

Policy Evaluation during a Pandemic*

Brantly Callaway[†]

Tong Li[‡]

May 17, 2021

Abstract

National and local governments have implemented a large number of policies, particularly non-pharmaceutical interventions, in response to the Covid-19 pandemic. Evaluating the effects of these policies, both on the number of Covid-19 cases as well on other economic outcomes is a key ingredient for policymakers to be able to determine which policies are most effective as well as the relative costs and benefits of particular policies. In this paper, we consider the relative merits of common identification strategies exploiting variation in policy choices made across different locations by checking whether the identification strategies are compatible with leading epidemic models in the epidemiology literature. We argue that unconfoundedness type approaches are likely to be more useful for evaluating policies than difference in differences type approaches due to the highly nonlinear spread of cases during a pandemic. For difference in differences, we further show that a version of this problem continues to exist even when one is interested in understanding the effect of a policy on other economic outcomes when those outcomes also depend on the number of Covid-19 cases. We propose alternative approaches that are able to circumvent these issues. We apply our proposed approach to study the effect of state level shelter-in-place orders early in the pandemic.

JEL Codes: C21, C23, I1

Keywords: Policy Evaluation, Difference in Differences, Unconfoundedness, Covid-19, Pandemic, Mediators

*We thank Andrew Goodman-Bacon, Ian Schmutte, and Meghan Skira for helpful comments.

[†]Department of Economics. University of Georgia. brantly.callaway@uga.edu

[‡]Department of Economics. Vanderbilt University. tong.li@vanderbilt.edu

1 Introduction


There have been a large number of policies implemented in order to decrease the spread of Covid-19. The most important of these policies have been non-pharmaceutical interventions such as requirements to wear masks, making Covid-19 tests widely available, contact tracing, school closures, lockdowns, and others. These policies are likely to come with a number of tradeoffs in terms of effectiveness in reducing the spread of Covid-19 as well as their effects on individuals' economic and psychological well-being. Thus, understanding the effects of different policies along a number of dimensions (both effects on number of cases as well as effects on other outcomes) is a key ingredient for researchers, policymakers, and governments to consider when evaluating Covid-19 related policies.

The main way that these policies have been studied by researchers is to compare outcomes in locations that implemented some policy to outcomes in another location that did not implement the policy. Researchers typically exploit having access to panel data — data on cases, testing, and economic variables is generally widely available for particular locations over multiple time periods — to try to understand these effects. This sort of setup is very familiar to many researchers in economics that have exploited policy variation across locations and time to identify policy effects. With this sort of data availability, an almost default strategy of empirical researchers is to use difference in differences. And, indeed, difference in differences has been widely used to study the effects of policies in response to Covid-19.

In the current paper, we argue that difference in differences has properties that make it relatively less attractive for conducting policy evaluation during a pandemic than it typically would be for most applications in economics. The intuition for our results is that difference in differences methods are typically motivated by a two-way fixed effects model for untreated potential outcomes (see, for example, Blundell and Costa Dias (2009)). The key feature of these models is that they include additively separable unit-level unobserved heterogeneity (i.e., a fixed effect). This sort of heterogeneity is very common in applications in economics; for example, a textbook example would be an application on the effect of some policy on individuals' earnings where the researcher is worried that "ability" is unobserved, affects earnings, and is distributed differently between the group of individuals that are affected by the policy and the group of individuals not affected by the policy. In this setup, the unobserved heterogeneity can be differenced out and paths of outcomes for the group of treated units and the group of untreated units can be compared to each other to deliver the effect of the policy.

However, this sort of motivation does not apply in the case of Covid-19. In particular, the main epidemiological models for Covid-19 transmission are highly nonlinear and depend on (i) the number of currently infected individuals in a particular location, (ii) the number of susceptible individuals in a location, and (iii) the transmission properties of Covid-19. In other words, the key challenge in identifying effects of Covid-19 related policies on the number of Covid-19 cases is not that different locations are different in terms of unobserved heterogeneity, but rather that the spread of cases may differ in systematic ways across locations that experience the policy and those that do not especially due to the timing of the first cases in a particular

Figure 1: A Simulated SIRD Model



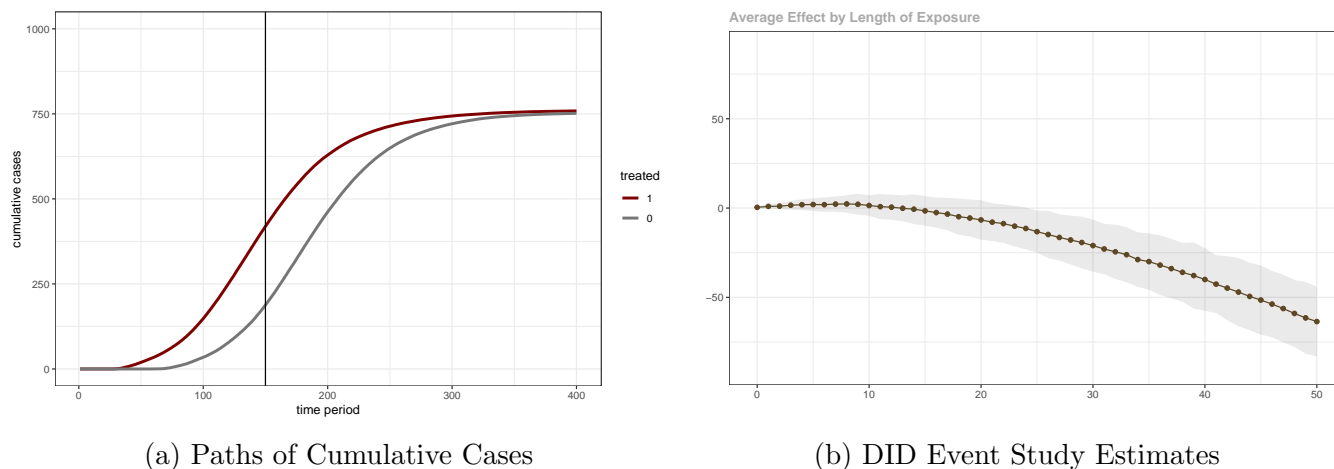
Notes: The figure shows paths of each variable in a stochastic SIRD model for one (treated) location in the simulated data. The vertical black line is placed at $t = 150$ when the policy was implemented. S stands for the number of susceptible individuals in the population, I stands for the number of currently infected individuals, R stands for the total number of recovered individuals in the population, D stands for the cumulative number of deaths, and C stands for the cumulative number of cases. In this example, the total population is 1000, there are 400 time periods, and the policy has no effect on the pandemic. The specific values for the parameters in this simulation are provided in Table 4 in Appendix A.

location.

In order to motivate our approach, we start by giving an example of the main types of issues that can confound policy analysis during a pandemic. Panel (a) of Figure 1 shows the paths of the key variables during a simulated pandemic coming from a stochastic SIRD model (SIRD models are a leading class of epidemic models, and we discuss this model in substantially more detail in the next section). The shapes of the path of each variable is typical of a SIRD model. In particular, at some point in time, some small number of cases shows up in a particular location. Then, the number of infections rise in early periods when there are a large number of susceptible individuals in that location combined with an increasing number of currently infected (which also implies contagious). As the number of susceptible decreases (i.e., as infected individuals recover or die), eventually the number of infected individuals decreases. Simultaneously, the cumulative number of cases, number of recovered individuals, and number of deaths all initially grow before eventually leveling off.

In this example, we consider the case where a new policy is implemented in some locations in period 150. For simplicity, we consider the case where the policy has no effect on Covid-19 cases. Locations that participate in the treatment and locations that do not participate in the treatment are alike in all ways except that treated locations tend to experience their first cases

Figure 2: Simulated Policy Effects on Cumulative Cases



Notes: Panel (a) plots simulated average paths of cumulative cases among treated and untreated locations. The vertical black line is at $t = 150$ when the policy is implemented for the treated locations. In this example, the policy is constructed so that it has no effect of Covid-19 cases. Panel (b) plots event study type estimates based on a parallel trends assumption for the effect of the policy on the number of cases by length of exposure to the treatment.

earlier than untreated locations. Panel (a) of Figure 2 shows plots of the average paths of cumulative cases for treated locations and untreated locations in this setup.

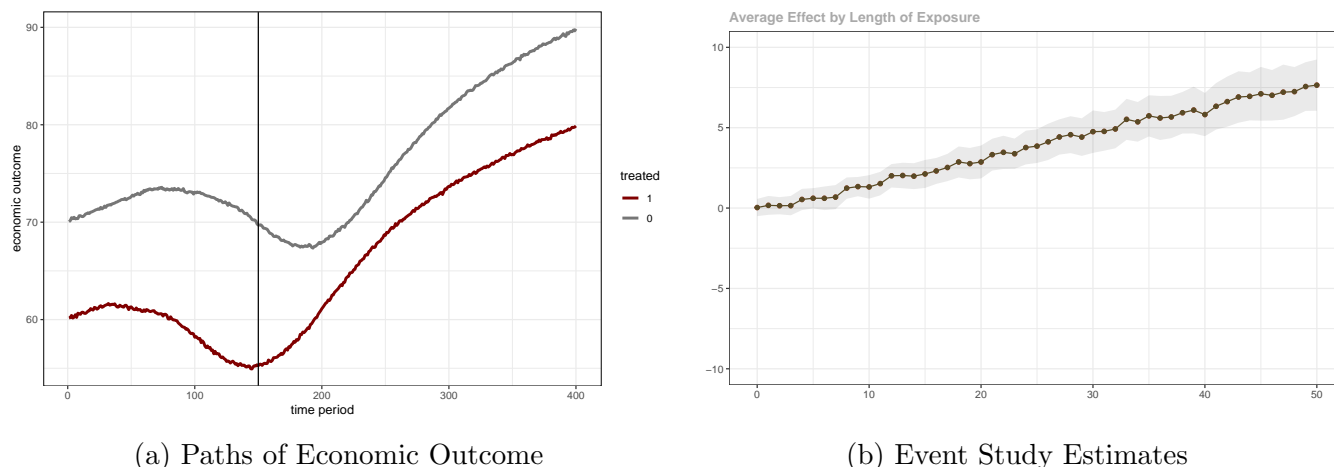
Panel (b) of Figure 2 shows event study-type estimates of the effect of the treatment on the number of cases. To be precise, these are difference in differences type estimates where the estimated effect comes from the average change in cases experienced by the treated group of locations relative to the change in cases experienced by the untreated group of locations over the same time periods. Taken at face value, the estimated effects in Panel (b) suggest that the policy decreased the number of Covid-19 cases in treated locations relative to what they would have been in the absence of the policy. However, recall that, in our simulation setup, the policy has no effect on Covid-19 cases. Thus, this example demonstrates that difference in differences can perform poorly in the context of trying to evaluate the effect of a policy on the number of Covid-19 cases. The key driver of this poor performance is (i) the nonlinearity of the model for Covid-19 transmission and (ii) differences in the timing of the first cases between locations that participate in the treatment and those that do not. The first of these is an inherent feature of trying to evaluate the effects of policies on Covid-19 cases. For the latter, generally, the bias of difference in differences approaches for policy evaluation becomes more severe as timing of first cases becomes more different between and untreated locations.¹

Another central issue in the economics literature is to understand the effect of Covid-19 related policies on various economic outcomes.² Understanding the effects of Covid-19 related

¹Interestingly, the best case for difference in differences is when the timing of first cases is the same across treated and untreated locations. However, this is also a case where there is no need to take a time difference at all and one could just make level comparisons of Covid-19 cases across locations.

²Throughout the text, we use the term “economic outcomes” but our results apply to any outcome of interest that is outside the epidemic model that we consider in the paper.

Figure 3: Simulated Policy Effects on Economic Outcome



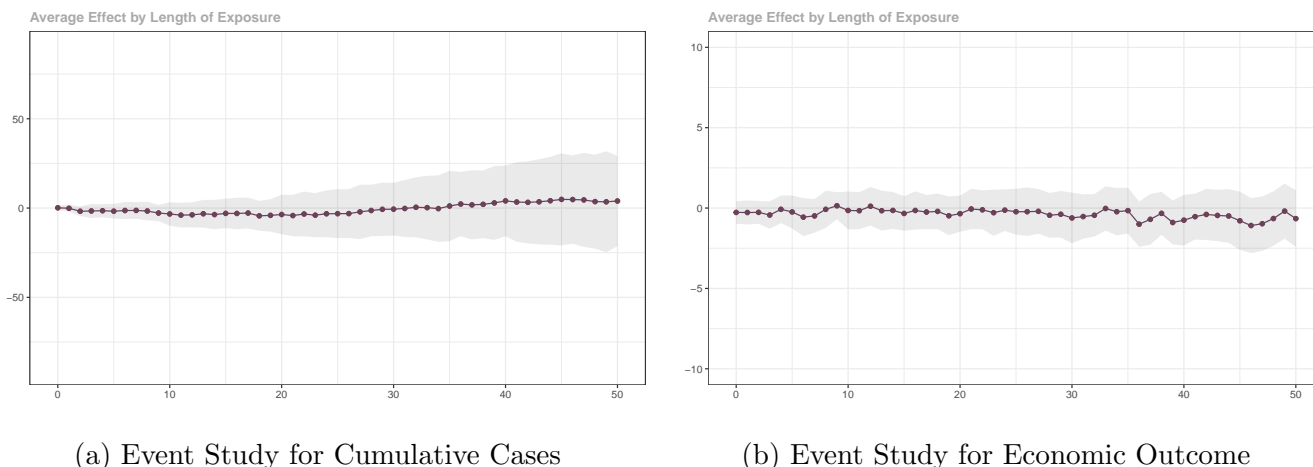
Notes: Panel (a) plots simulated average paths of outcomes among treated and untreated locations. The vertical black line is at $t = 150$ when the policy is implemented for treated locations. In this example, the policy is constructed so that it has no effect on economic outcomes (either directly or through its effect on Covid-19 cases). Panel (b) plots event study type estimates based on a parallel trends assumption for the effect of the policy on the number of cases by length of exposure to the treatment.

policies on economic outcomes is essential in order to understand the costs and benefits of various policies that are aimed at reducing the number of Covid-19 cases.

We focus on the simple leading case where untreated potential outcomes are generated by a two way fixed effects model that also depends on the current number of Covid-19 cases in a particular location. This setup allows both for the active number of Covid-19 cases to have an effect on the outcome of interest and for the active number of cases to themselves be affected by the policy. In this case, we consider (i) a “standard” version of difference in differences that directly compares paths of economic outcomes among treated and untreated locations, and (ii) difference in differences when the current number of Covid-19 cases is included as a regressor. We show that, generally, neither approach can deliver the average effect of the treatment on economic outcomes across treated locations. The first strategy breaks down because of differences in the current number of cases across treated and untreated locations. The second strategy breaks down when the policy affects the number of Covid-19 cases (which is the goal of the policy).

Figure 3 continues with the same simulated policy as above. As above, we consider the case where the policy has no effect on cases or on economic outcomes. However, in this simulation we allow for the economic outcome to depend on the number of active Covid-19 cases in a particular location (here, more active cases tends to decrease the economic outcome), but, otherwise, the economic outcome would follow parallel trends. Panel (a) shows average paths of outcomes for treated and untreated locations in this setup. Panel (b) shows event study type estimates (under the assumption of parallel trends). As before, and even in this very simple example, differences in the timing of first cases lead to violations of parallel trends that lead to poor estimates of the effect of the policy on the economic outcome of interest.

Figure 4: Estimated Policy Effects under Unconfoundedness



Notes: This figure plots event study-type estimates of the effect of a simulated policy on cumulative cases (in Panel (a)) and on an economic outcome (in Panel (b)) using the unconfoundedness-type identification arguments considered in the paper and using the doubly robust estimator discussed below. In this simulation, the policy has no effect on either outcome.

We propose alternative approaches that address both sets of issues mentioned above. In particular, for evaluating the effect of policies on Covid-19 cases, we first show that unconfoundedness-type identification strategies (i.e., strategies that compare locations that have the same pre-treatment values of key Covid-19 related variables) do not suffer from the same drawbacks as differences in differences approaches. For this case, we propose an estimation strategy that involves estimating (i) the propensity score (i.e., the probability of experiencing the policy conditional on pre-treatment values of Covid-19 related variables) and (ii) an outcome regression for Covid-19 cases in the absence of the policy that is related to the epidemic model. Our approach is doubly robust in the sense that it delivers consistent estimates of policy effects if either the propensity score or outcome regression are correctly specified. This is important as it implies that we can circumvent having to estimate a full SIRD model while still delivering an estimate of policy effects that are *compatible* with the SIRD model. Our approach also does not impose any model/restrictions on Covid-19 cases following participating in the treatment.

Second, for evaluating the effect of Covid-19 related policies on economic outcomes, we propose a two-step approach where the parameters of a two way fixed effects model that additionally allows for current cases to affect the outcome are identified using untreated locations in the first step. Then, in the second step, we recover the path of active cases that treated locations would have experienced on average if they had not been treated (this follows under similar unconfoundedness type arguments used for cumulative cases above). Using these two pieces of information, we are able to construct the average economic outcome that treated locations would have experienced if they had not participated in the policy — and, therefore, the average effect of the policy is identified for treated locations. Unlike other common approaches, this approach allows both for the policy to affect the current number of cases and for the current

number of cases to affect untreated potential outcomes.

We conclude the paper by studying the effect of state-level shelter-in-place orders (SIPOs) on the number of Covid-19 cases and on travel early in the pandemic. These are challenging policies to evaluate because states that implemented these policies also tended to have a large number of Covid-19 cases earlier than states that did not implement this type of policy (or implemented it later). This correlation mechanically leads to larger increases in Covid-19 cases among early treated states relative to untreated and later-treated states. This additionally implies violations of parallel trends and results in difference in differences estimates that these policies *increased* the number of Covid-19 cases (which is clearly unreasonable). In contrast, using our approach, our point estimate of the effect of SIPOs on Covid-19 cases is close to zero.³

Related Work There are a number of recent papers at the intersection of economics, Covid-19, and policy evaluation, and here we only briefly summarize some of the most related ones.

The most related papers to ours are Chernozhukov, Kasaha, and Schrimpf (2020) and Rowthorn and Maciejowski (2020). Like those papers, our paper considers policy evaluation motivated by SIRD models, but in practice, our approach is substantially different. Our approach generally imposes less structure than this strand of the literature. This has advantages and disadvantages. On the one hand, our approach generally requires fewer assumptions to evaluate policies that were actually enacted; on the other hand, our approach is less suitable for counterfactual policy analysis. Allcott et al. (2020) propose an event-study regression estimator that is motivated by SIRD models though their approach ends up being substantially different from ours.

Difference in differences has been widely used to study the effects of Covid-19 policies. Some examples of papers that consider Covid-19 related policies using difference in differences types of identification strategies include Bartik et al. (2020), Chetty, Friedman, Hendren, and Stepner (2020), Courtemanche et al. (2020), Dave et al. (2020b), Dave et al. (2020a), Gapen, Millar, Blerina, and Sriram (2020), Glaeser, Jin, Leyden, and Luca (2020), Goolsbee and Syverson (2021), Gupta et al. (2020), Juranek and Zoutman (2020), Kong and Prinz (2020), Villas-Boas, Sears, Villas-Boas, and Villas-Boas (2020), Wright, Sonin, Driscoll, and Wilson (2020), and Ziedan, Simon, and Wing (2020). Goodman-Bacon and Marcus (2020) provide an overview of differences in differences targeted to researchers studying Covid-19.

Finally, on the econometrics side, our paper is related to a large literature on unconfoundedness and difference in differences (see Imbens and Wooldridge (2009) for a survey of this literature). More notably, our contributions on allowing for infections to both be affected by the policy and to have a direct effect on economic outcomes appear to be conceptually new but related to work on mediation analysis (see Huber (2020) for a summary of this literature) and, in particular, see also Bonhomme and Sauder (2011) and Lechner (2011) for related discussion on time varying covariates that can potentially be affected by treatment participation.

³Most work on SIPOs that exploits state-level variation in policies has estimated that SIPOs decreased the number of Covid-19 cases (Courtemanche et al. (2020), Dave, Friedson, Matsuzawa, and Sabia (2020), Dave et al. (2020a), and Villas-Boas, Sears, Villas-Boas, and Villas-Boas (2020)). We discuss the reasons for these differences in Section 6.

2 A Baseline Stochastic SIRD Model

To start with, in this section, we briefly discuss a Stochastic SIRD model which is the workhorse model of epidemic spread in epidemiology and has been used extensively to forecast the spread of Covid-19 cases. SIRD models categorize individuals in a population into being S-Susceptible, I-Infected, R-Recovered, or D-Dead. SIRD models have a long history in epidemiology — a deterministic version of this kind of model was proposed by Kermack and McKendrick (1927). Stochastic SIRD models are discussed in Allen (2008) and Allen (2017) and have been considered by economists in Oka, Wei, and Zhu (2020), Fernández-Villaverde and Jones (2020), Ellison (2020), and Acemoglu, Chernozhukov, Werning, and Whinston (2020), among others.

Notation: Let N_l denote the number of individuals in location l . Let \mathcal{T} denote the total number of time periods. The number of susceptible individuals in location l in a particular time period t is denoted by S_{lt} , the number of currently infected individuals in location l at time period t is denoted by I_{lt} , the cumulative number of recovered individuals is denoted by R_{lt} , and the number of cumulative deaths is denoted by δ_{lt} . All individuals in the population are in exactly one of these states at a particular point in time so that

$$N_l = S_{lt} + I_{lt} + R_{lt} + \delta_{lt}$$

in all time periods. Later, we will be interested in the effect of the policy on the cumulative number of cases by time period t , and we denote this variable by C_{lt} and note that $C_{lt} = N_l - S_{lt}$.

In a SIRD model, the paths of each of these variables is governed by some transition equations. The transition equations have the Markov property; i.e., the path of each outcome over time only depends on the “state” of location l in the immediately preceding period. And, in particular, these transition equations are given by

$$\mathbb{E}[I_{lt} | \mathcal{F}_{lt-1}] = (1 - \lambda - \gamma)I_{lt-1} + \beta \frac{I_{lt-1}}{N_l} S_{lt-1} \quad (1)$$

$$\mathbb{E}[R_{lt} | \mathcal{F}_{lt-1}] = R_{lt-1} + \lambda I_{lt-1} \quad (2)$$

$$\mathbb{E}[\delta_{lt} | \mathcal{F}_{lt-1}] = \delta_{lt-1} + \gamma I_{lt-1} \quad (3)$$

$$\mathbb{E}[S_{lt} | \mathcal{F}_{lt-1}] = S_{lt-1} - \beta \frac{I_{lt-1}}{N_l} S_{lt-1} \quad (4)$$

$$\mathbb{E}[C_{lt} | \mathcal{F}_{lt-1}] = C_{lt-1} + \beta \frac{I_{lt-1}}{N_l} S_{lt-1} \quad (5)$$

where for some time period t , we define $\mathcal{F}_{lt} = (S_{lt}, I_{lt}, R_{lt}, \delta_{lt})$. It is worth considering each of these equations in some more detail. To start with, consider the term $\beta \frac{I_{lt-1}}{N_l} S_{lt-1}$ which shows up in Equations (1), (4) and (5). This is the expected number of new cases in time period t conditional on the state of location l in time period $t - 1$. The expected number of new cases from one period to the next depends on three things. First, it depends on I_{lt-1}/N_l which is the fraction of individuals that are infected in period $t - 1$. Holding other things constant, when more individuals are infected, it implies an expected larger increase in the number

of cases. Second, the expected number of new cases depends on the number of susceptible individuals in the population. Intuitively, when there are more susceptible individuals, the number of cases grows more rapidly (other things constant). The spread of a pandemic stops when the number of susceptible becomes small enough which can happen either through “herd immunity” or by decreasing the number of susceptible (for example, through the introduction of a vaccine). Finally, the change in the number of cases depends on the parameter β which is called the infection rate. Most non-pharmaceutical interventions are aimed at changing the infection rate — here, there are two potential benefits: (i) decreasing the infection rate through non-pharmaceutical interventions decreases the total number of cases that need to occur before reaching herd immunity,⁴ and (ii) if there is a vaccine on the horizon, it also would decrease the total number of cases that occur before herd immunity is reached through the vaccine.

Next, consider Equation (2). This transition equation says that, on average, the number of total recoveries in location l in time period t (conditional on the state in period $t - 1$) is equal to number of individuals in location l that have already recovered by time period $t - 1$ plus some fraction of infected individuals in period $t - 1$. This fraction is determined by the parameter λ which is the recovery rate from Covid-19. Equation (3) is the transition equation for deaths. The key parameter is γ which parameterizes the death rate from being infected with Covid-19.

Next, consider Equation (1). This is the transition equation for active Covid-19 cases. The expected number of infections in period t thus depends on (i) the remaining cases after accounting for recoveries and deaths (this is the first term in Equation (1)), and (ii) the expected number of new cases (this is the second term in Equation (1)). Finally, in Equation (4), the expected number of susceptible individuals is equal to the number of susceptible individuals in time period $t - 1$ minus the expected number of new cases; likewise, the expected number of cumulative cases by time period t is equal to the number of cumulative cases in time period $t - 1$ plus the expected number of new cases.

3 Identification Strategies for Policy Effects on Covid-19 Cases

The previous section presented a basic stochastic SIRD model. This section connects that sort of model with the treatment effects literature and considers the relative merits of difference in differences and unconfoundedness strategies for evaluating the effect of a policy on the number of Covid-19 cases.

The strategy of this section is to impose the stochastic SIRD model for untreated potential outcomes and to check if differences in differences and/or unconfoundedness are compatible with the stochastic SIRD model. This setup does not place restrictions on how treated potential outcomes (i.e., Covid-19 cases under the policy) are generated. In particular, this is consistent with Covid-19 cases under the policy continuing to follow a stochastic SIRD model but where

⁴This is also one explanation for repeated “waves” of Covid-19 cases. That is, the infection rate may be temporarily reduced by policy intervention or individual choices but then increases again once these interventions are relaxed.

the values of the parameters potentially change in response to the policy; but it is also more general than that in the sense that there are no substantive restrictions on treated potential outcomes.

Additional Treatment Effects Notation To make the connection with the treatment effects literature, we start by introducing some additional notation. First, we define D_l as a binary variable indicating whether or not location l participated in the treatment. We also define treated and untreated versions of all of the variables in the stochastic SIRD model. In particular, for generic time period t , $S_{lt}(0)$, $I_{lt}(0)$, $R_{lt}(0)$, and $\delta_{lt}(0)$ are the number of susceptible, infected, recovered, and dead individuals in location l in time period t if the policy had not been enacted. We also define $C_{lt}(0)$ as the cumulative number of cases in location l by time period t . Similarly, we define $S_{lt}(1)$, $I_{lt}(1)$, $R_{lt}(1)$, $\delta_{lt}(1)$, and $C_{lt}(1)$ to be the corresponding treated potential variables; i.e., the values of each of these if the policy had been enacted. Following a large literature on policy evaluation which exploits having access to panel data, we consider the case where the researcher has access to some pre-treatment periods. We suppose that the policy is implemented for treated locations in time period t^* where $1 < t^* \leq \mathcal{T}$.⁵ For random variables indexed by time periods, we define $\Delta X_t := (X_t - X_{t-1})$. Because we are also interested in how policy effects vary across time, some of our arguments involve “long differences” where, for $t_2 > t_1$, we define $\Delta^{(t_1, t_2)} X_t := X_{t_2} - X_{t_1}$. In Appendix A.1, we write the SIRD model given in the previous section in terms of untreated potential outcomes, and we refer to this model as the [Stochastic SIRD Model for Untreated Potential Outcomes](#) throughout the remainder of the paper.

Our main interest for this part of the paper is the effect of the policy on the cumulative number of Covid-19 cases. Typically, the main parameter of interest in DID applications (and the parameter that we focus on in the current paper) is the Average Treatment Effect on the Treated (ATT). It is given by

$$ATT_t^C = \mathbb{E}[C_t(1) - C_t(0) | D = 1] \tag{6}$$

where we index the ATT by C to indicate that we are considering the effect of the policy on the cumulative number of cases in time period t . ATT_t^C is the difference between cumulative cases under the policy relative to cumulative cases in the absence of the policy on average among locations that participated in the treatment. That this parameter is disaggregated by time period makes it straightforward to report across time periods (as in an event study), but it is also straightforward to, for example, average it across post-treatment time periods in order to report an overall average effect of participating in the treatment.

⁵In practice, the timing of implementing a particular policy may vary across different locations. Extending our arguments to this case is relatively straightforward, and, therefore, this section considers the case where the policy is implemented at the same time across all treated locations. See Remark 2 below for additional discussion on this point.

3.1 Using Difference in Differences to Evaluate the Policy Effects on Covid-19 Cases

The main underlying motivation for considering a DID approach is when a researcher thinks that untreated potential outcomes are generated from a two-way fixed effects model (see, for example, Blundell and Costa Dias (2009)). These sorts of models are attractive in many applications in economics where there are thought to be important unobserved differences between individuals (or firms, etc.) that are not observed by the researcher. In labor economics, these are often thought of as being unobserved skill; in industrial organization, these may be unobserved differences in productivity across firms; and, in health economics, these may be thought of as proneness to particular health conditions. However, there is an important difference of Covid-19 relative to all of these cases. In general, particular locations do not have time invariant unobservables that make them more or less likely to have a large number of cases; instead, the key differences between locations are (i) the timing of their initial case(s), and (ii) the pandemic response (both in terms of policies and in terms of actions taken by the populations in different locations).⁶

The main result in this section is that there are likely to be major drawbacks to using DID to evaluate the effects of Covid-19 related policies on the number of Covid-19 cases. The two primary reasons for this are (i) the highly nonlinear spread of Covid-19 cases during a pandemic and (ii) that the key difference between locations is the current number of Covid-19 cases rather than some fixed unobserved difference between locations in terms of “proneness” to having a large number of cases.

In this section, we consider whether difference in differences approaches are compatible with the stochastic SIRD model presented above. We begin by providing some background on using difference in differences to identify the effect of some policy. The key identifying assumption in a DID application is the following parallel trends assumption.

Parallel Trends Assumption. For all $t = 2, \dots, \mathcal{T}$

$$\mathbb{E}[\Delta C_t(0)|D = 1] = \mathbb{E}[\Delta C_t(0)|D = 0]$$

The parallel trends assumption says that the path of Covid-19 cases that locations in the treated group would have experienced if they had not participated in the treatment is the same as the path of Covid-19 cases that locations in the untreated group did experience. Under this assumption, it is straightforward to show that, for $t \geq t^*$,

$$ATT_t^C = \mathbb{E}[\Delta^{(t^*-1,t)} C_t | D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)} C_t | D = 0] \quad (7)$$

In other words, under the parallel trends assumption, ATT_t^C is equal to the path of Covid-19 cases that treated locations experienced adjusted by the path of Covid-19 cases that untreated location experienced; under the parallel trends assumption, the latter is the path of Covid-19

⁶One caveat to this is that different locations may have characteristics that are related to the parameters of the SIRD model discussed above. See Remark 3 below for more discussion along these lines.

cases that treated locations would have experienced on average if they had not experienced the policy. And, regardless of whether or not the parallel trends assumption holds, the estimand on the right hand side of Equation (7) is what is estimated in DID applications on Covid-19. Before providing our main result on using difference in differences to identify/estimate the effect of a policy on Covid-19 cases, it is also worth mentioning that the primary motivating model for differences in differences identification strategies is one where

$$C_{lt}(0) = \theta_t + \eta_l + v_{lt} \quad (8)$$

where θ_t is a time fixed effect, η_l is location-specific unobserved heterogeneity that can be distributed differently between the treated group and untreated group and v_{lt} is an idiosyncratic time varying unobservable. Comparing Equation (8) to the equation for cumulative Covid-19 cases in Equation (5), it is immediately clear that these are notably different. In the stochastic SIRD model, the important difference between treated and untreated locations is not unobserved heterogeneity, but rather differences in the current number of Covid-19 cases and the number of susceptible individuals across locations.

The next result makes explicit that the parallel trends assumption is generally violated in stochastic SIRD models.

Theorem 1. *In the stochastic SIRD model discussed above*

$$\mathbb{E}[\Delta^{(t^*-1,t)}C_t(0)|D = d] = \sum_{s=t^*}^t \mathbb{E}\left[\mathbb{E}[\Delta C_s(0)|\mathcal{F}_{t^*-1}, D = d]|D = d\right]$$

which implies that

- (i) *The parallel trends assumption does not generally hold*
- (ii) *Further, the bias from incorrectly imposing the parallel trends assumption is given by*

$$\begin{aligned} & \left(\mathbb{E}[\Delta^{(t^*-1,t)}C_t|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}C_t|D = 0] \right) - ATT_t^C \\ & = \left(\sum_{s=t^*}^t \mathbb{E}\left[\mathbb{E}[\Delta C_s(0)|\mathcal{F}_{t^*-1}, D = 1]|D = 1\right] - \sum_{s=t^*}^t \mathbb{E}\left[\mathbb{E}[\Delta C_s(0)|\mathcal{F}_{t^*-1}, D = 0]|D = 0\right] \right) \end{aligned}$$

The proof of Theorem 1 is provided in Appendix B. Theorem 1 shows that difference in differences generally delivers (potentially severely) biased estimates of the effect of a policy on cumulative Covid-19 cases. It is worth making a few additional comments before proceeding. First, the key reason why the difference in differences strategy breaks down is that, in general, the distribution of the “state” of pandemic related variables immediately before the policy (contained in \mathcal{F}_{t^*-1}) is not the same across treated and untreated locations. Due to the nonlinearity of the SIRD model, this leads to violations of the parallel trends assumption. Second, the sign of the bias cannot generally be determined from these expressions. For example, in Figure 2 above, difference in differences resulted in downward biased estimates of the effect of the policy, but the direction of the bias is sensitive to both (i) timing of first cases in treated and untreated

locations, and (ii) the timing of the policy itself (this can be clearly seen in Panel (a) of Figure 2 where setting the policy at an alternative time period could result in parallel trends being violated in the opposite direction).

Some of the expressions in Theorem 1 seem complicated. One special case of this result that is worth pointing out is when $t = t^*$ (so that we are considering the effect of the policy on Covid-19 cases “on impact”). In that case, the bias from using DID is given by

$$\mathbb{E} \left[\beta \frac{I_{t^*-1}(0)}{N} S_{t^*-1}(0) | D = 1 \right] - \mathbb{E} \left[\beta \frac{I_{t^*-1}(0)}{N} S_{t^*-1}(0) | D = 0 \right]$$

This bias is the difference between the expected number of new cases that treated locations would have experienced in the absence of the policy relative to the expected number of new cases for untreated locations. And, here, it is straightforward to see key reasons why difference in differences can perform poorly: if the joint distribution of currently infected and number of susceptible individuals is different among treated and untreated locations, then they would have experienced a different number of new Covid-19 cases *even if the policy had not been implemented*. In the context of Covid-19, there are some cases where these biases could be substantial. Perhaps the leading example is when the timing of initial Covid-19 cases varied across locations and Covid-19 related policies were implemented earlier in locations that tended to have cases earlier.

3.2 Unconfoundedness in SIRD Models

A main alternative to DID for evaluating the effects of policies is to assume some version of unconfoundedness. Unconfoundedness means that, after conditioning on some covariates, treatment assignment is as good as randomly assigned. In other words, in order to identify the effect of some policy on Covid-19 cases, one can compare Covid-19 cases in locations that experienced the treatment to Covid-19 cases in locations that did not participate in the treatment *and had the same characteristics related to the pandemic* as treated locations. In this section, we consider a particular version of unconfoundedness that does not suffer from the same limitations as difference in differences for evaluating the effects of policies on the number of Covid-19 cases.

Intuitively, the reason why an unconfoundedness strategy works better for studying policy effects of Covid-19 is that the key differences between locations are the current amount of cases and the current number of susceptible individuals rather than differences in location-specific unobserved heterogeneity. Therefore, conditioning on current cases and the current number of susceptible individuals is sufficient for comparisons of treated and untreated locations to deliver causal effects of policies on Covid-19 cases; while the differencing strategy of difference in differences is not able to do the same.

The next result is a main result on the validity of identifying policy effects under the assumption of unconfoundedness. Before stating this result, define the propensity score as

$$p(\mathcal{F}_{t^*-1}) := \mathbb{P}(D = 1 | \mathcal{F}_{t^*-1})$$

which is the probability of being treated conditional on pre-treatment characteristics \mathcal{F}_{t^*-1} and make the following assumption

Assumption 1 (Overlap). *There exists some $\epsilon > 0$ such that $P(D = 1) > \epsilon$ and $p(\mathcal{F}_{t^*-1}) < 1 - \epsilon$ almost surely.*

Assumption 1 is a standard assumption in the treatment effects literature. In the context of Covid-19 related policies, the first part says that there are some locations that participate in the treatment, and the second part says that, for all values of \mathcal{F}_{t^*-1} , one can find untreated locations that have those characteristics. This implies that, for all treated locations, there exists matching untreated locations with the same pre-treatment characteristics. In practice, if this condition is violated, one can identify treatment effects that are local to the region of common support (see, for example, Crump, Hotz, Imbens, and Mitnik (2009)).

Proposition 1. *In the [Stochastic SIRD Model for Untreated Potential Outcomes](#) and under Assumption 1, and for any $t \geq t^*$,*

$$\mathbb{E}[C_t(0)|\mathcal{F}_{t^*-1}, D = 1] = \mathbb{E}[C_t(0)|\mathcal{F}_{t^*-1}, D = 0]$$

The proof of Proposition 1, provided in Appendix B, is very straightforward and holds essentially immediately in the [Stochastic SIRD Model for Untreated Potential Outcomes](#). That said, it is an important result and implies that, on average, the unobserved number of cumulative cases that locations that participated in the treatment would have experienced if they had not participated in the treatment is the same as the cumulative number of cases that untreated locations actually did experience *among locations that had the same pre-treatment characteristics*.

Finally, in this section, we provide an identification result for ATT_t^C which is valid under the SIRD model for Covid-19 cases.

Theorem 2. *In the [Stochastic SIRD Model for Untreated Potential Outcomes](#) and under Assumption 1, and for any $t \geq t^*$,*

$$ATT_t^C = \mathbb{E} \left[\omega(D, \mathcal{F}_{t^*-1}) (C_t - m_{0,t}^C(\mathcal{F}_{t^*-1})) \right] \quad (9)$$

where

$$\omega(D, \mathcal{F}_{t^*-1}) := \frac{D}{\mathbb{E}[D]} - \frac{\frac{p(\mathcal{F}_{t^*-1})}{1-p(\mathcal{F}_{t^*-1})}(1-D)}{\mathbb{E} \left[\frac{p(\mathcal{F}_{t^*-1})}{1-p(\mathcal{F}_{t^*-1})}(1-D) \right]} \quad \text{and} \quad m_{0,t}^C(\mathcal{F}_{t^*-1}) := \mathbb{E}[C_t|\mathcal{F}_{t^*-1}, D = 0] \quad (10)$$

Theorem 2 says that, under a stochastic SIRD model, we can evaluate the effect of a policy using an unconfoundedness strategy that compares the number of cases in locations that participated in the treatment to the number of cases in locations that did not participate in the treatment *and* which had the same Covid-19 related characteristics in the period before the policy was implemented.

It is worth making several additional comments related to the result in Theorem 2. First, estimating ATT_t^C from the expression in Equation (9) involves estimating the propensity score, $p(\mathcal{F}_{t^*-1})$ and the outcome regression $m_{0,t}^C(\mathcal{F}_{t^*-1})$. It is also possible to derive alternative expressions for ATT_t^C that only require either estimating the propensity score (these would be similar to propensity score re-weighting estimators as in Hirano, Imbens, and Ridder (2003)) or estimating the outcome regression (these would be similar to regression adjustment estimators). However, the expression for ATT_t^C in Equation (9) possesses the double robustness property.⁷ A main advantage of a doubly robust estimator is that it provides consistent estimates of ATT_t^C if either the propensity score model or the outcome regression model are correctly specified (see, for example, Bang and Robins (2005) and Słoczyński and Wooldridge (2018)). Double robustness is particularly appealing in this context as it enables us to side-step the problem of estimating the full SIRD model and instead involves estimating a model of the treatment assignment process which is both familiar to economists and may be substantially more feasible to do with a simple parametric model. In unreported simulations, we found that imposing flexible parametric models for both the propensity score and the outcome regression performed notably better than either the pure outcome regression approach or the propensity score re-weighting approach.

Second, it is worth briefly mentioning that the weights in Equation (7) are normalized to have mean one in finite samples. This type of normalized weights is said to be of the Hájek-type (Hájek (1971)) and typically results in estimators with improved finite sample properties relative to its unnormalized counterpart (Busso, DiNardo, and McCrary (2014)). Finally, we provide the asymptotic properties of our estimator in the Supplementary Appendix. In order to conduct inference, we use a multiplier bootstrap procedure that involves perturbing the influence function of the estimator of ATT_t^C ; we also discuss how to conduct uniform inference across different time periods to account for multiple testing. These results primarily follow from recent results on doubly robust estimators with Hájek-type weights in Sant’Anna and Zhao (2020).

Remark 1. *The results in this section have focused on the effect of a policy on the number of cumulative cases. However, it is straightforward to show that the same sorts of arguments apply to other possible variables of interest such as the current number of infections (which we use in the next section).*

Remark 2. *The identification arguments in this section have been for the case where the timing of the policy does not vary across different locations. However, it is straightforward to extend these arguments to the case where the timing of the policy varies across locations (as is the case in our application). In this case, one can think of our identification arguments holding specifically for each “group” where a group is defined by the time period when a location first becomes treated. In this case, instead of identifying ATT-type parameters, one would identify group-time average treatment effects along the lines of Callaway and Sant’Anna (2020) and can*

⁷For completeness, we provide a proof in the Supplementary Appendix, but the arguments are identical to arguments for existing doubly robust estimators under unconfoundedness.

follow their approach to aggregating these sorts of parameters into an overall average effect of participating in the treatment or into an event study type result. This is the approach that we follow in the application.

Remark 3. *It is also worthwhile to consider extensions of the [Stochastic SIRD Model for Untreated Potential Outcomes](#) and which sorts of extensions are compatible with the unconfoundedness condition discussed in this section. Let $\nu = (\beta, \lambda, \gamma)$ denote the collection of parameters in the [Stochastic SIRD Model for Untreated Potential Outcomes](#) and consider the following extensions of the model:*

- *Time varying parameters (i.e., ν can vary across time but not across locations): Unconfoundedness is compatible with time varying SIRD parameters that are common across all locations.*
- *Parameters that vary arbitrarily by location (i.e., ν can vary arbitrarily across locations but not through time). Unconfoundedness would not hold in this case as each location would essentially be experiencing its own unique pandemic. An alternative approach that might be feasible here would be to estimate the parameters of the SIRD model separately for each location in pre-treatment periods and then use these estimates to impute the outcomes that a particular location would have experienced if the policy had not been implemented.*
- *Parameters that vary arbitrarily over time and with observed location characteristics (i.e., ν can vary arbitrarily across locations and time). Unconfoundedness will not hold in this case either, but when parameters can vary arbitrarily across locations and time, it is likely that there is not enough structure in this case for any policy evaluation strategy to work.*
- *Parameters that vary with observed location-specific characteristics and vary over time. This is a realistic setup as the parameters of the SIRD model likely can vary with location-specific characteristics such as demographics and population density. Unconfoundedness can hold in this type of SIRD model though the location-specific characteristics should be included among the conditioning variables in addition to the pandemic-related variables discussed above.*

As a final comment, all the restrictions discussed above are for the SIRD model for untreated potential outcomes. The approach considered in the paper does not put any structure on treated potential outcomes. In particular, they could follow any of the SIRD setups mentioned above but, more generally, do not have to be generated by a SIRD model at all).

4 Identification Strategies for Policy Effects on Economic Outcomes

Another interest of economists is studying the effect of Covid-19 related policies on other (particularly economic) outcomes. This is likely to be useful for thinking about a cost-benefit analysis of particular policies. Relative to textbook versions of difference in differences, what

is different in this section is that we allow for economic outcomes to depend on the current number of Covid-19 cases in a particular location.⁸ We denote the economic outcome of interest by Y_{lt} which is the observed economic outcome for location l in time period t . We also define treated potential outcomes, $Y_{lt}(1)$, and untreated potential outcomes, $Y_{lt}(0)$, and note that $Y_{lt} = D_l Y_{lt}(1) + (1 - D_l) Y_{lt}(0)$. The target parameter in this section is given by

$$ATT_t^Y = \mathbb{E}[Y_t(1) - Y_t(0)|D = 1]$$

which is the difference between treated potential outcomes and untreated potential outcomes on average, in time period t , and among treated locations. As discussed above, DID designs are closely related to TWFE models for untreated potential outcomes, and, in this section, we consider the following model for untreated potential outcomes

$$Y_{lt}(0) = \tau_t + \xi_l + \alpha I_{lt}(0) + v_{lt} \quad (11)$$

where τ_t is a common macro shock. For economic outcomes, there is clear evidence of common macroeconomic shocks which can be motivated by, for example, common information about the dangers of Covid-19 across locations. ξ_l is a location specific fixed effect allowing for time-invariant location-specific differences in economic outcomes, and v_{lt} are idiosyncratic, time varying unobservables.

What is different about this model from standard DID is the term involving $I_{lt}(0)$ where $I_{lt}(0)$ is the number of Covid-19 cases in location l in time period t if the policy were not implemented. It is likely to be very important to include this sort of term during the pandemic as it allows for economic outcomes to depend on the local spread of cases. In particular, this allows for current cases to directly affect outcomes as well as individuals and/or firms taking more Covid-19 related precautions when the number of local cases is high.

In this section, we propose an approach that is able to deliver consistent estimates of ATT_t^Y in the case when policies can have an effect on current Covid-19 cases and current Covid-19 cases can, in turn, have an effect on the outcome of interest. Throughout this section, we contrast our suggested approach with two very common DID-type approaches. First, we consider the case where a researcher compares the path of outcomes of treated locations to the path of outcomes among untreated locations. Throughout this section, we refer to this case as “standard DID”. Second, we consider a version of DID that includes the number of cases as a regressor. Throughout this section, we refer to this case as “regression DID”. We show that both of these approaches generally deliver biased estimates of ATT_t^Y under the model in Equation (11). In the standard DID case, biased estimates arise because the researcher does not account for current cases in a particular location having a direct effect on outcomes. In the regression DID case, biased estimates arise because the approach does not accommodate the possibility that the policy has an effect on the current number of cases (which in turn has an

⁸As above, because the target parameter is an ATT-type parameter, the setup in this section does not require assumptions on how treated potential outcomes are generated and, therefore, the discussion about the effect of current cases and SIRD models in this section need only apply for untreated potential outcomes.

effect on outcomes).

In light of this discussion, we propose an alternative approach that simultaneously addresses both of these issues. We call our approach “adjusted regression DID”. Our idea is to include an adjustment term that accounts for the possibility that the policy affects the current number of cases. This adjustment term is closely related to the arguments in the previous section; in particular, we can recover an estimate of the number of active cases that a treated location would experience in a particular time period by recovering the number of active cases in untreated locations with similar pre-treatment pandemic-related characteristics.

Before stating the main result in this section, it is helpful to notice that, in the model in Equation (11),

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D = d] = \tilde{\tau}_t + \alpha\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = d] \quad (12)$$

where we define $\tilde{\tau}_t := (\tau_t - \tau_{t^*-1})$. Also note that $\tilde{\tau}_t$ and α are both identified using the untreated group (in that case, untreated potential outcomes and untreated potential active cases are observed in all time periods which implies that the parameters are identified as this amounts to a simple linear regression of $\Delta^{(t^*-1,t)}Y_t$ on $\Delta^{(t^*-1,t)}I_t$ using untreated locations).

Next, to fix ideas, under standard DID, the estimator of ATT_t^Y is the sample analogue of

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 0]$$

Likewise, under regression DID, the estimator of ATT_t^Y is the sample analogue of

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D = 1] - (\tilde{\tau}_t + \alpha\mathbb{E}[\Delta^{(t^*-1,t)}I_t|D = 1])$$

Including covariates in this sort of way is a common strategy⁹ and would amount to comparing paths of outcomes for treated and untreated locations that experienced the same change in cases over time.

The next result provides an alternative identification result for ATT_t^Y as well as results for the bias of standard DID and regression DID.

Theorem 3. *In the model for untreated potential outcomes Equation (11) and [Stochastic SIRD Model for Untreated Potential Outcomes](#) and under [Assumption 1](#), and for any $t \geq t^*$,*

$$\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 1] = \mathbb{E}\left[\frac{D}{\mathbb{E}[D]}\Delta^{(t^*-1,t)}I_t - \omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^I(\mathcal{F}_{t^*-1}))\right] \quad (13)$$

⁹Notice that this estimand is similar in spirit, though not exactly the same, as two way fixed effects regressions that include a treatment dummy variable along with other time varying covariates. Besides the issues pointed out in this section (related to the covariates), those sorts of regressions do not generally deliver an interpretable treatment effect parameter in the case with multiple time periods and variation in treatment timing (see, for example, [Goodman-Bacon \(2019\)](#)). The estimand mentioned above avoids the issues related to multiple periods and variation in treatment timing but, as we point in this section, still suffers from issues stemming from actual cases in treated locations not being equal to what cases would have been if the policy had not been implemented.

where $m_{0,t}^I(\mathcal{F}_{t^*-1}) := \mathbb{E}[I_t | \mathcal{F}_{t^*-1}, D = 0]$ and ω are the same weights as in Theorem 2. Moreover,

$$ATT_t^Y = \mathbb{E}[\Delta^{(t^*-1,t)} Y_t | D = 1] - \left\{ \tilde{\tau}_t + \alpha \left(\mathbb{E}[\Delta^{(t^*-1,t)} I_t(0) | D = 1] \right) \right\} \quad (14)$$

and all the terms on the RHS of the expression for ATT_t^Y are identified. In addition, standard difference in differences is biased with bias given by

$$\mathbb{E}[\Delta^{(t^*-1,t)} Y_t | D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)} Y_t | D = 0] - ATT_t^Y = \alpha (\mathbb{E}[\Delta^{(t^*-1,t)} I_t(0) | D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)} I_t(0) | D = 0])$$

and the bias of regression DID (i.e., including current cases as a covariate) is given by

$$\mathbb{E}[\Delta^{(t^*-1,t)} Y_t | D = 1] - (\tilde{\tau}_t + \alpha \mathbb{E}[\Delta^{(t^*-1,t)} I_t | D = 1]) - ATT_t^Y = -\alpha ATT_t^I$$

where $ATT_t^I := \mathbb{E}[I_t(1) - I_t(0) | D = 1]$.

It is worth sketching the arguments underlying the result in Theorem 3. To start with, notice that

$$ATT_t^Y = \mathbb{E}[\Delta^{(t^*-1,t)} Y_t | D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)} Y_t(0) | D = 1]$$

The bias of standard DID arises from (incorrectly) setting $\mathbb{E}[\Delta^{(t^*-1,t)} Y_t(0) | D = 1] = \mathbb{E}[\Delta^{(t^*-1,t)} Y_t | D = 0]$. In general, this sort of substitution is not appropriate because the path of outcomes that treated locations would have experienced in the absence of participating in the treatment depends on the path of active cases (which is not accounted for here).

Next, based on the model in Equation (11), it follows from Equation (12) that

$$\mathbb{E}[\Delta^{(t^*-1,t)} Y_t(0) | D = 1] = \tilde{\tau}_t + \alpha \mathbb{E}[\Delta^{(t^*-1,t)} I_t(0) | D = 1] \quad (15)$$

The bias of regression DID (that directly includes current cases as a covariate) comes from (incorrectly) setting $\mathbb{E}[\Delta^{(t^*-1,t)} I_t(0) | D = 1] = \mathbb{E}[\Delta^{(t^*-1,t)} I_t | D = 1]$. This strategy is also not generally appropriate because the policy can change (and is likely targeted at changing) the path of active cases.

By contrast, our approach uses the expression in Equation (13) for $\mathbb{E}[\Delta^{(t^*-1,t)} I_t(0) | D = 1]$. This expression takes the observed path of active cases and subtracts from it the effect of the policy on active cases (which is the term $\mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^I(\mathcal{F}_{t^*-1}))]$ and holds under the [Stochastic SIRD Model for Untreated Potential Outcomes](#)). Notice that this term is analogous to the expression for the effect of the policy on cumulative cases in Theorem 2 in the previous section. The difference between the observed path of active cases among treated locations and the effect of the policy on active cases recovers the path of active cases that would have occurred if the policy had not been implemented. Given this expression, it can be plugged into Equation (15) to recover the the path of untreated potential outcomes and, hence, to recover ATT_t^Y .

Estimation: The above discussion suggests the following estimation strategy:

1. Run a regression of $\Delta^{(t^*-1,t)}Y_{lt}$ on $\Delta^{(t^*-1,t)}I_{lt}$ using untreated locations in order to estimate $\tilde{\tau}_t$ and α .
2. Estimate the propensity score, $p(\mathcal{F}_{t^*-1})$, using the entire sample. Plug these estimates into the sample analogue of Equation (10) in order to estimate the weights ω . Estimate $m_{0,t}^I(\mathcal{F}_{t^*-1})$ using untreated locations. Plug in the estimates of ω and $m_{0,t}^I$ into the sample analogue of Equation (13) to compute an estimate of $\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D = 1]$.
3. Plug in the preliminary estimators in Steps 1 and 2 to estimate ATT_t^Y directly using the expression in Theorem 3.

In the Supplementary Appendix, we provide the asymptotic distribution of our estimator of ATT_t^Y . The estimation procedure involves several steps, but each step is parametric and the limiting distribution of the estimate of ATT_t^Y can be obtained following well-known arguments about multiple step estimation procedures that account for estimation effects of each step. In particular, the term $\mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^I(\mathcal{F}_{t^*-1}))]$ can be handled using exactly the same arguments as in the previous section. The other steps in the estimation procedure only involve either running simple parametric regressions or directly calculating averages and are therefore straightforward to account for. As earlier, in practice, we use the multiplier bootstrap to conduct inference and discuss how to conduct uniform inference across different time periods.

Remark 4. *In general, it is not possible to sign the bias from using standard DID or regression DID (as an extreme example, over long enough time horizons, a policy that is effective at slowing the spread of Covid-19 could lead to higher current infections if untreated locations reach herd immunity). That said, over relatively short time horizons, it is possible to get a sense of the likely directions of bias. In particular, suppose that (i) $\alpha < 0$ so that more current cases lead to lower economic outcomes, (ii) that the time horizon is short and the policy decreases the number of active cases over a short horizon, (iii) that the pandemic is in its early stages and that treated locations tend to have earlier arrival times of their first cases (so that, in the absence of participating in the treatment, treated locations would have experienced larger increases in the number of active cases than untreated locations), and (iv) the policy has a negative effect on economic outcomes. In this case, both standard DID and regression DID (that adjusts for the actual number of cases) will both overstate the magnitude of the effect of the policy.*

5 Monte Carlo Simulations

In this section, we provide some Monte Carlo simulations to demonstrate the performance of the main estimation strategies considered in the paper. To begin with, we consider estimating the effect of a policy on cumulative Covid-19 cases. In order to generate the data, we consider the case where untreated potential outcomes are generated by the [Stochastic SIRD Model for Untreated Potential Outcomes](#). The values for the main parameters in the SIRD model are provided in Table 4 in Appendix A. We also suppose that the policy has no effect on the pandemic so that all treatment effects are equal to 0.

Table 1: Monte Carlo Simulations for Cumulative Cases

Policy Time				Unconfoundedness			DID		
				Bias	RMSE	Rej. Prob.	Bias	RMSE	Rej. Prob.
Vary Treated First Case Timing									
150	40	60	0.009	0.478	0.039	-3.044	4.169	0.162	
150	60	60	0.008	0.582	0.044	0.031	3.233	0.055	
150	80	60	-0.012	0.750	0.065	2.931	4.542	0.153	
Vary Policy Timing									
75	40	80	0.034	0.803	0.036	-12.829	14.416	0.469	
150	40	80	0.034	0.428	0.024	-5.593	6.464	0.438	
225	40	80	0.047	0.196	0.031	-1.133	1.389	0.323	
Vary Number of Locations, $n = 1000$									
150	40	80	0.031	0.194	0.044	-5.680	5.895	0.951	

Notes: The table provides Monte Carlo simulations for cumulative cases using the unconfoundedness approach suggested in the paper as well as difference in differences and with the simulation parameters discussed in the text. The column labeled “Policy Time” indicates the timing when the policy is implemented among treated locations; λ_D and λ_U are the mean timing of the first case for treated and untreated locations, respectively. The other columns report the bias, root mean squared error (RMSE), and rejection probabilities for each simulation setup.

Throughout this section, we consider the case where there are 250 locations (we vary this number in a few cases), where the probability of a location being treated is equal to 0.5, and where there are 1000 individuals in each location. We report bias, root mean squared error, median absolute deviation, and rejection probabilities for $H_0 : ATT = 0$ for the average effect of the policy across the first 50 post-treatment time periods (i.e., we compute event-study type estimates for 50 periods following the treatment, average them across event time to get an overall average treatment effect parameter, and compute the properties of this estimator). To implement our doubly robust estimator, we include a third order polynomial (also including all interactions) in the pre-treatment number of infected individuals and pre-treatment number of susceptible individuals both for the outcome regression and for the propensity score. Across simulations, we primarily focus on varying the timing of initial Covid-19 cases among treated and untreated locations, and on varying the treatment timing across treated and untreated locations.

The results for our first set of simulations are provided in Table 1. The high level takeaway from this table is that the unconfoundedness approach uniformly appears to perform better than difference in differences. Difference in differences is severely biased when the timing of initial cases is different between treated and untreated locations (this is in line with our earlier discussion). The magnitude of the bias of difference in differences is also sensitive to the timing of the policy (this holds because the direction/magnitude of violations of parallel trends depends

Table 2: Monte Carlo Simulations for Economic Outcomes

Policy Time	λ_D	λ_U	Adjusted Regression DID			Standard DID			
			Bias	RMSE	Rej. Prob.	Bias	RMSE	Rej. Prob.	
<i>n</i> = 250									
150	40	60	0.000	0.127	0.048	-0.134	0.227	0.092	
150	60	60	0.001	0.132	0.048	0.002	0.193	0.043	
150	80	60	-0.015	0.132	0.049	0.110	0.230	0.081	
<i>n</i> = 1000									
150	40	60	0.003	0.066	0.055	-0.129	0.159	0.263	
150	60	60	0.005	0.067	0.045	0.005	0.098	0.051	
150	80	60	0.001	0.071	0.068	0.127	0.165	0.240	

Notes: The table provides Monte Carlo simulations for economic outcomes using the adjusted regression DID approach suggested in the paper as well as standard DID and with the simulation parameters discussed in the text. The column labeled “Policy Time” indicates the timing when the policy is implemented among treated locations; λ_D and λ_U are the mean timing of the first case for treated and untreated locations, respectively. The other columns report the bias, root mean squared error (RMSE), and rejection probabilities for each simulation setup.

on the shape of pandemic related variables which are, in turn, dependent on how long ago the pandemic started). Across simulations, difference in differences also tends to over-reject.

On the other hand, the doubly robust unconfoundedness approach performs much better with good performance across each specification. Interestingly, even in the case where the first cases show up, on average, at the same time across treated and untreated locations (in this case, as expected, DID appears to be unbiased), the unconfoundedness approach suggested in the paper has notably smaller root mean squared error.

Next, we provide analogous results but for the effect of the policy on economic outcomes. For this part, we generate untreated potential outcomes according to Equation (11). We set $\alpha = -0.1$, $\tau_t = (50 + 20 \times t/\mathcal{T})$, and we set $\eta|D = d \sim N(\mu_d, 1)$ where $\mu_d = 20 - 10d$ for $d \in \{0, 1\}$. We also set the parameters for the pandemic related variables as in the baseline specification discussed above.

These results are provide in Table 2 where we vary the timing of initial cases as well as the number of locations across simulations. As in the previous case, the approach suggested in the paper (adjusted regression DID) performs well uniformly across DGPs. By contrast, standard DID performs less well particularly in cases where the timing of initial cases is systematically different across treated and untreated locations.

6 Application: Effects of Shelter-in-Place Orders on Covid-19 Cases and Travel

To conclude the paper, we briefly apply our approach to study the effect of shelter-in-place orders (SIPOs) on Covid-19 cases and travel. The effects of shelter-in-place orders have been studied by Courtemanche et al. (2020), Dave, Friedson, Matsuzawa, and Sabia (2020), Dave et al. (2020a), and Villas-Boas, Sears, Villas-Boas, and Villas-Boas (2020), among others.

6.1 Data

We consider a period early in the pandemic — March 8, 2020 to April 20, 2020 — when a large number of states implemented shelter-in-place orders. We follow Dave, Friedson, Matsuzawa, and Sabia (2020) in terms of definitions of shelter-in-place orders and the timing of implementation across states. In order to facilitate estimating conditional treatment assignment probabilities, we assign states into “groups” on the basis of the timing when they adopted a shelter-in-place order. And, in particular, we assign states that adopted a shelter-in-place order within a five day window, starting on March 18, into the same group. For example, California was the first state to implement a shelter-in-place order on March 19; Illinois and New Jersey followed on March 21, and New York on March 22. These form a group of states that we refer to as the March 18 group. Fifteen other states adopted shelter-in-place orders between March 23 and March 27 and form the group that we refer to as the March 23 group. We include five such groups total as well as an untreated group of ten states that did not adopt a shelter-in-place order over the time period that we consider.

Next, we obtained data on state-level Covid-19 cases and testing from the Covid Tracking Project (<https://covidtracking.com>). We also obtain 2019 state-level populations from the Census Bureau. Finally, we use travel data from Google’s Covid-19 Community Mobility Report (<https://www.google.com/covid19/mobility>). We focus on state-level retail and recreation travel (these are aggregated together) which is reported as a percentage change relative to pre-Covid travel. We dropped four states (Alaska, Massachusetts, Maryland, and Oregon) due to missing data either on cases or testing on some of the days during the time period that we consider. In order to deal with heterogeneity in terms of state populations, we use versions of pandemic related variables in terms of their number per million in a particular state (e.g., cumulative cases per million). In terms of the SIRD model, this amounts to dividing all variables by N_l and multiplying by one million; this transformation is compatible with the SIRD model. We construct the current number of active Covid-19 cases (and therefore contagious individuals) as the total number of newly reported cases over the past five days; as for the other pandemic related variables, we use the number of current cases per million individuals in a state.

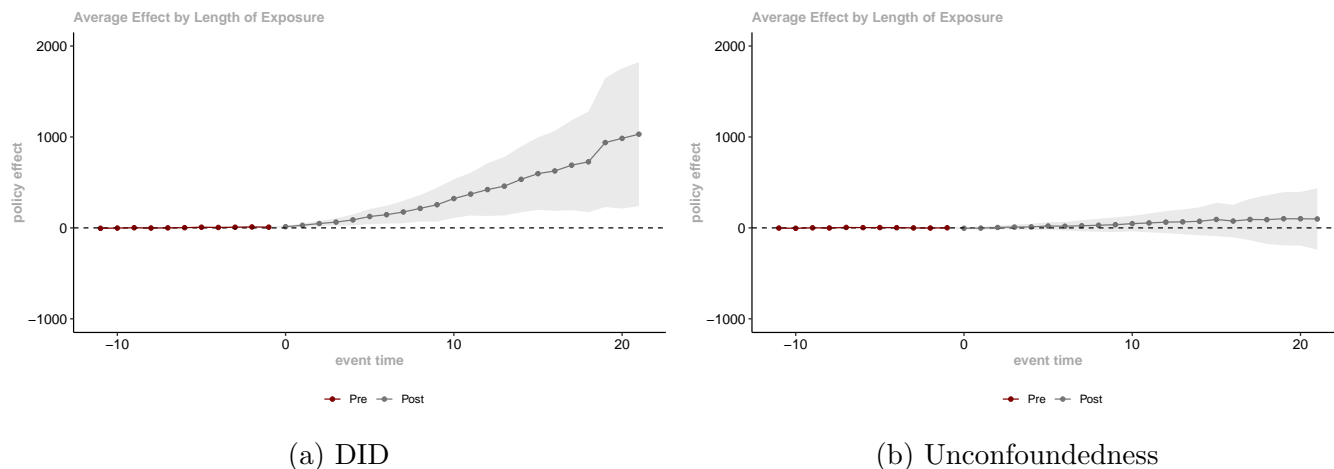
Summary statistics for the data that we use are provided in Table 3. There are some things that are immediately notable from the summary statistics. First, early in the pandemic, the number of cases were substantially different for states that adopted shelter-in-place orders earlier relative to states that adopted them later or that did not adopt them at all. This immediately

Table 3: Summary Statistics

	Group					
	Untreated	Mar 18	Mar 23	Mar 28	Apr 2	Apr 7
group size	10	4	14	12	7	1
pop (millions)	2.9	20.1	4.6	5.1	10.9	5.1
Cumulative Cases per million						
Mar 17	18.4	38.7	166.9	26.0	15.1	6.4
Mar 22	78.4	333.9	408.9	90.8	70.1	37.9
Mar 27	313.0	1024.0	863.2	282.0	273.3	88.6
Apr 1	719.0	2051.8	1525.3	644.8	636.4	251.1
Apr 6	1268.8	3358.6	2435.0	1190.3	1109.7	398.0
Apr 11	2039.7	4711.9	3262.8	1869.3	1739.0	622.9
Apr 16	2861.3	5913.5	3935.1	2411.5	2238.6	710.1
Cumulative Tests per million						
Mar 17	806.6	349.7	1380.2	725.4	453.7	66.8
Mar 22	2311.3	1444.4	3379.2	2009.9	1668.7	322.6
Mar 27	4912.5	3501.2	6764.2	4952.6	3945.4	536.6
Apr 1	8346.5	5816.3	10965.9	8537.1	6980.5	1228.7
Apr 6	12650.5	9974.8	15748.3	12887.0	11416.6	3685.6
Apr 11	18615.6	13966.3	20716.4	18009.4	16101.6	5844.8
Apr 16	23841.2	18099.2	25434.3	21773.1	21032.1	6746.0
Travel Relative to Baseline						
Mar 17	-15.2	-26.0	-19.4	-16.8	-12.1	-8.0
Mar 22	-40.2	-56.0	-40.6	-39.9	-39.6	-39.0
Mar 27	-37.6	-48.5	-47.3	-38.6	-37.1	-32.0
Apr 1	-32.0	-44.8	-39.5	-38.4	-28.6	-30.0
Apr 6	-35.2	-50.2	-40.4	-38.8	-40.3	-33.0
Apr 11	-42.2	-54.2	-45.9	-43.6	-42.1	-41.0
Apr 16	-33.2	-53.2	-37.9	-34.7	-31.4	-26.0

suggests that it will be challenging for DID to perform well to evaluate the effect of shelter-in-place orders on the number of cases. In addition, notice that early treated states (e.g., groups March 18 and March 23) experience very large increases in their number of cases relative to later- and never-adopters of the policy. The next panel shows the number of tests across different groups of states. These differences appear to be relatively smaller than the number of cases which suggests that our results will be driven by actual differences in cases rather than differences in testing rates between treated and untreated states. Finally, the third panel of the table shows changes in retail and recreation trips. The most notable feature of this part of the table is that there were drastic decreases in travel across all states regardless of their shelter-in-place policies.

Figure 5: Policy Effects of SIPOs on Covid-19 Cases



Notes: The figure contains event study type estimates of the effect of SIPOs on the number of cumulative Covid-19 cases. $e = 0$ corresponds to the time period when the policy was implemented. Negative values of e correspond to pre-treatment estimates of the effect of the policy (and can be thought of as pre-tests), and positive values of e correspond to estimates of the effect of the policy at different lengths of exposure to the treatment. Panel (a) contains estimates using a DID approach. Panel (b) contains estimates using an unconfoundedness approach. 95% uniform confidence bands are provided by the shaded areas in each panel.

6.2 Results

Our first set of results are for the effect of shelter-in-place orders on Covid-19 cases. Here, we focus on comparing estimates coming from a difference in differences strategy to estimates coming from an unconfoundedness strategy. Our main results are presented in event study form in Figure 5. Using DID, we estimate that shelter-in-place orders had *increased* Covid-19 cases by just over 1000 cases per million, on average, three weeks after the policies were implemented relative to what they would have been if the policies had not been implemented. These results are statistically significant. More importantly, it is hard to rationalize these sorts of results; in particular, it would seem that shelter-in-place orders could either have no effect or decrease Covid-19 cases, but it is hard to see how they could increase cases. However, even from the summary statistics, one can see that DID estimates are likely to be positive as early policy adopters were tending to experience larger increases in cases. In our view, a better explanation for these results is that Covid-19 was more prevalent earlier in locations that adopted the policy earlier and that the strong, early exponential growth of Covid-19 cases overwhelms any reduction in the infection rate due to the policy. It is exactly in this case where DID would be susceptible to attributing faster growth in Covid-19 cases to the policy rather than to simply a larger number of early cases in treated locations.

Next, we turn to estimates using our approach based on unconfoundedness. As discussed above, these estimates require estimating a model for treatment participation and an outcome regression model. For both of these models, we include a cubic polynomial in the current number of cases in a state; we do not include the number of susceptible individuals as this is very close

to the full population in all states during the period early in the pandemic that we consider; we additionally include a dummy variable for region of the country so that states are compared to other states in the same region; finally, we include the number of tests run per million individuals in the state in order to control for the possibility that some states were detecting Covid-19 cases better than others. When we estimate the propensity score, we find strong indications that the overlap condition is violated indicating that there are a substantial number of states that do not have reasonable comparisons among never-treated and late-treated states. From this step, we drop 20 states from our analysis.¹⁰ These states include some not surprising states such as New York, New Jersey, and California which had large numbers of early Covid-19 cases and all were early implementers of SIPOs. It is, therefore, hard to find reasonable comparisons for these states. Another large number of states from the South are dropped due to the timing of their policies being very similar which, as above, makes it challenging to find reasonable comparison states when region is included as a conditioning variable.

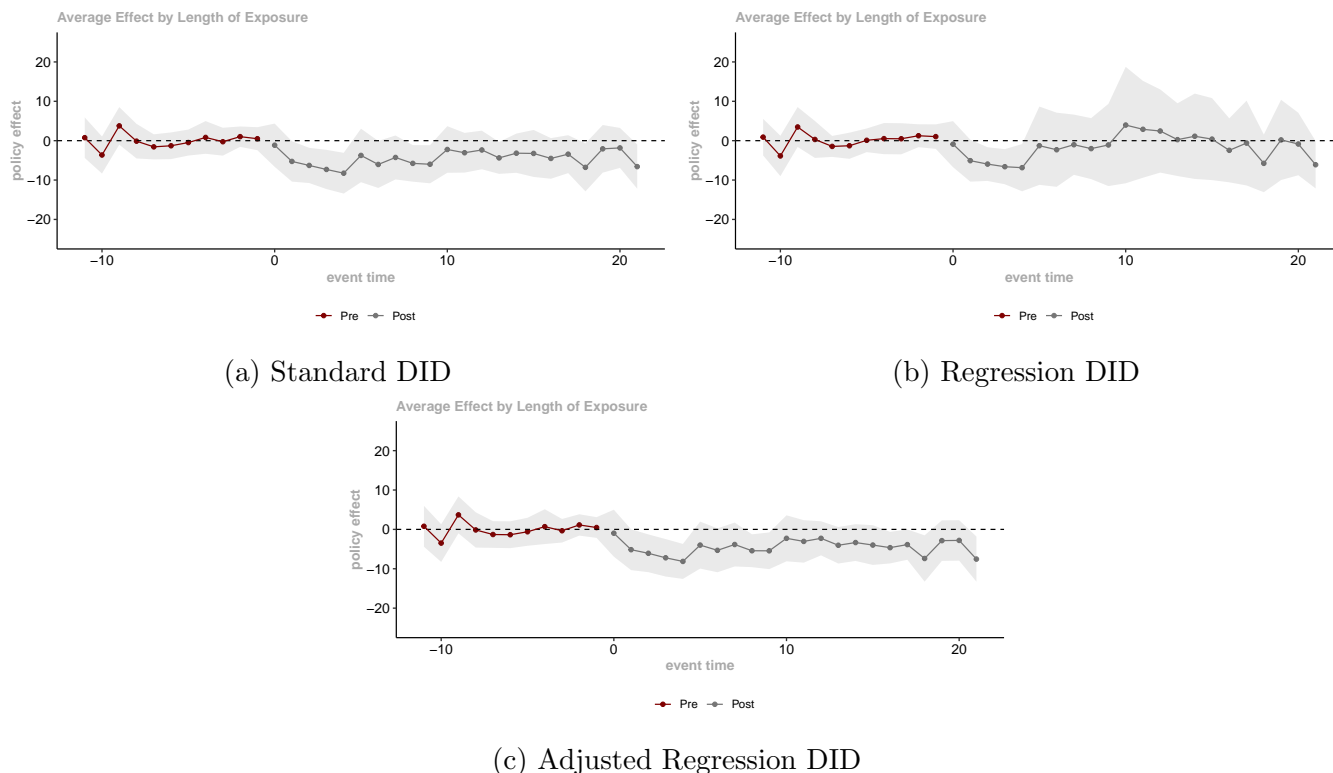
These results are presented in Panel (b) of Figure 5. These results are much different from the previous ones. In this case, we estimate a small and not statistically significant effect of SIPOs on Covid-19 cases.

Before moving on, one other notable feature of Figure 5 is that the pre-treatment estimates of policy effects are very small in both cases. One primary explanation for this in both cases is just that the number of Covid-19 cases is much smaller in early periods relative to later periods (due to the rapid growth of Covid-19 cases over the period that we consider). It is not possible to see in the figure, but, interestingly, the DID estimates of the effect of SIPOs on Covid-19 are positive (though small) and statistically significant when $e = -1$ and $e = -2$ (i.e., the two periods immediately before treatment). For an applied researcher, these results would perhaps provide an empirical suggestion that the DID approach may not be reasonable here; however, without the additional theory coming from the SIRD model, it would also be tempting to compare these “smaller” violations of parallel trends in pre-treatment periods to the larger effects in post-treatment periods and conclude that these violations can safely be ignored. On the other hand, the pre-treatment estimates of policy effects using the unconfoundedness approach suggested in the paper are all small and none are statistically different from zero. These differences in pre-treatment estimated policy effects between DID and unconfoundedness strategies are another piece of suggestive evidence that the unconfoundedness approach is more appropriate for evaluating the effect of the policy.

Finally, we consider the effect of SIPOs on travel. We focus on the percentage change in retail and recreation travel from a pre-Covid baseline, and, for all the results below, we use the group of states that appear to satisfy the overlap condition mentioned above. These results are available in Figure 6. The results in Panel (a) come from a standard DID approach that implicitly imposes that Covid-19 cases do not directly affect travel; the results in Panel (b) come from the regression DID approach that includes current cases as a covariate but not accounting

¹⁰These states include Alabama, California, Connecticut, Florida, Georgia, Illinois, Louisiana, Maine, Michigan, Missouri, Mississippi, New Hampshire, New Jersey, New York, Pennsylvania, Rhode Island, South Carolina, Texas, Vermont, and Washington.

Figure 6: Estimates of SIPO Orders on Travel



Notes: The figure contains event study type estimates of the effect of SIPOs on the percentage change in retail and recreation travel. $e = 0$ corresponds to the time period when the policy was implemented. Negative values of e correspond to pre-treatment estimates of the effect of the policy and can be thought of as pre-tests, and positive values of e correspond to estimates of the effect of the policy at different lengths of exposure to the treatment. Panel (a) provides estimates using standard DID (without accounting for cases), Panel (b) provides regression DID estimates (accounting for cases but not that the policy may have a direct effect on cases), and Panel (c) provides adjusted regression DID estimates (accounting for cases and allowing for the policy to have had an effect on cases as is proposed in the text). 95% uniform confidence bands are provided by the shaded areas in each panel.

for the possibility that SIPOs could have affected the number of cases directly; and the results in Panel (c) use the adjusted regression DID approach proposed in the current paper that allows for the policy to have had an effect on Covid-19 cases.

In this case, the estimates are more broadly similar than they were for cumulative Covid-19 cases previously. Standard DID estimates indicate a relatively small but persistent negative effect of SIPOs on retail and recreation travel.¹¹ Using standard DID, the overall estimate of the effect of SIPOs across post-treatment time periods is -4.44 ($s.e. = 1.51$); i.e., across the first twenty-one days of the SIPO, the policy reduced travel by about 4.5 percentage points relative to what travel would have been if the policy had not been implemented. The point estimates from regression DID (these estimates just include observed current cases as a covariate) are somewhat

¹¹Interestingly, the effect of SIPOs on travel seems to occur more immediately than the effect of SIPOs on Covid-19 cases. This is not surprising though as a SIPO should have an immediate effect on travel, but any effects on Covid-19 cases are likely to take longer to materialize.

smaller in magnitude but substantially noisier; in that case the estimated overall effect of SIPOs on travel is -1.71 ($s.e. = 2.79$). Finally, the adjusted regression DID overall estimated effect of SIPOs is -4.53 ($s.e. = 1.33$) when we allow for the policy to have had an effect on cases.

6.3 Discussion of Results

Our results, especially those about the effect of SIPOs on the number of Covid-19 cases, are substantially different from existing estimates, and it is worth making a few additional comments.

Most other papers that have studied the effect of SIPOs on Covid-19 cases have used variations of difference in differences such as taking the logarithm of the number of cumulative cases and/or including linear time trends. Using essentially the same data, some of these sorts of specifications do appear to be able to deliver an estimate that SIPOs reduced Covid-19 cases, but these variations on DID also do not appear to be compatible with SIRD models (see Allcott et al. (2020) for additional discussion along these lines). A related issue is that most estimates of the effects of SIPOs have been estimated using two-way fixed effects regressions. Recent work (for example, Goodman-Bacon (2019)) has shown that these sorts of regressions are often not able to deliver reasonable estimates of policy effects especially in the case where there is variation in treatment timing (which occurs here due to states adopting SIPOs at different points in time) and treatment effect dynamics (which is especially likely due to the highly nonlinear spread of Covid-19). The approach proposed in the current paper does not suffer from these issues either.

On the other hand, our results on the effect of SIPOs on travel are more similar to existing estimates (Goolsbee and Syverson (2021)). Since the primary channel through which SIPOs would likely reduce Covid-19 cases is through reducing travel/contact with other individuals, it seems reasonable to simultaneously estimate relatively small (but statistically significant) effects of SIPOs on travel coinciding with small (and not statistically significant) effects of SIPOs on the number of Covid-19 cases; but harder to rationalize SIPOs strongly decreasing Covid-19 cases while having only a small effect on travel.

That being said, we hesitate to interpret our results as providing strong evidence that SIPOs did not reduce Covid-19 cases. For one thing, the interpretation of our treatment effect parameters is somewhat subtle. Untreated potential outcomes here do not correspond to particular states not reacting at all to Covid-19 but to outcomes that would have occurred if the state had not implemented the policy (but other things about the state remained the same). This can be seen to be clearly relevant from the summary statistics in Table 3 where all states (not just those who implemented the policy) were experiencing massive decreases in travel over the period that we consider. Importantly, this indicates that our results are not at all saying that staying at home did not have an effect on Covid-19 cases. Second, our standard errors, while not systematically larger than standard errors from DID, are large enough to be compatible with a wide variety of possible effects of SIPOs on Covid-19 cases.¹² Instead, we interpret the results

¹²To give an example, the lower end of a 95% confidence interval for our estimated effect of the policy on Covid-19

from the application as indicating that these sorts of policies are likely to be very challenging to precisely evaluate — our analysis involved data only from 46 (and sometimes fewer) states while trying to deal with a highly nonlinear outcome and a policy that was adopted at different times by states whose exposure to the pandemic also varied widely and was correlated with the timing of the policy being adopted.

7 Conclusion

In this paper, we have considered several different policy evaluation strategies during a pandemic. For identifying the direct effects of policies on the number of Covid-19 cases, our results suggest that strategies based on unconfoundedness type assumptions are likely to perform better than difference in differences type strategies due to the highly nonlinear nature of the spread of Covid-19.

Our second main set of results were about evaluating the effects of policies on other economic outcomes when (i) the policy can affect the number of Covid-19 cases and (ii) the number of Covid-19 cases can have a direct effect on the economic outcome of interest. For this case, we also showed that two of the most common ways to evaluate these policies (difference in differences directly or including the number of cases as a covariate in a DID setup) do not generally deliver an average effect of the policy. We proposed an alternative estimator that is valid in this case.

We applied our approach to study the effects of shelter-in-place orders on Covid-19 cases and travel early in the pandemic. We showed that our theoretical arguments were indeed relevant in this context and led to notably different estimates (particularly for the number of Covid-19 cases) relative to the most common approaches used in applications.

There remain a number of interesting possible extensions to this work, and we conclude by mentioning two of them. First, evaluating the effects of various policies (especially early in the pandemic) is complicated by nonrandomly missing testing data (see, for example, Callaway and Li (2020) and Manski and Molinari (2020)), and it would be interesting to extend our results along these dimensions. Second, another common policy evaluation approach in the context of Covid-19 related policies is synthetic controls (examples include Cho (2020), Dave et al. (2020b), Friedson, McNichols, Sabia, and Dave (2020), and Mitze, Kosfeld, Rode, and Wälde (2020)). It appears that, most often, the researcher’s decision between difference in differences or synthetic controls is driven by whether the number of treated locations is large or small. However, it is less clear if the synthetic control type “interpolations” are compatible with the sorts of epidemiological models that we considered in the current paper (see, for example, Kellogg, Mogstad, Pouliot, and Torgovitsky (2020) for related discussion on synthetic control and matching).

cases is that it reduced cases by 215 per million individuals. If we divide this by the estimated number of Covid-19 cases that treated locations would have experienced in the same period if they had not implemented the policy (i.e., $\mathbb{E}[C_t(0)|D = 1]$ in Equation (6) which is identified and can be recovered in our setup), we would estimate that SIPOs decreased Covid-19 cases by 31%. In other words, our estimates do not necessarily rule out the possibility that SIPOs may have had quite large effects on Covid-19 cases.

References

- [1] Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D Whinston. “Optimal targeted lockdowns in a multi-group SIR model”. Working Paper. 2020.
- [2] Allcott, Hunt, Levi Boxell, Jacob Conway, Billy Ferguson, Matthew Gentzkow, and Benny Goldman. “Economic and health impacts of social distancing policies during the coronavirus pandemic”. Working Paper. 2020.
- [3] Allen, Linda JS. “An introduction to stochastic epidemic models”. *Mathematical Epidemiology*. Springer, 2008, pp. 81–130.
- [4] Allen, Linda JS. “A primer on stochastic epidemic models: Formulation, numerical simulation, and analysis”. *Infectious Disease Modelling* 2.2 (2017), pp. 128–142.
- [5] Bang, Heejung and James M Robins. “Doubly robust estimation in missing data and causal inference models”. *Biometrics* 61.4 (2005), pp. 962–973.
- [6] Bartik, Alexander W, Marianne Bertrand, Feng Lin, Jesse Rothstein, and Matt Unrath. “Measuring the labor market at the onset of the COVID-19 crisis”. Working Paper. 2020.
- [7] Blundell, Richard and Monica Costa Dias. “Alternative approaches to evaluation in empirical microeconomics”. *Journal of Human Resources* 44.3 (2009), pp. 565–640.
- [8] Bonhomme, Stephane and Ulrich Sauder. “Recovering distributions in difference-in-differences models: A comparison of selective and comprehensive schooling”. *Review of Economics and Statistics* 93.2 (2011), pp. 479–494.
- [9] Busso, Matias, John DiNardo, and Justin McCrary. “New evidence on the finite sample properties of propensity score reweighting and matching estimators”. *Review of Economics and Statistics* 96.5 (2014), pp. 885–897.
- [10] Callaway, Brantly and Tong Li. “Understanding the effects of Tennessee’s open Covid-19 testing policy: Bounding policy effects with nonrandomly missing data”. Working Paper. 2020.
- [11] Callaway, Brantly and Pedro HC Sant’Anna. “Difference-in-differences with multiple time periods”. Working Paper. 2020.
- [12] Chernozhukov, Victor, Hiroyuki Kasaha, and Paul Schrimpf. “Causal impact of masks, policies, behavior on early COVID-19 pandemic in the US”. 2020.
- [13] Chetty, Raj, J Friedman, Nathaniel Hendren, and Michael Stepner. “The economic impacts of COVID-19: Evidence from a new public database built from private sector data”. *Opportunity Insights* (2020).
- [14] Cho, Sang-Wook. “Quantifying the impact of nonpharmaceutical interventions during the COVID-19 outbreak: The case of Sweden”. *The Econometrics Journal* 23.3 (2020), pp. 323–344.
- [15] Courtemanche, Charles, Joseph Garuccio, Anh Le, Joshua Pinkston, and Aaron Yelowitz. “Strong social distancing measures in the United States reduced the COVID-19 growth rate:” *Health Affairs* 39.7 (2020), pp. 1237–1246.
- [16] Crump, Richard K, V Joseph Hotz, Guido W Imbens, and Oscar A Mitnik. “Dealing with limited overlap in estimation of average treatment effects”. *Biometrika* 96.1 (2009), pp. 187–199.

- [17] Dave, Dhaval, Andrew Friedson, Kyutaro Matsuzawa, Joseph J Sabia, and Samuel Safford. “Were urban cowboys enough to control COVID-19? Local shelter-in-place orders and coronavirus case growth”. *Journal of Urban Economics* (2020), p. 103294.
- [18] Dave, Dhaval M, Andrew I Friedson, Kyutaro Matsuzawa, Drew McNichols, and Joseph J Sabia. “Did the Wisconsin Supreme Court restart a COVID-19 epidemic? Evidence from a natural experiment”. 2020.
- [19] Dave, Dhaval M, Andrew I Friedson, Kyutaro Matsuzawa, and Joseph J Sabia. “When do shelter-in-place orders fight COVID-19 best? Policy heterogeneity across states and adoption time”. Working Paper. 2020.
- [20] Ellison, Glenn. “Implications of heterogeneous SIR models for analyses of COVID-19”. Working Paper. 2020.
- [21] Fernández-Villaverde, Jesús and Charles I Jones. “Estimating and simulating a SIRD model of COVID-19 for many countries, states, and cities”. Working Paper. 2020.
- [22] Friedson, Andrew I, Drew McNichols, Joseph J Sabia, and Dhaval Dave. “Did California’s shelter-in-place order work? Early coronavirus-related public health effects”. Working Paper. 2020.
- [23] Gapen, Michael, Jonathan Millar, U Blerina, and Pooja Sriram. “Assessing the effectiveness of alternative measures to slow the spread of COVID-19 in the United States”. *Covid Economics* 40 (2020), pp. 46–75.
- [24] Glaeser, Edward L, Ginger Zhe Jin, Benjamin T Leyden, and Michael Luca. “Learning from deregulation: The asymmetric impact of lockdown and reopening on risky behavior during COVID-19”. Working Paper. 2020.
- [25] Goodman-Bacon, Andrew. “Difference-in-differences with variation in treatment timing”. Working Paper. 2019.
- [26] Goodman-Bacon, Andrew and Jan Marcus. “Using difference-in-differences to identify causal effects of COVID-19 policies”. *Survey Research Methods*. Vol. 14. 2. 2020, pp. 153–158.
- [27] Goolsbee, Austan and Chad Syverson. “Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020”. *Journal of Public Economics* 193 (2021), p. 104311.
- [28] Gupta, Sumedha, Laura Montenegro, Thuy Dieu Nguyen, Felipe Lozano-Rojas, Ian M Schmutte, Kosali Ilayperuma Simon, Bruce A Weinberg, and Coady Wing. “Effects of social distancing policy on labor market outcomes”. Working Paper. 2020.
- [29] Hájek, Jaroslav. “Discussion of ‘An essay on the logical foundations of survey sampling, Part I’, by D. Basu”. *Foundations of statistical inference* (1971), p. 326.
- [30] Hirano, Keisuke, Guido Imbens, and Geert Ridder. “Efficient estimation of average treatment effects using the estimated propensity score”. *Econometrica* 71.4 (2003), pp. 1161–1189.
- [31] Huber, Martin. “Mediation analysis”. *Handbook of Labor, Human Resources and Population Economics* (2020), pp. 1–38.
- [32] Imbens, Guido and Jeffrey Wooldridge. “Recent developments in the econometrics of program evaluation”. *Journal of Economic Literature* 47.1 (2009), pp. 5–86.

- [33] Juranek, Steffen and Floris Zoutman. “The effect of social distancing measures on the demand for intensive care: Evidence on Covid-19 in Scandinavia”. Working Paper. 2020.
- [34] Kellogg, Maxwell, Magne Mogstad, Guillaume Pouliot, and Alexander Torgovitsky. “Combining matching and synthetic controls to trade off biases from extrapolation and interpolation”. Working Paper. 2020.
- [35] Kermack, William Ogilvy and Anderson G McKendrick. “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 (1927), pp. 700–721.
- [36] Kong, Edward and Daniel Prinz. “Disentangling policy effects using proxy data: Which shutdown policies affected unemployment during the COVID-19 pandemic?” *Journal of Public Economics* 189 (2020), p. 104257.
- [37] Lechner, Michael. “The estimation of causal effects by difference-in-difference methods”. *Foundations and Trends in Econometrics* 4.3 (2011), pp. 165–224.
- [38] Manski, Charles F and Francesca Molinari. “Estimating the COVID-19 infection rate: Anatomy of an inference problem”. *Journal of Econometrics* (2020).
- [39] Mitze, Timo, Reinhold Kosfeld, Johannes Rode, and Klaus Wälde. “Face masks considerably reduce COVID-19 cases in Germany”. *Proceedings of the National Academy of Sciences* 117.51 (2020), pp. 32293–32301.
- [40] Oka, Tatsushi, Wei Wei, and Dan Zhu. “A spatial stochastic SIR model for transmission networks with application to COVID-19 epidemic in China”. Working Paper. 2020.
- [41] Rowthorn, Robert and Jan Maciejowski. “A cost–benefit analysis of the COVID-19 disease”. *Oxford Review of Economic Policy* 36. Supplement 1 (2020), S38–S55.
- [42] Sant’Anna, Pedro HC and Jun Zhao. “Doubly robust difference-in-differences estimators”. *Journal of Econometrics* 219.1 (2020), pp. 101–122.
- [43] Słoczyński, Tymon and Jeffrey M Wooldridge. “A general double robustness result for estimating average treatment effects”. *Econometric Theory* 34.1 (2018), pp. 112–133.
- [44] Villas-Boas, Sofia B, James Sears, Miguel Villas-Boas, and Vasco Villas-Boas. “Are we #StayingHome to flatten the curve?” Working Paper. 2020.
- [45] Wright, Austin L, Konstantin Sonin, Jesse Driscoll, and Jarnickae Wilson. “Poverty and economic dislocation reduce compliance with Covid-19 shelter-in-place protocols”. *Journal of Economic Behavior & Organization* 180 (2020), pp. 544–554.
- [46] Ziedan, Engy, Kosali I Simon, and Coady Wing. “Effects of state COVID-19 closure policy on non-COVID-19 health care utilization”. Working Paper. 2020.

A More Details on Stochastic SIRD Models

A.1 Stochastic SIRD Model for Untreated Potential Outcomes

In this section, we write down a Stochastic SIRD model along the lines of Equations 1 - 5 but for untreated potential outcomes and written in an error form.

Stochastic SIRD Model for Untreated Potential Outcomes. For all $t = 2, \dots, \mathcal{T}$

$$\begin{aligned}\Delta I_{lt}(0) &= \left(\beta \frac{I_{lt-1}(0)}{N_l} S_{lt-1}(0) + u_{lt}^I \right) - \Delta R_{lt}(0) - \Delta \delta_{lt}(0) \\ \Delta R_{lt}(0) &= \lambda I_{lt-1}(0) + u_{lt}^R \\ \Delta \delta_{lt}(0) &= \gamma I_{lt-1}(0) + u_{lt}^\delta \\ N_l &= S_{lt}(0) + I_{lt}(0) + R_{lt}(0) + \delta_{lt}(0) \\ C_{lt}(0) &= N_l - S_{lt}(0)\end{aligned}$$

where $\mathbb{E}[u_t | \mathcal{F}_{t-1}(0), \dots, \mathcal{F}_1(0)] = 0$ where $u_{lt} := (u_{lt}^I, u_{lt}^R, u_{lt}^\delta)'$. In addition, for some $t^* \leq t \leq \mathcal{T}$,

$$u^{(t^*, t)} \perp\!\!\!\perp D | \mathcal{F}_{t^*-1}$$

where $u_l^{(t^*, t)} := \text{vec}(u_{lt^*}, u_{lt^*+1}, \dots, u_{lt})$

The additional condition in [Stochastic SIRD Model for Untreated Potential Outcomes](#) concerns the error terms in the [Stochastic SIRD Model for Untreated Potential Outcomes](#). It says that, conditional on \mathcal{F}_{t^*-1} , the distribution of the error terms is the same across treated and untreated locations. Relative to typical Stochastic SIRD models, this is a mild assumption. For example, it is typical to impose that, $\left(\beta \frac{I_{lt-1}(0)}{N_l} S_{lt-1}(0) + u_{lt}^I \right)$, which is the number of new infections, follows a Poisson distribution with parameter $\beta \frac{I_{lt-1}(0)}{N_l} S_{lt-1}(0)$, and that the number of new recoveries $(\lambda I_{lt-1}(0) + u_{lt}^R)$, new deaths $(\gamma I_{lt-1}(0) + u_{lt}^\delta)$, and continued infections follows a multinomial distribution with parameters $(\lambda I_{lt-1}(0), \gamma I_{lt-1}(0), (1 - \lambda - \gamma) I_{lt-1}(0))$. In this case, this condition holds by construction.

A.2 Simulation Details

Table 4: Simulation Parameters

parameter	notation	value
infection rate	β	0.08
post policy β	β_{pol}	0.08
recovery rate	λ	0.04
death rate	γ	0.003
time periods	T	400
population size	N	1000
treatment probability	p	0.5
initial case period (treated)		Poisson($\lambda = 40$)
initial case period (untreated)		Poisson($\lambda = 80$)
initial # cases		10
# locations		250
policy start period		150

Table 4 provides the values of the parameters that we use in the simulations. The most important things to notice are that (i) the policy has no effect on Covid-19 infections rates (i.e., $\beta = \beta_{pol}$) and (ii) the distribution of the timing of the first case is different across locations that participate in the treatment and those that do not. This is roughly analogous to the idea that locations that were exposed to Covid-19 earlier tended to implement policies earlier.

B Proofs

Proof of Theorem 1

Proof. For the first part, notice that

$$\begin{aligned}
\mathbb{E}[\Delta^{(t^*-1,t)}C_t(0)|D = d] &= \mathbb{E}[C_t(0) - C_{t^*-1}(0)|D = d] \\
&= \sum_{s=t^*}^t \mathbb{E}[\Delta C_s(0)|D = d] \\
&= \sum_{s=t^*}^t \mathbb{E}\left[\mathbb{E}[\Delta C_s(0)|\mathcal{F}_{t^*-1}, D = d]|D = d\right] \tag{16}
\end{aligned}$$

where the first equality holds by the definition of $\Delta^{(t^*-1,t)}C_t(0)$, the second equality holds by adding and subtracting $\mathbb{E}[C_s(0)|D = d]$ for all $s = t^*, \dots, (t - 1)$, and the last equality holds by the law of iterated expectations.

For the second part,

$$\begin{aligned}
ATT_t^C &= \mathbb{E}[C_t(1) - C_t(0)|D = 1] \\
&= \mathbb{E}[C_t(1) - C_{t^*-1}(0)|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}C_t(0)|D = 1] \\
&= \mathbb{E}[C_t(1) - C_{t^*-1}(0)|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}C_t(0)|D = 0] \\
&\quad + \left(\mathbb{E}[\Delta^{(t^*-1,t)}C_t(0)|D = 0] - \mathbb{E}[\Delta^{(t^*-1,t)}C_t(0)|D = 1]\right) \\
&= \mathbb{E}[\Delta^{(t^*-1,t)}C_t|D = 1] - \mathbb{E}[\Delta^{(t^*-1,t)}C_t|D = 0] \\
&\quad + \left(\sum_{s=t^*}^t \mathbb{E}\left[\mathbb{E}[\Delta C_s(0)|\mathcal{F}_{t^*-1}, D = 0]|D = 0\right] - \sum_{s=t^*}^t \mathbb{E}\left[\mathbb{E}[\Delta C_s(0)|\mathcal{F}_{t^*-1}, D = 1]|D = 1\right]\right)
\end{aligned}$$

where the first equality comes from the definition of ATT_t^C , the second equality holds by adding and subtracting $\mathbb{E}[C_{t^*-1}(0)|D = 1]$ (which is the average number of Covid-19 cases across treated locations in the pre-treatment period), the third equality holds by adding and subtracting $\mathbb{E}[\Delta^{(t^*-1,t)}C_t(0)|D = 0]$ (which is the average path of Covid-19 cases from period $t^* - 1$ to t among untreated locations), and the fourth equality holds by rewriting potential outcomes in terms of their observed counterparts and by Equation (16). This implies the second result in the theorem. \square

Proof of Proposition 1

Proof. First, notice that [Stochastic SIRD Model for Untreated Potential Outcomes](#) implies that, for any time period $s \geq t^*$,

$$\mathcal{F}_s(0) \perp\!\!\!\perp D | \mathcal{F}_{s-1}(0), \mathcal{F}_{t^*-1}(0)$$

Applying this argument recursively, starting with period t , implies the result in the proposition. \square

Proof of Theorem 2

Proof. Given the result in Proposition 1, the proof of this result holds under standard arguments for unconfoundedness. We provide them here for completeness. First, notice that

$$\begin{aligned} ATT_t^C &= \mathbb{E}[C_t(1) - C_t(0)|D = 1] \\ &= \mathbb{E}[C_t|D = 1] - \mathbb{E}[\mathbb{E}[C_t|\mathcal{F}_{t^*-1}, D = 0]|D = 1] \end{aligned} \quad (17)$$

where the first equality holds from the definition of ATT_t^C , and the second equality holds by Proposition 1. This completes the first part of the proof. For the second part, notice that continuing from Equation (17),

$$\begin{aligned} ATT_t^C &= \mathbb{E}\left[\frac{D}{\mathbb{E}[D]}C_t\right] - \mathbb{E}\left[\frac{\mathbb{E}[(1-D)C_t|\mathcal{F}_{t^*-1}]}{1-p(\mathcal{F}_{t^*-1})}\Big|D = 1\right] \\ &= \mathbb{E}\left[\frac{D}{\mathbb{E}[D]}C_t\right] - \mathbb{E}\left[\frac{p(\mathcal{F}_{t^*-1})}{\mathbb{E}[D](1-p(\mathcal{F}_{t^*-1}))}\mathbb{E}[(1-D)C_t|\mathcal{F}_{t^*-1}]\right] \\ &= \mathbb{E}\left[\frac{D}{\mathbb{E}[D]}C_t\right] - \mathbb{E}\left[\frac{p(\mathcal{F}_{t^*-1})(1-D)}{\mathbb{E}[D](1-p(\mathcal{F}_{t^*-1}))}C_t\right] \end{aligned} \quad (18)$$

Further, notice that

$$\begin{aligned} \mathbb{E}\left[\frac{p(\mathcal{F}_{t^*-1})(1-D)}{\mathbb{E}[D](1-p(\mathcal{F}_{t^*-1}))}\right] &= \mathbb{E}\left[\mathbb{E}\left[\frac{p(\mathcal{F}_{t^*-1})(1-D)}{\mathbb{E}[D](1-p(\mathcal{F}_{t^*-1}))}\Big|\mathcal{F}_{t^*-1}\right]\right] \\ &= \frac{\mathbb{E}[p(\mathcal{F}_{t^*-1})]}{\mathbb{E}[D]} \\ &= 1 \end{aligned} \quad (19)$$

Combining the results from Equations (18) and (19) implies that

$$ATT_t^C = \mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})C_t] \quad (20)$$

To conclude the proof, notice that

$$\begin{aligned} \mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})|\mathcal{F}_{t^*-1}] &= \frac{1}{\mathbb{E}[D]}\left(\mathbb{E}[D|\mathcal{F}_{t^*-1}] - \mathbb{E}\left[\frac{p(\mathcal{F}_{t^*-1})(1-D)}{(1-p(\mathcal{F}_{t^*-1}))}\Big|\mathcal{F}_{t^*-1}\right]\right) \\ &= 0 \end{aligned} \quad (21)$$

which holds from the definition of ω and also uses the argument in Equation (19). This implies that

$$\begin{aligned} \mathbb{E}\left[\omega(D, \mathcal{F}_{t^*-1})m_{0,t}^C(\mathcal{F}_{t^*-1})\right] &= \mathbb{E}\left[m_{0,t}^C(\mathcal{F}_{t^*-1})\mathbb{E}[\omega(D, \mathcal{F}_{t^*-1})|\mathcal{F}_{t^*-1}]\right] \\ &= 0 \end{aligned} \quad (22)$$

Combining Equation (18) and Equation (22) implies the second part of the result. \square

Proof of Theorem 3

Proof. For the first part, notice that

$$\begin{aligned}
\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=1] &= \mathbb{E}[I_t(0)|D=1] - \mathbb{E}[I_{t^*-1}|D=1] \\
&= \mathbb{E}[I_t(1)|D=1] - (\mathbb{E}[I_t(1)|D=1] - \mathbb{E}[I_t(0)|D=1]) - \mathbb{E}[I_{t^*-1}|D=1] \\
&= \mathbb{E}[\Delta^{(t^*-1,t)}I_t|D=1] - ATT_t^I \\
&= \mathbb{E}\left[\frac{D}{\mathbb{E}[D]}\Delta^{(t^*-1,t)}I_t - \omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^C(\mathcal{F}_{t^*-1}))\right]
\end{aligned}$$

where the first equality holds by splitting the difference, the second equality adds and subtracts $\mathbb{E}[I_t(1)|D=1]$, the third equality holds by combining terms and the definition of ATT_t^I , and the last equality holds by using the same arguments for ATT_t^I as were used for ATT_t^C in Theorem 2.

For the second part, to start with, notice that

$$\begin{aligned}
ATT_t^Y &= \mathbb{E}[Y_t(1) - Y_t(0)|D=1] \\
&= \mathbb{E}[Y_t(1) - Y_{t^*-1}(0)|D=1] - \mathbb{E}[Y_t(0) - Y_{t^*-1}(0)|D=1] \\
&= \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D=1]
\end{aligned} \tag{23}$$

where the first equality holds by the definition of ATT_t^Y , the second equality holds by adding and subtracting $\mathbb{E}[Y_{t^*-1}(0)|D=1]$, and the third equality holds by replacing the potential outcomes in the first term with their observed counterparts. In Equation (23), the first term is directly identified (it is the observed path of outcomes for treated locations) while the second term is not directly identified, and we consider it in detail below. As a preliminary step, notice that

$$\mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D=0] = \tilde{\tau}_t + \alpha\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=0]$$

Since, $\Delta^{(t^*-1,t)}Y_{t^*-1}(0)$ and $\Delta^{(t^*-1,t)}I_{t^*-1}(0)$ are both observed outcomes for locations in the untreated group, this implies that $\tilde{\tau}_t$ and α are identified and can be recovered from the regression of $\Delta^{(t^*-1,t)}Y_{t^*-1}(0)$ on $\Delta^{(t^*-1,t)}I_{t^*-1}(0)$ among the untreated group. Next,

$$\begin{aligned}
\mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D=1] &= \tilde{\tau}_t + \alpha\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=1] \\
&= \tilde{\tau}_t + \alpha\left(\mathbb{E}\left[\frac{D}{\mathbb{E}[D]}\Delta^{(t^*-1,t)}I_t - \omega(D, \mathcal{F}_{t^*-1})(I_t - m_{0,t}^C(\mathcal{F}_{t^*-1}))\right]\right)
\end{aligned}$$

and all terms are identified. Plugging this expression into Equation (23) implies the result.

Next, we consider the bias coming from using standard DID. This is given by

$$\begin{aligned}
&\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D=0] - (\mathbb{E}[Y_t(1) - Y_t(0)|D=1]) \\
&= \mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}Y_t(0)|D=0] \\
&= \alpha(\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=0])
\end{aligned}$$

where the second equality holds by adding and subtracting $\mathbb{E}[Y_{t^*-1}(0)|D=1]$ and canceling terms and the last equality holds by the model in Equation (11).

Finally, we consider the bias coming from DID including actual cases (rather than cases in the absence of the policy) as a covariate. This bias is given by

$$\begin{aligned}
&\mathbb{E}[\Delta^{(t^*-1,t)}Y_t|D=1] - (\tilde{\tau}_t + \alpha\mathbb{E}[\Delta^{(t^*-1,t)}I_t|D=1]) - (\mathbb{E}[Y_t(1) - Y_t(0)|D=1]) \\
&= \alpha(\mathbb{E}[\Delta^{(t^*-1,t)}I_t(0)|D=1] - \mathbb{E}[\Delta^{(t^*-1,t)}I_t|D=1]) \\
&= -\alpha ATT_t^I
\end{aligned}$$

where the first equality holds by adding and subtracting $\mathbb{E}[Y_{t^*-1}(0)|D = 1]$ and from the model in Equation (11) (and by canceling terms), and the last equality holds by the definition of ATT_t^I after canceling terms. □