COVID-19 and Stay-At-Home Orders: Identifying Event Study Designs with Imperfect Testing

Jaedo Choi Elird Haxhiu Thomas Helgerman Nishaad Rao Taeuk Seo*

April 16, 2021

Abstract

This paper estimates the dynamic effect of Stay-At-Home (SAH) orders on the transmission of COVID-19 in the United States. Identification in this setting is challenging due to differences between real and reported case data given the imperfect testing environment, as well as the clearly non-random adoption of treatment. We extend a Susceptible-Infected-Recovered (SIR) model from Epidemiology to account for endogenous testing at the county level, and exploit this additional structure to recover identification. With the inclusion of model-derived sufficient statistics and fixed effects, SAH orders have a large and sustained negative effect on the growth of cases under plausible assumptions about the progression of testing. Point estimates range from a 44% to 54% reduction in the growth rate of cases one month after a SAH order. We conclude with a discussion on extending the methodology to later phases of the pandemic.

*Department of Economics, University of Michigan, Ann Arbor, MI. Contact: Jaedo Choi (jaedohi@umich.edu), Elird Haxhiu (haxhiu@umich.edu), Thomas Helgerman (tehelg@umich.edu), Nishaad Rao (nsrao@umich.edu), and Taeuk Seo (taeuks@umich.edu). We thank SafeGraph for making data available for research. We thank seminar participants at the Causal Inference Reading Group, Labor Lunch, and Business Economics Lunch at the University of Michigan as well as Charles Brown, Andrew Goodman-Bacon, Florian Gunsilius, and Justin Wolfers for valuable comments and feedback.

1 Introduction

It is difficult to exaggerate the severity of the COVID-19 pandemic in the United States. After "flattening the curve" three separate times since April, new daily cases are once again on the rise (see Figure 1). Even as we approach a spring season with cautious optimism and a set of vaccines on hand, it is important to understand how effective existing policy tools are at stopping the virus given uncertainty surrounding both vaccine take-up and protection offered by the shot against newer variants of COVID-19.[1] One policy used widely during the initial propagation of the disease was the Stay-At-Home (SAH) order, where states require citizens to remain at home unless they need to leave for an essential activity.[2] While the aggregate time-series evidence suggests that these policies were effective at arresting the growth rate of new cases, credible causal estimation of treatment effects on disease spread remains challenging due to a number of identification challenges.

The most difficult of these is related to non-classical measurement error: we do not expect *actual* COVID-19 cases to match those *reported* by local health authorities. We could use standard econometric methods if this measurement error is classical, but as infections are almost surely systematically under-reported, we cannot reasonably maintain that this error is mean zero. Overcoming this issue therefore requires a theoretical framework relating observed positive cases to actual infections.[3] We augment a standard Susceptible-Infected-Recovered (SIR) model from Epidemiology with endogenous testing and derive assumptions about variation in testing capacity over time and across space that must be made in order to conduct valid inference. The SIR model governs the evolution of cases and disciplines all of our estimating equations, while the theory on testing reveals how different fixed effects amount to substantive assumptions about testing capacity.

The second challenge to credibly identifying the effect of SAH orders on disease spread is the non-random adoption of treatment. All else equal, states with higher rates of spread are more likely to implement the policy on a given day. Simple comparisons then suggest that SAH orders *increase*

[1] This is especially true given the slow roll-out of the vaccine in countries other than the United States.

[2] Starting on March 19th with California and concluding with South Carolina on April 7th, 42 states enacted a SAH order during the sample period ending on April 30th. We maintain an absorbing treatment assumption throughout, and show robustness to dropping Montana and Colorado which removed their respective SAH orders a number of days before the last day in our sample.

[3] Yang et al. (2020) estimates the active prevalence of COVID-19 based on an SIR model that allows for the fact that testing is skewed towards symptomatic individuals. They find that accounting for this reveals that the prevelance of COVID-19 could be up to three times higher in the United States, highlighting the importance of considering testing and how it relates to caseload data. However, the paper focusses on the extent of spread, not the effect that a policy such as stay-at-home orders would have on this prevalence.

the spread of COVID-19, an implausible conclusion given our knowledge of disease transmission. The recursive structure of the SIR model implies a sufficient statistic that captures this underlying heterogeneity: the lag of cumulative cases is sufficient to determine the current evolution of daily cases, in the absence of policy. A parallel trends assumption on the evolution of daily cases across different cohorts of SAH adopters then identifies potential dynamic effect of SAH orders. Pre-trends tests in our model-derived estimating equations can then be used to assess the validity of this parallel trends assumption in the *post*-period, which we find strong support for. [4]

Figure 1: New Cases of COVID-19 Reported to the CDC

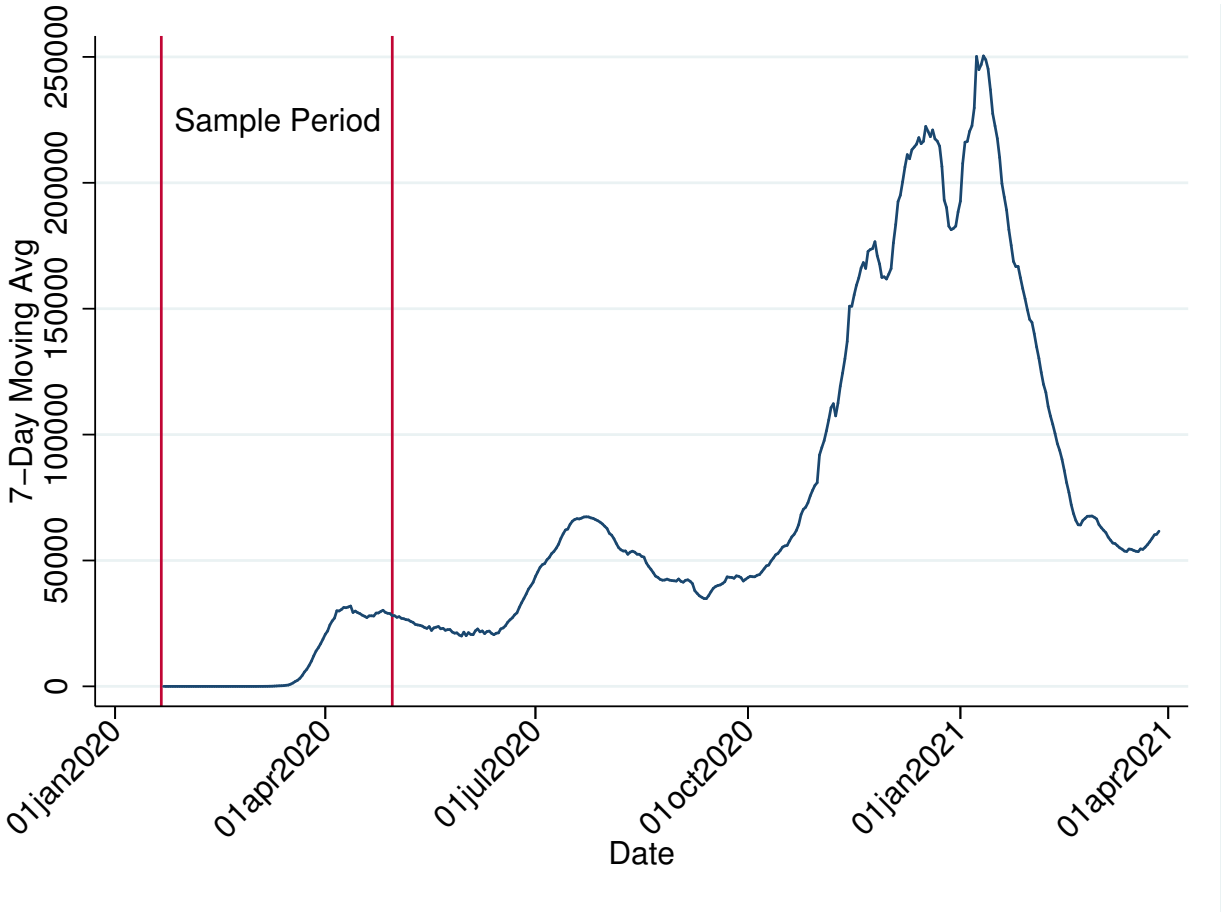

This figure plots the 7-Day Moving Average of daily new cases of COVID-19 reported in the U.S. Data was collected from the CDC on 3/30/2021 at

`https://www.cdc.gov/coronavirus/2019-ncov/cases-updates/cases-in-us.html`

Despite these unique challenges to identification, there has been an explosion of economic re-

[4]This is subject to the additional assumption that dynamic treatment effects associated with a SAH do not vary *too much* across different cohorts adopting treatment (Sun and Abraham (2020)). We believe this is reasonable given the short time period between the first and last adopter (20 days), but estimates assessing robustness are forthcoming.

search focusing on the effectiveness of policy responses: as of the time of writing, more than 380 pandemic related working papers have been uploaded to NBER; and *COVID Economics* has published 468 papers on the economics of COVID-19, with at least 20 focusing specifically on the effects of lock-down measures. Early work based on US data includes Lyu and Wehby (2020), Dave et al. (2020) and Fowler et al. (2020), which all focus on the treatment effect of SAH orders, but each use specifications with differing outcome variables and fixed effects. Each find significant effects with lock-downs, but notable pre-trends suggest the parallel trends identifying assumption is suspect, at least as specified.[5] Similar specifications have been used to compare policies across countries using cross-country regression analysis (Alfano and Ercolano (2020), Bonardi et al. (2020)). Research since then has expanded to look at the effects of different policies; for example, Isphording et al. (2020) study the impact of public health informed school re-openings in Germany on the spread of COVID-19. Further, Schlosser et al. (2020) focus on how geographic mobility interacts with SAH orders to reduce spread.[6]

Related work to ours includes Allcott et al. (2020), which derives event study specifications from the SIR model in a similar fashion. However, we explicitly deal with the problem that observed data is the endogenous outcome of testing and show how different sets of fixed effects are structurally related to different assumptions on testing capacity. Some practitioners focus on simulating the SIR model to derive estimates of effect size (as in Giordano et al. (2020)), while we estimate it using a specification from the model and minimal data requirements. We derive the *additional* assumptions on the evolution of testing capacity over time and across space needed to properly interpret fixed effects estimators of this effect.

Finally, Chernozhukov et al. (2020) construct an SIR model that is similar in spirit to ours but written in continuous time. They find an expression for the growth rate of confirmed cases, which they approximate using a discrete time difference equation. They proceed by assuming that the growth of testing capacity[7] is a linear function of the growth of the number of tests administered. Their equation (10) restricts the growth rate in the number of confirmed cases to be positively

[5] Their estimates naturally vary in size and interpretation: for instance, Fowler et al. (2020) find that the infections declined by 30% in the first week after the lock-down, while Lyu and Wehby (2020) estimate that there is a difference of 0.51 per 10,000 resident in cases after imposing lock-down in a state compared to its neighbor.

[6] They find that counties become more disconnected once they impose lock-downs, and there is a significant reduction in mobility, which also leads to a decline in disease spread.

[7] Defined here as the percentage of true cases reported by local public health officials.

related to the corresponding growth rate in tests administered in a linear fashion. We pursue an alternative approach in our main specifications by making assumptions on the *progression* of testing capacity itself; this amounts to restricting the variation in confirmed cases we permit to identify the effect of SAH on disease transmission. As such, these approaches are complementary ways of solving the challenge of making inference about real COVID-19 cases from reported ones.

We proceed as follows. Section 2 documents our data sources and implements the standard two-way fixed effects (TWFE) estimator, while section 3 derives alternative estimating equations based on the SIR model augmented with endogenous testing. This section also presents and discusses our main finding, that SAH orders adopted early in the pandemic had a strong and significant effect on curbing the spread of disease, as well as the identifying assumptions on testing and parallel trends needed to believe it. We conclude in section 4 with some extensions and robustness checks, as well as a discussion of how to interpret and adapt our framework during later phases of the pandemic.

2 Standard Event Study

At its core, this paper is concerned with the research question: "What is the effect of Stay At Home Orders on the spread of the COVID-19?" On face, this is a policy evaluation that seems to be easily answerable with standard econometric tools. This section attempts to evaluate the effect of SAH orders using a TWFE event study approach and illustrates its pitfalls in this context.

2.1 Data

To measure the spread of SARS-CoV-2[8], we collect county level data on both the number of positive cases of and deaths attributable to COVID-19 compiled by *The New York Times*. The data begins on January 21st, 2020, with the first reported case in the United States, and includes the cumulative number of cases and deaths for each county on each day through April 30th, at the time of writing. To compile this dataset, *The New York Times* collected historical information from local and state governments and health departments; as a result, these data are subject to the same limitations as this source material. First, there are cases in which state reports do not report cases separately by county, or the county of residence of an individual is simply reported as "unknown." Second,

[8]Throughout, we interchangeably use the terms: COVID-19, SARS-CoV-2, the coronavirus, and the pandemic.

and more importantly, these data should be interpreted as the incidence of COVID-19 conditional on the level of testing at the county level. Preliminary estimates suggest that a large fraction of positive cases remain undetected (Jagodnik et al. (2020)), which we tackle explicitly in Section 4.

We collect the date each state implemented a SAH order from the *New York Times*.[9] In our data set, we count 42 states as implementing SAH orders (as well as Washington, D.C.); the eight that do not are Arkansas, Iowa, Nebraska, North Dakota, Oklahoma, South Dakota, Utah and Wyoming. Importantly, we do not count Oklahoma's "Safer-At-Home" order as an SAH order, as it only applied to the older and more vulnerable to serious infection. In addition, while several localities within Oklahoma, Utah and Wyoming issued their own SAH orders, we do not consider them as treated, as these policies were not implemented state-wide.

2.2 Empirical Specification

We observe C counties for T days in our data set. We define the Stay-At-Home event for county c, E_c, as the day on when the state in which that county resides imposed a stay at home order. Then, we can utilize the TWFE event study specification:

$$Y_{c,t} = \sum_{l=-7, l\neq -1}^{l=28} \mu_l \cdot \mathbb{1}\{t - E_c = l\} + \mu_{29+} \cdot \mathbb{1}\{t - E_c \geq 29\} + \alpha_c + \gamma_t + \varepsilon_{c,t} \tag{1}$$

$Y_{c,t}$ is an unspecified outcome variable, as it is *ex-ante* unclear what transformation of positive cases is appropriate to use. In each county c, we count the total number of people who tested positive *by* day t, $T_{c,t}$. Then, to measure the daily number of new cases, we take first differences $\triangle T_{c,t} := T_{c,t} - T_{c,t-1}$. We set $Y_{c,t} = \ln\{\triangle T_{c,t} + 1\}$, noting that $\triangle T_{c,t}$ grows in an exponential manner during the beginning of a pandemic.

We use an event window of 1 week before the SAH order through 4 weeks after, omitting the day before the event. In addition, we include a term that bins together all days after 4 weeks to capture "long run" effects of treatment, though in practice for most states the data do not extend far beyond our event window. Finally, we include county and time fixed-effects to account for fixed differences across counties in spread and national trends in cases, respectively.

[9] Data collected at `https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html`

Figure 2: Standard TWFE Event Study Results

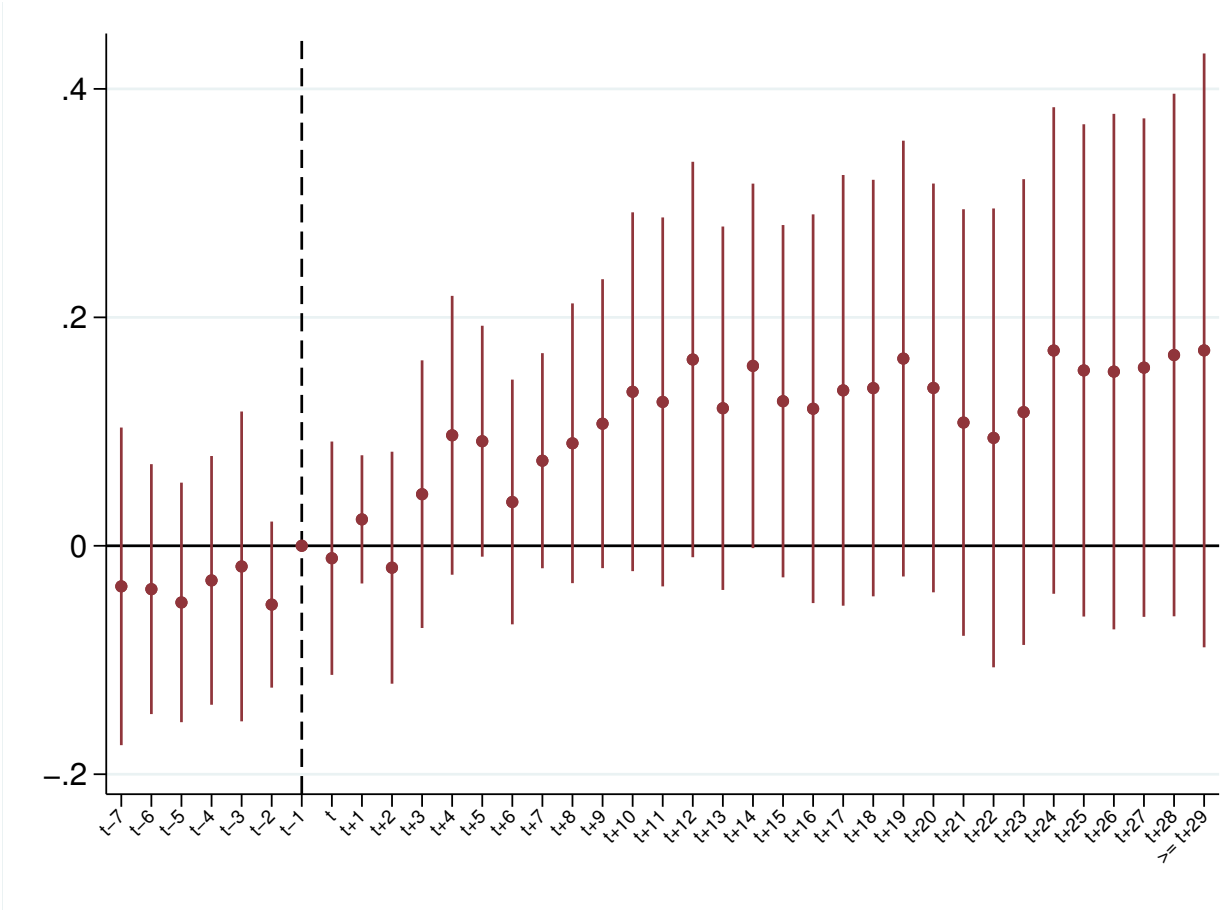

Figure plots point estimates and 95 percentile confidence interval of Equation (1). The dependent variable is the daily number of new infections. The regression model includes time and county fixed effects. Standard errors are clustered at the state level.

2.3 Results & Discussion

Figure 2 reports the results from estimating equation (1). The point estimates are noisy, especially around a month following the SAH order, but it is clear that this model estimates a weakly positive impact of SAH orders on positive cases.

These results strike us as suspect for two reasons. First, it seems extremely unlikely *ex-ante* that SAH orders increase the spread of SARS-CoV-2. A null estimate seems perfectly plausible, as it is unsettlingly easy to imagine Americans roundly ignoring this public health directive, reducing SAH orders to a placebo treatment in effect. A positive coefficient estimate seems more consistent with the interpretation that treatment is not randomly assigned conditional on the covariates included; rather, it is likely that the imposition of an SAH order is correlated with the growth rate of positive

cases in a state. This interpretation is bolstered by the second problem with these results: there is a clear pre-trend in the event coefficients preceding the SAH orders. This suggests that counties receiving earlier treatment are experiencing faster case growth than those receiving later treatment (or no treatment at all), which casts serious doubt on interpreting these results as causal.

In light of these issues, we focus on correctly modeling the expected spread of SARS-CoV-2 in the absence of treatment. Adding structure to our estimation strategy reduces dependence on the faulty assumption of random timing of SAH orders. To achieve this, we augment an epidemiological model with endogenous testing to inform which variables explain the unmitigated spread of the virus and how these variables enter the conditional expectation. Then, under the assumption that the restrictions imposed by our model are correct, we are able to recover causal identification.

3 Causal Evidence

3.1 SIR Model

To understand how we expect COVID-19 to spread, we utilize a simple discrete time SIR model.[10] An SIR model divides the population in each county c at time t into three groups: Susceptible ($S_{c,t}$), Infected ($I_{c,t}$) and Recovered ($R_{c,t}$). It then describes how the stock of these groups will change over time as a function of two key parameters. The model is described completely by the following set of difference equations:

$$\text{Susceptible} \quad : \quad S_{c,t} = S_{c,t-1} - \beta_{c,t} \cdot I_{c,t-1} \frac{S_{c,t-1}}{N_c} \tag{2}$$

$$\text{Infected} \quad : \quad I_{c,t} = (1-\gamma) \cdot I_{c,t-1} + \beta_{c,t} \cdot I_{c,t-1} \frac{S_{c,t-1}}{N_c} \tag{3}$$

$$\text{Recovered} \quad : \quad R_{c,t} = R_{c,t-1} + \gamma \cdot I_{c,t-1} \tag{4}$$

The dynamics of the model are determined by the parameters $(\beta_{c,t}, \gamma)$. $\beta_{c,t}$ describes the rate of infection in county c at time t under the assumption of random mixing: in period t, every infected person from last period $I_{c,t-1}$ spreads the disease to $\beta_{c,t}$ people at random, of which the portion $\frac{S_{c,t-1}}{N_c}$ are susceptible, leading to $m_{c,t} := \beta_{c,t} I_{c,t-1} \frac{S_{c,t-1}}{N_c}$ new infections in period t. At the same time, some proportion γ of $I_{c,t-1}$ recover, while $(1-\gamma) I_{c,t-1}$ individuals stay infected in period t.

[10] Atkeson (2020) provides an excellent introduction to a continuous time SIR model for economists.

We can now consider how to identify policy effects separately from the county-specific growth rate $\beta_{c,t}$. We do this by taking logs and linearizing the model. We show that without making further restrictions on $\beta_{c,t}$ we cannot identify policy effects; in light of this, we outline two identification strategies under different restrictions on the progression of $\beta_{c,t}$.

3.2 Identifying Policy Effects

To begin, suppose that all of the model variables are observable. In this case, we can follow the derivations in Gupta et al. (2020). Early in an epidemic, it is reasonable to assume that almost all of the population is susceptible to the disease, i.e. $S_{c,t}/N_c \approx 1$. Taking the log of new cases yields

$$\ln(m_{c,t}) = \ln(\beta_{c,t}) + \ln(I_{c,t-1}) \tag{5}$$

Note that we can do this in a general sense only because of the recursive nature of the problem, i.e. because this relationship holds for all t within the framework of our model. New cases today are *only* a function of the growth rate of cases in a county, $\beta_{c,t}$, and the number of infected people in that county yesterday, $I_{c,t-1}$. The former is a structural parameter, and the latter is a variable that summarizes the state of the world going into the current period. This is the sense in which we consider infections yesterday a "sufficient statistic" to determine cases today.[11]

Of course, we do not believe that this equation holds exactly in practice; rather, we assume it is modeling the conditional mean of $\ln(m_{c,t})$ in the absence of policy. Formally, $\mathbb{E}[\ln(m_{c,t})|I_{c,t-1}] = \ln(\beta_{c,t}) + \ln(I_{c,t-1})$, and we define $\varepsilon_{c,t} := \ln(m_{c,t}) - \mathbb{E}[\ln(m_{c,t})|I_{c,t-1}]$, allowing us to rewrite 5 as

$$\ln(m_{c,t}) = \ln(\beta_{c,t}) + \ln(I_{c,t-1}) + \varepsilon_{c,t}$$

Now, let $\mu_{s(c),t}$ denote the impact of some state-level SAH order on the spread of the virus. Specifically, we model this treatment as lowering the growth rate of the virus by a fixed proportion in every county c in state s at time t. Then, the effective growth rate in county c at time t will be

[11] A parallel can be drawn here with the motivation for using lagged variables as IVs in the macroeconomics literature. Current variables are only affected by structural parameters through state variables, making the exclusion restriction clear, at least in theory. State variables are also uncorrelated with unobservables that might affect the control variables by design, satisfying exogeneity. Likewise, the control variable in our case, i.e. the number of cases today, is determined completely by the number of infected cases yesterday and a structural parameter.

$\beta_{c,t}\mu_{s(c),t}$, and the progression of daily cases will be given by

$$\ln(m_{c,t}) = \ln(\beta_{c,t}) + \ln(\mu_{s(c),t}) + \ln(I_{c,t-1}) + \varepsilon_{c,t}$$

Proposition: State-level policy effects are not identified in the fully general SIR model.

Proof: Let $\{\beta_{c,t}, \mu_{s(c),t}, \gamma, F^e_{c,t}\}$ describe the progression of cases, where $\varepsilon_{c,t} \sim F^e_{c,t}$. Define $\beta'_{c,t}$ such that $\ln(\beta'_{c,t}) = \ln(\beta_{c,t}) + \ln(\mu_{s(c),t})$ and let $\mu'_{s(c),t} = 1$. Then, $\{\beta'_{c,t}, \mu'_{s(c),t}, \gamma, \varepsilon_{c,t}\}$ leads to the same distribution of observed daily cases.

In light of proposition 1, we know that we need to impose structure on the structural parameters $\beta_{c,t}$ in order to identify policy effects. We first consider the simplest assumption - that this growth rate does not vary over time:

$$\beta_{c,t} = \beta_c \qquad \text{(Assumption 1B)}$$

Under Assumption 1B, we can identify the treatment effect of a SAH order using a panel event study approach with only county fixed effects $\alpha_c = \ln(\beta_c)$:[12]

$$\ln(m_{c,t}) = \sum_{l=-7, l\neq -1}^{l=28} \mu_l \cdot \mathbb{1}\{t - E_c = l\} + \alpha_c + \ln(I_{c,t-1}) + \varepsilon_{c,t}$$

Unfortunately, this estimation is infeasible, as we are unable to observe neither $m_{c,t}$ nor $I_{c,t-1}$. We do not observe $m_{c,t}$ as we are limited in each period by existing testing capacity and infrastructure. We miss $I_{c,t-1}$ for two reasons. It is clear that total infections yesterday is a function of all previous new daily cases; that is, $I_{c,t-1} = f(m_{c,t-1}, ..., m_{c,t_0}; \beta, \gamma)$. As a result, any issues with testing in the past will impact our estimate of the stock of infections today. In addition, this variable is not equivalent to the stock of positive cases yesterday, even if testing is complete, as it does not include people who have recovered and can no longer pass the disease along to others. Therefore, to the extent that recovery data is incomplete, we will have issues measuring the number of individuals who are currently infected.

[12] This section will omit the binned terms from our specifications for brevity.

3.3 Testing

In light of this issue, we account for testing within our model.[13] Specifically, we add the variable $T_{c,t}$, the total number number of individuals who have tested positive in county c at time period t, to the model. Then, by making assumptions on the relationship between the underlying mechanisms of the SIR model and what we observe, we can recover the event study coefficients. We assume that, in each period, we are able to estimate some fraction $\tau_{c,t}$ of daily new cases in each county ($m_{c,t}$). In the general form of this model, $\tau_{c,t}$ can vary both across counties and over time. Using this new parameter, we can write $T_{c,t}$ using the recursion $T_{c,t} = T_{c,t-1} + \tau_{c,t} m_{c,t}$, which is equivalent to $\triangle T_{c,t} = \tau_{c,t} m_{c,t}$. The left hand side of this equation is now the number of daily new cases, which we observe. Taking logs of this relationship and utilizing (5) gives us that

$$\ln(\triangle T_{c,t}) = \ln(\tau_{c,t}) + \ln(\beta_c) + \ln(I_{c,t-1}) \tag{6}$$

We now need only find an expression for $\ln(I_{c,t-1})$ in terms of variables we observe. Adding together (3) and (4) and utilizing the definition of $m_{c,t}$ gives us that

$$I_{c,t} + R_{c,t} = I_{c,t-1} + R_{c,t-1} + m_{c,t} = \sum_{i=t_0}^{t} m_{c,i}$$

A similar recursion using positive tests allows us to write

$$T_{c,t} = T_{c,t-1} + \tau_{c,t} m_{c,t} = \sum_{i=t_0}^{t} \tau_{c,i} m_{c,i}$$

Unless we are willing to make the assumption that testing capacity is not evolving over time, we cannot solve for the general relationship between $T_{c,t}$ and $I_{c,t} + R_{c,t}$. Nevertheless, it is clear from these derivations that we can write down the reduced form relationship between these variables as $T_{c,t} = \tau_{c,t}^{RF} \cdot (I_{c,t} + R_{c,t})$, where $\tau_{c,t}^{RF}(\tau_{c,t}, ..., \tau_{c,t_0}; \beta_c, \gamma)$ is some unknown function of the evolution of testing capacity in county c. Taking logs of this relationship and rearranging gives us that

[13] Berger et al. (2020) extends the basic SEIR model to include testing as well. They are able to show that a mixture of higher levels of testing and targetted quarantining can both reduce transmission and dampen the impact of the virus on the economy. However, theirs is a model that calibrates certain facts about COVID-19 (such as the infection rate or the quarantine rate) and therefore does not have en empirical strategy per se – which makes sense because they are attempting to recommend policy. However, we take the SEIR model as a starting point to inform our empirical analysis to assess the effectiveness of the stay-at-home orders.

$\ln(I_{c,t} + R_{c,t}) = \ln(T_{c,t}) - \ln(\tau_{c,t}^{RF})$. We can utilize this expression to rewrite (6) as

$$\begin{aligned}
\ln(\triangle T_{c,t}) &= \ln(\tau_{c,t}) + \ln(\beta_c) + \ln(I_{c,t-1}) \\
&= \ln(\tau_{c,t}) + \ln(\beta_c) + \ln(I_{c,t-1}) - \ln(I_{c,t-1} + R_{c,t-1}) + \ln(I_{c,t-1} + R_{c,t-1}) \\
&= \ln(\tau_{c,t}) - \ln(\tau_{c,t-1}^{RF}) + \ln(\beta_c) + \ln(T_{c,t-1}) + \{\ln(I_{c,t-1}) - \ln(I_{c,t-1} + R_{c,t-1})\}
\end{aligned}$$

Early in a pandemic, we would expect $\ln(I_{c,t-1}) - \ln(I_{c,t-1} + R_{c,t-1}) = \ln(\frac{I_{c,t-1}}{I_{c,t-1}+R_{c,t-1}})$ to be small. The size of this fraction depends on how fast the disease is spreading relative to the rate of recovery at the beginning of the pandemic - this is referred to as the basic reproduction rate $R_0 = \beta_c/\gamma$.[14] Intuitively, as the disease spreads initially with a high R_0, the stock of infected individuals grows much faster than the stock of recovered individuals. Further, since $\frac{\partial}{\partial x \partial y} \ln(\frac{x}{x+y}) < 0$, even as the stock of recovered individuals inevitably rises, this difference remains small. Utilizing the approximation $\ln(I_{c,t-1}) - \ln(I_{c,t-1} + R_{c,t-1}) \approx 0$, we can write:

$$\ln(\triangle T_{c,t}) = \ln(\tau_{c,t}/\tau_{c,t-1}^{RF}) + \ln(\beta_c) + \ln(T_{c,t-1}) \tag{7}$$

In principle, this motivates the event study design

$$\ln(m_{c,t}) = \sum_{l=-7, l\neq -1}^{l=28} \mu_l \cdot \mathbb{1}\{t - E_c = l\} + \ln(\tau_{c,t}/\tau_{c,t-1}^{RF}) + \alpha_c + \ln(T_{c,t-1}) + \varepsilon_{c,t}$$

Since this empirical design is informed directly by an SIR model, we can interpret μ_l in the same way as before. Notice that $\alpha_c = \ln(\beta_c)$ and μ_l both enter additively into the log of daily cases; as a result, properly transformed, μ_l gives the average change in β_c on day l after the implementation of a SAH order. Formally, $\ln(\beta_c) + \mu_l = \ln(\beta_c e^{\mu_l})$, so that $100 \cdot (1 - e^{\mu_l})$ gives the average percentage change in the growth rate across counties induced by SAH orders.

3.3.1 Relationship to Literature

Before continuing, it is worthwhile to briefly discuss the relationship between this model of imperfect testing and what has been developed in the empirical literature. As noted earlier, Chernozhukov

[14] A medical metastudy (Liu et al. (2020)) finds the number to be within 1.4 to 6.49 with a mean of 3.28.

et al. (2020) derive a similar model: their equation (8) describing the evolution of confirmed cases is the continuous time analog of our equation $\triangle T_{c,t} = \tau_{c,t} m_{c,t}$. We derive their estimating equation within our discrete time framework to illustrate how our focus differs from theirs.

Taking the first difference of the natural log of our identity $\triangle T_{c,t} = \tau_{c,t} m_{c,t}$ for confirmed cases decomposes the growth rate of this variable into two sources, growth in testing capacity and spread of disease: $\triangle \ln(\triangle T_{c,t}) = \triangle \ln(\tau_{c,t}) + \triangle \ln(m_{c,t})$. Taking the first difference of equation 6 (and adding that $\beta_{c,t} = \beta_c$) reveals that $\triangle \ln(m_{c,t}) = \triangle \ln(I_{c,t-1})$, allowing us to write $\triangle \ln(\triangle T_{c,t}) = \triangle \ln(\tau_{c,t}) + \triangle \ln(I_{c,t-1})$. Maintaining our earlier assumption that $S_{c,t}/N_c \approx 1$, we can rearrange equation 3 from our SIR model to describe the progression of infections as $I_{c,t} = (1 + \beta_c - \gamma) I_{c,t-1}$. Taking logs and rearranging slightly yields $\triangle \ln(I_{c,t}) = \ln(1 + \beta_c - \gamma)$, which we can substitute into our expression for the growth rate of confirmed cases:

$$\triangle \ln(\triangle T_{c,t}) = \triangle \ln(\tau_{c,t}) + \ln(1 + \beta_c - \gamma) \approx \triangle \ln(\tau_{c,t}) + (\beta_c - \gamma)$$

This is the discrete time version of equation (10) in Chernozhukov et al. (2020), which provides a theoretical basis for their estimating equation. They proceed by assuming that the growth of testing capacity $\triangle \ln(\tau_{c,t})$ is a linear function of the growth of the number of tests administered, which is observable, and focus on modeling how $\beta_c - \gamma$ evolves in response to changes in behavior and information over time. In contrast, we assume that $\beta_c - \gamma$ is stationary and focus on making inferences when testing, $\triangle \ln(\tau_{c,t})$, evolves in a complicated manner over time. Equation 7 is better suited for this approach, so we return to our model of how daily positive cases evolve over time.

3.4 Identification

Unfortunately, at this point, we are still unable to identify μ_l. This stems from the fact that $\tau_{c,t}$ varies at the county-day level, so we cannot differentiate a rise in cases over time due to spread of the disease from changes in testing. We view this as the fundamental identification challenge stemming from endogenous testing. At this point we need to impose more structure in order to recover a useful estimating equation. We consider several different assumptions that allow us to identify μ_l, and proceed from the most to least restrictive.

3.4.1 Time-Invariant Testing

First, we could impose that

$$\tau_{c,t} = \tau_c \tag{Assumption 1T}$$

is constant over time.[15] As a result, $T_{c,t} = \tau_c \sum_{i=t_0}^{t} m_{c,i} = \tau_c(I_{c,t}+R_{c,t})$, which gives us the intuitive result that $\tau_c = \tau_c^{RF}$, since we are simply reporting some fixed fraction of cases in every period. Then, 7 will simplify to $\ln(\triangle T_{c,t}) = \ln(\beta_c) + \ln(T_{c,t-1})$, which is the observed analog of (5). This result is important, as it establishes a sufficient condition under which we can estimate β_c as if we are observing the true number of cases.

Under (Assumption 1B) & (Assumption 1T), we can utilize an event study approach to identify the impact of a stay at home order. Our SIR model implies the following specification:

$$\ln(\triangle T_{c,t}) = \sum_{l=-7, l\neq -1}^{l=28} \mu_l \cdot \mathbb{1}\{t - E_c = l\} + \delta_{\text{lag}} \cdot \ln(T_{c,t-1}) + \alpha_c + \text{DOTW}_t + \varepsilon_{c,t} \tag{8}$$

We use county level fixed effects to difference out the county specific growth rate, and we include cumulative cases yesterday to control for the expected number of cases today. Intuitively, this functions as a sufficient statistic (along with α_c) for what we expect daily new cases to be in the absence of a public health intervention. Lastly, we include 7 day of the week fixed effects DOTW_t to account for systematic differences in case reporting throughout the week.[16]

Figure 3(a) plots the results of this specification. The contrast with Figure 2 is immediately clear - we no longer have a pre-trend in daily new cases, and the effect of a SAH order is negative and significant. One month after the SAH order is implemented, our event coefficient $\mu_{28} = -.661$, implying a 48% average reduction in the spread of COVID-19. As noted previously, Liu et al. (2020) find that the average estimated R_0 was 3.28; our estimates imply that the reproductive rate at time $t + 28$ is $R_{t+28} = 1.70$, implying a substantial reduction but not elimination of spread.[17]

[15] It is important to note that this assumption does *not* impose that testing capacity remains constant over time. Rather, we assume that we are identifying a fixed fraction of infected individuals, which implies that testing capacity is expanding at the same rate as the virus is spreading. As a result, this assumption is not inconsistent with the commonsense observation that more tests are being performed over time, though it might still not be appropriate.

[16] Including day-of-the-week fixed effects doesn't change our results much, but we notice that it shrinks the standard error bands around our estimates. This is to be expected given that we are essentially controlling for the variation that is happening because of the cyclicality in the reporting – it is well documented that there is a big drop in cases reported on weekends compared to the middle of the week, for instance.

[17] R_t would need to drop below 1 for spread to eventually stop, as this reflects a scenario in which each infected person is spreading the disease to less than 1 other person.

Figure 3: Causal Event Study Results

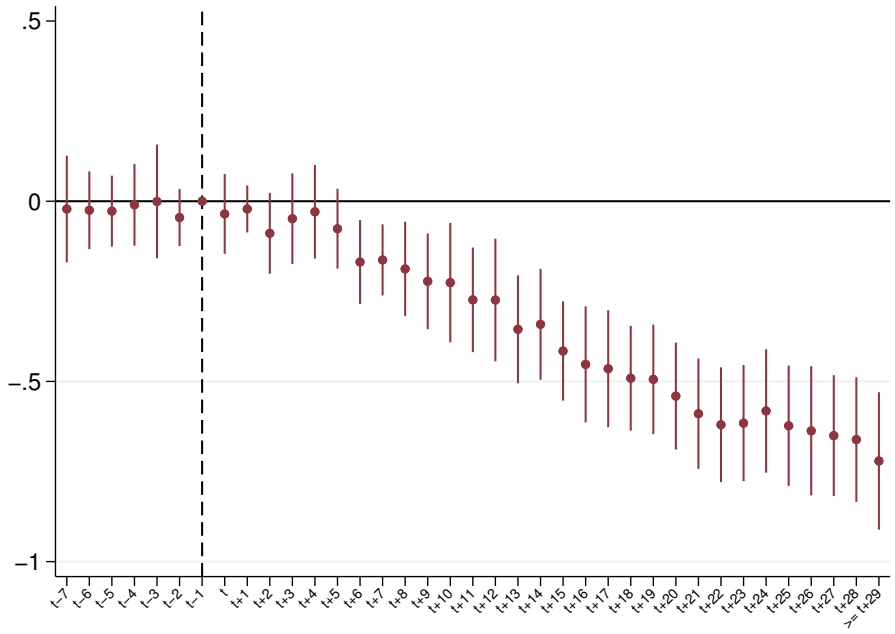

(a) Assumption 1B & Assumption 1T

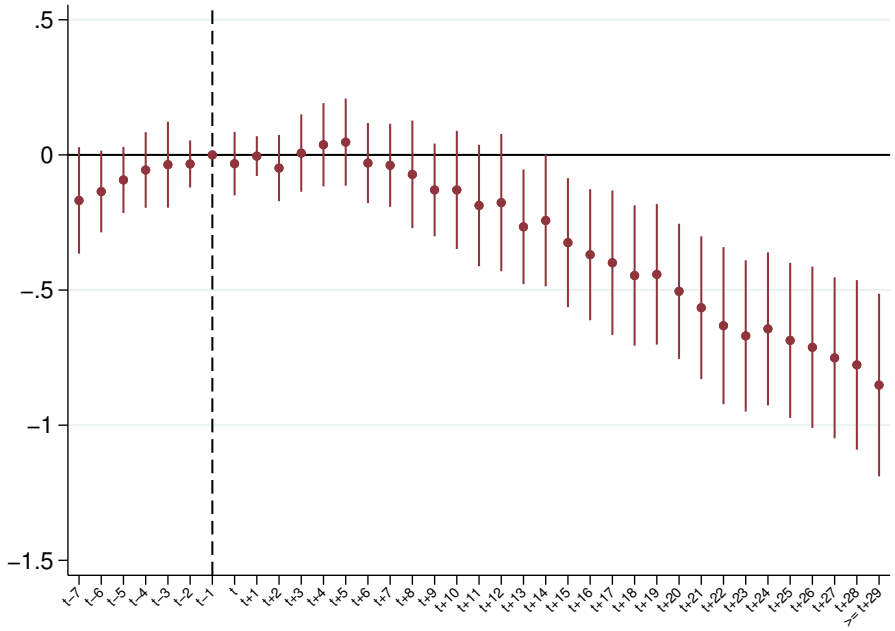

(b) Assumption 2B & Assumption 1T

Panel (a) plots the estimates and 95 percent confidence intervals for Equation (8). In all models, the dependent variable is the daily number of new infections, and county fixed effects and lagged stock of positive cases are controlled for. Panel (b) plots the estimates and 95 percent confidence intervals for Equation (10). Standard errors are clustered at the state level.

We view this simple model as having two major lessons. First, $\ln(T_{c,t-1})$ should be included as a control to correctly adjust for what cases are expected to be this period (as well as unit fixed effects). Second, if the determinants of case growth are specified properly, there is no need to include daily fixed effects, which are standard in the TWFE design. In practice, we find that including these fixed effects only serves to increase noise in the model, as a lot of important variation is soaked up, without a clear theoretical justification. The assumptions needed for this model to be correct are restrictive, so we now turn to estimation when testing can change differentially over time.

3.4.2 Identification Issues: Changing Testing Over Time

To begin, we could impose the similar but less restrictive assumption that $\tau_{c,t}$ varies over time but does not change *too* quickly. To illustrate, let $w(t)$ denote the week that day t is a part of and assume that $\tau_{c,t} = \tau_{c,w(t)}$ is determined entirely by $w(t)$ rather than t[18]. On face, this would seem to have solved the identification challenge, as we can now contribute intraweek variation in cases to spread, not changes in testing. However, even when the structural testing parameter does not vary at the daily level, the reduced form parameter will.

To see this, let $m_{c,w}$ denote the total number of new cases in week w,[19] and let $I_{c,w}$ and $R_{c,w}$ denote the total stock of infected and recovered individuals at the end of week w, respectively. Notice that we can rewrite $T_{c,t}$ as the weighted sum of weekly averages:

$$\begin{aligned} T_{c,t} &= \sum_{w=0}^{w(t)} \tau_{c,w} m_{c,w} = \tau_{c,w(t)} m_{c,w(t)} + \sum_{w=0}^{w(t)-1} \tau_{c,w} m_{c,w} \\ &= \tau_{c,w(t)} \{(I_{c,t} + R_{c,t}) - (I_{c,w(t)-1} + R_{c,w(t)-1})\} + \tau^{INT}_{c,w(t)-1} (I_{c,w(t)-1} + R_{c,w(t)-1}) \end{aligned}$$

Where the last equality utilizes the identity $m_{c,w(t)} = (I_{c,t} + R_{c,t}) - (I_{c,w(t)-1} + R_{c,w(t)-1})$ and implicitly defines $\tau^{INT}_{c,w} = T_{c,w}/(I_{c,w} + R_{c,w})$, which are fixed for all t' after week w. A bit of rearranging shows us that

$$T_{c,t} = \tau_{c,w(t)} (I_{c,t} + R_{c,t}) + \{\tau^{INT}_{c,w(t)-1} - \tau_{c,w(t)}\} (I_{c,w(t)-1} + R_{c,w(t)-1}) \tag{9}$$

[18] This analysis remains unchanged for any different level of aggregation.

[19] Strictly speaking, $m_{c,w}$ is a function of t, as it gives the number of cases in week $w(t)$ on day t. We write $m_{c,w}(t) = m_{c,w}$ in the derivations in this section.

What this tells us is that testing today has an *affine* relationship with the stock of infected and recovered individuals, with coefficients that vary at the weekly level. As a result, since $\tau_{c,t}^{RF}$ imposes a *linear* relationship between these variables, it must vary at the daily level.

Unfortunately, we cannot escape this problem by assuming that testing does not vary too much across space, our other dimension of variation. Suppose testing capacity does not vary across state so $\tau_{c,t} = \tau_{s,w(t)}$. We can still write $T_{c,t}$ as the weighted sum of weekly averages, but since each county has a different path of daily new cases, the intercept still varies at the county-week level:

$$\begin{aligned} T_{c,t} &= \tau_{s,w(t)} m_{c,w(t)} + \sum_{w=0}^{w(t)-1} \tau_{s,w} m_{c,w} \\ &= \tau_{s,w(t)}\{(I_{c,t} + R_{c,t}) - (I_{c,w(t)-1} + R_{c,w(t)-1})\} + \tau_{c,w(t)-1}^{INT}(I_{c,w(t)-1} + R_{c,w(t)-1}) \end{aligned}$$

This illustrates exactly how assumptions on testing impact the theoretical relationship between total population of ever-infected individuals ($I_{c,t} + R_{c,t}$) and the number of these individuals that we have identified ($T_{c,t}$). Any level of testing aggregation across time will be reflected in the intercept and slope terms, but any level of testing aggregation across space will only be reflected in the slope term. However, neither assumption or its combination, will give us that the relationship between these variables is linear *and* varies at a level higher than our level of observation.

3.4.3 Modeling the Affine Relationship

In light of this, we deal directly with the affine relationship identified earlier. Since assuming that testing only varies at the state level does not simplify the model, we maintain the previous assumption that testing capacity progresses at the county-week level, i.e. $\tau_{c,t} = \tau_{c,w(t)}$. We can rearrange (9) to obtain the following expression for $I_{c,t} + R_{c,t}$:

$$\begin{aligned} \tau_{c,w(t)}(I_{c,t} + R_{c,t}) &= T_{c,t} - \{\tau_{c,w(t)-1}^{INT} - \tau_{c,w(t)}\}(I_{c,w(t)-1} + R_{c,w(t)-1}) \\ &= T_{c,t} - \frac{\tau_{c,w(t)} - \tau_{c,w(t)-1}^{INT}}{\tau_{c,w(t)-1}^{INT}} T_{c,w(t)-1} \end{aligned}$$

To simplify notation, let $\zeta_{c,w(t)} := -\frac{\tau_{c,w(t)} - \tau_{c,w(t)-1}^{INT}}{\tau_{c,w(t)-1}^{INT}}$, and define $T_{c,t}^{*} := T_{c,t} + \zeta_{c,w(t)} T_{c,w(t)-1}$. Substituting into the expression above, we now have that $\tau_{c,w(t)}(I_{c,t} + R_{c,t}) = T_{c,t}^{*}$.

Once again recalling (6), we can perform a similar derivation to obtain that

$$\begin{aligned}\ln(\triangle T_{c,t}) = \ln(\tau_{c,w(t)}) + \ln(\beta_c) + \ln(I_{c,t-1}) &\approx \ln(\beta_c) + \ln(T^*_{c,t-1}) \\ &= \ln(\beta_c) + \ln(T_{c,t-1}) + \ln\left(1 + \frac{\zeta_{c,w(t-1)} T_{c,w(t-1)-1}}{T_{c,t-1}}\right)\end{aligned}$$

Finally, we linearize the last term around 1 so that $\ln\left(1 + \frac{\zeta_{c,w(t-1)} T_{c,w(t-1)-1}}{T_{c,t-1}}\right) \approx \frac{\zeta_{c,w(t-1)} T_{c,w(t-1)-1}}{T_{c,t-1}}$. This is appropriate when $\zeta_{c,w(t-1)} T_{c,w(t-1)-1} \ll T_{c,t-1}$, which relies on $\zeta_{c,w(t-1)}$ and $T_{c,w(t-1)-1}/T_{c,t-1}$ being small. Notice that $T_{c,w(t-1)-1}/T_{c,t-1} \leq 1$ by construction, so we need to assume like before that testing capacity does not change drastically from week to week; formally

$$\tau_{c,t} = \tau_{c,w(t)} \text{ and } \tau_{c,w(t)} - \tau^{INT}_{c,w(t)-1} \approx 0 \qquad \text{(Assumption 2T)}$$

Under Assumption 2T, our model implies that

$$\ln(\triangle T_{c,t}) = \ln(\beta_c) + \ln(T_{c,t}) + \zeta_{c,w(t-1)-1} \frac{T_{c,w(t-1)-1}}{T_{c,t-1}}$$

Accordingly, we estimate policy effects using the event study specification

$$\ln(\triangle T_{c,t}) = \sum_{l=-7, l\neq -1}^{l=28} \mu_l \cdot \mathbb{1}\{t - E_c = l\} + \alpha_c + \delta_{\text{lag}} \ln(T_{c,t-1}) + \alpha_{c,w(t-1)-1} \frac{T_{c,w(t-1)-1}}{T_{c,t-1}} + \text{DOTW}_t + \varepsilon_{c,t} \tag{10}$$

We interact county-by-week fixed effects $\alpha_{c,w}$ with $T_{c,w(t-1)-1}/T_{c,t-1}$ to estimate the coefficient on this term $\zeta_{c,w(t-1)-1}$, which varies at the county-by-week level.

Figure 3(b) plots the result of this specification. The pattern of event coefficients is different here - there is a clear pre-trend leading up to and continuing past the event date, suggesting that the growth rate of cases in treated counties is increasing relative to control areas. Nevertheless, despite this, after a month counties with a SAH order have significantly decreased spread of COVID-19. Around one month after the SAH order is implemented, our event coefficient $\mu_{28} = -.777$, implying a 54% average reduction in the growth rate β_c and effective reproductive rate of $R_{t+28} = 1.50$.

3.5 Adding Mobility

3.5.1 Theory

So far, we have maintained Assumption 1B and held that the rate of infection β_c can differ across county but must remain constant over time. Clearly, we need to impose structure on $\beta_{c,t}$, as we could then not distinguish changes in the underlying spread from policy effects by the proposition established earlier. However, we can break $\beta_{c,t}$ into its component parts and introduce data to measure some of these components to estimate a model with weaker assumptions on the components we cannot observe.

To begin, we expand our basline model and now allow $\beta_c = \beta_{c,t}$ to change over time. Recall that β gives the expected number of susceptible individuals one infected person will pass the virus to. Following Arnon et al. (2020), we can decompose this rate into the product of the expected number of susceptible individuals one infected person will come in contact with and the probability that the virus is transmitted, conditional on contact. We define the former as the *contact rate* $\kappa_{c,t}$ and the latter as the *infection rate* $\theta_{c,t}$, giving us that

$$\beta_{c,t} = \kappa_{c,t} \cdot \theta_{c,t} \tag{11}$$

Taking logs of this equation reveals that $\ln(\beta_{c,t}) = \ln(\kappa_{c,t}) + \ln(\theta_{c,t})$, which enter additively into our estimating equation.

We assume that the contact rate can change over time, and we introduce new data in the next subsection to proxy for these changes. To recover identification, we assume that the infection rate is constant across time; that is, $\theta_{c,t} = \theta_c$. This rate is likely a function of two factors: biological factors specific to the disease, and the demographic make-up of the population. The former are constant across all counties,[20], but the latter need to be accounted for. Let $\mathbf{X}_c$ denote all demographic factors that determine the spread of COVID-19. To the extent that these determinants do not change over time, we can write $\theta_c(\mathbf{X}_c)$ as their aggregate impact on the infection rate, which we can account for using fixed effects estimation.

[20] While there are documented mutations of SARS-CoV-2, it is unlikely that new strains are changing over time, within county, during our sample frame, in a manner that would seriously impact our results.

3.5.2 Data

We use the patterns data from SafeGraph[21] to measure mobility. Using anonymized data from mobile devices, the company has compiled daily visits to about 7 million points of interest (POI) located across the United States. The POI include a wide range of physical locations such as restaurants, retail and grocery stores.

Our mobility measure is very similar to weekly estimates used in Allcott et al. (2020), except that ours is at a daily frequency. On each calendar day, we aggregate visits to each POI using its FIPS code into a county level panel. For simplicity, we do not differentiate across types of establishments and focus on capturing widest range of foot traffic possible. While this increased frequency allows for greater time variation, it also creates strong day-of-the-week seasonality. To address the issue, we use a 7-day moving average of the POI visit counts as our baseline measure.

Our mobility measure closely captures decreased mobility across counties around March. To illustrate, we plot the time series of total POI visits in Los Angeles County of California and Washtenaw County of Michigan on Figures 4a and 4b. On March 13th, the White House declared the pandemic a National Emergency (US President, Proclamation (2020)) - as these figures illustrate, mobility declines markedly upon declaration of national emergency by the federal government ('NE') and continue to drift downwards following the announcement of stay-at-home orders ('SAH') in their respective states.

To match this empirical fact, we model the March 13th emergency declaration as a national-level information event that is responded to differentially by county. Formally, we modify equation (11) such that

$$\ln(\beta_{c,t}) = \ln(\kappa_c) + \mathbb{1}(t \geq \text{March 13}) \ln(\hat{\kappa}_{c,t}) + \ln(\theta_c) \qquad \text{(Assumption 2B)}$$

To be clear, we are assuming that the 7-day average contact rate is constant within county in the period leading up the March 13 emergency declaration. Then, following this date, we assume the contact rate is measured by our POI metric $\hat{\kappa}_{c,t}$, which we are able to include as a control. This

[21] `https://docs.safegraph.com/docs/weekly-patterns`

Figure 4: Mobility Patterns in Selected Counties

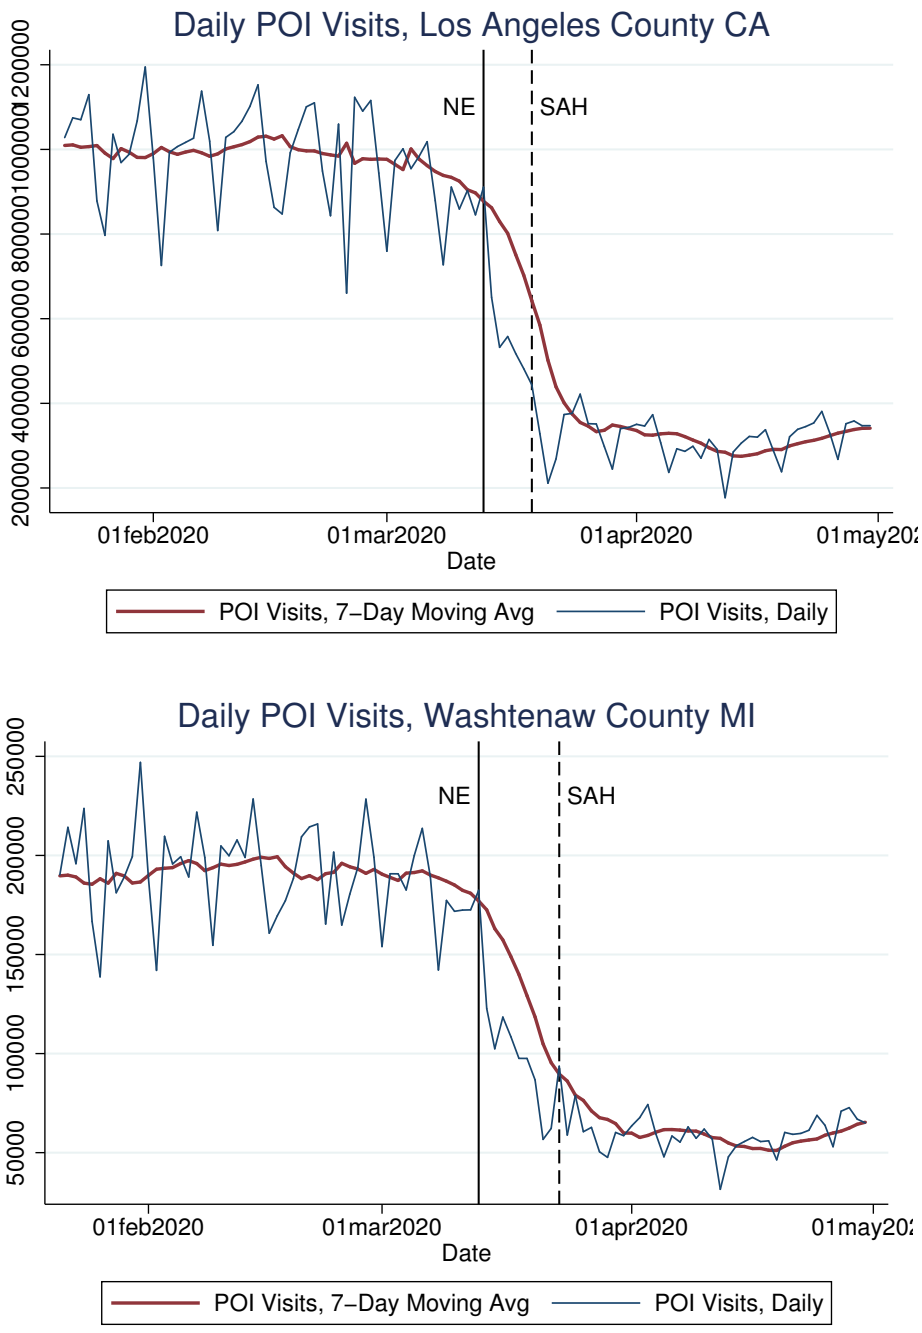

This figure plots the time series of total visits to points of interest (POIs) located in Los Angeles County of California and Washtenaw Coiunty of Michigan according to SafeGraph. 7-day moving averages of daily visits are in solid red lines, and the raw series are plotted in thinner blue lines. Solid vertical line is drawn on March 13 of 2020, when the federal government declared national emergency. Dashed vertical lines are denoted on dates when the governors of Michigan and California issued Stay-at-Home orders.

motivates the event study design:

$$\ln(\triangle T_{c,t}) = \sum_{l=-7, l\neq -1}^{l=28} \mu_l \cdot \mathbb{1}\{t - E_c = l\} + \alpha_c + \delta_{\text{lag}} \ln(T_{c,t-1}) + \delta_{\text{mob}} \mathbb{1}(t \geq \text{March 13}) \ln(\hat{\kappa}_{c,t}) + \text{DOTW}_t + \varepsilon_{c,t} \tag{12}$$

Where $\alpha_c = \ln(\beta_c) + \ln(\kappa_c) + \ln(\theta_c)$, maintaining Assumption 1T.

3.5.3 Results

Figure 5(a) plots our results when we control for POI visit mobility and impose Assumption 1T. The results are similar in Figure 5(b), which utilizes controls for both mobility (under Assumption 2B) and testing (under Assumption 2T). These results look very similar qualitatively to those without mobility controls. We still have the same upward pre-trend leading up to lockdown measures in counties under Assumption 2T, but there is a significant negative effect of SAH orders in both specifications. Our main results are therefore robust to accounting for a time varying $\beta_{c,t}$, at least as far as the time variation is captured by changes in mobility, which the literature and theory both suggest it should be. Quantitatively, our $\mu_{28} = -.589$ & $\mu_{28} = -.705$ in each respective specification, which implies a 44-50% reduction in the spread of cases.

4 Conclusion

We build a novel SIR model with endogenous testing to identify the effect of Stay-at-Home orders on COVID-19 spread. One insight is that the number of infections in the previous period (or a proxy for it) is a key control variable that must be included in the event study in order to properly estimate the impact of the policy. This is because the number of infections yesterday summarizes the state of the world going into the current period, and together with some structural parameters determines the spread of the virus today. Another key contribution is to show how different sets of fixed effects amount to different assumptions about the progression of a county's ability to detect the virus over time. Taking into account these two issues, and assuming a conditional parallel trends assumption holds, we find that SAH orders have a strong sustained negative effect on the growth of cases under various assumptions about the progression of testing, with point estimates varying from 44% to 54%.

Figure 5: Causal Event Study Results

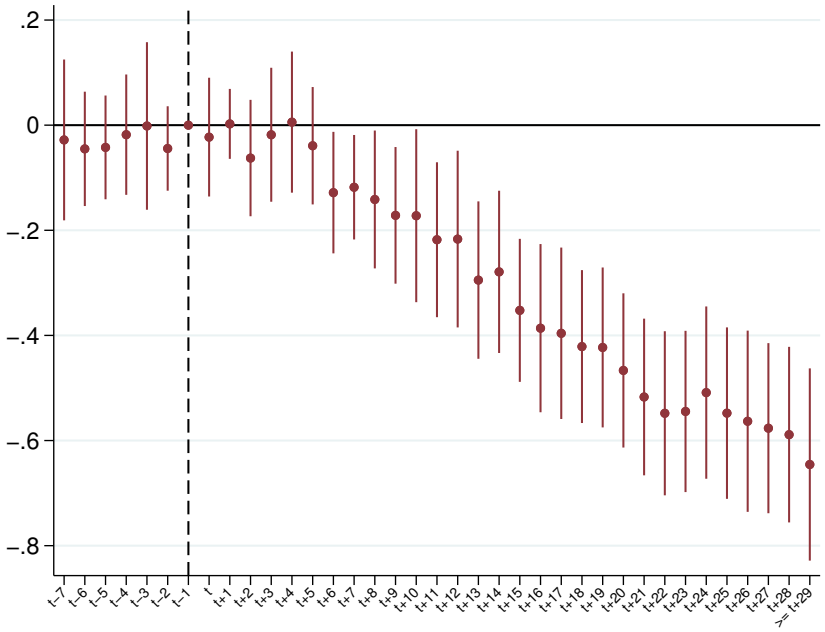

(a) Assumption 1B & Assumption 2T

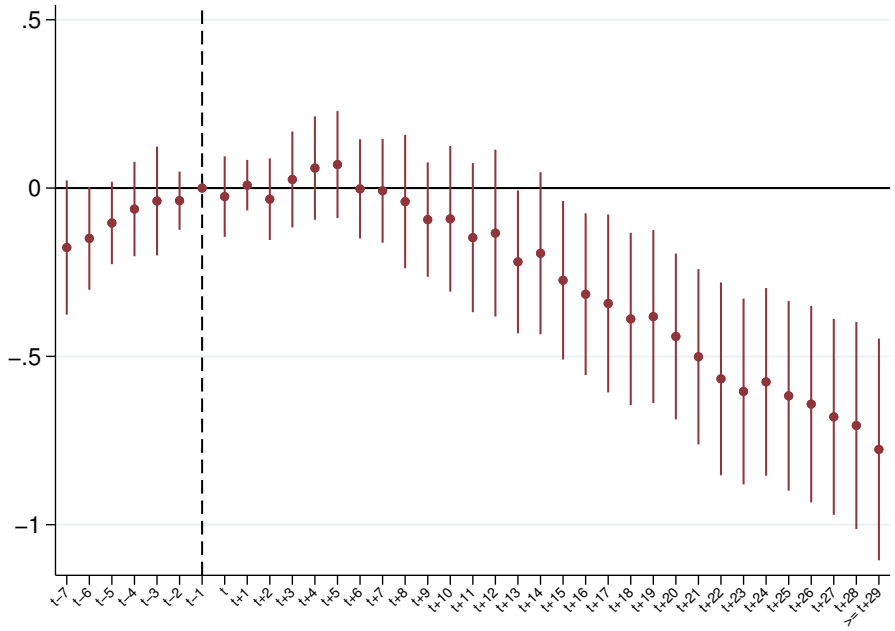

(b) Assumption 2B & Assumption 2T

Panel (a) plots the estimates and 95 percent confidence intervals for Equation (12). In all models, the dependent variable is the daily number of new infections, and county fixed effects and lagged stock of positive cases are controlled for. Panel (b) plots the estimates and 95 percent confidence intervals for Equation (12) with additional testing controls from Equation (10). Standard errors are clustered at the state level.

These estimates imply that SAH orders are a strong policy tool to eliminate the spread of COVID-19. However, it is important to note that initial estimates of R_0 suggest that reductions of $\beta_{c,t}$ of the magnitude we find will not eventually drive daily cases to zero. This suggests that Stay-at-Home Orders are best used in combination with other policy tools, such as mask mandates and contact tracing. Not only are these likely effective approaches in their own right, but they might augment the effects that we estimate as well. There is also a growing literature studying the costs and benefits of lockdown orders (e.g. Kaplan et al. (2020)), and while an economic model is needed to trade these off, we believe that our estimates add to this debate by precisely estimating the *benefits* of a lockdown from a public health perspective.

While our results are encouraging, it is unreasonable to expect that we can capture the full dynamics of an infectious disease with a simple SIR model. There is a nascent but deep literature exploring several important dimensions of the spread of COVID-19. We hope to augment this work by giving researchers the tools to inform their empirical approach using insights from Epidemiology rather than a standard econometric design. In particular, future research could expand on this basic model in two important directions.

First, while we note that $I_{c,t-1}$ (along with the county-specific growth rate) is a sufficient statistic for the future evolution of cases, this is only true in the context of our model. If the state policymaker has information outside of this model that informs their SAH timing decision (for example, that the disease is spreading among undergraduate students in a college town but not circulating outside of this subgroup), this timing might fail to be random conditional on knowing cumulative cases last period. Accordingly, future work ought implement a more complicated multi-group SIR model (see Acemoglu et al. (2020)), as well as a political economy model to explicitly address the non-random adoption of SAH orders.

Finally, it is important to note that the model we construct here is only appropriate at the beginning of a pandemic. As a result, any study interested in the current spread of COVID-19 should reconsider the assumptions made here. We sketch an outline of these assumptions and their plausibility in the present. First, we assume that almost every individual in the population is susceptible to disease, i.e. $S_{c,t}/N_c \approx 1$. This allows us to ignore the progression of the susceptible population in an (S)IR model. However, as this population drops, this assumption begins to

introduce important bias into estimates. If $S_{c,t}/N_c < 1$, then equation 5 becomes

$$\ln(m_{c,t}) = \ln(S_{c,t}/N_c) + \ln(\beta_c) + \ln(I_{c,t-1})$$

Due to the shape of the natural log function, significant drops in $S_{c,t}$ lead to a large negative bias in estimates of event coefficients. Second, we assume that the population of recovered individuals $R_{c,t}$ is small relative to to the population of infected people $I_{c,t}$. This obviously becomes inappropriate later in a pandemic; however, it is in principle possible to adjust for this with data on the number of recovered individuals or additional assumptions about the recovery rate. Third, recent months have seen the spread of newer mutations of COVID-19. If the infection rate $\theta_{c,t}$ differs across variants, additional assumptions are needed to model the change in this rate over time. For example, suppose that variant B overtakes variant A over the course of a month. Even if we assume that the infection rate is constant across the country and over time for each variant so that the series $\{\theta^A, \theta^B\}_{t=1}^{30}$, the true infection rate will be $\alpha_t\theta^A + (1-\alpha_t)\theta^B$ where α_t measures the prevalence of each variant in the population. This setup would motivate the include of daily fixed effects to control for this changing composition.

References

Acemoglu, D., I. W. Chernozhukov, and M. D. Whinston (2020). Optimal targeted lockdowns in a multi-group sir model. Technical report, National Bureau of Economic Research.

Alfano, V. and S. Ercolano (2020). The efficacy of lockdown against covid-19: a cross-country panel analysis. *Applied health economics and health policy 18*, 509–517.

Allcott, H., L. Boxell, J. Conway, M. Gentzkow, M. Thaler, and D. Yang (2020, nov). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics 191*, 104254.

Allcott, H., J. C. Conway, B. A. Ferguson, H. Allcott, L. Boxell, J. C. Conway, B. A. Ferguson, M. Gentzkow, and B. Goldman (2020). What Explains Temporal and Geographic Variation in the Early US Coronavirus Pandemic? *NBER Working Paper*.

Arnon, A., J. Ricco, and K. Smetters (2020). Epidemiological and economic effects of lockdown. Technical report, Brookings Papers on Economic Activity.

Atkeson, A. (2020). What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios. *NBER Working Paper*.

Berger, D., K. Herkenhoff, and S. Mongey (2020). An SEIR infectious disease model with testing and conditional quarantine. *COVID Economics* (13), 1–30.

Bonardi, J.-P., Q. Gallea, D. Kalanoski, and R. Lalive (2020). Fast and local: How did lockdown policies affect the spread and severity of the covid-19. *Covid Economics 23*, 325–351.

Chernozhukov, V., H. Kasaha, and P. Schrimpf (2020). Causal impact of masks, policies, behavior on early covid-19 pandemic in the us. *arXiv preprint arXiv:2005.14168*.

Dave, D. M., A. I. Friedson, K. Matsuzawa, and J. J. Sabia (2020). When do shelter-in-place orders fight covid-19 best? policy heterogeneity across states and adoption time. Technical report, National Bureau of Economic Research.

Fowler, J. H., S. J. Hill, N. Obradovich, and R. Levin (2020). The effect of stay-at-home orders on covid-19 cases and fatalities in the united states. *medRxiv*.

Giordano, G., F. Blanchini, R. Bruno, P. Colaneri, A. Di Filippo, A. Di Matteo, and M. Colaneri (2020). Modelling the covid-19 epidemic and implementation of population-wide interventions in italy. *Nature medicine 26*(6), 855–860.

Gupta, S., T. D. Nguyen, F. L. Rojas, S. Raman, B. Lee, A. Bento, K. I. Simon, and C. Wing (2020). Tracking public and private response to the covid-19 epidemic: evidence from state and local government actions. Technical report, National Bureau of Economic Research.

Isphording, I. E., M. Lipfert, N. Pestel, et al. (2020). School re-openings after summer breaks in germany did not increase sars-cov-2 cases. Technical report, Institute of Labor Economics (IZA).

Jagodnik, K. M., F. Ray, F. M. Giorgi, and A. Lachmann (2020). Correcting under-reported covid-19 case numbers: estimating the true scale of the pandemic. Technical report, medRxiv.

Kaplan, G., B. Moll, and G. L. Violante (2020). The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the us. Technical report, National Bureau of Economic Research.

Liu, Y., A. A. Gayle, A. Wilder-Smith, and J. Rocklöv (2020). The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine 27*(2), 1–4.

Lyu, W. and G. L. Wehby (2020). Comparison of estimated rates of coronavirus disease 2019 (covid-19) in border counties in iowa without a stay-at-home order and border counties in illinois with a stay-at-home order. *JAMA Network Open 3*(5), e2011102–e2011102.

Schlosser, F., B. F. Maier, O. Jack, D. Hinrichs, A. Zachariae, and D. Brockmann (2020). Covid-19 lockdown induces disease-mitigating structural changes in mobility networks. *Proceedings of the National Academy of Sciences 117*(52), 32883–32890.

Sun, L. and S. Abraham (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.

US President, Proclamation (2020, March 18). Declaring a national emergency concerning the novel coronavirus disease (covid-19) outbreak, proclamation 9994 of march 13, 2020. *Federal Register 85*(53), 15337–15338. https://www.govinfo.gov/content/pkg/FR-2020-03-18/pdf/2020-05794.pdf.

Yang, M.-J., N. Seegert, M. Gaulin, A. Looney, B. Orleans, A. Pavia, K. Stratford, M. Samore, and S. Alder (2020). What is the active prevalence of covid-19? *Available at SSRN 3734463*.