but fewer black and older residents. Even within a state, locales that issued earlier lockdowns may differ systematically in ways that may or may not be observable. For example, the Associated Press reports that the Bay Area lockdown had its roots in an association of local health officials that formed during the AIDS epidemic and has met regularly to discuss prior epidemics like Ebola and swine flu (Rodriguez, 19 April 2020). The presence of such an institution may have had other impacts on the local response to COVID-19 beyond the lockdown, making it difficult to isolate the effect of early social distancing. The problem of selection bias is compounded by the problem of measurement. It is possible that the states and counties that responded more quickly are also more active in testing for the disease, creating non-classical measurement error.

We sidestep these challenges by exploiting within-state variation in early social distancing induced by rainfall. We measure county-level rainfall on the last weekend before the county's home state went into mandatory lockdown. This key weekend is the last day that people had wide discretion in leaving home for reasons unrelated to work (dining at restaurants, for example). Focusing on this weekend creates a natural experiment for a marginally longer period of social distancing. After controlling for average historical rainfall, temperature, and state fixed effects, rainfall on this specific weekend is plausibly exogenous. Counties that had heavy rainfall were exogenously induced to exercise a marginal degree of extra social distancing just a few days before counties that had less rainfall. We measure whether these counties had fewer COVID-19 cases and deaths in the weeks after the statewide lockdown.

Figure 1
States that Lock Down Earlier are Systematically
Different on Baseline Characteristics



Note: The size of each circle is proportional to the number of states that shut down on that date. Demographics are from the 2014-2018 American Community Survey (5-year estimates). Dates of state-wide lockdown orders come from the Institute of Health Metrics and Evaluation. See Section 2.1 for details about the data.

We detect highly significant effects even two weeks after the statewide lockdown, many days after the crucial weekend. The two-stage least squares estimates imply that a 1 percentage point increase in the number of people leaving home causes an additional 14 cases and 1.3 deaths per 100,000 residents. These effects are all the more remarkable because the variation in social distancing induced by rainfall, though precise, is relatively small. But the impact of the initial reduction is propagated over time. We measure growing impacts that have not leveled off even 18 days after the lockdown, nearly 3 weeks after the crucial weekend. These effects appear to be driven by the right tail of the distribution. Counties where more people left home on the pre-shutdown weekend are no more likely to have a marginally higher case count, but are slightly more likely to have a big outbreak. This result is what might be expected given that differences in the number of infections on the eve of a statewide lockdown will either vanish or be drastically amplified depending on whether the county lowers the viral reproduction rate below 1 and avoids “superspreader” events.

Our paper joins a small but growing number of papers that study the impact of social distancing on COVID-19 transmission. Our research question is most similar to Pei et al. (2020), who use an epidemiological model to simulate COVID-19 trajectories in a counterfactual world where lockdowns had begun a few weeks sooner. Our study approaches this question using a natural experiment rather than a model. A few recent studies (Courtemanche et al., 2020; Fowler et al., 2020) use difference-in-differences designs to study the impact of statewide closures and lockdowns on transmission. Aside from exploiting an orthogonal source of variation, our study aims to answer a different question: whether marginal improvements in early distancing can affect medium-run outcomes.

Meanwhile, Brzezinski et al. (2020) use state-level rainfall and temperature as exogenous variation in non-mandated social distancing to study whether state governments are less likely to mandate social distancing where it is already being practiced.¹ Methodologically our study is most similar to Madestam et al. (2013), which measures the impact of rainfall on a single pivotal date (Tax Day 2010) to measure the long-run impacts of Tea Party protests. One major advantage to studying a one-time shock rather than panel variation is that we can fully trace differential tra-

¹ Since we exploit only within-state variation, their result is not a threat to our design. We verify in Section 3.4 and Appendix A.7 that county-level policy responses do not bias our results.

jectories across counties. And since that shock is on the weekend before statewide lockdown, it is the closest possible counterfactual to having a longer policy of social distancing.

Our results suggest that even small differences in the extent of early social distancing can have sizable impacts on the scale of the outbreak. As states begin to loosen their stay-at-home orders, health officials are considering plans to potentially return to lockdown if there are signs of a resurgence. Subject to caveats discussed in the final section, our results suggest moving even a few days more quickly could make a measurable difference. Our results also suggest that completely random events early in the course of a local outbreak can have surprisingly persistent effects on its size. Finally, the estimates from our atheoretic natural experiment may serve as targets for calibrating and validating epidemiological models.²

2 Research Design

2.1 Data

Weather : We measure rainfall by spatially merging weather stations from the Global Historical Climatology Network-Daily Database (Menne et al., 2012) to U.S. counties based on 2012 Census TIGER/Line shapefiles. We calculate county-level average precipitation and daily maximum temperatures. For each day in 2020 we calculate the average precipitation and max temperature for that same day-of-year from 2015—2019. We then take the inverse hyperbolic sine of all of these quantities.³ This transformation is a standard way to approximate a log transformation without having to discard zero-rainfall days. As long as rainfall itself is exogenous, the transformed quantity is also exogenous and, as we show in Figure 2 below, has a roughly linear relationship with our primary measure of social distancing. We apply the same transformation to temperature to maintain consistency. From here on we refer to these transformed quantities as simply current or historical rainfall and temperature.

Social Distancing: Our primary measure of social distancing is the percentage of people

² A modeler would do so by simulating our natural experiment and comparing the results against our estimates. We have posted the details needed for this exercise at <https://docs.google.com/spreadsheets/d/1RUHWFS85QmZl1oOuyA74tvuQREq2DEnWqNFOVZVVe7s/edit?usp=sharing>

³ The inverse hyperbolic sine transformation $\log(x + \sqrt{x^2 + 1})$ is a convenient approximation to the natural logarithm that is well-defined when $x = 0$ and converges to $\log 2 + \log x$ as $x \rightarrow \infty$.

that leave home, calculated using aggregated mobile phone GPS data provided by SafeGraph (SafeGraph, 2020a). The data report the total devices in SafeGraph’s sample by block group, and the number that leave their home.⁴ We aggregate these two counts by county and calculate the percentage leaving home.

Leaving home is our first-stage regressor because keeping people home is the primary impact of rain on social distancing, and keeping people at home for an extra weekend is the most natural analogy to locking down a few days sooner. But to better understand what activities people are deterred from doing when they stay home—and whether those who do leave change where they go—we draw on several other measures of social distancing. We use two measures of indoor exposure. The first is the Device Exposure Index (Couture et al., 2020a), which represents the number of people (cell phones) an average individual was exposed to in small commercial venues within the county. We also use SafeGraph’s Weekly Patterns data to compute a measure of “gatherings” based on whether more than 5 devices ping within a single indoor non-residential location within one hour (SafeGraph, 2020b). Since the SafeGraph sample represents roughly 6% of a typical county, 5 devices represent a large number of people. We rescale both measures by their daily average on the first full weekend in March, meaning a value of 100 denotes the same exposure or number of gatherings as the first weekend of March (which was before any local or state lockdown).

We also use several measures of long-distance travel. Using SafeGraph’s data we measure the percentage of devices that travel greater or less than 16 kilometers from home (among those that leave home). We also measure cross-county travel using the Location Exposure Index (Couture et al., 2020b). We measure the fraction of people in a county who were not present on any of the prior 14 days.⁵

COVID-19 Cases and Deaths : We measure daily (cumulative) COVID-19 cases/deaths by combining data from Johns Hopkins University and the CoronaDataScraper project (Center for Systems Science and Engineering (Johns Hopkins University); Corona Data Scraper (2020)). As described in detail in Appendix B, we manually corrected missing values by consulting county

⁴ SafeGraph defines “home” as the “common nighttime location of each mobile device over a 6 week period to a Geohash-7 granularity (153m x 153m).” Leaving home is defined as leaving that square.

⁵ For more information on the Device Exposure Index and the Location Exposure Index see Appendix B.1.

public health departments and local newspapers. All of these measures are cumulative cases and deaths rather than new cases and deaths. Our primary outcomes are the number of cases and deaths per 100,000 population, measured 14 days after the statewide lockdown.⁶

Demographics : We measure demographic characteristics (such as population size, median income, age profiles of the population) using the 2014-2018 five-year estimates from the American Community Survey (Manson et al., 2019).

Lockdowns: Finally, we measure statewide lockdown dates using the Institute of Health Metrics and Evaluation’s record of state policies as of 17 April 2020 (Institute for Health Metrics and Evaluation (2020)). The dataset has all shutdown dates up to 7 April. Any state that had not shut down by that date (or was not recorded as doing so by the Institute) is excluded from our study.

2.2 Instrument and Specifications

Why the Last Weekend?: Our ideal experiment would be to randomly assign some counties to begin social distancing sooner than others. Since such an experiment is not feasible, our natural experiment focuses on rainfall-induced social distancing on the weekend just prior to the statewide lockdown. People in rainy counties began a marginal degree of additional social distancing a few days sooner than other counties.⁷

Defining the Instrument: We identify the last Saturday and Sunday before the day of the shutdown order. If the shutdown was announced on a Sunday we take only the Saturday of that weekend as the “weekend before.” If it is announced on a Saturday we take the prior weekend. We average rainfall and temperature (both current and historical) as well as social distancing across the days of this weekend. We compute baseline cases and deaths as those recorded for the day before this last weekend, and baseline growth in these measures as the average change in the inverse hyperbolic sine of each in the prior 7 days.

Specification, Identification, and Inference: We estimate first-stage, reduced form, and

⁶ We choose these measures both because they are the measures most commonly used by policymakers to gauge the severity of an outbreak, and because they give the most accurate reflection of the number of infections relative to the number who could potentially be infected. We choose 14 days as our default horizon because this is the typical quarantine period for the disease, though Section 3.2 shows the impact at every horizon.

⁷ The gap between the last weekend (as defined above) and the the shutdown is 3 days for the median county.

second-stage regressions of the form

$$D_i = \alpha_s + \gamma R_i + \tau_1 \bar{R}_i + \tau_2 T_i + \tau_3 \bar{T}_i + X_i \omega + u_i \quad (1)$$

$$Y_i = \zeta_s + \rho R_i + \xi_1 \bar{R}_i + \xi_2 T_i + \xi_3 \bar{T}_i + X_i \theta + v_i \quad (2)$$

$$Y_i = \kappa_s + \beta \hat{D}_i + \phi_1 \bar{R}_i + \phi_2 T_i + \phi_3 \bar{T}_i + X_i \vartheta + z_i \quad (3)$$

where i and s index counties and states, D is the percentage of people leaving home, Y is the outcome, α_s and κ_s are state fixed-effects, R and \bar{R} are current and historical rainfall, T and \bar{T} are current and historical temperature, and X is a vector of baseline and demographic control variables that vary across specifications, with the most basic specification having no controls.

We must control for historical rainfall because even within a state, counties that are typically rainy in March and April may be systematically different from those that are not (e.g. Santa Cruz versus San Diego in California). The instrument R_i is thus excess or unexpected rainfall, which is plausibly uncorrelated with historical demographic characteristics. We control for temperature because some experts and politicians have hypothesized that it may directly impact COVID-19 transmission.⁸ The identification assumption is that, after controlling for state fixed-effects, historical rainfall, and temperature, rainfall on the pre-shutdown weekend only affects endline case counts through its impact on the number of people leaving home. We show in Appendix A.1 that, as expected, rainfall is uncorrelated with baseline cases, deaths, and a host of demographic characteristics.

In all of these regressions β is the two-stage least squares estimate of the impact on the outcome of having 1 percentage point more people leave home on the weekend before the lockdown.⁹ Since there is spatial correlation in both rainfall and COVID-19 infections, we cluster standard errors using a 3°x 3° latitude-longitude grid.¹⁰

Additional Control Variables: Since rainfall is exogenous, the control variables X_i will not affect the consistency of the estimates. But they can make the estimates more precise by reducing the unexplained variation in social distancing and COVID-19 cases and deaths. Our basic specification includes nothing in X_i . Our preferred specification adds controls for baseline COVID-19

⁸ Chin et al. (2020), for example, find that temperature affects virus stability in lab samples.

⁹ Since there is a single endogenous regressor and a single excluded instrument, $\hat{\beta} = \hat{\rho}/\hat{\gamma}$.

¹⁰ To be precise, we generate a grid and assign each county to the cell that contains its centroid.

prevalence. We include the number of cases per 100,000 at baseline, the raw number of cases at baseline, and the growth rate of cases in the week prior to the pre-lockdown weekend.¹¹ Our most comprehensive specification includes baseline controls as well as demographic characteristics.¹²

3 Results

3.1 Basic Estimates

First-Stage—Impact of Rainfall on Social Distancing: Column 1 in Panel A of Table 1 shows estimates of the first-stage (Equation 1). After controlling for historical rainfall and temperature, a one-unit increase in our measure of rainfall causes a 0.4 percentage point decrease in the number of people who leave home. The F-statistic is 11.68, well above conventional measures of instrument strength.

Columns 2—6 explore what activities become less prevalent because of rainfall and because people are staying home. One concern might be that although some people stay home because of the rain, those who do leave will pack into bars and restaurants instead of visiting the outdoors. Column 2 shows that the average exposure, based on how many people visit small indoor venues, declines by 0.87 percentage points relative to its level the first weekend of March (prior to any lockdown). Column 3 shows that our measure of large gatherings declines by 1.7 percentage points relative to early March.

Is the impact of rainfall on the prevalence of COVID-19 driven more by reducing local transmission, or by reducing the spread of the virus over long distances and across counties? Columns 4 and 5 measure the impact on the percentage of people leaving home and traveling a short or long distance (based on whether they traveled more than 16 kilometers from home). The esti-

¹¹ We control for both cases per 100,000 and raw case counts at baseline because both are independently informative about social distancing and endline outcomes. That is likely because while the one measures the baseline rate of prevalence, the other drives initial local media coverage. It is also likely that a greater raw number of cases lowers the probability that the infection dies out because all initially infected self-isolate. The case growth rate, which we calculate as the average change in the inverse hyperbolic sine of case counts, is informative about the trajectory prior to the pre-shutdown weekend.

¹² Total population; fraction of population in the bins 60-69, 70-79, and over 80; fraction African American; and median household income.

Table 1
Two-Stage Least Squares Estimates

Panel A: Interpreting the First-Stage

	First-Stage		Activities Averted by Staying Home			
	(1)	(2)	(3)	(4)	(5)	(6)
	% Leaving Home	Exposure	Gatherings	Travel Near	Travel Far	Non-Locals
Rainfall	-0.432*** (0.126)	-0.876*** (0.327)	-1.670** (0.694)	-0.217* (0.131)	-0.336** (0.146)	-0.267** (0.122)
Counties	1946	1397	1757	1946	1946	1397
Clusters	139	113	124	139	139	113
Outcome Mean	64.77	37.35	38.34	41.41	21.42	9.13
F-stat: Rainfall	11.68	7.18	5.79	2.75	5.27	4.77
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X

Panel B: Reduced-Form


	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-6.776** (3.160)	-6.132*** (1.705)	-5.921*** (1.670)	-0.717 (0.463)	-0.581*** (0.216)	-0.537*** (0.176)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
Outcome Mean	58.12	58.12	58.12	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Panel C: Two-Stage Least Squares

	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
% Leaving Home	15.686 (9.653)	14.596*** (4.852)	14.824*** (5.130)	1.660 (1.274)	1.383** (0.556)	1.344** (0.517)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	11.68	16.54	17.80	11.68	16.54	17.80
Outcome Mean	58.12	58.12	58.12	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: All standard errors are clustered using a 3°x 3° latitude-longitude grid to adjust for spatial correlation.
Panel A: “Exposure” refers to the Device Exposure Index, a measure of the number of devices (cell phones) visiting small indoor venues. “Gatherings” measures the number of times more than 5 devices ping in a single indoor venue within the span of an hour. Both of these measures are rescaled as a percentage of their level on the weekend 7–8 March. “Travel Near” and “Travel Far” give the percentage of devices that leave home and travel less than versus more than 16 kilometers. “Non-Locals” gives the percentage of devices in the county that were not present on any of the prior 14 days.
Panels B and C: “Baseline Case Controls” are the number of COVID-19 cases the day before the pre-shutdown weekend (both the raw count and the number per 100,000), and the average growth (change in the inverse hyperbolic sine) of cases in the week preceding the last weekend. “Demographic Controls” are total population; fraction of population in the bins 60-69, 70-79, and over 80; fraction African American; and median household income.
 *p=0.10 **p=0.05 ***p=0.01

Figure 2
First-Stage and Reduced Form



Note: Each panel shows a partial correlation plot of rainfall on the weekend before the statewide lockdown against either the percentage of people leaving home on that weekend (left-hand panel) or total cases per 100,000 as of 14 days after the lockdown. We calculate residuals from a regression of both X and Y variable on state fixed-effects, historical rainfall, current and historical temperature, and baseline case controls. We define bins based on residualized rainfall. Each dot shows the average residualized outcome within the bin, and the line shows the linear prediction. The histogram shows the number of observations that fall into each bin.

mates suggest a larger impact on long distance travel (especially compared to the mean). Column 6 shows that a one-unit increase in rainfall causes a 0.27 percentage point decrease in the fraction of people in the county who had not been there in the previous two weeks, suggesting a sizable decline in cross-county travel.

Reduced-Form and Two-Stage Least Squares: Panel B of Table 1 shows estimates of the reduced-form impact of rainfall on COVID-19 cases and deaths per 100,000 at endline, which these regressions define as 14 days after the statewide lockdown. Column 1 shows that a 1 unit increase in rainfall on the weekend before lockdown lowers the number of cases at endline by 6.7 per 100,000. Columns 2 and 3 show that controlling for baseline prevalence and demographics tightens the standard errors without substantially changing the estimates. Columns 4–6 imply that the reduction in cases translates to a reduction in deaths, as well. A 1 unit increase in rainfall causes a 0.5 to 0.7 per 100,000 reduction in the death rate.

Figure 2 shows a partial correlation plot of the first-stage and reduced form of the regression in Column 2 (which includes baseline case controls). The plot illustrates how rainfall on the last weekend before the state-wide lockdown lowers both the percentage of people leaving home (left-hand panel) and the number of cases at endline (right-hand panel). The plot shows that our estimates are not driven by outliers, and that both relationships are approximately linear.

Under the assumption that rainfall only affects disease transmission through its impact on early social distancing, the two-stage least squares estimate—the ratio of the reduced-form and first-stage coefficients—gives the causal impact of early social distancing on COVID-19 cases and deaths. Panel C of Table 1 presents these estimates. All three specifications have a strong first-stage, with the F-statistic on the excluded instrument (weekend rainfall) varying from 11 to 18. The basic specification, which has no controls, is relatively noisy and statistically insignificant.

But after controlling for baseline case controls the standard errors become tight enough to make the estimates highly significant (Columns 2 and 5). The final specification additionally controls for county demographics, which makes little difference in size or significance of the estimates (Columns 3 and 6). Indeed, all three specifications produce near-identical estimates. A 1 percentage point increase in the number of people leaving home on the weekend before the shutdown causes an additional 13 cases and 1.3 deaths per 100,000.

The size of these estimates relative to the mean of the outcome may seem surprising. As we discuss in Section 3.3, the average impact represented by these estimates is misleading because it is in large part driven by changes in the probability of large outbreaks. Clearly it is not the case that a 4 percentage point decrease in people leaving home would have eradicated the disease across the country. It is more accurate to say that it would marginally reduce the probability of a catastrophic outbreak.¹³ Finally, it is not clear whether the proper benchmark is the population that is infected at endline or the population that is susceptible to infection. In the latter case the reference group is roughly the entire population, in which case our results imply that an additional 1 percentage point of people leaving home causes an additional 0.013 percentage

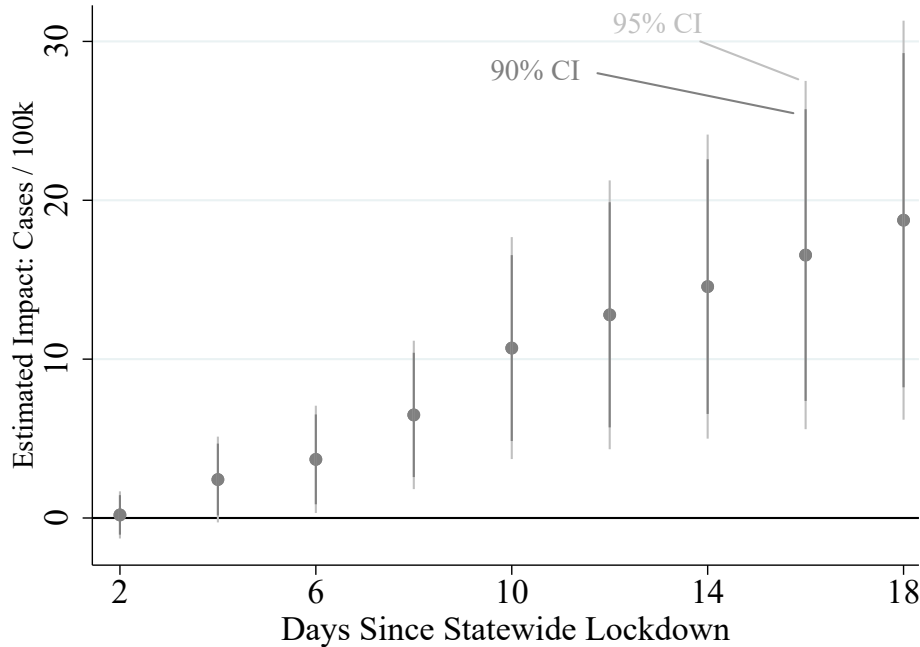
¹³ Another possibility is that the types of activities deterred by rainfall are the riskiest—for example, visits by family members to skilled nursing facilities. It is also possible that there are substantial spillovers across counties, and that each person staying home actually reduces the risk of spreading cases across several counties.

points of the population to become infected.

3.2 Comparative Dynamics in Counties with Less Early Social Distancing

Figure 3

The Excess Case Count in Counties with Less Early Distancing Continues to Increase Even 18 Days after Lockdown



Note: Using total cases per 100,000 at each horizon $h = 2, 4, 6, \dots, 18$ we estimate the two-stage least squares coefficient controlling for baseline case controls (analogous to Column 2 of Panel C, Table 1). Each coefficient is from a separate regression (and the regression at $h = 14$ is identical to that reported in Table 1).

Table 2 gives a relatively limited picture of the trajectory of cases because all outcomes are measured at the fixed horizon of 14 days after the statewide lockdown. One advantage of our research design is that we can estimate the comparative dynamics of case rates between counties that quasi-randomly practiced different levels of early social distancing. Using the same specification as Column 2 of Table 1, we estimate the impact on cases per 100,000 2 days after the lockdown, 4 days after, and so on for every horizon $h = 2, 4, 6, \dots, 18$. Figure 3 plots each coefficient against h . The estimated impact appears to increase linearly over time with no sign of leveling off within the horizon available to us.¹⁴ The figure suggests the impact of a one-time

¹⁴ At longer horizons we would start to lose states because our case count data ends 18 days after the last state in


difference in early social distancing is surprisingly long-lived.

We find no evidence, however, that the growth *rate* of cases increases because of more people leaving home on the last weekend (see Appendix A.4). That is not surprising because the natural experiment induces some counties to begin early social distancing just before all counties go uniformly into lockdown. The effect is analogous to quasi-randomly inducing some counties to begin lockdown with a larger infected population. As long as this difference in initial population does not affect how carefully the lockdown is observed, it will rescale the case count without affecting the transmission rate.¹⁵

3.3 Distributional Impact: Early Social Distancing Lowers the Chance of Right-Tail Outcomes

Figure 4

Counties with Less Social Distancing are More Likely to Have Very Large (Right-Tail) Outbreaks



Note: We estimate the impact across the distribution of outcomes. Each point and confidence interval is the two-stage least squares estimate of the impact of early social distancing on the probability of having endline cases per 100,000 greater than the percentile or absolute number indicated on the horizontal axis. Each estimate controls for baseline case rate, count, and growth (analogous to Column 2 of Panel C, Table 1).

our sample to go on lockdown.

¹⁵ If endline case count is $Y_T = Y_0 \exp(gT)$, our natural experiment is analogous to increasing Y_0 .

Given the nature of exponential growth, local COVID-19 outbreaks may quickly die down or rapidly spiral out of control. That feature of transmission dynamics suggests early social distancing may have extended rather than shifting the distribution. We test for the impact on the full distribution by defining dummies for whether the endline number of cases per 100,000 is greater than each decile of the distribution. We estimate Equation 3 using these dummies as the outcomes (using the specification with baseline case controls). This procedure is analogous to testing how the inverse cumulative distribution function is shifted by a 1 percentage point reduction in early social distancing.

The left-hand panel of Figure 4 plots the estimates with their 90 and 95 percent confidence intervals. The figure suggests that although the estimated impact becomes positive around 0.4 (meaning less early distancing increases the probability of being above the 40th percentile), the effect only becomes significant at 0.7. That suggests early distancing is lowering the probability of a right-tail outbreak. The most precise estimate is the last. A 1 percentage point increase in the number of people leaving home on the weekend before lockdown causes a 2 percentage point increase in the probability of an outbreak that puts the county in the top 10 percent of the distribution. The right-hand panel clarifies just how large these right-tail events are. This panel is analogous to the first one, but it defines dummies based on having an endline case rate above some absolute cutoff. The size and significance peaks at 100 cases per 100,000, a very large case count.

The results suggest early social distancing worked less by causing a moderate reduction in cases than by reducing the chance of a big outbreak. This result may be consistent with several recent studies that find that COVID-19 has a very low dispersion factor, meaning small groups of “superspreaders” are responsible for the vast majority of cases (Kupferschmidt, 19 May 2020). Endo et al. (2020) estimate using a mathematical model that as few as 10% of initially infected people may be responsible for as much as 80% of subsequent cases. Miller et al. (2020) find a similar result when they use genome sequencing to trace the virus’s spread across Israel. If early social distancing marginally reduces the probability a superspreader begins a transmission chain, it could explain why our estimates are driven by changes in the number of large outbreaks. Regardless of the cause, our estimates imply that most counties that began distancing sooner had little benefit, but those that did benefit did so tremendously.

3.4 Robustness and Threats to Validity

In the appendix we run several other tests:

Balance: One concern is that rainfall, even after controlling for state fixed effects, historical rainfall, and current and historical temperature, is not truly exogenous. We show in Appendix A.1 that rainfall is uncorrelated with baseline measures of COVID-19 prevalence and county demographic characteristics.

Heterogeneity: We show in Appendix A.2 that there is little evidence of heterogeneous impacts by baseline case levels, baseline case growth, the time between the last weekend and the start of the statewide lockdown, and a host of demographic characteristics. This seems largely a consequence of not having enough data to generate a strong first stage when splitting the sample or identifying an interaction as well as a direct effect. There is some slight evidence that early social distancing has less of an impact in counties with an older population, though the mechanism for that result is uncertain.

Outliers: Given that Section 3.3 shows the effect comes largely from changes in the likelihood of right-tail events, one may worry that the entire estimate is driven by a few outliers. Appendix A.3 shows that Winsorizing the very largest outcomes still yields significant effects. Although the top of the distribution does drive the result, it is a genuine distributional impact rather than a handful of fluke outliers.

Other Outcomes: Though endline cases and deaths per 100,000 is the most logical outcome (see Section 2.1), we show in Appendix A.5 the results are qualitatively similar if we instead use raw counts and the log of endline cases and deaths per 100,000.¹⁶

Measurement Error in COVID-19 Prevalence: One inevitable challenge to any study of COVID-19 is that the true number of cases far exceeds reported cases. One strength of our design is that rainfall is unlikely to be correlated with local testing capacity, making it unlikely that our result is spuriously driven by non-classical measurement error. However, we cannot rule out that counties with larger outbreaks are more aggressive in testing. Then *any* variation that reduces COVID-19 cases rates, be it rainfall or a hypothetical randomized controlled trial, would find

¹⁶ To be precise, we estimate a Poisson Maximum Likelihood estimator using Equation 2 as the link function. Unlike simply taking the log, the Poisson estimator is consistent even though endline cases and deaths equal zero in many counties (Silva and Tenreiro, 2006).

accentuated impacts. We acknowledge that this caveat applies to our study as it does to any other.

Local Policy Response: One concern is that even if rainfall is exogenous, local governments might respond to either social distancing or (more likely) rising numbers of cases by instituting their own emergency orders or lockdowns. Our estimates might reflect not just the initial shock to social distancing but the policy response triggered by that shock. Although such a response is possible, it is likely to be a countervailing response. Local officials would likely loosen restrictions wherever case counts are low and vice-versa.¹⁷ That would, if anything, bias our estimates towards zero. Nevertheless we show in Appendix A.7 that controlling for a dummy for whether the county has any policy restriction by the end of the 14 day horizon of our regressions does not change the results.

Direct Impact of Weather: Some news reports and health experts have observed that warmer countries (e.g. Singapore and South Korea) have been more successful in controlling outbreaks than more temperate ones (e.g. the U.S. and Western Europe). That has led to a theory that temperature may directly affect virus transmission (e.g. Sajadi et al., 2020). If the weather directly affects transmission it could violate the single-channel assumption needed for a valid instrument.

We find no evidence for a link between transmission and temperature on the last weekend in our county-level results. Regardless, all of our specifications control for temperature, making it unlikely to be driving our results. Some reports have also suggested humidity may separately affect transmission.¹⁸ Though the evidence for this is limited, we test for whether humidity is driving the results. If the impact of rainfall on cases and deaths were through its correlation with humidity rather than its impact on social distancing, we would expect that the reduced-form impact of rainfall on cases and deaths would vanish after controlling for humidity. But we show in Appendix A.6 that the reduced-form coefficient is essentially unchanged.¹⁹ Other links are possible but not yet well substantiated. It is possible that sunlight, through ultraviolet radiation,

¹⁷ Brzezinski et al. (2020) find that states where people are already social distancing of their own accord are less likely to impose a lockdown.

¹⁸ Luo et al. (2020) is one example, though they actually find that humidity predicts *lower* transmission.

¹⁹ Since we only have humidity data for 60% of the sample, controlling for it directly in all specifications (as we do with temperature) would be too costly for precision.

reduces virus spread. If that is true it would bias our estimates towards rainfall *increasing* the number of COVID-19 cases.

That said, we cannot categorically rule out that rainfall has some unanticipated impact or interaction with the environment. Given what is currently known about the virus and the nature of our own results, we believe these effects to be second-order compared to the direct impact on human behavior.

4 Directions for Future Research

Our results suggest that a marginal increase in social distancing a few days before a statewide lockdown has persistent effects two to three weeks later. One interpretation is that policy makers wishing to (re)institute a lockdown would reap surprisingly large gains from moving more quickly.

Our results come with a few caveats. First, as noted above we cannot categorically rule out that rainfall directly affects COVID-19 transmission through some as-yet unknown mechanism. Second, the type of social distancing induced by rainfall may differ from that induced by a government order.²⁰ Finally, the context of our natural experiment—the weekend before a statewide lockdown—was one in which many people were already voluntarily social distancing. Policy-makers may face a different context when deciding on whether to begin a future lockdown. We leave disentangling these mediating factors to future research.

²⁰ For example, state and county lockdowns have sparked protests and political opposition, while a rainy weekend presumably would not.

References

- Bell, Michael**, “How do I calculate dew point when I know the temperature and the relative humidity?,” Accessed 17 May 2020. <https://iridl.ldeo.columbia.edu/dochelp/QA/Basic/dewpoint.html>.
- Brzezinski, Adam, Guido Deiana, Valentin Kecht, and David Van Dijcke**, “COVID Economics,” *Covid Economics*, 2020, 7, 115.
- Center for Systems Science and Engineering (Johns Hopkins University)**, “COVID-19 Data Repository,” 2020.
- Chin, Alex W H, Julie T S Chu, Mahen R A Perera, Kenrie P Y Hui, Hui-Ling Yen, Michael C W Chan, Malik Peiris, and Leo L M Poon**, “Stability of SARS-CoV-2 in different environmental conditions,” *The Lancet Microbe*, 2020, 1 (1), e10.
- Corona Data Scraper**, “Corona Data Scraper,” 2020.
- Courtemanche, Charles, Joseph Garuccio, Anh Le, Joshua Pinkston, and Aaron Yelowitz**, “Strong Social Distancing Measures In The United States Reduced The COVID-19 Growth Rate,” *Health Affairs*, 2020.
- Couture, Victor, Jonathan Dingel, Allison Green, Jessie Handbury, and Kevin Williams**, “Device Exposure Indices,” 2020. <https://github.com/COVIDExposureIndices/COVIDExposureIndices>.
- , —, —, —, and —, “Location exposure indices,” 2020. <https://github.com/COVIDExposureIndices/COVIDExposureIndices>.
- Endo, Akira, Sam Abbott, Adam J Kucharski, Sebastian Funk et al.**, “Estimating the overdispersion in COVID-19 transmission using outbreak sizes outside China,” *Wellcome Open Research*, 2020, 5 (67), 67.
- Fowler, James H., Seth J. Hill, Nick Obradovich, and Remy Levin**, “The Effect of Stay-at-Home Orders on COVID-19 Cases and Fatalities in the United States,” *medRxiv*, 2020.
- Institute for Health Metrics and Evaluation**, “COVID-19 Projections,” 2020. <https://covid19.healthdata.org/united-states-of-america> . Accessed 17 April 2020.
- Kupferschmidt, Kai**, “Why Do Some COVID-19 Patients Infect Many Others, Whereas Most Don’t Spread the Virus at All?,” *Science*, 19 May 2020.

- Luo, Wei, Maimuna S Majumder, Dianbo Liu, Canelle Poirier, Kenneth D Mandl, Marc Lipsitch, and Mauricio Santillana**, “The Role of Absolute Humidity on Transmission Rates of the COVID-19 Outbreak,” *medRxiv*, 2020.
- Madestam, Andreas, Daniel Shoag, Stan Veuger, and David Yanagizawa-Drott**, “Do Political Protests Matter? Evidence from the Tea Party Movement*,” *The Quarterly Journal of Economics*, 09 2013, *128* (4), 1633–1685.
- Manson, Steven, Jonathan Schroeder, David Van Riper, and Steven Ruggles**, “IPUMS National Historical Geographic Information System: Version 14.0 [Database].,” 2019. Minneapolis, MN: IPUMS. <http://doi.org/10.18128/D050.V14.0>.
- Menne, Matthew J, Imke Durre, Russell S Vose, Byron E Gleason, and Tamara G Houston**, “An Overview of the Global Historical Climatology Network-Daily Database,” *Journal of Atmospheric and Oceanic Technology*, 2012, *29* (7), 897–910.
- Miller, Danielle, Michael A Martin, Noam Harel, Talia Kustin, Omer Tirosh, Moran Meir, Nadav Sorek, Shiraz Gefen-Halevi, Sharon Amit, Olesya Vorontsov, Dana Wolf, Avi Peretz, Yonat Shemer-Avni, Diana Roif-Kaminsky, Na’ama Kopelman, Amit Huppert, Katia Koelle, and Adi Stern**, “Full Genome Viral Sequences Inform Patterns of SARS-CoV-2 Spread Into and Within Israel,” *medRxiv*, 2020.
- Pei, Sen, Sasikiran Kandula, and Jeffrey Shaman**, “Differential Effects of Intervention Timing on COVID-19 Spread in the United States,” *medRxiv*, 2020.
- Rodriguez, Olga R.**, “Fast Decisions in Bay Area Helped Slow Virus Spread,” *Associated Press*, 19 April 2020. <https://apnews.com/10c4e38a0d2241daf29a6cd69d8d7b43>.
- SafeGraph**, “SafeGraph Social Distancing Metrics, Version 2,” 2020. <https://docs.safegraph.com/docs/social-distancing-metrics>.
- , “SafeGraph Weekly Patterns Metrics, Version 1,” 2020. <https://docs.safegraph.com/docs/weekly-patterns>.
- Sajadi, Mohammad M, Parham Habibzadeh, Augustin Vintzileos, Shervin Shokouhi, Fernando Miralles-Wilhelm, and Anthony Amoroso**, “Temperature and Latitude Analysis to Predict Potential Spread and Seasonality for COVID-19,” *Preprint, Available at SSRN 3550308*, 2020.

Silva, JMC Santos and Silvana Tenreyro, “The Log of Gravity,” *The Review of Economics and Statistics*, 2006, 88 (4), 641–658.

The National Association of Counties, “County Explorer,” Accessed 22 May 2020.
<https://ce.naco.org/?dset=COVID-19&ind=State%20Declaration%20Types>.

A Empirical Appendix

A.1 Balance Tests

Table 2
First Stage and Balance

Panel A				
	(1)	(2)	(3)	(4)
	% Leaving Home	Baseline Cases	Baseline Cases/100k	Baseline Case Growth
Rainfall	-0.432*** (0.126)	-3.135 (4.377)	-0.054 (0.324)	0.005 (0.004)
Counties	1946	1946	1946	1946
Clusters	139	139	139	139
F-stat: Rainfall	11.68	0.51	0.03	1.44
State FEs	X	X	X	X
Avg. Rain	X	X	X	X
Temperature	X	X	X	X
Baseline Case Controls				
Demographic Controls				

Panel B				
	(1)	(2)	(3)	(4)
	Baseline Deaths	Baseline Deaths/100k	Baseline Death Growth	Population
Rainfall	-0.174 (0.172)	-0.033 (0.032)	0.005 (0.004)	4258.716 (11014.234)
Counties	1946	1946	1946	1946
Clusters	139	139	139	139
F-stat: Rainfall	1.03	1.04	1.44	0.15
State FEs	X	X	X	X
Avg. Rain	X	X	X	X
Temperature	X	X	X	X
Baseline Case Controls				
Demographic Controls				

Panel C					
	(1)	(2)	(3)	(4)	(5)
	Median HH Income	Fraction 60-69	Fraction 70-79	Fraction over 80	Fraction Black
Rainfall	3148.256 (8802.852)	0.001* (0.001)	0.001 (0.001)	-0.000 (0.000)	-0.001 (0.002)
Counties	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139
F-stat: Rainfall	0.13	2.79	1.10	0.01	0.22
State FEs	X	X	X	X	X
Avg. Rain	X	X	X	X	X
Temperature	X	X	X	X	X
Baseline Case Controls					
Demographic Controls					

Note: We estimate Equation 1 using the basic specification on each outcome. Standard errors are clustered as in Table 1.
*p=0.10 **p=0.05 ***p=0.01

A.2 Heterogeneity

Table 3
Heterogeneity By Interaction Terms

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline Cases	Baseline Case Growth	Days Until Lockdown	Fraction over 80	Fraction Black	Median HH Income
Main Effect	16.482** (8.300)	16.981** (6.890)	8.614 (17.789)	15.350*** (5.102)	11.772** (5.133)	18.513* (9.741)
Interaction	-0.373 (0.812)	-19.173 (27.065)	1.587 (4.471)	-16.891*** (5.432)	27.187 (22.896)	-0.000 (0.000)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
K-P Stat.	0.14	6.49	1.73	9.64	5.77	1.60
State FES	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls	X	X	X	X	X	X

Note: Each regression adds $C_i \times D_i$ to Equation 3, instrumenting for it with $C_i \times R_i$. The column header gives the variable used for C_i . The outcome in all regressions is endline cases per 100,000.
*p=0.10 **p=0.05 ***p=0.01

Table 4
Heterogeneity by Splitting the Sample

Panel A

	Baseline Cases		Baseline Case Growth		Days Until Lockdown	
	(1)	(2)	(3)	(4)	(5)	(6)
	Below	Above	Below	Above	Below	Above
% Leaving Home	7.824 (5.830)	22.356* (12.145)	8.257* (4.655)	32.642 (21.240)	16.916** (8.328)	13.614** (6.646)
Counties	998	948	1450	496	1055	891
Clusters	123	105	133	85	79	93
First-Stage F	3.97	8.54	6.37	6.24	8.37	6.23
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls	X	X	X	X	X	X

Panel B

	Fraction over 80		Fraction Black		Median HH Income	
	(1)	(2)	(3)	(4)	(5)	(6)
	Below	Above	Below	Above	Below	Above
% Leaving Home	24.445** (12.197)	7.941* (4.688)	6.141 (7.875)	18.321*** (6.835)	11.647 (9.222)	62.631 (44.183)
Counties	973	973	973	973	973	973
Clusters	119	112	118	89	118	107
First-Stage F	6.28	9.65	3.46	21.11	3.25	2.29
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls	X	X	X	X	X	X

Note: The sample is split based on whether a county is above or below the median value of the variable given in the header. “Days Until Lockdown” is the difference between the date of statewide lockdown and the first day of the final pre-shutdown weekend.

*p=0.10 **p=0.05 ***p=0.01

Table 5
Winsorized Outcomes

Panel A: Reduced-Form						
	Endline Cases/100k			Endline Deaths/100k		
	(1) .01	(2) .02	(3) .04	(4) .01	(5) .02	(6) .04
Rainfall	-4.088*** (1.143)	-3.027*** (0.891)	-2.199*** (0.712)	-0.160*** (0.059)	-0.130** (0.052)	-0.077** (0.038)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
Outcome Mean	54.18	51.93	48.91	1.66	1.61	1.45
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls						

Panel B: Two-Stage Least Squares						
	Endline Cases/100k			Endline Deaths/100k		
	(1) .01	(2) .02	(3) .04	(4) .01	(5) .02	(6) .04
% Leaving Home	9.730*** (3.407)	7.204*** (2.677)	5.234** (2.155)	0.381** (0.157)	0.310** (0.136)	0.184* (0.097)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	16.54	16.54	16.54	16.54	16.54	16.54
Outcome Mean	54.18	51.93	48.91	1.66	1.61	1.45
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls						

Note: Outcomes are Winsorized at the percentiles shown in the column header.
*p=0.10 **p=0.05 ***p=0.01

A.3 Winsorized Outcomes

Table 6
Growth Rates

Panel A: Reduced-Form

	Average Growth Rate in Cases			Average Growth Rate in Deaths		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-0.000 (0.002)	-0.000 (0.002)	0.000 (0.002)	-0.000 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
Outcome Mean	0.10	0.10	0.10	0.04	0.04	0.04
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Panel B: Two-Stage Least Squares

	Average Growth Rate in Cases			Average Growth Rate in Deaths		
	(1)	(2)	(3)	(4)	(5)	(6)
% Leaving Home	0.000 (0.004)	0.001 (0.004)	-0.000 (0.004)	0.000 (0.003)	0.001 (0.002)	0.001 (0.002)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	11.68	16.54	17.80	11.68	16.54	17.80
Outcome Mean	0.10	0.10	0.10	0.04	0.04	0.04
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: We calculate the growth rate as the average of the day-to-day change in the inverse hyperbolic sine of cases and deaths from the pre-shutdown weekend through 14 days after the statewide lock-down.

*p=0.10 **p=0.05 ***p=0.01

A.4 Growth Rates

Table 7
Alternative Outcome: Raw Endline Counts of Cases and Deaths

Panel A: Reduced-Form						
	Endline Cases			Endline Deaths		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-58.417 (53.157)	-31.696** (12.158)	-34.127*** (11.987)	-7.818 (7.393)	-4.886* (2.833)	-4.293* (2.175)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
Outcome Mean	164.71	164.71	164.71	7.21	7.21	7.21
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Panel B: Two-Stage Least Squares						
	Endline Cases			Endline Deaths		
	(1)	(2)	(3)	(4)	(5)	(6)
% Leaving Home	135.222 (141.248)	75.448** (33.577)	85.438** (35.271)	18.097 (19.153)	11.631* (6.924)	10.749* (5.850)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	11.68	16.54	17.80	11.68	16.54	17.80
Outcome Mean	164.71	164.71	164.71	7.21	7.21	7.21
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: The outcomes are endline cases and deaths without adjustment for county population.
*p=0.10 **p=0.05 ***p=0.01

A.5 Other Outcomes

Table 8
 Alternative Outcome: “Log” of Cases and Deaths per 100,000

	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-0.110*** (0.033)	-0.078*** (0.023)	-0.076*** (0.023)	-0.306*** (0.090)	-0.233*** (0.084)	-0.164*** (0.047)
Counties	1946	1946	1946	1942	1942	1942
Clusters	139	139	139	139	139	139
Outcome Mean	58.12	58.12	58.12	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: We estimate a Poisson Maximum Likelihood model that assumes the outcome equals the exponential of the specifications in the main text. This is in concept similar to regressing the log of the outcome on each specification, but the Poisson estimate is consistent even though the outcome equals zero for many counties. We are unable to estimate second-stage IV coefficients because the GMM estimator is unable to converge to estimates of so many state fixed-effects.

*p=0.10 **p=0.05 ***p=0.01

Table 9

The Impact of Rainfall on Cases/Deaths Is Unchanged When We Control for Humidity

	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-6.132*** (1.705)	-7.307*** (2.069)	-6.870*** (2.040)	-0.707** (0.353)	-0.707** (0.353)	-0.703* (0.373)
Rel. Humidity			-10.416 (17.458)			-0.113 (0.993)
Counties	1946	1131	1131	1131	1131	1131
Clusters	139	135	135	135	135	135
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls						
Sample	Full	Humidity	Humidity	Full	Humidity	Humidity

Note: The “Full” sample is the sample used in the main text. The “Humidity” sample is the subsample of counties for which we have data on dew point.

*p=0.10 **p=0.05 ***p=0.01

A.6 Humidity

We use data from the Global Surface Summary of Day. The dataset does not record humidity but does record dew point temperature. We calculate relative humidity using an approximation of the Clausius-Clapeyron equation (Bell, Accessed 17 May 2020).²¹

$$\begin{aligned}
 E &= E_0 \exp \left\{ \frac{L}{R_v} \left(\frac{1}{T_0} - \frac{1}{T_d} \right) \right\} \\
 E_s &= E_0 \exp \left\{ \frac{L}{R_v} \left(\frac{1}{T_0} - \frac{1}{T} \right) \right\} \\
 H_R &= 100\% \times \frac{E}{E_s} = 100 \exp \left\{ \frac{L}{R_v} \left(\frac{1}{T} - \frac{1}{T_d} \right) \right\}
 \end{aligned} \tag{4}$$

where the terms in (4) are

- H_R : relative humidity
- T : Temperature (in Kelvin)
- T_d : Dew Point Temperature (in Kelvin)
- $\frac{L}{R_v} = 5423K$

²¹ In a few cases the calculation gives a number greater than 100%, likely because a measurement error in the

We average dew point for all stations within a county and calculate the inverse hyperbolic sine of the dew point on the last weekend before statewide lockdown.

We estimate the reduced-form of our specification

$$Y_i = \kappa_s + \omega R_i + \phi_1 \bar{R}_i + \phi_2 T_i + \phi_3 \bar{T}_i + X_i \vartheta + v_i$$

which gives the direct impact of rainfall on the last weekend on cases and deaths. We see if the reduced-form coefficient $\hat{\omega}$ changes when we add dewpoint to the set of controls X_i . The specifications in Table 9 first show the reduced form coefficient for the entire sample. Since we only have humidity data for a subset of this sample, the next specification estimates the same reduced-form coefficient using the restricted sample. The final specification adds relative humidity. The reduced-form coefficient is essentially unchanged when we control for humidity.

Table 10
Controlling for Local Policy Response Does not Change the Results

Panel A: Reduced-Form						
	Endline Cases/100k			Endline Cases/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-6.253** (2.506)	-6.004*** (1.705)	-5.884*** (1.682)	-0.761 (0.502)	-0.660** (0.293)	-0.626** (0.252)
Counties	1904	1904	1904	1904	1904	1904
Clusters	134	134	134	134	134	134
Outcome Mean	57.11	57.11	57.11	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Panel B: Two-Stage Least Squares						
	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
% Leaving Home	15.686 (9.653)	14.596*** (4.852)	14.824*** (5.130)	1.660 (1.274)	1.383** (0.556)	1.344** (0.517)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	11.68	16.54	17.80	11.68	16.54	17.80
Outcome Mean	58.12	58.12	58.12	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: We define a dummy equal to 1 if the county has adopted some measure (emergency declaration, safer-at-home instruction, shutting down businesses) by the end of the horizon for our outcome, 14 days after the statewide lockdown. All regressions control for this dummy (in addition to the controls discussed in the main text).

*p=0.10 **p=0.05 ***p=0.01

A.7 Policy Response

Using dates on county-level policy responses from The National Association of Counties (Accessed 22 May 2020), we define a dummy for whether the county has put any social distancing measure (emergency declaration, safer-at-home instruction, shutting down businesses) before the date at which we measure the outcome (14 days after the statewide lockdown). Table 10 reports our reduced-form and two-stage least squares estimates after controlling for the policy response.

B Data Appendix

B.1 Measures of Social Distancing

Device-exposure index (DEX): The index is computed using cellular data from PlaceIQ. Daily exposure of a device is defined as the number of distinct devices that visit the commercial venues that the particular device visits that day. DEX is then calculated by averaging the exposure values for all devices in the sample in the geographical unit (e.g. county) on a particular day. The set of devices included in the calculation of DEX are those that pinged on at least 11 days over any 14-day period from November 1, 2019 through the date in question. The venues covered are mainly commercial venues (with the largest category being restaurants). The set of venues is restricted to those “small enough such that visiting devices are indeed exposed to each other.” The set excludes Nature and Outdoor, Theme Parks, Airports, Universities, as well as any location whose category is unidentified by PlaceIQ.

Location-exposure index: The LEX dataset is a daily matrix of 2018 counties in which each cell $[i, j]$ reports, among devices that pinged on a particular day in county j and pinged anywhere in the previous 14 days, the share of devices that pinged in county i at least once during the previous 14 days. The dataset is restricted to counties with reasonably large device samples. We assume that diagonal elements of the matrix represent the fraction of cellphones pinged in a particular county that belong to that county itself, and hence $1 - lex[i, i]$ represents the total fraction of devices pinged in county i that had not been in i during the prior 14 days.

B.2 COVID-19 Cases and Deaths

The data for county-level COVID-19 cases and deaths was extracted from two sources: (I) CoronaDataScraper project, and (ii) JHU COVID-19 daily cases and deaths repository. Both the sources are updated daily. While the JHU dataset is more comprehensive of the two, we identified several county-date combinations for which:

- There were missing observations, or
- Data was discontinued for subsequent time periods

There were 136 such counties identified for confirmed cases and 63 for number of deaths.²² For the counties we check, we also corrected two additional errors:

- The cumulative number of cases (or deaths) decrease after the particular date, implying a negative growth rate in cases (or deaths)
- The cases were reported with a lag of more than 1 day

We start with confirming the first reported case for the aforementioned 136 counties. This is important, since in some cases a presumed case was erroneously reported as the first confirmed case, or an administrative error assigned a case from another county or state to the county in question (or person was a temporary resident). What we observe is one (incorrect) entry in the number of cases on a particular day and then no observations for multiple days after that.

While some counties have regular press releases or a daily updated dashboard to check the numbers for a particular day, for the others we rely on multiple news reports. We follow the same procedure for other cases in the panel where the cumulative numbers on the subsequent dates mysteriously reduce only to increase again. Links to the rectification provided by the County Public Health Department as well as the news reports have been provided in the dataset. In cases where the county corrected the numbers but an associated press release was not found, we rely on multiple local news reports for the dates in question. We follow the same procedure for counties which did not have any confirmed cases but the dataset recorded one.

For randomly missing observation on particular dates, we look at the county public health department daily releases and dashboard charts, or the state public health department daily status updates for counties, and finally if there is a lack of information from both sources, we look at reports from the local media. Some state public health departments also provide a disclaimer attributing missing data for certain counties to lag in time between testing and reporting (e.g. Jeff Davis County, Georgia). For these county-date pairs, we rely solely on multiple local news reports that confirm the number of cases on that date.

We follow the above steps for correcting the cumulative number of deaths decreasing over time.

²² A subsequent release of the JHU data corrected 33 of the case count errors.