# What Explains Temporal and Geographic Variation in the Early US COVID-19 Pandemic?

Hunt Allcott, *Stanford University and NBER**
Levi Boxell, *Independent*
Jacob Conway, *University of Chicago*
Billy Ferguson, *Stanford University*
Matthew Gentzkow, *Stanford University and NBER*
Benny Goldman, *Cornell University*

November 2024

**Abstract**

We provide new evidence on the drivers of the early US COVID-19 pandemic and develop a methodology that future researchers can use to similarly analyze the outbreaks of new diseases. We combine an epidemiological model of disease transmission with quasi-random variation arising from the timing of stay-at-home-orders to estimate the causal roles of policy interventions and voluntary social distancing. We then relate the residual variation in disease transmission rates to observable features of cities. We estimate significant impacts of policy and social distancing responses, but we show that the magnitude of policy effects was modest, and most social distancing was driven by voluntary responses. Moreover, we show that neither policy nor rates of voluntary social distancing explained a meaningful share of geographic variation. The most important predictors of which cities were hardest hit by the pandemic were exogenous characteristics such as population and density.

*E-mail: allcott@stanford.edu, leviboxell@gmail.com, jacob.conway@chicagobooth.edu, billyf@stanford.edu, gentzkow@stanford.edu, bgoldman@cornell.edu. We thank Zane Kashner and Peng Zhang for their excellent research assistance. We thank SafeGraph, Facteus, and Homebase for providing access to the data. We also thank the SafeGraph and Homebase COVID-19 response communities respectively for helpful input. We also thank Christopher Avery, Christopher Meissner, David McAdams, an anonymous referee, and seminar participants at Harvard University and Stanford University for their feedback and comments. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1656518. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation We also acknowledge funding from the Stanford Institute for Economic Policy Research (SIEPR), the John S. and James L. Knight Foundation, the Sloan Foundation, and the Institute for Humane Studies. Previously circulated beginning in May 2020 as "Economic and Health Impacts of Social Distancing Policies during the Coronavirus Pandemic."

## 1 Introduction

The course of the COVID-19 pandemic in the US varied dramatically across space and time. In the early months, the epicenter was primarily large cities on the West Coast and in the Northeast. Slow growing outbreaks in Seattle and San Francisco were followed by rapid surges of cases in cities such as New York and Boston. Other large cities such as Dallas, Atlanta, Phoenix, Miami, and Denver were largely spared.

Popular discussion and prior work has pointed to a number of plausible drivers of this variation, with policy differences as one leading explanation.[1] However, evidence below and in prior work clearly shows that substantial behavioral changes preceded the first stay-at-home orders, and that mobility levels remained significantly depressed even after many of these policies were lifted—suggesting voluntary social distancing behaviors may have been important. Other potential drivers include physical characteristics of cities such as population size and density that might have affected virus transmission rates, or connections via international flights to early overseas epicenters such as China and Italy.

In this paper, we provide new evidence on the importance of policy relative to these other factors during the first months of the US pandemic (through April 2020). Our empirical strategy exploits short-run changes around the onset of policy in event-study and difference-in-difference specifications within the context of an SIRD (Susceptible, Infected, Recovered, Deceased) model of disease transmission. By taking the ratio between the estimated effects of policy on disease transmission and the estimated effects of policy on social distancing behaviors, we are able to provide causal estimates for the change in disease transmission rates per unit change in social distancing—allowing us to decompose the role that policy and social distancing behaviors played in driving temporal and geographic variation during the early pandemic period. We then relate the residual variation in disease transmission rates to observable features of cities.

We begin by using event-study designs to estimate the short-run effects of stay-at-home orders on social distancing. We measure social distancing using visits to points of interest (POIs) such as shops, parks, hospitals, and other places measured in cell phone location data aggregated by a company called SafeGraph. We combine POI visits with data on the implementation date of stay-at-home orders at the Combined Statistical Area (CSA) level. Our estimates suggest that POI visits

[1] For example, New York Governor Andrew Cuomo was frequently blamed for the severity of the crisis in New York, with critics citing his slow response to the pandemic and his derision of Mayor De Blasio's earlier suggestion of closing down New York City (Gold and Robinson 2020).

dropped by 18 percent on the day after a stay-at-home order was implemented. We complement our GPS data with information on consumer spending and small business employment—finding that consumer spending dropped by 7 percent and employment dropped by 12 percent.

We next estimate the effects of stay-at-home orders on disease transmission in the context of an SIRD epidemiological model (Kermack and McKendrick 1927). Infected people infect Susceptible people at rate $\beta_t$ and recover at rate $\gamma$. The reproduction number $\mathscr{R}$—the number of people to whom one Infected person transmits the virus—is $\beta_t/\gamma$. If a stay-at-home order is modeled as a proportional effect $\tau$ on the contact rate, then we can estimate $\tau$ in a linear regression framework with the natural log of the number of new cases on the left-hand side. We find that stay-at-home orders reduced the contact rate by 9 to 14 percent for different plausible values of $\gamma$. We show that our estimator performs well when estimated on data simulated from a SIRD model.

We use these estimates to examine the share of the temporal variation in health, social distancing, and economic outcomes that can be attributed to stay-at-home orders versus voluntary and other policy responses. Consistent with other work released around the same time as our initial May 2020 working paper, we find that much of the reduction in POI visits pre-dated stay-at-home orders (Brzezinski et al. 2020; Chetty et al. 2024; Gupta et al. 2021; Sears et al. 2023). We calculate that by mid-April 2020, the short run effects of stay-at-home orders accounted for only 16 percent of observed social distancing, 16 percent of observed reductions in economic activity (measured by small business employment), and 13 percent of the reduction in contact rate.

Next, we use our previous estimates to compute the counterfactual transmission rates if policy, and subsequently social distancing rates, were equalized across all CSAs. We find that neither policy nor social distancing rates explain the geographic variation in transmission rates. Rather, fixed differences across CSAs were the primary drivers. Using a lasso model to select features of the data that are highly predictive of differential transmission rates, we find that population and population density explain nearly half of the average differences across high and low transmission CSAs, with racial composition and partisanship explaining a smaller share. Our model accurately predicts transmission rates in epicenters, such as New York City. These results suggest that much of the observed variation across CSAs was not driven by different policy or voluntary behavioral responses, but was driven by invariant characteristics of CSAs.

Lastly, we use our estimates to examine the effect of counterfactual policies on the overall prevalence of the virus in the United States. While policy explains a small proportion of the temporal variation in case growth and an even smaller proportion of the geographic variation, policy still led to an important reduction in cases. Absent the observed policy response, there would have been 494,000 more confirmed cases by April 30th, 2020 and 14,800 more deaths. In contrast, a uniform stay-at-home order implemented on March 17th, 2020 (the effective date of the first county-level stay-at-home order) would have resulted in 154,000 fewer cases by April 30th.

We emphasize a number of important caveats. Our GPS-based social distancing measure captures overall movement patterns without distinguishing activity with a high vs. low risk of virus transmission. Our measures of economic cost only consider two dimensions of policies' overall economic impact, and these are captured imperfectly. Our measure of health impact relies on the assumptions of our SIRD model, is overall relatively imprecise, and may be biased by factors such as endogenous reporting of COVID-19 cases. In each case, we are able to capture only short-term, on-impact effects. We provide a set of data points that speak to the benefits and costs of social distancing policies but stop well short of a comprehensive welfare analysis. Moreover, our analysis focuses on the early onset of the pandemic. Predictors and drivers of temporal or geographic variation later in the pandemic may have been different.[2] Our analysis also does not distinguish whether policy- and social distancing-induced reductions in cases during the onset of the pandemic permanently avoided these negative health outcomes, or whether they instead delayed these infections (potentially still "flattening the curve" and avoiding health care capacity constraints and unvaccinated infections).[3]

Our work connects to several research areas. First, a series of recent papers used GPS data from SafeGraph or similar providers to quantify social distancing and estimate the effects of stay-at-home orders and other policies (Alexander and Karger 2023; Allcott et al. 2020; Chen et al. 2020a; Engle et al. 2020; Goolsbee and Syverson 2021; Painter and Qui 2021; Sears et al. 2023). Second, several recent papers have studied the effects of stay-at-home policies on economic outcomes (Baker et al. 2020; Bartik et al. 2020; Chen et al. 2020b; Chetty et al. 2024; Kong and Prinz 2020; Lin and Meissner 2020). Third, another set of papers quantified the effects of regulation on health outcomes (Childs et al. 2021; Flaxman et al. 2020; Fowler et al. 2020; Friedson et al. 2021; Greenstone and Nigam 2020; Lasry et al. 2020; Lin and Meissner 2020). Collectively, these contemporaneous literatures have reached a growing consensus (supported by our analysis) that the majority of social distancing was voluntary rather than policy-induced. However, these

[2]For example, Wallace et al. (2023) show that excess death rates during the COVID-19 pandemic were higher for Republican voters after COVID-19 vaccines were made available to all adults (May 2021), but not before.

[3]See Budish (2024) for analysis of mitigation policies sufficient to keep the effective reproduction number below one, resulting in falling infections until a vaccine or cure can be developed. See Rachel (2024) for analysis of how temporary mitigation policies might increase the likelihood of subsequent waves in the absence of a vaccine or cure.

literatures have not fully pinned down the drivers of pandemic outbreaks or the causal effect of policy and social distancing on health outcomes. In the epidemiological literature, there is a set of what economists might call "structural" models that use Bayesian techniques to estimate the reproduction number $\mathscr{R}$; these estimates often pay less attention to identifying the causal effect of policies on $\mathscr{R}$ (Cori et al. 2013; Thompson et al. 2013). In the economics literature, there is a set of papers that use reduced form event-study approaches to estimate the effects of policies on some measure of disease transmission, but many of these papers are not closely tied to structural models of disease transmission. Our paper forms a bridge between these two lines of work by deriving simple linear estimating equations (which are useful for standard quasi-experimental analysis) from structural epidemiological models and using these estimates to decompose the role of various drivers in explaining the temporal and geographic variation of virus transmission in the early months of a pandemic.[4] Our methodology may also serve as a blueprint for future researchers seeking to analyze the outbreaks of new diseases.

# 2 Data

## 2.1 Policy Data

We explore both stay-at-home and business closure policies in this paper. Due to the decentralized policy response of states, cities, and counties, there is no single resource documenting non-pharmaceutical interventions (NPIs) in the United States. To get the best coverage of these NPIs, we combine data from four sources and define both our stay-at-home and business closure policies by sequentially assigning enforcement dates by data source. We prioritize the data sources in the following way (first to last priority): the New York Times (Mervosh et al. 2020a), Keystone Strategy, a crowdsourcing effort from Stanford and University of Virginia, and Hikma Health. Once a state enacted a policy, the counties inherited the policy of the state. In Appendix Table A2, we provide summary statistics reporting the share of county policies from each source. We visualize the distribution of county stay-at-home order implementation dates over time in Appendix Figure A1. Additional detail on data construction can be found in Appendix Section A.1.

[4]Desmet and Wacziarg (2021), Knittel and Ozaltun (2020), and Kuchler et al. (2022) examine spatial variation in COVID cases and/or deaths in the United States without employing the framework of an SIRD model.

### 2.2 Social Distancing Data

The data on social distancing behaviors come from SafeGraph, a data company that aggregates anonymized location data from about 45 million mobile devices and numerous applications in order to provide insights about physical points of interest (POIs). POIs include restaurants, coffee shops, grocery stores, retail outlets, hospitals, and many other business establishments. For each POI, SafeGraph reports the daily number of unique device visits along with information on the POI's industry and location.[5] For each CSA, we construct the total number of visits to POIs in that CSA for a given day.

### 2.3 Economic Data

Our analysis uses economic data from two sources. We incorporate spending on approximately 10 million debit cards in data from Facteus, a financial data provider that directly partners with banks. This sample consists of traditional debit cards issued by banks, general purpose debit cards issued by merchants, payroll cards issued by employers, and government alimony disbursement cards. Lower- and middle-income individuals are represented more heavily in this data than in the US population. We construct the total number of transactions and dollar amount spent by cards from a given home CSA on a given day.[6]

We also source information on employment from Homebase, a company providing scheduling and time tracking software to over 60,000 small businesses.[7] For each day, we analyze the number of work hours and individuals employed by Homebase partner firms in a given CSA.

### 2.4 Health Data

We pull case and death counts by day at the county level from a continually-updated repository by the New York Times that aggregates reports from state and local health agencies. For all dates up to the first available data, we assume no cases nor deaths. We collect state-level testing and hospitalization data from the Covid Tracking Project.

[5] See https://web.archive.org/web/20201024234927/https://docs.safegraph.com/docs/weekly-patterns for additional information on the data's construction.

[6] To protect privacy, Facteus injects a small amount of mathematical noise into key record attributes. This has very minimal impact on aggregate data. More information on this differential privacy procedure can be found at https://web.archive.org/web/20210519180241/https://www.facteus.com/products/data-products/.

[7] Additional information regarding Homebase's data can be found at https://joinhomebase.com/data/.

### 2.5 Demographic Data

We supplement the policy and outcome data with data on CSA characteristics. For our measure of partisanship, we use the Republican vote share in the 2016 presidential election (MIT Election Data and Science Lab 2018). Following Allcott et al. (2020), we use the SafeGraph OpenCensus data to assign demographic variables such as race, income, occupation, and commuting to the various geographies analyzed. The SafeGraph OpenCensus data is derived from the 2016 5-year ACS at the census block group level. We add the urban share of the population from the 2010 Census. We also use average seasonal temperatures by geography from Wu et al. (2020), which is ultimately sourced from gridMET (Abatzoglou 2013). To characterize potential global migration flows for transmission, we use flight data from the OpenSky Network (Schäfer et al. 2014; Olive 2019).

## 3 Effects on Social Distancing and Economic Outcomes

Our main results are at the Combined Statistical Area (CSA) by order date level (i.e., we group counties within a CSA who received a stay-at-home order on the same day together) using data from February 1, 2020 to April 30, 2020.[8] We call our unit of observation a "CSA" for simplicity. To estimate the causal effect of these stay-at-home orders, we estimate the following event-study specification

$$Y_{it} = \mu_i + \delta_t + \sum_{k=-21, k\neq -1}^{k=21} \omega_k \mathbf{1}_{\{t-T_i=k\}} + \varepsilon_{it} \tag{1}$$

where $Y_{it}$ is the outcome of interest in CSA $i$ during time $t$, $\mu_i$ is a CSA fixed effect, $\delta_t$ is a date fixed effect, and $\mathbf{1}_{\{t-T_i=k\}}$ is an indicator for the days relative to the first stay-at-home order $T_i$.[9] Standard errors are clustered at the CSA level irrespective of order timing. Earlier and later time periods are pooled in the $k = -21$ and $k = 21$ time indicators respectively.

Panels A–C of Figure 1 show clear effects of stay-at-home orders on social distancing and economic outcomes. Panel A shows that a stay-at-home order decreased POI visits by 17.8 percent (se = 1.3) by the day after the order's effective start date ($k = 1$). This decrease persisted and was relatively stable throughout the window of analysis. Panel B estimates that a stay-at-home order decreased consumer debit spending (in total \$) by 7.1 percent (se = 0.9) by the day after an

[8] See Appendix A.1 for detail regarding CSAs and the benefits of conducting our analysis at the CSA level.

[9] For CSAs without a stay-at-home order in our sample, $\mathbf{1}_{\{t-T_i=k\}}$ is always set to zero.

order's effective start date. Panel C estimates a 12.3 percent (se = 1.5) reduction in the number of employees working in our Homebase sample on the day following an order's implementation.

In Appendix Figure A4 Panel A, we find no evidence of heterogeneous policy effects on mobility by implementation date. In Panel B, we also find similar policy effects on mobility in majority Democrat vs. majority Republican CSAs. Appendix Figure A5 analyzes the effects of mandatory business closure orders, which induced less negative mobility and spending responses than stay-at-home orders but which caused a comparable reduction in employment.

# 4 Effects on Health Outcomes

## 4.1 SIRD Model

We start with a discrete-time SIRD model (Kermack and McKendrick 1927), suppressing notation for different geographies $i$. In outlining this model, we make the assumption that there are no health spillovers across geographies. Furthermore, we abstract from issues around testing and the endogeneity of stay-at-home order timing.

The population is defined by

$$S_t + I_t + R_t + D_t = N \tag{2}$$

where $S_t$, $I_t$, $R_t$, and $D_t$ are the number of susceptible, infected, recovered, and deceased individuals at time $t$. Dynamics in the SIRD model are defined by the transition probabilities between states. The laws of motion are given by:

$$S_{t+1} - S_t = -\beta_t S_t \frac{I_t}{N} \tag{3}$$

$$I_{t+1} - I_t = \beta_t S_t \frac{I_t}{N} - \gamma I_t \tag{4}$$

$$R_{t+1} - R_t = (1-\kappa)\gamma I_t \tag{5}$$

$$D_{t+1} - D_t = \kappa\gamma I_t \tag{6}$$

where $\beta_t$ is the contact rate that governs the speed at which new infections propagate, $\gamma$ is the rate at which infected individuals recover, and $\kappa$ is the proportion of recovered individuals that die. We treat the recovery rate $\gamma$ and death rate $\kappa$ as fixed parameters during the time period we analyze.

Defining the total number of cases to be $C_t = I_t + R_t + D_t$ and combining equations (4), (5), and (6), we get that new cases evolve as

$$C_{t+1} - C_t = \beta_t S_t \frac{I_t}{N} \tag{7}$$

We make the simplifying assumption that $S_t \approx N_t$ so that we can treat the ratio $\frac{S_t}{N_t} = 1$. As of May 1, less than 0.5 percent of the US population had a confirmed case on or before this date—making this approximation reasonable—though the true case count may have been greater. This allows us to replace equations (3) and (4) with

$$C_{t+1} - C_t = \beta_t I_t. \tag{8}$$

Furthermore, we can write

$$I_t = (C_t - C_{t-1}) + (1-\gamma) I_{t-1}. \tag{9}$$

Given initial conditions $C_0$, $I_0$, the contact rate $\beta_t$, and the recovery rate $\gamma$, equations (8) and (9) define the dynamics of cases over time.

### 4.2 Estimation Framework

A key parameter of interest for policymakers is the contact rate $\beta_t$. As social distancing increases, the contact rate decreases—yielding fewer new cases. The contact rate $\beta_t$ is proportionally related to the reproduction number $\mathscr{R}_{0t}$ as $\mathscr{R}_{0t} = \beta_t/\gamma$. A proportional effect on the contact rate $\beta_t$ will have the same proportional effect on the reproduction number $\mathscr{R}_{0t}$.

We suggest it is natural to think of stay-at-home orders as a proportional effect ($\tau$) on the contact rate that would occur in the absence of policy interventions ($\tilde{\beta}_t$). That is, the contact rate is $\beta_t = \tilde{\beta}_t(1+\tau)$ under a stay-at-home order as opposed to $\tilde{\beta}_t$.

Applying this substitution and taking logs of equation (8), we get

$$\log(C_{t+1} - C_t) = \log(I_t) + \log(\tilde{\beta}_t) + \log(1 + \tau \mathbf{1}_{\{t \geq T_i\}}).$$

where $\mathbf{1}_{\{t \geq T_i\}}$ denotes an indicator for being under a stay-at-home order. We then use an event-study framework to estimate the impact of treatment

$$\log(C_{i,t+1} - C_{it}) = \alpha \log(I_{it}) + \delta_t + \xi_{it} + \sum_{k=-14, k\neq -1}^{k=21} \omega_k \mathbf{1}_{\{t-T_i=k\}} + \varepsilon_{it}, \tag{10}$$

where, relative to equation (1), $\xi_{it}$ is an indicator for the non-binned event-study window.[10]

The $\omega_k$ coefficients can be interpreted as estimates of $\tau$ while $\delta_t$ and $\xi_{it}$ control for variation in $\beta_t$ over time unrelated to the policy. Even though $I_t$ is not directly observed, given initial conditions $C_0 = I_0$ and $\gamma$, no additional data is required to construct the time-path for $I_t$ beyond the growth in cases. Below, we use values for $\gamma$ suggested by the epidemiology literature and examine robustness to alternative values. Note that one test of the model and its assumptions is whether $\hat{\alpha} = 1$.

In Appendix Figure A6, we show that this estimator performs well when estimated on data simulated from an SIRD model, and contrast this with the poor performance of other specifications used in the literature.

### 4.3 Results

A key input into the estimation process is $\gamma$ which is the inverse of the average infectious period for COVID-19. We report estimates using a range of values for $\gamma$. On one extreme, we set $\gamma = 0$ which implies $I_{it} = C_{it}$ or an infinite infectious period. On the other extreme, we set $\gamma = 1$ which implies an average infectious period of 1 day. Early indications in the literature suggested an infectious period of 4.4 to 7.5 days (Anderson et al. 2020). As of May 8, 2020, the CDC website recommended home isolation until at least 10 days have passed since symptoms first appeared, whereas the UK NHS recommended a minimum of 7 days.[11] We view the range of $\gamma = 1/3$ (an infectious period of three days on average) and $\gamma = 1/12$ (an infectious period of twelve days on average) as limits to the range of likely values.

Table 1 reports our estimates using case data and stay-at-home orders. To reduce instances where $\log(C_{i,t+1} - C_{it})$ is undefined, we group counties by the interaction between their cumulative statistical area and the timing of their stay-at-home order.[12] We restrict attention to CSA-days with

[10]The event-study window indicator is required for normalization when geography fixed effects are excluded. We exclude geography fixed effects because they bias estimates in our simulations (see Appendix Figure A7).

[11]https://web.archive.org/web/20200508003242/https://www.cdc.gov/coronavirus/2019-ncov/if-you-are-sick/steps-when-sick.html and https://web.archive.org/web/20200522075042/https://www.nhs.uk/conditions/coronavirus-covid-19/what-to-do-if-you-or-someone-you-live-with-has-coronavirus-symptoms/staying-at-home-if-you-or-someone-you-live-with-has-coronavirus-symptoms/

[12]For simplicity, we subsequently refer to these CSA-timing groups as CSAs.

at least 10 cases, the set of CSAs to either never treated CSAs or CSAs which are observed for at least 8 days before and 20 days after the order, and the time period prior to May 1, 2020. Because of the imprecision of the estimates, we estimate an aggregated event study specification (see table notes). In general, $\gamma$ is positively correlated with the estimated treatment effects of stay-at-home orders on case prevalence.

Using likely values of $\gamma$, we find a negative estimated effect of stay-at-home orders on case prevalence though these estimated effects have wide confidence intervals. Setting $\gamma = 1/6$, which implies an average infectious period of 6 days, our baseline estimates suggest that stay-at-home orders decreased the contact rate $\beta_t$ (i.e., the rate of new cases) by 9.1 percent (se = 4.8) relative to their pre-order levels. Consistent with the data coming from an SIRD data generating process, estimates for $\alpha$ are close to 1.

Panels D and E of Figure 1 report the full event-study plot for $\gamma = 0$ which sets $I_t = C_t$, along with our preferred value of $\gamma = 1/6$. Appendix Figure A9 reports the full event studies for the other values of $\gamma$.

## 5 Explaining Variation in Outcomes

### 5.1 Temporal Variation

In this section, we compute the share of the overall change in each outcome that is attributable to stay-at-home and business closure orders. Secular trends in the outcomes were prominent over our time period as individuals made voluntary behavioral changes (e.g., Appendix Figure A2).

To examine the share of aggregate changes explained by policy, we first compute the average total percent reduction in the outcome as

$$\text{Total}\Delta = \frac{Y_T - Y_0}{Y_0} \tag{11}$$

where $Y_t$ is the weighted average of the level of the outcome in week $t$ taken over geographies that enacted the corresponding order during our time period. $t = 0$ is the first week of February, and $t = T$ is the third week of April. We average across days in the week when computing $Y_t$ to remove any day-of-week effects.[13]

[13] Most CSAs did not have any cases in early February, so the pre-period levels cannot be estimated from the data when examining the overall change in the contact rate $\beta_t$. Since we must choose $\gamma$ for each specification and the reproduction number $\mathscr{R}_0$ is $\beta_t/\gamma$, we can recover $\beta_t$ in the pre-period by specifying $\mathscr{R}_0$. Anderson et al. (2020) and

Next, we compute the average policy-induced change relative to baseline levels as

$$\text{Policy}\Delta = \frac{1}{N}\sum_i w_i \frac{\omega_k \times Y_{\mathscr{T}(i)}}{Y_0} \tag{12}$$

where $\omega_k$ is the estimated treatment effect from Sections 3 and 4, $w_i$ are geography population weights that sum to $N$, and $\mathscr{T}(i)$ is the period prior to the order's implementation for geography $i$. We account for uncertainty induced by the estimation of $\omega_k$ in the standard errors, and treat the values of $Y_t$ as given. We set $k = 1$ in our baseline specification and examine robustness to alternative assumptions.

Figure 2 presents the ratio Policy$\Delta$/Total$\Delta$ for our social distancing, employment, and health outcomes.[14] We compute this estimate separately for stay-at-home orders and mandatory business closure orders. Appendix Table A3 reports Total$\Delta$ and Policy$\Delta$ individually for the stay-at-home orders, the business closure orders, and the simultaneous implementation of both orders.

We estimate that stay-at-home orders explain 16.2 percent of the change in POI visits, 15.6 percent of the change in total wages, 16.0 percent of the change in total employment, and 13.1 percent of the change in the contact rate $\beta_t$ when setting $\gamma = 1/6$ and $\mathscr{R}_0 = 3.0$. Overall, while stay-at-home orders only explain a small proportion of the overall changes in social distancing and case growth, they do not appear inefficient. Stay-at-home orders have a wage cost per unit of social distancing that is 0.96 times as large as the average cost across voluntary behavioral changes and other government orders and has a relative employment cost that is 0.99 times as large.[15]

In contrast to the stay-at-home orders, the employment cost per unit of social distancing for business closures is relatively high. Our estimates suggest that business closure orders have a wage cost per unit of social distancing that is 2.25 times larger than the average cost across voluntary behavioral changes and other government orders and has a relative employment cost that is 1.91 times greater.[16]

Appendix Table A4 considers alternative estimators and provides qualitatively similar conclu-

D'Arienzo and Coniglio (2020) provide an overview of estimates of the initial reproduction rate $\mathscr{R}_0$, ranging from 2.5 to 3.5. Our preferred reproduction rate is $\mathscr{R}_0 = 3.0$. We then compute $Y_T$ and $Y_{\mathscr{T}(i)}$ using equation (8) and the assumed $\gamma$ value; we use the assumed $\mathscr{R}_0$ when (8) is undefined in our data.

[14] Debit card transactions and total spending, according to the Facteus data, did not have a similarly strong decrease between February and April. As a result, the decomposition of this small reduction (or increase in the case of total spending) into a voluntary and mandatory portion is more difficult to interpret. We note that Farrell et al. (2020) find a decrease in consumer spending when using data sources other than Facteus' debit card panel.

[15] That is, $\frac{.156}{.162} / \frac{1-.156}{1-.162} = 0.96$ and $\frac{.160}{.162} / \frac{1-.160}{1-.162} = 0.99$.

[16] That is, $\frac{.182}{.090} / \frac{1-.182}{1-.090} = 2.25$ and $\frac{.159}{.090} / \frac{1-.159}{1-.090} = 1.91$.

sions.

### 5.2 Geographic Variation

We next examine factors that explain the geographic variation in health outcomes.

To formalize the comparison across CSAs, we extend our parameterization of the contact rate $\beta_{it}$ as a function of social distancing behaviors along with date and CSA fixed effects

$$\log(\beta_{it}) = \Lambda_0 \Delta \log(POI_{it}) + \theta_i + \rho_t + \nu_{it} \tag{13}$$

where $\Delta \log(POI_{it})$ is the change in the log of POI visits between March 1, 2020 and $t$, $\theta_i$ are CSA fixed effects, $\rho_t$ are date fixed effects, and $\nu_{it}$ are residuals unobserved by the econometrician. We use our constructed series for $I_{it}$ and the case growth $C_{i,t+1} - C_{it}$ to define $\beta_{it}$ as in equation (8) assuming $\gamma = 1/6$. Note that we can use the ratio of the treatment effects of stay-at-home orders on social distancing behaviors in Section 3 and on the contact rate from Table 1 to provide a *causal* estimate $\hat{\Lambda}_0$ for the change in the contact rate per change in social distancing. Constraining $\Lambda_0$ to be the estimated ratio from our event-study specifications, we estimate equation (13) via a fixed effects estimator.[17]

Table 2 reports the average difference in log contact rates $\log(\beta_{it})$ across various groups of CSAs with high and low average case growth. To understand the role of social distancing and policy in explaining differences between these groups, we use equation (13) and our previous event-study estimates to obtain the predicted log contact rates if (i) stay-at-home policies had been equated across all CSAs and (ii) if changes in social distancing had been equated across all CSAs. Our estimates suggest that little to none of the average geographic variation in log contact rates were explained by differences in social distancing behaviors or policy. The timing of virus onset, accounted for by $\rho_t$, also did little to explain these average differences.[18]

These results suggest the vast majority of differences across CSAs were driven by the estimated fixed effects $\theta_i$.[19] We next turn to examining which observable factors help explain the variation

[17] To implement the regression constraint and avoid taking the log of zero, we set $\beta_{ik} = \frac{1}{1{,}000{,}000}$ when $\beta_{ik} < \frac{1}{1{,}000{,}000}$ and we estimate

$$\log(\beta_{it}) - \hat{\Lambda}_0 \Delta \log(POI_{it}) = \theta_i + \rho_t + \nu_{it}$$

where we use the same estimating sample as in Section 4 restricted to the period between March 15 and April 30, 2020, and we set $\hat{\Lambda}_0 = \frac{.091}{.178} = .51$. We do not use population weights.

[18] Appendix Table A5 reports the share of the cross-CSA variance in log contact rates $\log(\beta_{it})$ explained by each set of covariates and provides similar conclusions as the additive decomposition in Table 2.

[19] Appendix Table A1 reports summary statistics for all CSAs, CSAs with above median average contact rates, and

in the estimated CSA fixed effects.

To examine the determinants of these fixed differences across CSAs, we first perform the descriptive exercise of regressing the CSA fixed effects $\theta_i$ on CSA-level covariates using OLS. Panel A of Figure 3 shows the standardized coefficient estimates from univariate regressions. Panel B of Figure 3 shows the corresponding coefficients when controlling for log population. While many variables were significant predictors in the univariate regressions, only racial composition (share Black, Asian, and Other) and partisanship (Republican vote share) were significant predictors after controlling for log population.

Next, we formalize the prediction exercise using the full set of covariates and lasso to select and penalize coefficients.[20] We choose the lasso penalty to maximize out-of-sample fit in a 10-fold cross-validation without using population weights. Lasso selects variables that cover population (log population and log population density), racial composition (share Black and Other), and partisanship (Republican vote share).[21] Appendix Figure A10 shows the fit of the predicted fixed effects from the lasso regression and the estimated fixed effects $\hat{\theta}_i$. Overall, the model performs well—accurately predicting the contact rate fixed effect for NYC and other major CSAs but performing less well on smaller CSAs.

Using the lasso estimates, we return to the additive decomposition in Table 2. We recompute the log contact rates after equating various subsets of covariates using the estimated coefficients from our lasso model. Overall, the covariates in our lasso model explained 54.8 percent of the difference between CSAs with above-median case growth and CSAs with below-median case growth. The population variables, log population and log population density, were the primary predictors—explaining 48.2 percent of the above-below median difference in CSAs by themselves with partisanship explaining smaller shares.

These results suggest fixed differences across locations played a larger role in explaining differential case growth early on in the pandemic than policy or observed social distancing behaviors. Predictors and drivers of temporal or geographic variation later in the pandemic may have been different. The $S = N$ assumption was also less tenable during later periods in the pandemic as the recovered population grew—thus, complicating an analysis of this later period.

CSAs with below median average contact rates; and provides analogous conclusions.

[20] We exclude the log number of tests as of April 30th from the lasso exercise for endogeneity reasons.

[21] This analysis of the strongest predictors of differences in contact rates across geographies might help to identify the most critical characteristics to incorporate into heterogeneous SIR models (see Ellison 2024 for detail regarding these models).

### 5.3 Counterfactuals

Lastly, we conduct several counterfactual exercises. Our exercises take the form of constructing alternate sequences of contact rates $\{\beta_{ik}\}_{k=0}^{t}$ and using the SIRD model outlined above to compute counterfactual cases given observed initial conditions $I_0 = C_0$.[22]

In Panel C of Figure 2, we compute various counterfactual contact rate sequences $\{\beta_{ik}^{c}\}_{k=0}^{t}$ and examine how these alternative sequences would have shaped the evolution of total cases in our sample. Assuming a proportional impact of stay-at-home orders on social distancing behaviors as estimated in Section 3, a uniform stay-at-home order implemented on March 17 (when the San Francisco Bay Area implemented their stay-at-home order) would have resulted in 154,000 fewer cases by April 30th, or a 19.5 percent reduction in cases.[23] If, instead, we assume that stay-at-home orders caused social distancing behaviors to fall to a fixed level at 35 percent of baseline levels, a uniform stay-at-home order implemented on March 17 would have resulted in 494,000 fewer cases, or a 62.5 percent reduction in cases. If all CSAs followed the social distancing behaviors of the counties in the San Francisco Bay Area that were the first to initiate stay-at-home orders, there would have been 349,000 fewer cases in our sample, or a 44.1 percent reduction. Lastly, removing the proportional effect of all policy would have resulted in 494,000 more cases, or a 62.4 percent increase in cases.

While policy explains a small proportion of the temporal or geographic variation in case growth, policy still has led to a non-trivial reduction in cases over time. As of September 4, 2020, the observed case fatality rate in the United States was 3 percent.[24] Based on this observed case fatality rate, the stay-at-home policies saved 14,800 lives while a further 10,500 lives could have been saved if all CSAs followed the social distancing behavior of the counties in the San Francisco Bay Area that were the first to initiate stay-at-home orders during the initial period of the pandemic.

This counterfactual analysis comes with three important caveats. First, we have assumed that there were no spillover effects from stay-at-home orders to geographies outside of the implementing CSA, as we do not have information on the matrix of potential cross-CSA spillovers. While this limitation is mitigated by our use of CSAs (which combine areas with significant cross-geography employment), if stay-at-home orders reduced contact rates in other geographies, then our counter-

[22] We restrict $\beta_{ik} \geq \frac{1}{1{,}000{,}000}$.

[23] To implement, we counterfactually reduce the log of POI visits by .178 after March 17 for CSAs not under an order at that time.

[24] See https://coronavirus.jhu.edu/data/mortality accessed on September 4, 2020.

factuals are likely to underestimate the impacts of counterfactual policies. A second key limitation of our analysis is that the wide confidence intervals in our estimated impacts on health outcomes (see Section 4) similarly lead to significant imprecision in our health counterfactuals. Finally, our analysis assumes that stay-at-home orders cause a proportional reduction in the log of POI visits that is uniform across geographies and time periods. While Appendix Figure A4 Panel A provides empirical support for the assumption of uniform effects across different implementation dates during this period, it is plausible that effects may be heterogeneous across geographies. The existing literature provides mixes evidence regarding potential heterogeneity. Painter and Qui (2021) estimate that stay-at-home orders have larger effects on mobility in Democratic counties (which were more likely to implement these orders) than in Republican counties. If observed stay-at-home orders are positively selected on their potentially heterogeneous treatment effects, this would lead our counterfactuals to overestimate magnitudes from implementing these policies in geographies that did not in fact issue such orders.[25] In contrast and consistent with our homogeneity assumption, Alexander and Karger (2023) estimate fairly uniform responses to stay-at home orders across the country (including by county-level political leanings, and particularly after controlling for differences in the timing of stay-at-home orders). In Appendix Figure A4 Panel B, we estimate similar impacts on mobility in majority Democrat vs. majority Republican CSAs.[26]

## 6 Conclusion

We use event studies and a model-driven regression framework to provide new decompositions on the role of policy in driving the spatial and temporal variation in case transmission during the early months of the COVID-19 pandemic. We find that policy was responsible for roughly 13 percent of the change in virus contact rates between early March and mid-April 2020.[27] Moreover, policy explained little-to-none of the differences in average contact rates between CSAs with high- and low-contact rates; in other words, policy and social distancing choices were not the primary reason that areas like New York City experienced relatively severe outbreaks during the early pandemic period. Rather, the most important predictors of which cities were hardest hit by the pandemic

[25] We thank an anonymous referee for contributing this point and suggesting heterogeneity tests.

[26] Estimated treatment effect magnitudes are slightly larger in majority Democrat CSAs, but this difference is not statistically significant at standard significance levels for the day following the implementation of the stay-at-home order.

[27] This finding is consistent with a growing consensus that most social distancing during this time period was voluntary rather than policy-driven (e.g., Goolsbee and Syverson 2021), and further extends this analysis to quantify impacts on subsequent health outcomes.

were exogenous characteristics such as population and density.

These results inform the debate about the role policy played in the initial months of the COVID pandemic. While our results suggest that policy explained a small proportion of the overall spatial and temporal variation, this does not mean that (a) policy was inefficient or (b) policy was not important. We find evidence that policy was no less costly in the short-run than voluntary social distancing behaviors. Additionally, our counterfactual estimates suggest a uniform stay-at-home policy implemented on March 17th would have resulted in a 20 percent reduction in cases by April 30th.

The empirical strategy that we have developed can similarly be used to study potential future pandemics, as researchers can use this methodology to rapidly identify the causal effects of policy and voluntary responses, and to analyze the drivers of initial outbreaks.

## References


Abatzoglou, John T. 2013. Development of gridded surface meteorological data for ecological applications and modelling. *International Journal of Climatology*. http://www.climatologylab.org/gridmet.html

Alexander, Diane and Ezra Karger. 2023. Do stay-at-home orders cause people to stay at home? Effects of stay-at-home orders on consumer behavior. *Review of Economics and Statistics.* 105(4): 1017-1027.

Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Yang. 2020. Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics.* 191: 104254.

Anderson, Roy M., Hans Heesterbeek, Don Klinkenberg, T. Deirdre Hollingsworth. 2020. How will country-based mitigation measures influence the course of the COVID-19 epidemic? *The Lancet*. 395(10228): 931–934.

Baker, Scott R., Robert A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. 2020. How does household spending respond to an epidemic? Consumption during the 2020 COVID-19 pandemic. *Review of Asset Pricing Studies.* 10(4): 834–862.

Bartik, Alexander W., Marianne Bertrand, Zoë B. Cullen, Edward L. Glaeser, Michael Luca, and Christopher T. Stanton. 2020. How are small businesses adjusting to COVID-19? Early evidence from a survey. *Proceedings of the National Academy of Sciences.* 117(30): 17656–17666.

Brzezinski, Adam, Guido Deiana, Valentin Kecht, and David Van Dijcke. 2020. The COVID-19 pandemic: Government vs. community action across the United States. *Covid Economics: Vetted and Real-Time Papers* 7: 115–156.

Budish, Eric. 2024. $R < 1$ as an Economic Constraint. *Review of Economic Design.* Forthcoming.

Chen, M. Keith, Yilin Zhuo, Malena de la Fuente, Ryne Rohla, and Elisa F. Long. 2020a. Causal estimation of stay-at-home orders on SARS-CoV-2 transmission. *Working Paper*.

Chen, Sophia, Deniz Igan, Nicola Pierri, and Andrea F. Presbitero. 2020b. Tracking the economic impact of COVID-19 and mitigation policies in Europe and the United States. *Covid Economics: Vetted and Real-Time Papers* 36: 1–24.

Chetty, Raj, John N. Friedman, Michael Stepner, and the Opportunity Insights Team. 2024. The economic impacts of COVID-19: Evidence from a new public database built using private sector data. *The Quarterly Journal of Economics* 139(2): 829-889.

Childs, L. Marissa, Morgan P. Kain, Devin Kirk, Mallory Harris, Lisa Couper, Nicole Nova, Isabel Delwel, Jacob Ritchie, and Erin A. Mordecai. 2021. The impact of long-term non-pharmaceutical interventions on COVID-19 epidemic dynamics and control. *Proceedings of the Royal Society B* 288(1957): 20210811.

Cori, Anne, Neil M. Ferguson, Christophe Fraser, and Simon Cauchemez. 2013. A new framework and software to estimate time-varying reproduction numbers during epidemics. *American Journal of Epidemiology* 178(9): 1505–1512.

Courtemanche, Charles, Joseph Garuccio, Anh Le, Joshua Pinkston, and Aaron Yelowitz. 2020. Strong social distancing measures in the United States reduced the COVID-19 growth rate. *Health Affairs.* 39(7): 1237–1246.

Dave, Dhaval M., Andrew I. Friedson, Kyutaro Matsuzawa, and Joseph J. Sabia. 2021. When do shelter-in-place orders fight COVID-19 best? Policy heterogeneity across states and adoption time. *Economic Inquiry.* 59(1): 29–52.

Desmet, Klaus, and Romain Wacziarg. 2021. Understanding spatial variation in COVID-19 across the United States. *Journal of Urban Economics (Insights)*. 103332.

D'Arienzo, Marco and Angela Coniglio. 2020. Assessment of the SARS-CoV-2 basic reproduction number, R0, based on the early phase of COVID-19 outbreak in Italy. *Biosafety and Health.*

Ellison, Glenn. 2024. Implications of heterogeneous SIR models for analyses of COVID-19. *Review of Economic Design*. 1-37.

Engle, Samuel, John Stromme, and Anson Zhou. 2020. Staying at home: mobility effects of COVID-19. *Covid Economics: Vetted and Real-Time Papers* 4: 86–102.

Farrell, Diana, Fiona Greig, Natalie Cox, Peter Ganong, and Pascal Noel. 2020. The initial household spending response to COVID-19: Evidence from credit card transactions. *JPMorgan Chase Institute Report.* Accessed at https://institute.jpmorganchase.com/institute/research/household-income-spending/initial-household-spending-response-to-covid-19.

Fernández-Villaverde, Jesús, and Charles I. Jones. 2022. Estimating and simulating a SIRD model of COVID-19 for many countries, states, and cities. *Journal of Economic Dynamics and Control* 140: 104318.

Flaxman, Seth, Swapnil Mishra, Axel Gandy, H. Unwin, Helen Coupland, T. Mellan, Harisson Zhu et al. 2020. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature. 584:257-261.*

Fowler, James H., Seth J. Hill, Remy Levin, and Nick Obradovich. 2021. The effect of stay-at-home orders on COVID-19 cases and fatalities in the United States. *PLOS ONE*. 16(6): 1–15.

Friedson, Andrew I., Drew McNichols, Joseph J. Sabia, and Dhaval Dave. 2021. Shelter-in-place orders and public health: Evidence from California during the Covid-19 pandemic. *Journal of Policy Analysis and Management*. 40(1): 258–283.

Gold, Lyta and Nathan Robinson. 2020. Andrew Cuomo is no hero. He's to blame for New York's coronavirus catastrophe. *The Guardian.* Accessed at https://www.theguardian.com/commentisfree/2020/may cuomo-new-york-coronavirus-catastrophe.

Goolsbee, Austan, and Chad Syverson. 2021. Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline. *Journal of Public Economics*. 193: 104311.

Greenstone, Michael, and Vishan Nigam. 2020. Does social distancing matter? *Covid Economics: Vetted and Real-Time Papers.* 1(7): 1–22.

Gupta, Sumedha, Thuy Nguyen, Shyam Raman, Byungkyu Lee, Felipe Lozano-Rojas, Ana Bento, Kosali Simon, and Coady Wing. 2021. Tracking public and private response to the covid-19 epidemic: Evidence from state and local government actions. *American Journal of Health Economics* 7(4): 361-404.

Kermack, William Ogilvy, and Anderson G. McKendrick. 1927. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*. 115(772): 700–721.

Knittel, Christopher R., and Bora Ozaltun. 2020. What does and does not correlate with COVID-19 death rates. *NBER Working Paper 27391.*

Kong, Edward, and Daniel Prinz. 2020. The impact of non-pharmaceutical interventions on unemployment during a pandemic. *Journal of Public Economics*. 189. 104257.

Kuchler, Theresa, Dominic Russel, and Johannes Stroebel. 2022. JUE Insight: The geographic spread of COVID-19 correlates with the structure of social networks as measured by Facebook. *Journal of Urban Economics.* 127: 103314.

Lasry, Arielle, Daniel Kidder, Marisa Hast, Jason Poovey, Gregory Sunshine, Nicole Zviedrite, Faruque Ahmed, and Kathleen A. Ethier. 2020. Timing of community mitigation and changes in reported COVID-19 and community mobility—four US metropolitan areas, February 26–April 1, 2020. *CDC Morbidity and Mortality Weekly Report.* 69: 451-457.

Lin, Zhixian and Christopher M. Meissner. 2020. Health vs. wealth? Public health policies and the economy during COVID-19. *NBER Working Paper 27099.*

Mervosh, Sarah, Denise Lu, and Vanessa Swales. 2020a. See which states and cities have told residents to stay at home. *The New York Times.* Accessed at https://www.nytimes.com/interactive/2020/us/corona
stay-at-home-order.html.

Mervosh, Sarah, Jasmine C. Lee, Lazaro Gamio, and Nadja Popovich. 2020b. See how all 50 states are reopening. *The New York Times.* Accessed at https://www.nytimes.com/interactive/2020/us/states-reopen-map-coronavirus.html.

MIT Election Data and Science Lab. 2018. County presidential election returns 2000-2016. Harvard Dataverse. https://doi.org/10.7910/DVN/VOQCHQ.

Olive, Xavier. 2019. Traffic, a toolbox for processing and analysing air traffic data. *Journal of Open Source Software.* 4(39).

Painter, Marcus, and Tian Qiu. 2021. Political beliefs affect compliance with COVID-19 social distancing orders. *Journal of Economic Behavior & Organization.* 185: 688–701.

Rachel, Lukasz. 2024. The second wave. *Review of Economic Design.* Forthcoming.

Schäfer, Matthias, Martin Strohmeier, Vincent Lenders, Ivan Martinovic, and Matthias Wilhelm. 2014. Bringing up OpenSky: A large-scale ADS-B sensor network for research. *Proceedings of the 13th IEEE/ACM International Symposium on Information Processing in Sensor Networks.* 83–94.

Sears, James, J. Miguel Villas-Boas, Sofia Berto Villas-Boas,and Vasco Villas-Boas. 2023. Are we #stayinghome to flatten the curve? *American Journal of Health Economics* 9(1): 71-95.

Thompson, Kimberly M., Mark A. Pallansch, Radboud J. Duintjer Tebbens, Steve G. Wassilak, Jong-Hoon Kim, and Stephen L. Cochi. 2013. Preeradication vaccine policy options for poliovirus infection and disease control. *Risk Analysis.* 33(4): 516–543.

Vissat, Ludovica Luisa, Nir Horvitz, Rachael V. Phillips, Zhongqi Miao, Whitney Mgbara, Yue You, Richard Salter, Alan E. Hubbard, and Wayne M. Getz. 2022. A comparison of COVID-19 outbreaks across US Combined Statistical Areas using new methods for estimating R0 and social distancing behaviour. *Epidemics*. 41: 100640.

Wallace, Jacob, Paul Goldsmith-Pinkham, and Jason L. Schwartz. 2023. Excess death rates for republican and democratic registered voters in Florida and Ohio during the COVID-19 pandemic. *JAMA Internal Medicine.*183(9): 916-923.

Wu, Xiao, Rachel C. Nethery, Benjamin M. Sabath, Danielle Braun, and Francesca Dominici. 2020. Air pollution and COVID-19 mortality in the United States: Strengths and limitations of an ecological regression analysis. *Science Advances.* 6(45): eabd4049.

Figure 1: Effect of Stay-at-Home Orders

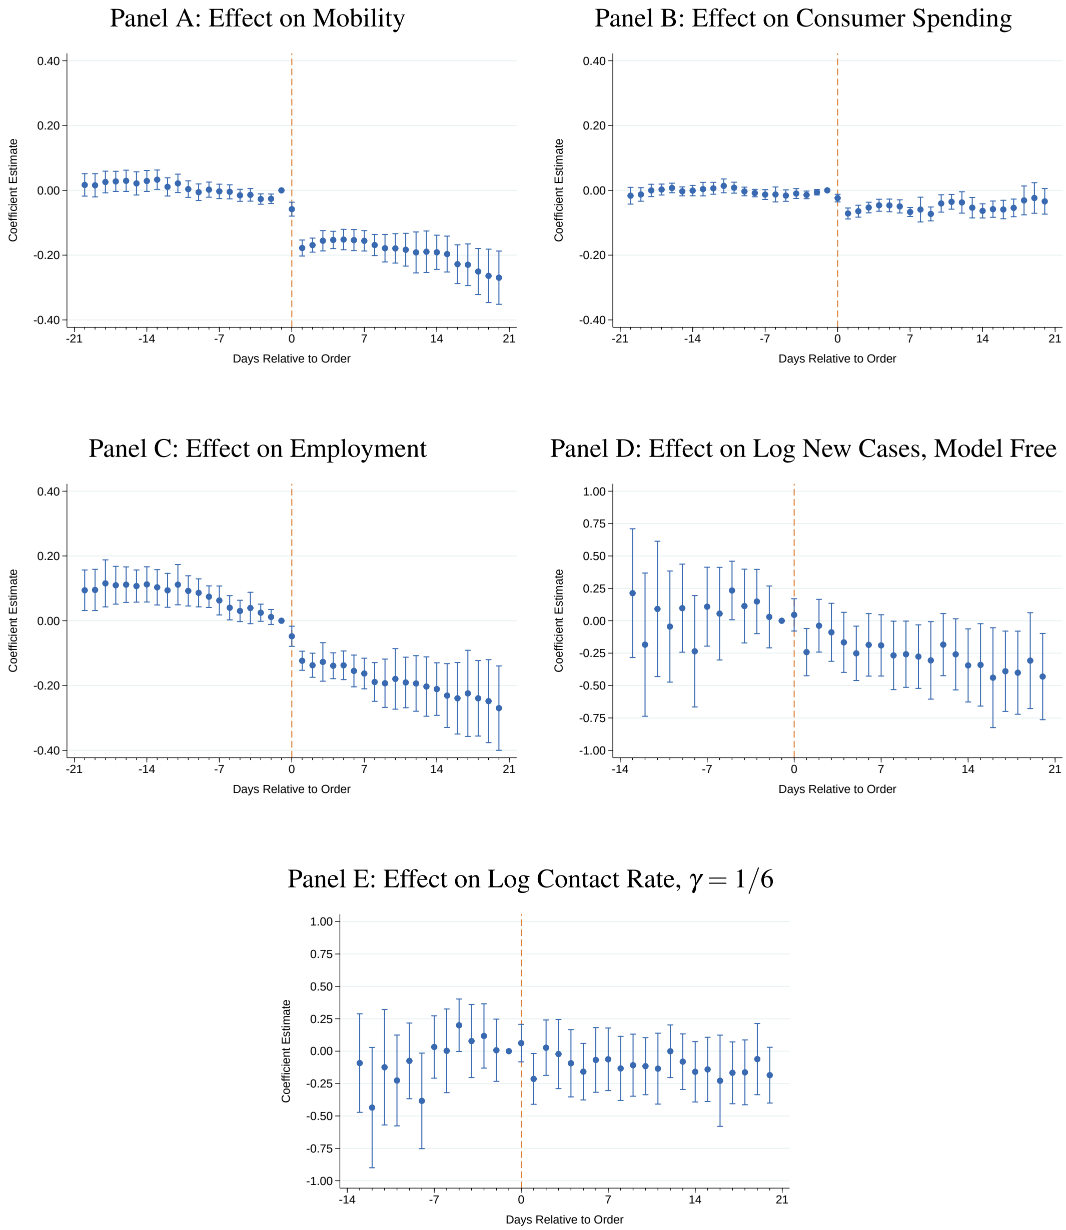


Note: Figure plots estimated treatment effects $\omega_k$ of stay-at-home orders on different outcomes, using the event-study specification at the CSA-day level outlined in equation (1) for the mobility and economic outcomes and equation (14) for the health outcomes. Panel A shows the effect on mobility, using the log of total POI visits in the SafeGraph data. Panel B shows the effect on consumer spending, using the log of total spending in Facteus' debit card sample. Panel C shows the effect on employment, using the log number of individuals with positive work hours from the Homebase sample. Panel D examines log new cases as the outcome and sets $\gamma = 0$ which implies $\log(I_{it}) = \log(C_{it})$. Panel E is the same as Panel D, except that it sets $\gamma = 1/6$. All regressions include date fixed effects $\delta_t$. Panels A-C include CSA fixed effects $\mu_i$; Panels D-E include the event window indicator $\xi_{it}$. CSAs are weighted by population in the regression. Standard errors are clustered at the CSA level irrespective of order timing.

Figure 2: Temporal Variation Explained by Policy

*Panel A: Share of Temporal Variation in Social Distancing and Employment Outcomes*

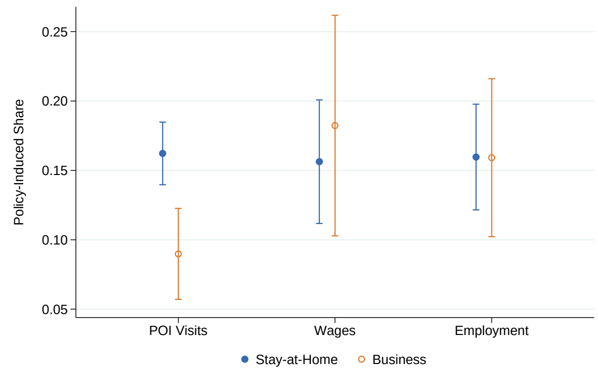


*Panel B: Share of Temporal Variation in Contact Rate $\beta_t$*

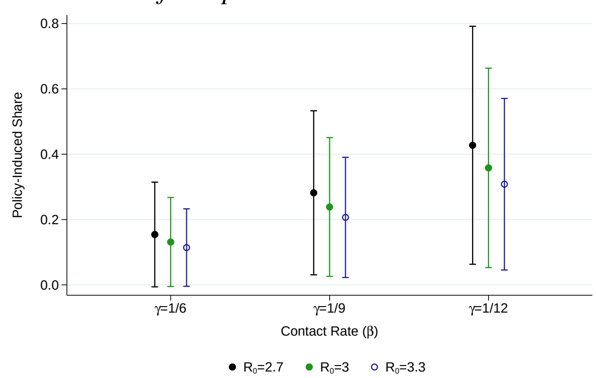


*Panel C: Counterfactual Case Growth*

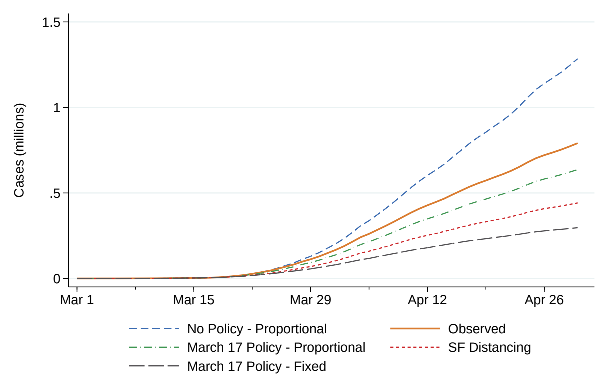


Note: Figure plots the share of the total change in each outcome that is attributable to a given policy Policy$\Delta$/Total$\Delta$ following Section 5.1. Panel A reports estimates using stay-at-home orders and business closure orders for POI visits from SafeGraph, total wages from Homebase, and employment from Homebase. For each policy treatment, we restrict attention to the CSAs treated by the given policy. Panel B reports estimates of the policy-induced change in the contact rate $\beta_t$ from stay-at-home orders, varying the assumed basic reproduction number $\mathscr{R}_0$. In both panels, the bars depict 95 percent confidence intervals. See Table A3 for additional details on Panels A and B. Panel C plots observed and counterfactual cases following the methodology outlined in Section 5.3 for our sample of cumulative statistical areas (CSAs) using $\gamma = 1/6$. Panel C reports the observed number of cases along with the estimated number of cases if a uniform stay-at-home order was implemented on March 17 with proportional effect on social distancing behaviors, a uniform stay-at-home order was implemented on March 17 that caused social distancing behaviors to fall to a fixed level of 35 percent of March 1 levels, and the estimated number of cases if all areas followed the same social distancing behavior as the San Francisco Bay Area (defined to be the counties in the San Francisco CSA that were first to implement a stay-at-home order).

Figure 3: Determinants of Geographic Variation in Log Contact Rates

*Panel A: Univariate Regression of $\hat{\theta}_i$ on Covariates*

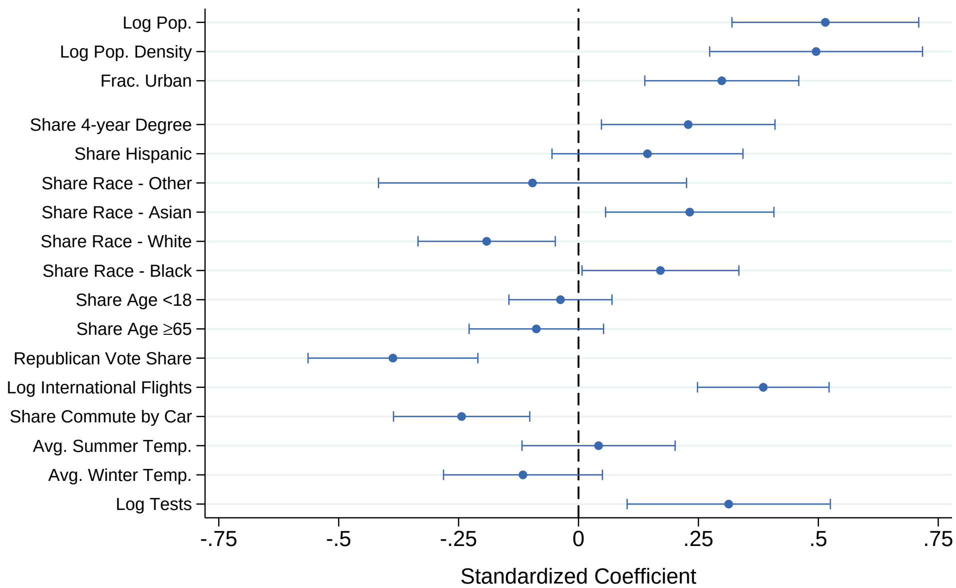


*Panel B: Controlling for Log Population*

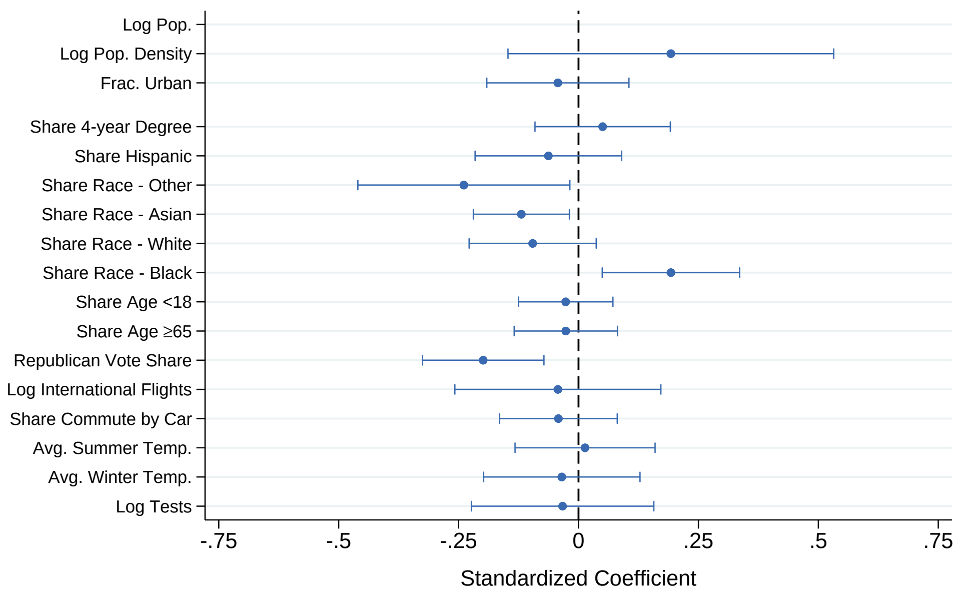


Note: Figure plots the coefficients of regressing the CSA log contact rate fixed effect estimates $\hat{\theta}_i$ from equation (13) on CSA-level determinants. The $\hat{\theta}_i$ and all covariates have been standardized to have a mean 0 and a standard deviation of 1. Panel A plots the standardized coefficients and 95 percent confidence intervals from univariate regressions. Panel B repeats Panel A but the regressions also control for the log of population. Population weights are not used. Robust standard errors are used to compute the confidence intervals.

Table 1: Estimated Effects of Stay-at-Home Orders on Contact Rate

| | (1) $\gamma = 0$ | (2) $\gamma = 1/12$ | (3) $\gamma = 1/9$ | (4) $\gamma = 1/6$ | (5) $\gamma = 1/3$ |
|---|---|---|---|---|---|
| Post-order | -0.203 | -0.139 | -0.121 | -0.091 | -0.023 |
| | (0.086) | (0.060) | (0.055) | (0.048) | (0.045) |
| $\log(I_{it})$ | 1.053 | 1.048 | 1.044 | 1.037 | 1.013 |
| | (0.027) | (0.019) | (0.017) | (0.012) | (0.006) |
| Clusters | 76 | 76 | 76 | 76 | 76 |
| Obs. | 5240 | 5240 | 5240 | 5240 | 5240 |

Note: Table shows estimated coefficients from estimating an aggregated version of the event study in equation (1):

$$\log(C_{i,t+1} - C_{it}) = \alpha \log(I_{it}) + \delta_t + \omega_0 E_{it} + \omega_1 \text{Post}_{it} + \xi_{it}^0 + \xi_{it}^1 + \varepsilon_{it} \tag{14}$$

where $E_{it} = \mathbf{1}_{\{-9<t-T_i<21\}}$, $\text{Post}_{it} = \mathbf{1}_{\{-1<t-T_i<21\}}$, $\xi_{it}^0 = \mathbf{1}_{\{t-T_i<-8\}}$, and $\xi_{it}^1 = \mathbf{1}_{\{t-T_i>20\}}$. 'Post-order' reports $\hat{\omega}_1$, and '$\log(I_{it})$' reports $\hat{\alpha}$. Each column reports estimates for a different value of $\gamma$ as reported in the header. Observations are weighted by population. Standard errors clustered by CSA are reported in parenetheses.

Table 2: Additive Decomposition for Geographic Differences in Log Contact Rates

| | Above/Below Med. | Top/Bot. Quart. | Top/Bot. Dec. |
|---|---|---|---|
| Overall Difference | 1.228 | 2.171 | 3.565 |
| *Share of difference explained by:* | | | |
| Social Distancing | -0.009 | -0.010 | -0.023 |
| Policy | -0.006 | -0.007 | -0.007 |
| Timing of Virus | -0.000 | -0.003 | -0.003 |
| Observed Covariates | 0.548 | 0.508 | 0.424 |
| Population | 0.482 | 0.423 | 0.336 |
| Climate | 0.000 | 0.000 | 0.000 |
| Transport | 0.000 | 0.000 | 0.000 |
| Race | -0.036 | -0.015 | -0.023 |
| Partisanship | 0.103 | 0.100 | 0.110 |
| College Degrees | 0.000 | 0.000 | 0.000 |
| Age Demographics | 0.000 | 0.000 | 0.000 |

Note: Table reports the difference in the average log contact rate $\log(\beta_{it})$ between March 15 and April 30, 2020 for each group of CSAs. It also reports the counterfactual share of the overall difference explained by each set of determinants as outlined in Section 5.2. The following sets of covariates were considered for possible inclusion, but only bolded variables were selected by a Lasso model and included in this variance decomposition:

- Population: **Log Population Density**; **Log Population**; Share Urban
- Climate: Average Summer Temperature; Average Winter Temperature
- Transport: Share Commute Auto; Log International Flights
- Race: **Share Black**; Share White; Share Asian; **Share Other Race**; Share Hispanic
- Partisanship: **Republican Vote Share**
- College Degrees: Share with Bachelor's or More
- Age Demographics: Share Age $\geq 65$; Share $< 18$

# A Appendix

See the replication code for exact details on data construction and estimation.

## A.1 Data Construction Procedures

We construct the datasets used in our analysis as follows.

1. We begin by matching SafeGraph POIs to the counties in which they are located. We use latitude and longitude from SafeGraph's August 2020 Core POI dataset, along with the 2010 TIGER county shapefile.[28] We successfully assign 99.9 percent of the POIs to a county.

2. We then merge the POI-county mapping from (1) onto SafeGraph's Patterns data using the safegraph-place-id variable. We sum visits by county for a given day, aggregating across POIs. Our SafeGraph series ranges from January 1, 2020 to August 30, 2020.

3. We then merge onto the output from (2) a dataset of county-level demographic information constructed as follows. We use the Open Census data from SafeGraph, aggregating up the data given at the census block group level to the county level. We combine this with data on county 2016 Presidential votes shares (MIT Election Data and Science Lab 2018). We define the Republican vote share to be the share of votes received by the Republican candidate over the sum of votes across all candidates. We exclude counties without valid vote data, which drops Alaska and two additional counties (FIPS: 15005, 51515). We also merge on the urban population share from the 2010 Census.[29] We also use average seasonal temperatures by geography from Wu et al. (2020), which is ultimately sourced from gridMET. Averages for a given county and season are taken across the years 2010-2016.

4. We then merge onto the output from (3) the number of incoming international flights for each US county from December 2019 to February 2020 made available by the OpenSky Network (Schäfer et al. 2014; Olive 2019).

5. We then merge data on Covid-19 health outcomes onto the output from (4). We source confirmed Covid-19 cases and deaths by county and day from The New York Times. We assume zero cases and deaths for the observations not observed in The New York Times data. We drop the four counties which overlap with Kansas City (MO), because The New York Times lists these as geographic exceptions where it either does not assign cases to these counties or excludes cases occurring within the city. For the five boroughs of New York City,

[28] Downloaded from https://www.census.gov/geo/maps-data/data/cbf/cbf_counties.html on July 24, 2018.

[29] Downloaded from https://www.census.gov/programs-surveys/geography/technical-documentation/records-layout/2010-urban-lists-record-layout.html on June 25, 2020.

we assign cases and deaths based on their population share. We also merge data on testing and hospitalizations from the Covid Tracking Project by state and day.

6. We then merge a dataset of county-level shelter-in-place and business closure order start dates onto the output from (5) and construct an indicator for whether a county had been subject to a shelter-in-place and/or business closure order by a given date. It is ultimately sourced from The New York Times, Keystone Strategy, a crowdsourcing effort from Stanford University and the University of Virginia, and Hikma Health. The New York Times has been tracking "shelter-in-place", "stay-at-home", "healthy-at-home", etc. policies enacted at the city, county and state level (Mervosh et al. 2020a). We use the dates reported by the NYT for our stay-at-home policy. The relevant NPIs from Keystone's data are shelter-in-place (SIP) and closure of public venues (CPV) policies. Keystone considers an "order indicating that people should shelter in their homes except for essential reasons" as a SIP intervention and a "government order closing gathering venues for in-person service" as a CPV intervention. We will use Keystone's SIP and CPV dates for our stay-at-home and business closure policies respectively. The crowdsourced data collected by a group from Stanford and University of Virginia solicits policy and personal information from survey participants in an online form. We use the "lockdown" and "business closed" dates for counties from this data. Hikma Health, a non-profit working on data systems and analysis for healthcare providers, has carried out their own crowdsourcing effort to document NPIs. We use the county "shelter date" and "work date" from Hikma Health in our construction of stay-at-home and business closure policies respectively. Given that none of the sources have entirely overlapping policy data, we define both our stay-at-home and business closure policies by sequentially assigning enforcement dates when data is available in the order: NYT, Keystone, Stanford/Virginia crowdsource, and Hikma Health. Once a state enacts a policy, the counties inherit the policy of the state. We then merge on a dataset of reopening dates at the state level collected by the NYT (Mervosh et al. 2020b) and curated by Rearc.

7. We then merge debit card transactions and spending totals from Facteus onto the output from (6). Prior to this merge, Facteus data is aggregated from the zip code to the county level using a 5-digit zip code to county crosswalk.[30] Facteus data ranges from Jan 1, 2020 to August 27, 2020, with missing data on August 7, 2020.

8. We then merge employment data from Homebase onto the output from (7). Homebase data is aggregated to the county-day level prior to this merge. This step completes the construction of the dataset used in our county-level analysis. Homebase data ranges from Jan 1, 2020 to

[30]Downloaded from HUD (https://www.huduser.gov/portal/datasets/usps_crosswalk.html) on April 12, 2020

August 31, 2020.

9. For analysis at the level of a CSA, order date, and day, we then aggregate the output from (8) to this level. We use a county to CSA crosswalk.[31] We sum countable variables such as POI visits or Covid-19 cases. For other variables, we take a population-weighted average across counties in a CSA-day with the same social distancing policy start date. In our CSA-level analysis, we exclude counties not assigned to a CSA.

CSAs consist of groups of counties that have significant commuting ties across counties within the group. More precisely (per https://www.census.gov/programs-surveys/metro-micro/about/glossary.html), each CSA consists of two or more adjacent core based statistical areas (CBSAs) that have an employment interchange measure of 15 or more. CBSAs in turn consist of the county or counties (or equivalent entities) associated with at least one core (urban area) of at least 10,000 population, plus adjacent counties having a high degree of social and economic integration with the core as measured through commuting ties. Employment interchange is defined as the sum of the percentage of workers living in the smaller entity who work in the larger entity and the percentage of employment in the smaller entity that is accounted for by workers who reside in the larger entity. We use CSAs as our unit of analysis in order to partially capture policy and health spillovers across geographies that have significant commuting ties, to avoid instances in which the log of new COVID-19 cases is undefined at the county level, and to focus our analysis on metropolitan and micropolitan areas (excluding particularly rural areas that likely had more limited COVID-19 testing capacity). This differs from the focus on county-level analysis in much of the existing literature, with Vissat et al. (2022) being one exception that also analyzes variation across CSAs (without analyzing causal effects or incorporating data on policy and social distancing). 80 percent of the US population resided within CSAs as of April 1, 2020, per data from the US Census (https://www2.census.gov/programs-surveys/popest/tables/2020-2023/metro/totals/csa-est2023-pop.xlsx, downloaded on August 20, 2024).

## A.2 SIRD Simulations and Estimation Details

### *A.2.1 Simulation*

To generate data from an SIRD data generating process, we assume a death rate $\kappa = .008$ and an average infectious period of 10 days ($\gamma = .1$). We assume $\beta_{it}$ evolves as $\beta_{it} = \gamma(\mathscr{R}_{i0}e^{\lambda_i t} + \mathscr{R}_1(1 - e^{\lambda_i t})) \times (1 + \tau T_{it})$ where $\mathscr{R}_{i0}$ is drawn from a normal distribution with mean 2.4 and standard deviation 0.1, $\mathscr{R}_1 = .95$, $\tau$ is either 0 or -0.1 depending on the simulation, $T_{it}$ is drawn from a

[31] Downloaded on May 29, 2020 from the NBER (http://data.nber.org/cbsa-csa-fips-county-crosswalk/cbsa2fipsxw.csv), which uses delineation files originally sourced from the Census

binomial distribution with size 150 and probability 4/15, and $\lambda$ is drawn from a normal distribution with mean $-0.08$ and standard deviation $0.01$. The exponential decay process for $\beta_t$ follows Fernández-Villaverde and Jones (2022). The initial share of the population infected is drawn from an exponential distribution with rate $10,000$. The size of the population is drawn from an exponential distribution with rate $1/100,000$.

We then follow the laws of motion outlined in the main text, updating $\beta_t$ each period. Note that in constructing the simulations, we do *not* use the $S/N \approx 1$ assumption but simulate data from the complete model. We simulate data for 200 geographies with 150 time periods. We then follow the estimation process outlined in the main text and assume that $\gamma$ is the true value used in simulations. We drop the event window indicator in equation (14) when including geography fixed effects. We also show robustness to incorrectly specifying $\tau$ as half its true value and twice its true value.

### *A.2.2 Dave et al. (2021)*

Dave et al. (2021; Figure 4) use an event-study specification with the log of confirmed cases on the left-hand side and geography-specific linear time trends, e.g.,

$$\log(C_{it}) = \mu_i + \mu_i \times t + \delta_t + \sum_{k=-7, k\neq -1}^{k=21} \omega_k \mathbf{1}_{\{t-T_i=k\}} + \varepsilon_{it} \tag{15}$$

where $\log(C_{it})$ is the log of confirmed cases in geography $i$ during time $t$, $\mu_i$ is a geography fixed effect, $\delta_t$ is a date fixed effect, $\mathbf{1}_{\{t-T_i=k\}}$ is an indicator for the days relative to the first stay-at-home order $T_i$, and $\mu_i \times t$ is a geography-specific linear time trend.[32] Following Dave et al. (2021), we exclude data after April 20, 2020.

### *A.2.3 Estimation Details*

To implement our SIRD estimator, we proceed as follows:

1. We assume zero cases in counties for which none have been reported, aggregate county data to the CSA-order timing level (subsequently "CSA"), and constrain cumulative cases to be non-decreasing at the CSA level. We then make the following restrictions:

    (a) CSA-days with at least 10 cases,

    (b) CSAs that either never receive a stay-at-home order or are observed at least 8 days prior and 20 days after the implementation of the order, and

[32] In their event-study specification, Dave et al. (2021) aggregate multiple post-treatment periods into a single treatment effect and include other control variables.

(c) CSA-days between February 1, 2020 and April 30, 2020.

2. We set $C_{i0}$ to be the number of confirmed cases during the period in which at least 10 cases are confirmed, and we define $C_{i0} = I_{i0}$.

3. We then use equation (9) to define the full time path of $I_{it}$ for each geography given $\gamma$.

4. We then estimate equation (14), using $\log(\frac{1}{2} + C_{t+1} - C_t)$ to avoid taking log of zero when no new cases are observed.

### *A.2.4 Other Methods of Estimation in the Literature*

Several previous attempts at estimating the effect of stay-at-home orders have not been model-driven. For example, Dave et al. (2021) use the log of confirmed cases as the outcome in an event-study framework with state-specific linear trends.[33] These estimators can produce unexpected results when the data comes from an SIRD data-generating process.

In Appendix Figure A8, we show that the Dave et al. (2021) estimator exhibits substantial pretrends and fails to recover the estimated treatment effect when estimated on simulated data. We also apply the estimator to real, state-level data as in Dave et al (2021). We qualitatively replicate their results when using the 7-day pre-period event window. When using a more complete 14-day or 21-day pre-period event window, the estimator produces null results with substantial pretrends—mirroring the estimates from this estimator when using data simulated from an SIRD model. To gain intuition for the poor performance of these estimators, equation (8) implies log cases follow

$$\log(C_{i,t+1}) = \log(\beta_{it} I_{it} + C_{it}). \tag{16}$$

Therefore, the $\omega_k$ coefficients from an event study with log cases on the left-hand side are going to pick up differential trends in a nonlinear function of $\beta_{it}$, $I_{it}$, and $C_{it}$ across treated and non-treated units rather than differential trends in $\log(\beta_{it})$ alone.[34]

[33] Friedson et al. (2021) use a synthetic control estimator with log cases on the left-hand side to estimate the effect of California's stay-at-home order. Lin and Meissner (2020), Fowler et al. (2021), and Courtemanche et al. (2020) use various difference-in-difference specifications with $\log(C_{i,t+1}) - \log(C_{it})$ on the left-hand side, which gives $\log(C_{i,t+1}) - \log(C_{it}) = \log(\beta_{it} I_{it} + C_{it}) - \log(\beta_{i,t-1} I_{i,t-1} + C_{i,t-1})$. Lin and Meissner (2020) also perform a matching exercise of counties across state borders. Given the nonlinear dynamics of the SIRD model, synthetic control or matching estimates may perform better when the model structure is not accounted for parametrically.

[34] Note that even under the simplifying assumption that $C_{it} = I_{it}$ (which implies $\gamma = 0$), rewriting equation (16) still gives

$$\log(C_{i,t+1}) = \log(1 + \beta_{it}) + \log(C_{it}).$$

Appendix Figure A1: Distribution of Timing of First Government Order

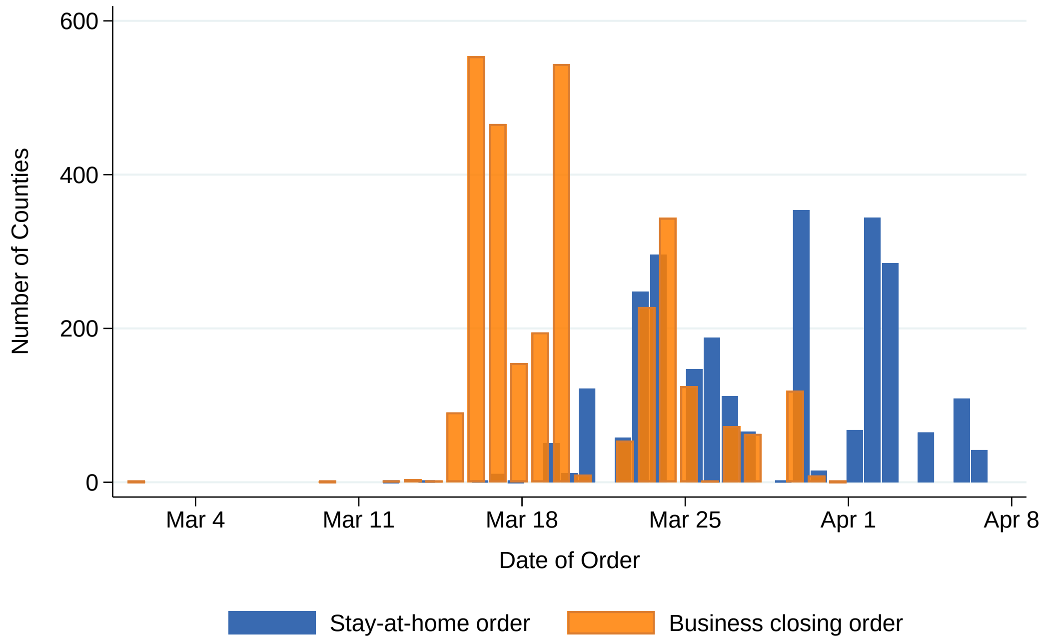


Note: Figure shows the distribution of government order effective start dates over time and across counties. Each bar represents the number of counties (y-axis) for which the first order of a given type went into effect on the date specified (x-axis). Stay-at-home and business closing orders are shown in blue and orange bars respectively. See Section 2.1 for detail on data sources and processing.

Appendix Figure A2: Trends in Average Mobility and Health

*Panel A: Mobility*

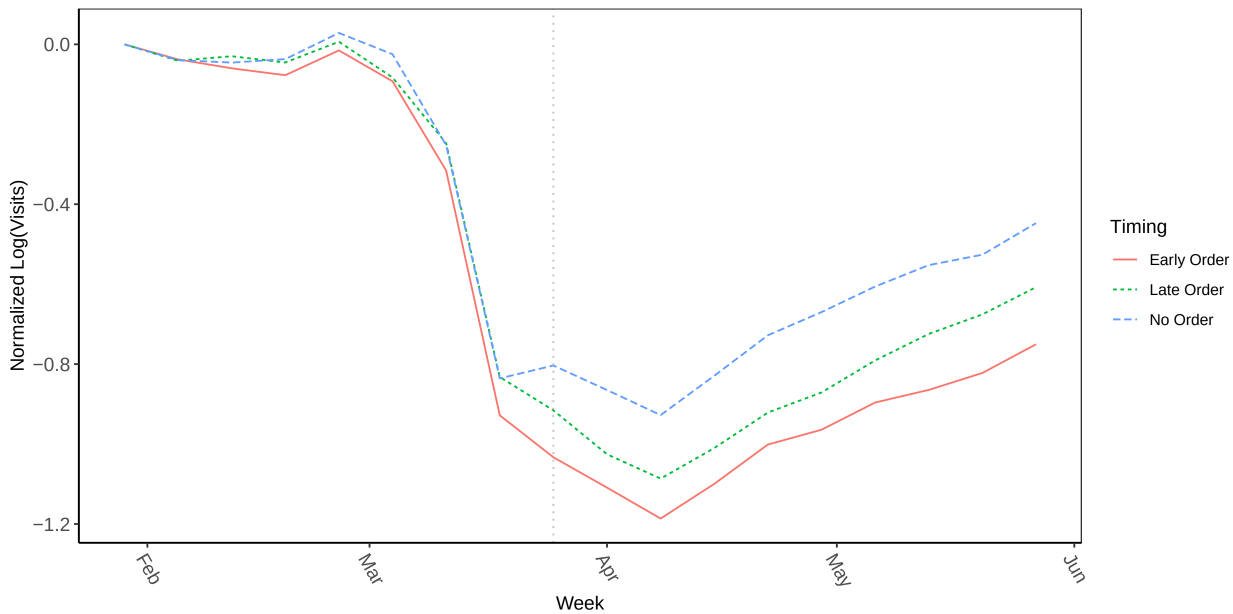


*Panel B: Health*

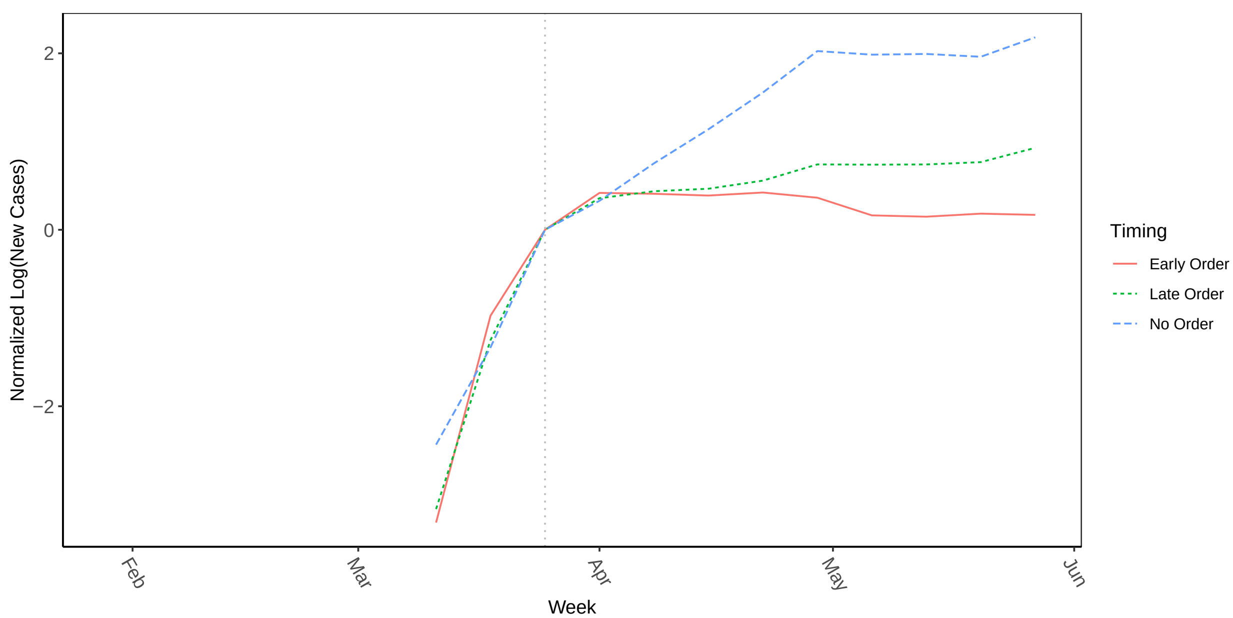


Note: Figure reports trends in average mobility and health by county order timing and week. Panel A plots the log of daily average POI visits, normalizing relative to the week starting January 29, 2020. Panel B plots the log of daily new COVID-19 cases, normalized to the week starting March 25, 2020. In both panels, averages are weighted by population and taken across counties and days prior to taking logs or normalization. 'Early Order' indicates counties with a stay-at-home order on or before March 25, 2020. 'Late Order' indicates counties with a stay-at-home order after this point. 'No Order' indicates counties which did not issue a stay-at-home order during this sample period. The dashed vertical line indicates the week starting March 25, 2020.

Appendix Figure A3: Geographic Variation in Social Distancing and Public Policy

*Panel A: % Change in SafeGraph Visits*

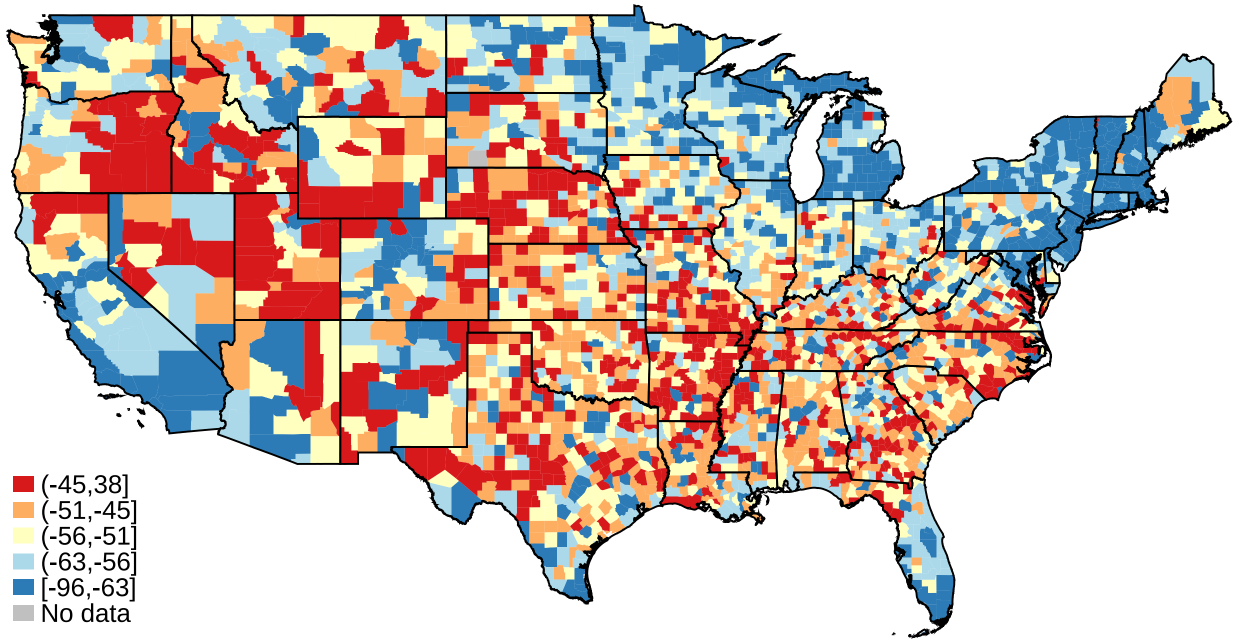


*Panel B: Shelter-in-Place Order Effective Start Date*

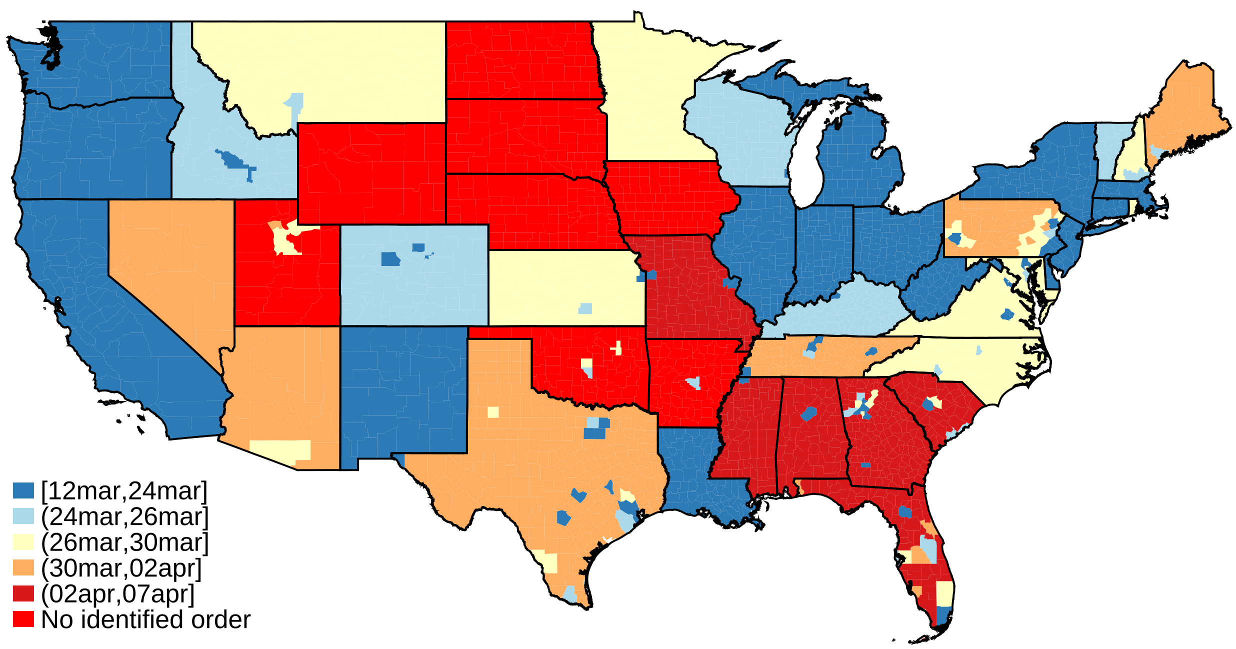


Note: This figure shows the U.S. geographic distribution of social distancing and public policy responses during the early COVID-19 pandemic. Panel A shows for each county the percent change in aggregate visits between the week beginning January 29, 2020 and the week beginning April 8, 2020. Blue shading denotes a more negative percent change in visits during the latter week relative to the former. Red shading indicates an increase or a smaller decrease in visits. These visits are sourced from SafeGraph's mobile device location data. Panel B shades U.S. counties by the effective start date for the earliest shelter-in-place order issued (see Section 2.1 for sources). Blue shading indicates an earlier order, while red shading indicates that an order was issued later or was never issued.

Appendix Figure A4: Heterogeneity in Effect of Stay-at-Home Orders on Social Distancing

*Panel A: Vs. Order Date*

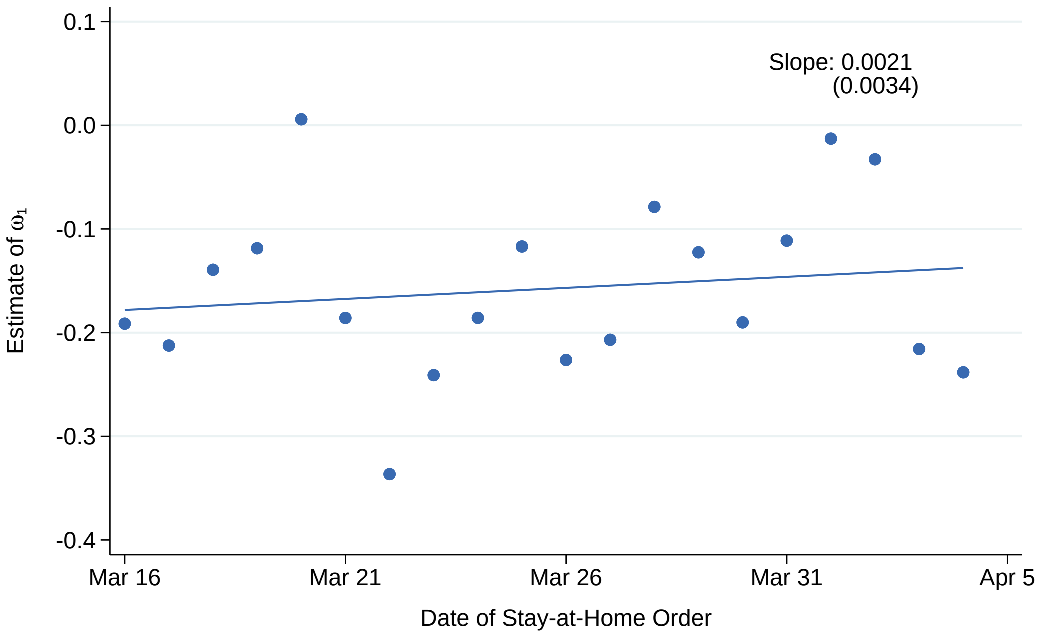


*Panel B: By CSA's Vote Share Majority*

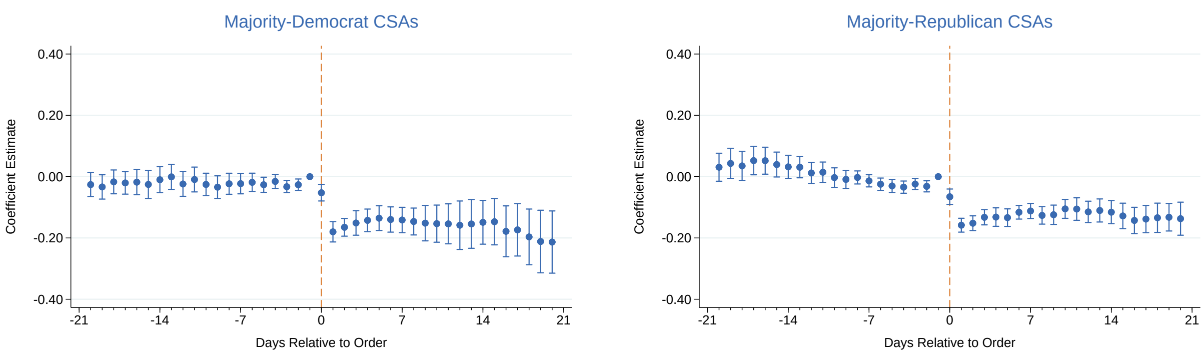


Note: Figure plots heterogeneity in estimated treatment effects across implementation dates or across geographies. Panel A plots estimated treatment effects $\omega_1$ from the event-study specification outlined in equation (1), run separately by date of the stay-at-home order, using the log of total POI visits in a CSA for a given day. Each regression is run on a sample of CSAs that either had an order issued on the day of interest, never had an order issued, or had an order issued later than April 6. As a result, each $\omega_1$ is estimated relative to the set of CSAs that never had orders or had orders after April 6. The regressions include CSA fixed effects $\mu_i$ and date fixed effects $\delta_t$. CSAs are weighted by population in the regression. Robust standard errors are used for the line of best fit. Panel B follows Figure 1 Panel A in plotting estimated treatment effects on mobility outcomes for all leads and lags, but does so separately for CSAs that consist of a majority of Democrats (left) vs. a majority of Republicans (right).

Appendix Figure A5: Effect of Mandatory Business Closure or Simultaneous Orders

*Panel A: Effect on Mobility*

Business Closure

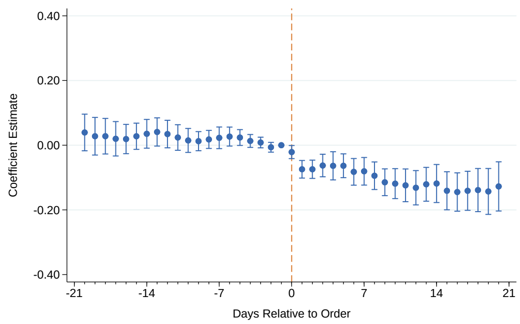


Combined Business Closure and Stay-at-Home

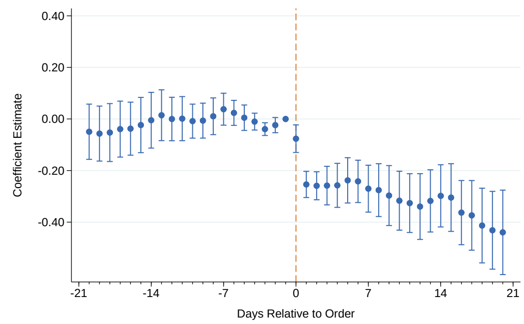


*Panel B: Effect on Consumer Spending*

Business Closure

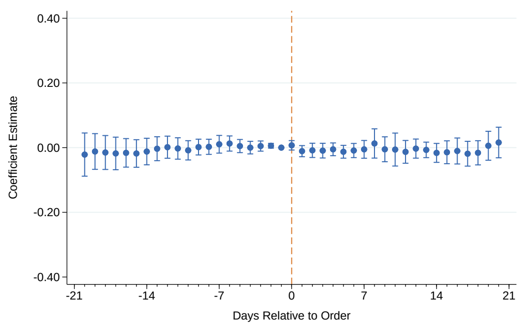


Combined Business Closure and Stay-at-Home

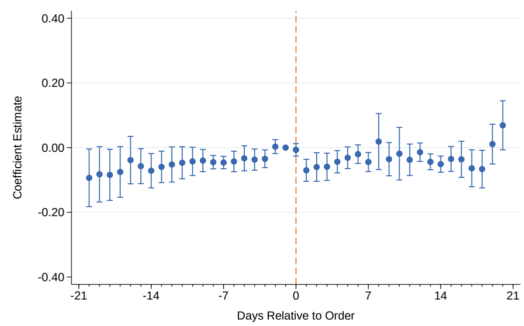


*Panel C: Effect on Employment*

Business Closure

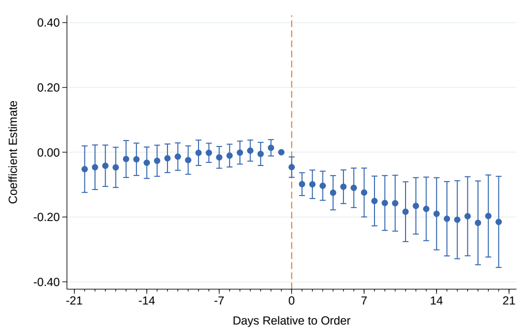


Combined Business Closure and Stay-at-Home

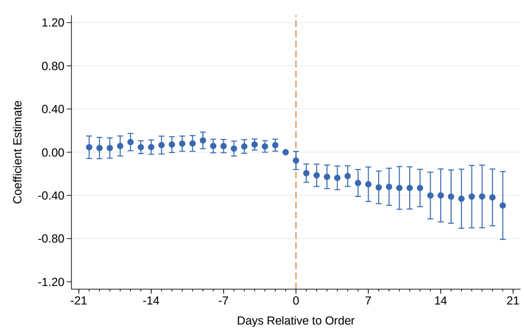


Note: Figure plots estimated treatment effects $\omega_k$ of policy orders on different outcomes, using the event-study specification outlined in equation (1) at the CSA-day level. 'Business Closure' subfigures analyze event studies of business closure orders. 'Combined Business Closure and Stay-at-Home' restricts to CSAs that implemented their business closure and stay-at home orders at the same time. Panel A shows the effect on mobility, using the log of total POI visits in the SafeGraph data. Panel B shows the effect on consumer spending, using the log of total spending in Facteus' debit card sample. Panel C shows the effect on employment, using the log number of individuals with positive work hours from the Homebase sample. All regressions include CSA fixed effects $\mu_i$ and date fixed effects $\delta_t$. CSAs are weighted by population in the regression. Standard errors are clustered at the CSA level irrespective of order timing.

Appendix Figure A6: Simulations

*Panel A: Preferred Estimator*

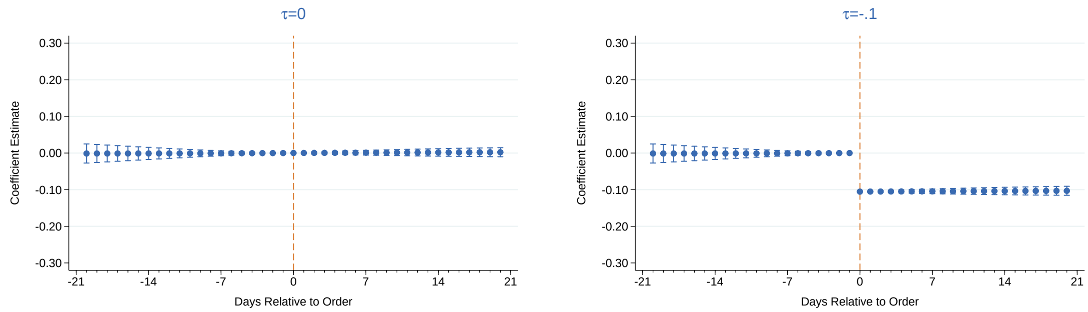


*Panel B: Preferred Estimator but not Controlling for* $\log(I_{it})$

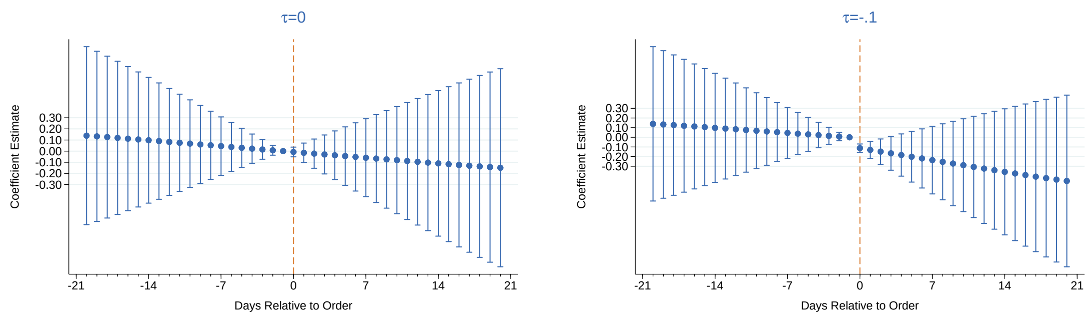


*Panel C: Log Cases Event Study with Linear Time Trends (Dave et al. 2021)*

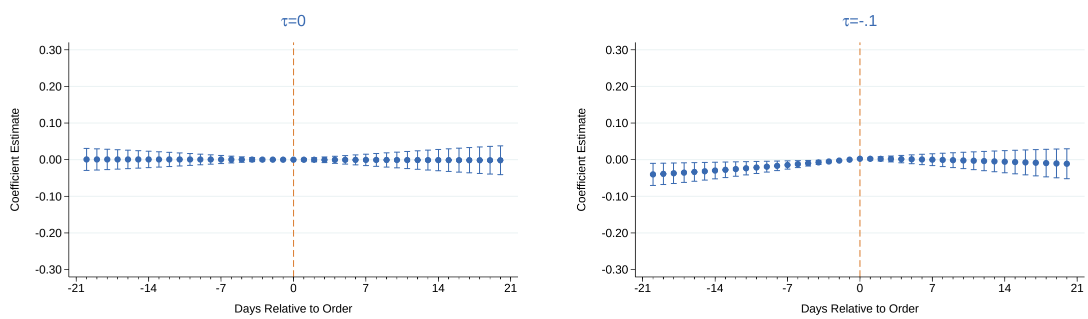


Note: Figure plots estimated treatment effects $\omega_k$ from the event-study specification outlined in equation (14) using the data simulated from an SIRD model. The first column uses data simulated from an SIRD model where the true treatment effect of the stay-at-home order is $\tau = 0$. The second column uses data simulated from an SIRD model where the true treatment effect of the stay-at-home order is $\tau = -.1$. Panel A is our preferred estimator. Panel B is the preferred estimator, but does not control for $\log(I_{it})$. Panel C is the Dave et al. (2021) estimator that uses log cases on the left-hand side and controls for geography-specific linear time trends. Geographies are weighted by population in the regression. Standard errors are clustered at the geography level. See Appendix Section A.2 for detail regarding this simulation exercise.

Appendix Figure A7: Simulations with Alternative Specifications

*Panel A: Preferred Estimator but Adding Geography FEs*

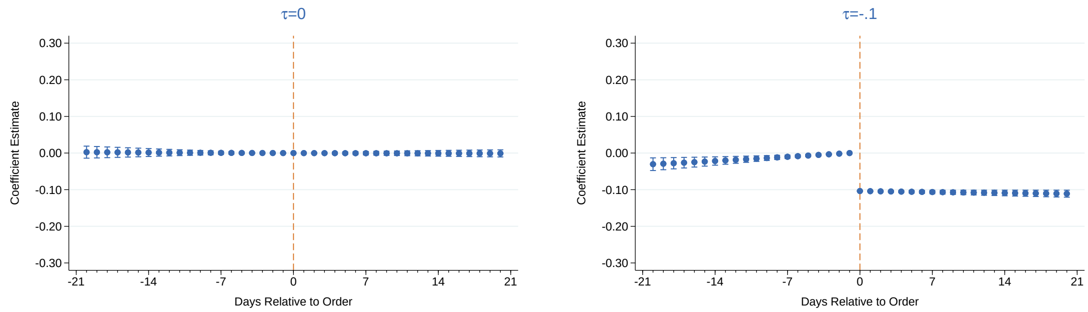


*Panel B: Preferred Estimator but Setting $\gamma$ at Half Its True Value*

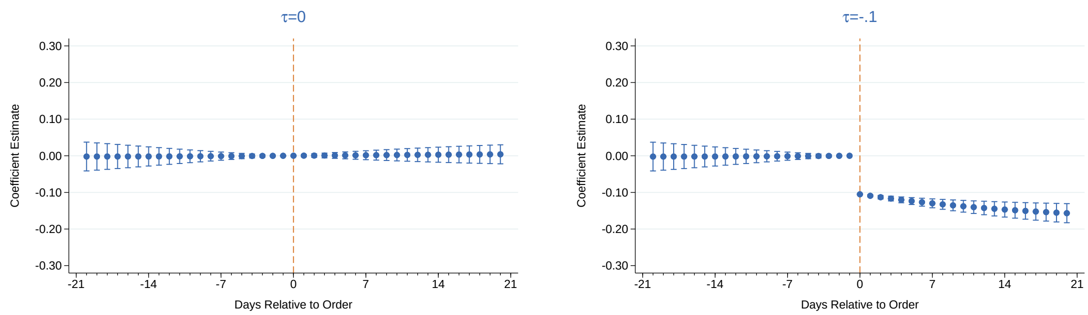


*Panel C: Preferred Estimator but Setting $\gamma$ at Twice Its True Value*

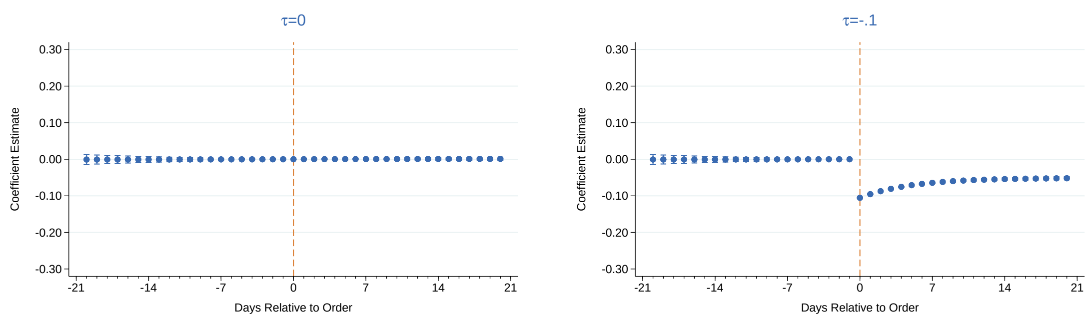


Note: Figure plots estimated treatment effects $\omega_k$ from the event-study specification outlined in equation (14) using the data simulated from an SIRD model. The first column uses data simulated from an SIRD model where the true treatment effect of the stay-at-home order is $\tau = 0$. The second column uses data simulated from an SIRD model where the true treatment effect of the stay-at-home order is $\tau = -.1$. Panel A is the preferred estimator, but adds geography fixed effects. Panel B is the preferred estimator, but sets $\gamma$ to half its true value. Panel C is the preferred estimator, but sets $\gamma$ to twice its true value. Geographies are weighted by population in the regression. Standard errors are clustered at the geography level. See Appendix Section A.2 for detail regarding this simulation exercise.

Appendix Figure A8: Dave et al. (2021) Estimators for Effect of Stay-at-Home Orders on COVID Cases

*Panel A: 7-day Preperiod as in Dave et al. (2021)*

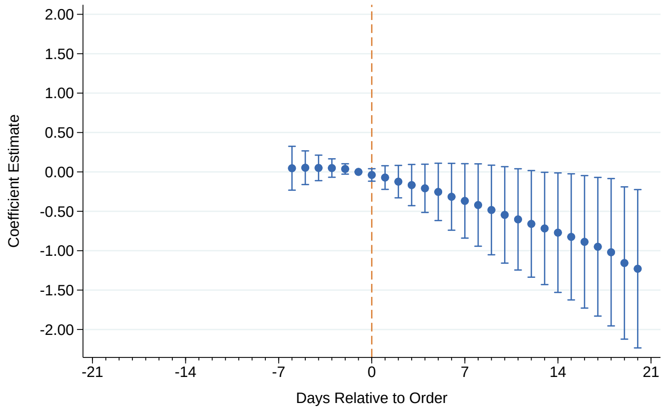


*Panel B: 14-day Preperiod*

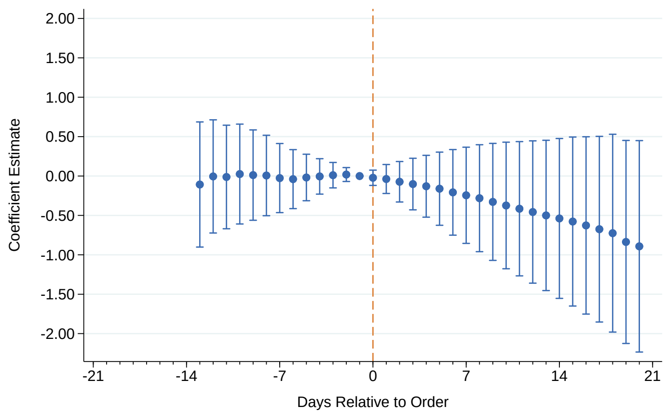


*Panel C: 21-day Preperiod*

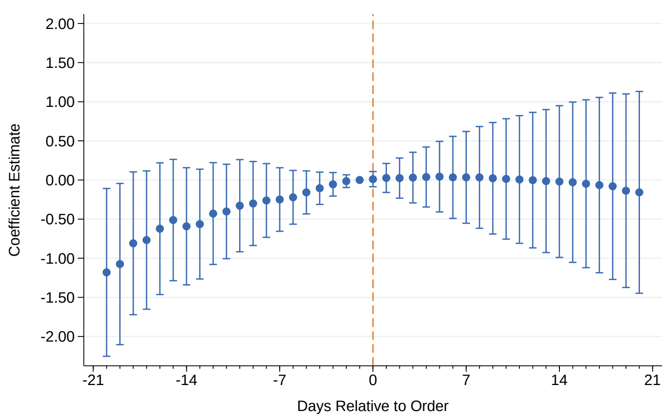


Note: Figure plots estimated treatment effects $\omega_k$ from the event-study specification outlined in equation (15) using the log of cases in a state for a given day and including state-specific linear trends. Panel A restricts the preperiod to 7 days. Panels B and C are the same as Panel A, except they use a 14- and 21-day preperiod. States are not balanced across the event window. States are weighted by population in the regression. Data is restricted to dates on or before April 20, 2020. Standard errors are clustered at the state level. See Appendix Section A.2 for detail regarding this simulation exercise.

Appendix Figure A9: Effect of Stay-at-Home Orders on Contact Rate, All $\gamma$ Values

Panel A: $\gamma = 0$

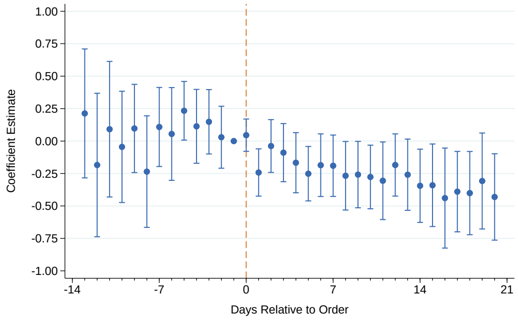


Panel B: $\gamma = 1/12$

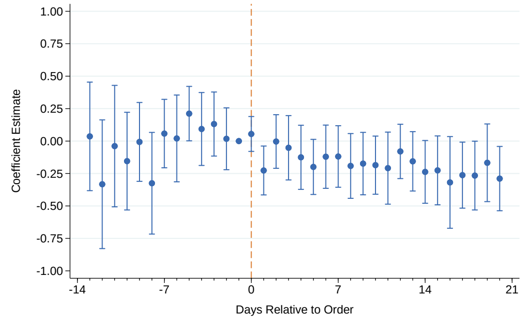


Panel C: $\gamma = 1/9$

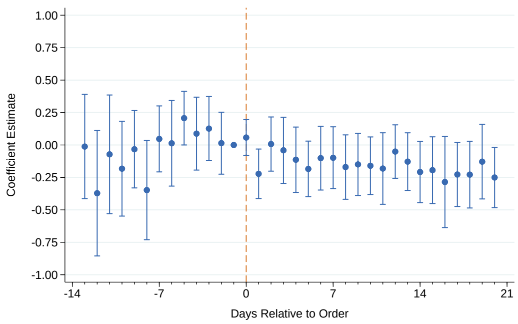


Panel D: $\gamma = 1/6$

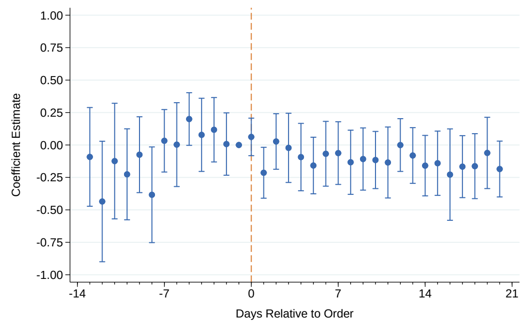


Panel E: $\gamma = 1/3$

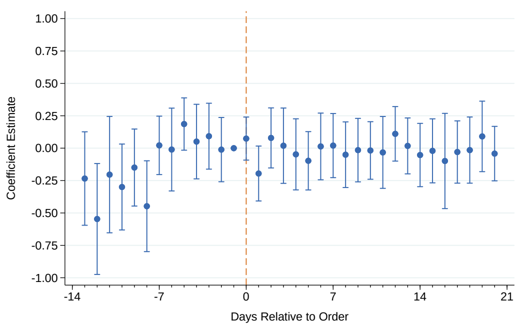


Panel F: $\gamma = 1$

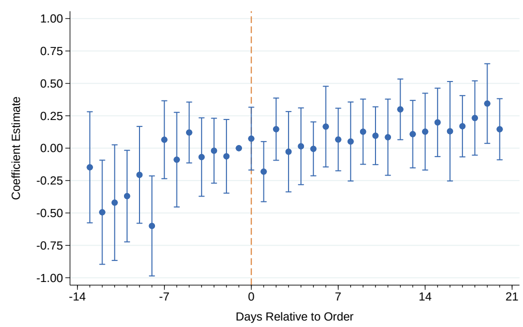


Note: Figure plots estimated treatment effects $\omega_k$ from the event-study specification outlined in equation (14) using the log of new cases in a state for a given day as in Figure 1. Each panel uses a different value of $\gamma$. CSAs are weighted by population in the regression. Standard errors are clustered at the CSA level irrespective of order timing.

Appendix Figure A10: Model Fit for Average Log Contact Rates, Full Lasso Model

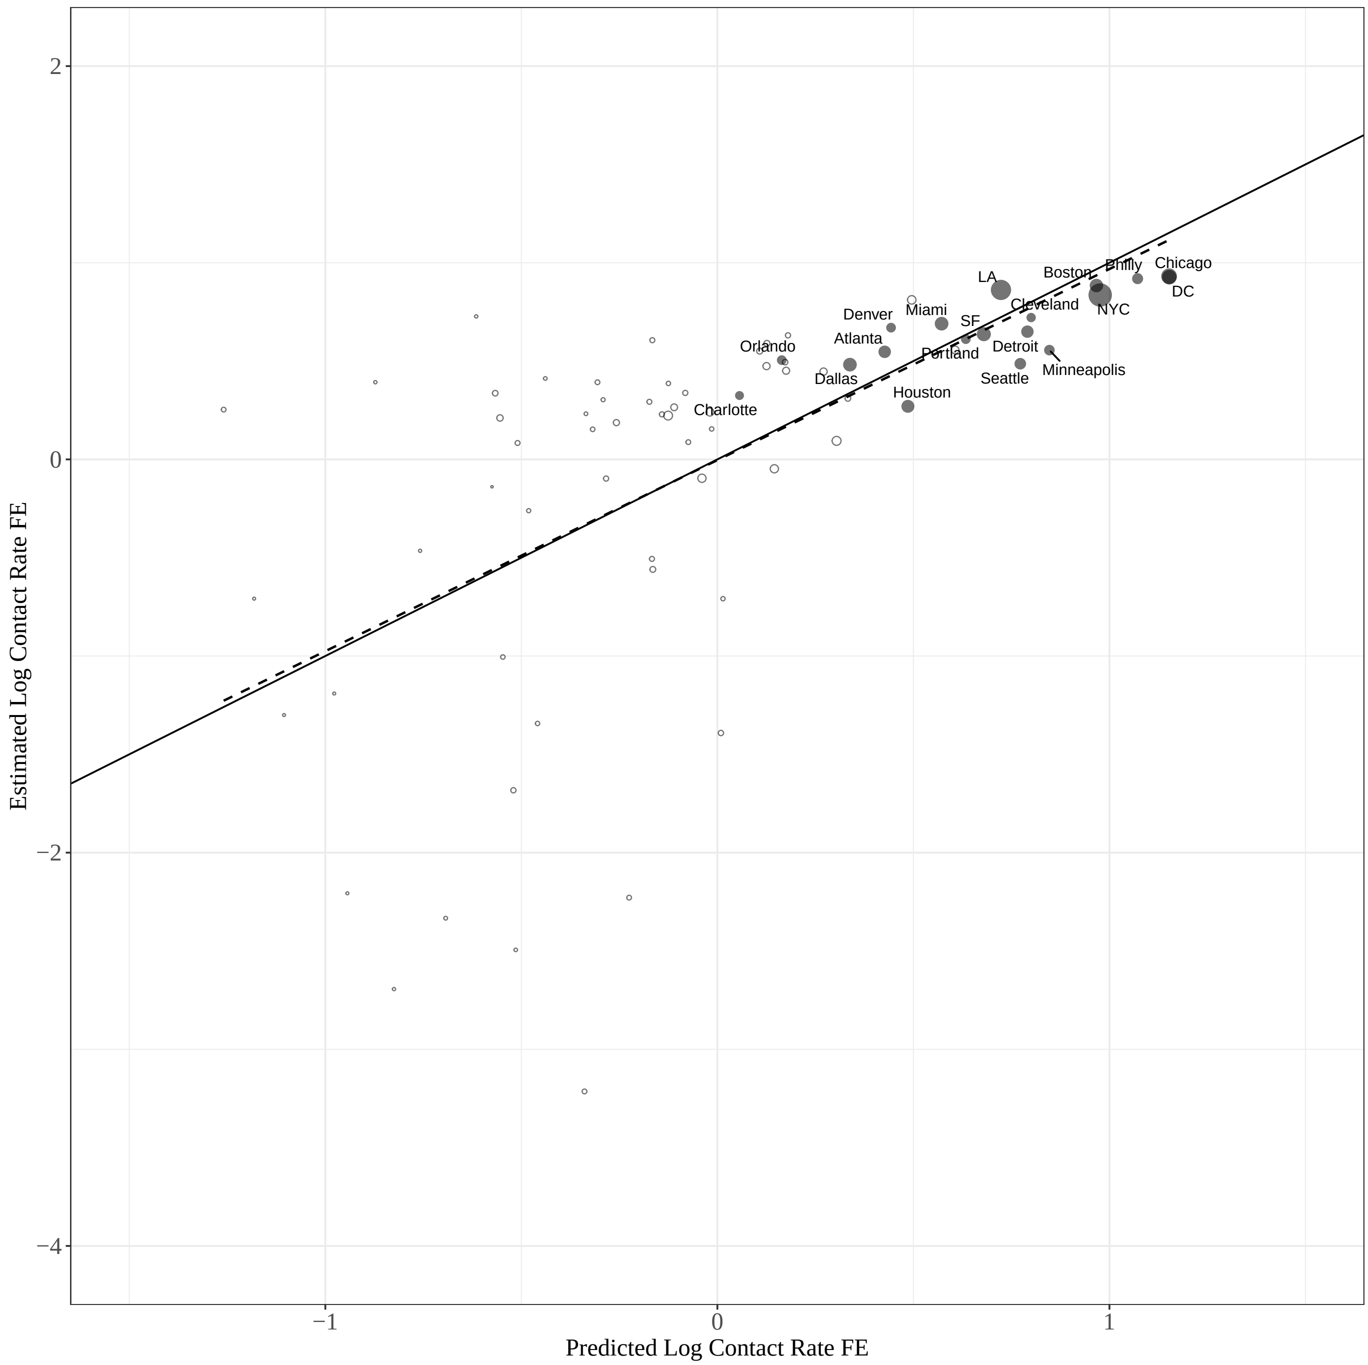


Note: Figure plots predictions of the estimated fixed differences in CSA log contact rates $\hat{\theta}_i$ in OLS regressions using $\gamma = 1/6$. The plotted points are sized proportional to CSA population and the top 25 most populous CSAs are filled in and labeled. In this figure, the fixed effects are averaged using population weights and the population is summed across order timings when a CSA has multiple order timings. The solid line is a 45 degree line indicating perfect prediction and the dashed line is a linear fit of the estimated fixed effects on their predictions for the CSAs plotted. We drop two observations when plotting to focus on the variation across the majority of CSAs.

Appendix Table A1: CSA Summary Statistics

| | All | Above Med. Contact Rate | Below Med. Contact Rate | Difference |
|---|---|---|---|---|
| Log Contact Rate ($\beta_t$, $\gamma = 1/6$) | -2.259 | -1.617 | -2.901 | **1.285** |
| | (1.230) | (0.138) | (1.482) | **[0.902, 1.710]** |
| Log Cases per 100,000 on April 30 | 5.100 | 5.403 | 4.798 | **0.605** |
| | (0.894) | (0.820) | (0.869) | **[0.298, 0.911]** |
| Log Deaths per 100,000 on April 30 | 1.738 | 2.154 | 1.308 | **0.846** |
| | (1.203) | (1.178) | (1.078) | **[0.386, 1.210]** |
| Log POI Visits | 10.545 | 11.137 | 9.953 | **1.184** |
| | (1.125) | (0.947) | (0.971) | **[0.807, 1.542]** |
| $\Delta$Log POI Visits | -0.642 | -0.658 | -0.626 | -0.032 |
| | (0.219) | (0.193) | (0.242) | [-0.106, 0.055] |
| Share of Days with Order | 0.650 | 0.697 | 0.603 | 0.094 |
| | (0.271) | (0.232) | (0.299) | [-0.004, 0.189] |
| Log Pop Density | 6.044 | 6.836 | 5.252 | **1.584** |
| | (1.501) | (1.447) | (1.083) | **[1.113, 2.056]** |
| Log Population | 13.436 | 14.114 | 12.758 | **1.356** |
| | (1.192) | (1.018) | (0.947) | **[0.967, 1.711]** |
| Share Urban | 0.821 | 0.881 | 0.761 | **0.120** |
| | (0.181) | (0.153) | (0.188) | **[0.065, 0.196]** |
| Average Summer Temperature | 88.440 | 87.480 | 89.400 | -1.920 |
| | (5.215) | (5.413) | (4.867) | [-3.634, 0.036] |
| Average Winter Temperature | 50.292 | 50.386 | 50.198 | 0.187 |
| | (12.432) | (13.712) | (11.125) | [-4.172, 4.822] |
| Share Commute Auto | 0.887 | 0.864 | 0.910 | **-0.045** |
| | (0.076) | (0.096) | (0.038) | **[-0.073, -0.019]** |
| Log International Flights | 2.759 | 4.297 | 1.222 | **3.075** |
| | (3.084) | (3.280) | (1.901) | **[1.892, 4.128]** |
| Republican Vote Share | 0.491 | 0.445 | 0.538 | **-0.093** |
| | (0.152) | (0.140) | (0.151) | **[-0.149, -0.039]** |
| Share Age $\geq 65$ | 0.145 | 0.142 | 0.147 | -0.005 |
| | (0.039) | (0.041) | (0.037) | [-0.021, 0.009] |
| Share Black | 0.157 | 0.163 | 0.150 | 0.014 |
| | (0.143) | (0.129) | (0.156) | [-0.038, 0.066] |
| Share White | 0.741 | 0.721 | 0.762 | -0.040 |
| | (0.148) | (0.135) | (0.158) | [-0.097, 0.010] |
| Share with Bachelor's or More | 0.304 | 0.328 | 0.280 | **0.047** |
| | (0.090) | (0.077) | (0.095) | **[0.015, 0.080]** |

Note: Table reports summary statistics of the average CSA covariate values between March 15 and April 30, except for cases and deaths for which the April 30th values are used. The means across all CSAs and the grouped CSAs by average log contact rate $\log(\beta_t)$ using $\gamma = 1/6$ are reported along with bootstrapped standard errors. The difference-in-means between the top and bottom contact rate groups are reported along with bootstrapped 95 percent confidence intervals using 100,000 random draws. Statistically significant difference-in-means are bolded. '$\Delta$Log Visits' reports the change in log POI visits relative to March 1, 2020. 'Share of Days with Order' reports the share of the March 15–April 30 time period in which a CSA is under a stay-at-home order.

Appendix Table A2: Sources of Non-pharmaceutical Interventions

| | Stay-at-home | | Business closures | |
|---|---|---|---|---|
| | Unweighted | Weighted | Unweighted | Weighted |
| Inherited from State | 0.898 | 0.464 | 0.923 | 0.629 |
| NYT | 0.895 | 0.638 | 0.000 | 0.000 |
| Keystone Strategy | 0.016 | 0.078 | 0.949 | 0.822 |
| Crowdsourced | 0.038 | 0.216 | 0.033 | 0.139 |
| Hikma Health | 0.018 | 0.057 | 0.016 | 0.038 |
| Manual Entry | 0.034 | 0.012 | 0.000 | 0.000 |

Note: Table summarizes the source of county stay-at-home and business closure policies. We report shares of county policies both unweighted and weighted by county population. The first row indiciates the share of county policies that are inheritied from the state; that is, counties did not enact the corresponding policy before the state took action. The remaing columns indicate the share of county policies coming from each of the our sources. The manual entry source is reserved for corrections to state policies which we hand checked. We only had to recode the Tennessee state policy.

Appendix Table A3: Role of Policy in Explaining Aggregate Temporal Changes in Distancing, Economic, and Health Outcomes

| *Panel A: Social Distancing and Economic Outcomes* | | | | | | |
|---|---|---|---|---|---|---|
| | (1) Stay-at-Home | | (2) Business | | (3) Both Orders | |
| | Total | Policy | Total | Policy | Total | Policy |
| POI Visits | -0.672 | -0.109 | -0.665 | -0.060 | -0.700 | -0.178 |
| | | (0.008) | | (0.011) | | (0.018) |
| Homebase Wages | -0.590 | -0.092 | -0.577 | -0.105 | -0.582 | -0.183 |
| | | (0.013) | | (0.023) | | (0.039) |
| Homebase Employment | -0.592 | -0.095 | -0.584 | -0.093 | -0.620 | -0.170 |
| | | (0.012) | | (0.017) | | (0.038) |
| Facteus Debit Transactions | 0.048 | -0.055 | 0.056 | -0.017 | -0.003 | -0.057 |
| | | (0.007) | | (0.007) | | (0.013) |
| Facteus Total Spending | 0.284 | -0.071 | 0.302 | -0.012 | 0.195 | -0.069 |
| | | (0.009) | | (0.010) | | (0.017) |

| *Panel B: Contact Rate $\beta_t$* | | | | | | |
|---|---|---|---|---|---|---|
| | (1) $\gamma = 1/12$ | | (2) $\gamma = 1/9$ | | (3) $\gamma = 1/6$ | |
| | Total | Policy | Total | Policy | Total | Policy |
| $\mathscr{R}_0 = 2.7$ | -0.574 | -0.245 | -0.610 | -0.172 | -0.634 | -0.098 |
| | | (0.107) | | (0.078) | | (0.052) |
| $\mathscr{R}_0 = 3.0$ | -0.617 | -0.221 | -0.649 | -0.155 | -0.670 | -0.088 |
| | | (0.096) | | (0.070) | | (0.047) |
| $\mathscr{R}_0 = 3.3$ | -0.652 | -0.201 | -0.681 | -0.141 | -0.700 | -0.080 |
| | | (0.087) | | (0.064) | | (0.042) |

Note: Table reports the total and policy-induced changes in various outcomes as outlined in Section 5.1. In Panel A, we consider the policy-induced changes of stay-at-home orders, business closures, and simultaneous stay-at-home and business closures. Each column restricts attention to the set of CSAs that receive a given treatment, with Column (3) restricting to counties in which business closure and stay-at-home orders went into effect on the same day. The treatment effects $\omega_k$ used in Panel A set $k = 1$. In Panel B, all specifications estimate the effect of stay-at-home orders using estimates from the CSA level and use the estimated treatment effect $\omega_k$ from Table 1.

Appendix Table A4: Decomposing Changes in Distancing, Economic, and Health Outcomes

*Panel A: Stay-at-Home Orders*

| | | (1) | (2) | (3) |
|---|---|---|---|---|
| | | $k=1$ | $k=20$ | Pre-Trend |
| | Total Δ | Policy-Induced Δ | | |
| POI Visits | -0.672 | -0.109 | -0.165 | -0.102 |
| | | (0.008) | (0.026) | (0.028) |
| Homebase Wages | -0.590 | -0.092 | -0.197 | -0.039 |
| | | (0.013) | (0.047) | (0.049) |
| Homebase Employment | -0.592 | -0.095 | -0.206 | -0.067 |
| | | (0.012) | (0.051) | (0.053) |
| Facteus Debit Transactions | 0.048 | -0.055 | -0.071 | -0.023 |
| | | (0.007) | (0.012) | (0.014) |
| Facteus Total Spending | 0.284 | -0.071 | -0.034 | 0.006 |
| | | (0.009) | (0.020) | (0.024) |

*Panel B: Business Closure Orders*

| | | (1) | (2) | (3) |
|---|---|---|---|---|
| | | $k=1$ | $k=20$ | Pre-Trend |
| | Total Δ | Policy-Induced Δ | | |
| POI Visits | -0.665 | -0.060 | -0.102 | -0.060 |
| | | (0.011) | (0.031) | (0.033) |
| Homebase Wages | -0.577 | -0.105 | -0.273 | -0.353 |
| | | (0.023) | (0.093) | (0.096) |
| Homebase Employment | -0.584 | -0.093 | -0.203 | -0.262 |
| | | (0.017) | (0.068) | (0.069) |
| Facteus Debit Transactions | 0.056 | -0.017 | -0.016 | -0.006 |
| | | (0.007) | (0.021) | (0.023) |
| Facteus Total Spending | 0.302 | -0.012 | 0.017 | -0.012 |
| | | (0.010) | (0.026) | (0.029) |

Note: Table reports the total and policy-induced changes in various outcomes as outlined in Section 5.1 using alternative estimators of the treatment effect. In Panel A, we consider the policy-induced changes of stay-at-home orders. In Panel B, we consider the policy-induced changes of business closure orders. For $k = 1$ and $k = 20$, we use the treatment effects $\omega_k$ with corresponding value $k$. The 'Pre-Trend' estimator uses the trend in estimates in the two weeks leading up to treatment to adjust the treatment effect $\omega_k$ for $k = 20$.

Appendix Table A5: Variance Decomposition for Differences in Log Contact Rates

| | |
|---|---|
| Cross-CSA Variance of Log Contact Rate | 1.514 |
| *Share of variance explained by:* | |
| Social Distancing | 0.008 |
| Policy | 0.000 |
| Timing of Virus | 0.001 |
| Observed Covariates | 0.264 |
| Population | 0.182 |
| Climate | 0.000 |
| Transport | 0.000 |
| Race | 0.032 |
| Partisanship | 0.031 |
| College Degrees | 0.000 |
| Age Demographics | 0.000 |

Note: Table reports the cross-CSA variance of the average log contact rate $\log(\beta_{it})$ between March 15 and April 30, 2020 in the first row. We calculate the cross-CSA variance of each explanatory variable and report the share of the log contact rate variance accounted for by the variation in each set of explanatory variables using the estimated coefficients from our lasso model in Section 5.2.