UNIVERSITY OF GOTHENBURG
SCHOOL OF BUSINESS, ECONOMICS AND LAW

Evaluating the Effect of School Closures on Societal Transmission of Covid-19: *Evidence from Lower Secondary School Closures in Sweden*

Sebastian Shaqiri

Supervisor: Mikael Lindahl

Master’s thesis in Economics, 30 hec

Spring 2022

Graduate School, School of Business, Economics and Law, University of Gothenburg, Sweden

Evaluating the Effect of School Closures on Societal Transmission of Covid-19: Evidence from Lower Secondary School Closures in Sweden*

Sebastian Shaqiri Johansson†

June 17, 2022

Abstract

In light of the Covid-19 pandemic, one of the most common interventions used by countries was to shut down schools in the hope of mitigating virus transmission. However, the effectiveness of school closures still lacks consensus in the literature due to the seemingly difficult task of estimating the causal effect of closing schools. This paper contributes to the literature by providing both empirical evidence of the effect of school closures and methodological guidance to isolate the effect of school closures. We exploit variations in school closures across municipalities by utilizing a novel dataset on Swedish lower secondary school closures to identify the causal effect using a synthetic control method. We find evidence that closing schools had a small and significant effect on the municipal transmission of Covid-19. In addition, we show that only holding all students at home simultaneously yielded mitigating effects on the transmission.

Keywords: Covid-19, Sweden, school closures, synthetic control method.

*Firstly, I want to thank my supervisor Mikael Lindahl for his time, encouragement and generally nice and giving discussions. Secondly, I would also like to thank Annika Lindskog and Elisabeth Olme, without them this paper would not have been possible.

†Department of Economics, School of Business, Economics, and Law

Contents

List of Figures

List of Tables

1 Introduction

In the wake of the SARS-CoV-2 (henceforth Covid-19) virus, governments worldwide have used various strategies to try and mitigate the virus transmission. In addition to societal lockdowns, school closures have been one of the most popular non-pharmaceutical interventions (NPIs), with more than 80 percent of the world's student population having to do their studies from home (Castex et al. 2020).

Although there are reasonably some cross-country differences in the policy implementation of school closures in how and to which extent schools have been closed, the purpose of school closures has been to reduce societal transmission or mitigate contagion clusters in schools to reduce the virus transmission. Nevertheless, school closures have broad impacts, not only affecting the virus transmission but also affecting the children in ways such as worsened learning outcomes (Andrew et al. 2020, Engzell et al. 2021), physical and mental health effects (Takaku & Yokoyama 2021, Yamamura & Tsustsui 2021), mothers labor force participation (Collins et al. 2021) and increased discrepancies in social and economic inequality due to seemingly disproportionate effects on less affluent families (Andrew et al. 2020, Engzell et al. 2021, Hanushek & Woessmann 2020).

Therefore, it is essential to examine if the expected positive effects of school closures outweigh the potential adverse effects documented in the literature. This paper tries to contribute to the insights on this topic by empirically assessing the effect of school closures on the Covid-19 transmission in Sweden. In contrast to most other countries, Sweden did not implement universal school closures for lower secondary schools. Instead, the decision to close schools was decided on a municipal level — or in some cases, school level when there are a large number of schools in the municipality — which has induced a variation across municipalities in both the intensity[1] of school closures and in absolute municipal school closures. In this paper, we collect data on lower secondary school (grade 7-9) closures for 40 municipalities in Sweden with the purpose of exploiting the variation in school closures across municipalities to compare the number of Covid-19 cases after school closures.

By utilizing a synthetic control method (SCM) approach, we are able to construct a synthetic control group from the municipalities that did not close their schools and thereby get around the problem of municipality pre-treatment heterogeneity in covariates when comparing outcomes. Using SCM, this paper aims to measure the causal effect of school closures on the number of Covid-19 cases by comparing the outcome between the treated municipalities and the

[1] With intensity of school closures, we refer to the number of grades that are closed simultaneously in a municipality. It is not necessarily the case that schools in the same municipality have had the same intensity of school closures.

synthetic control group.

Controlling for pre-trend covariates, we are able to construct a synthetic control group for the treated unit that resembles the average treatment unit before treatment in a way that makes it possible to evaluate the treatment effect of school closures and simultaneously deal with the endogeneity of school closures. We show that school closures among lower secondary schools in Sweden had a small negative effect on municipal transmission rates across municipalities that closed their schools. In line with Vlachos et al. (2021), our results suggest that closing schools had a mitigating effect on transmission; however, the minuscule effect shown by both this paper and Vlachos et al. (2021) suggest that school closures had no greater effect on transmission rates on a societal and school level. In addition, we also show that variations in the intensity of school closures directly affect the mitigating effects induced by the closures of schools, and only keeping all lower secondary school students at home simultaneously, on a municipal level, provided enough contagion management to reduce the societal transmission rate.

The rest of this paper is structured as follows. Section 2 provides a literature review and illustrates potential gaps in the literature that this paper aims to contribute to; section 3 gives a theoretical insight into why school closures would decrease the number of Covid-19 cases using a modified SIER model; section 4 gives an extensive explanation of the dataset and provides a comprehensive discussion of the empirical approach and how to reach identification using a synthetic control approach; section 5 presents the result of this paper along with robustness and heterogeneity analysis; lastly, section 6 gives a discussion on the results, limitations, and a conclusion of the paper.

2 Literature Review

This paper's primary concern is measuring the effect of school closures on the transmission of Covid-19 in society. In light of the empirical evidence suggesting adverse effects on students having to do their studies from home, it becomes increasingly important to evaluate the actual effect of school closures on societal transmission in order to disentangle the true cost of school closures. For the case of Sweden, Vlachos et al. (2021) exploit the fact that in the initial response to Covid-19, the Swedish government only required upper secondary schools to close while all lower secondary schools still had on-site learning. By comparing the outcome differences in the infection rate between children's parents, teachers, and teachers' partners among lower and upper secondary schools, Vlachos et al. (2021) finds that lower secondary teachers, their partners, and the children's parents were significantly more infected by Covid-19 than their upper secondary school counterparts. Although there is a statistically significant difference, the difference is relatively small, which according to Vlachos et al. (2021) is an indication that there only was a minimal effect of school closures on societal transmission.

This paper builds on what is initially found by Vlachos et al. (2021). First, as lower secondary schools were allowed to close, one could expect that the effect on societal transmission would be more significant. Second, we capture the actual cases of Covid-19 instead of using the parent's infection rate to measure societal transmission. Moreover, the detailed nature of the data used in this paper allows for measuring the effect of closing schools more directly by isolating the effect of school closures and utilizing variations across municipalities.

Svaleryd & Vlachos (2022) provides an extensive summary of the current research insights on school closures and the transmission of Covid-19. It is clear from Svaleryd & Vlachos (2022) that the literature on this topic provides quite substantial heterogeneity in the estimated effect of school closures. For example, using cross-country data for 130 countries, Liu et al. (2021) finds that school closures are one of several NPIs that reduce societal transmission of Covid-19. Similarly, Askitas et al. (2021) finds that school closures were an effective measurement to mitigate transmission. In contrast, both Chernozhukov et al. (2021) and Courtemanche et al. (2020) find no impact on school closures on transmission rates.

The ambiguous effect of school closures may result from unsatisfactory identification methods for both simulation and empirical studies: simulation methods using the SIR framework[2] are mainly a tool for predicting future outcomes and provide little help in estimating the causal effect. In contrast, empirical methods provide a better framework for estimating the causal effect of school closures. Although the empirical work is interested in the effect, most studies devote little focus to causal inference often lacking either (1) inclusion of sufficient control variables to isolate the effect of school closures and (2) seldom acknowledges the problem of reverse causality; that is, schools tend to close in times of high transmission. Therefore, failing to account for these potential problems would likely induce bias in the estimates.

The most straightforward approach in dealing with the endogeneity of school closures would be to control for other NPIs; however, school closures are often accompanied by other NPIs, making it difficult to isolate the effect of school closures. A few studies try to circumvent these problems by exploiting variations within family infection rates among students, close family members, and teachers

[2]SIR models are epidemiological models that explain how viruses transmit in society. In the most simple SIR model, everyone is susceptible to the virus, infected, and eventually recovered. Additional model extensions provide more dynamics, such as vaccinations and school closures. See section 3 for a theoretical explanation of the SIER model.

(Vlachos et al. 2021, Bravata et al. 2021). Since students living in close proximity to each other face similar NPIs, variations in school closures could be used to identify the effect on transmission by comparing the differences in family, student, or teacher infection rates. However, within-household transmission is insufficient to capture the full societal impact of school closures.

In the empirical literature examining the societal effects Fukumoto et al. (2021) and Rauscher & Burns (2021) tries to account for the endogeneity problems by utilizing matching methods. By using municipal variation in school closures across Japan, Fukumoto et al. (2021) matches municipalities that closed their schools with ones that did not close their schools. They found no evidence that closing schools had any impact on societal transmission. Rauscher & Burns (2021) estimates the timing of school closures and how it affects the death rates in the US and finds that each day of delay in school closures leads to a 1.5-2.4 percent higher death rate.

Another way of dealing with the pronounced endogeneity used in the literature is by examining school reopenings rather than school closures. Opening schools is more likely to be exogenous since it often corresponds to a particular set of pre-defined dates rather than being an effect of high transmission rates. Again, school openings show fairly large heterogeneity in the estimated effect. Using a synthetic control method approach, Reinbold (2021) studies the after-summer school reopenings in Illinois and shows that counties with high distance teaching or hybrid teaching experienced significantly fewer new cases than counties that had a majority of schools on on-site learning. However, both Goldhaber et al. (2020) and Harris et al. (2021) find no effect of school reopenings on the societal transmission and number of hospitalizations.

As can be seen from the equivocal results from the empirical literature, the effect of school closures and school openings on societal transmission is seemingly still uncertain. As noted above, the results seem to be highly dependent on the setting and how well the endogeneity of school closures, confounding, and collinearity with other NPIs are handled. In addition, the usefulness of school closures depends on the societal transmission as a whole, and the empirical evidence suggests that reopening schools at low or modest levels of societal transmission do not significantly affect any further spread of Covid-19. However, as the level of societal transmission likely affects both precautionary measures and endogenous behavioral responses, this conclusion should be regarded with caution. Furthermore, the usefulness of school closures also depends on the strategy of other NPIs; if, for example, students substitute their in-school interactions with out-of-school interactions, the impact of closing schools will crucially depend on the mitigating restrictions placed by other NPIs (Svaleryd & Vlachos 2022). Thus, there still exists a gap in the literature to provide both methodological and empirical evidence on this topic. This paper aims to shed some light on both these issues by providing a generalized synthetic control methodology to empirically assess the effect of school closures in a relatively unique setting, as Sweden's strategy to mitigate the spread of Covid-19 through school closures differs from most other countries.

3 Theoretical Effects of School Closures

To explain why school closures, at least theoretically, are a viable option to mitigate virus transmissions, we develop an extension of the standard SEIR (susceptible, exposed, infected, and recovered)[3] model by allowing for school closures as a measurement of social distancing and vaccinations.

Suppose that there are N agents that constitute the whole population. The population can then be decomposed into five non-overlapping groups that correspond to a different state of Covid-19. Firstly, let $S(t)$ denote the part of the population that is susceptible to the virus and let $E(t)$ be the share of the population that is exposed, that is, infected by the virus but does not yet transmit the virus. Individuals infected by the virus can spread the virus and are denoted by $I(t)$. The share of the population recovered from the virus is denoted by $R(t)$. Lastly, the vaccinated share of the population is denoted $V(t)$. Figure 1 shows the expected dynamics of Covid-19 in society. In the beginning,

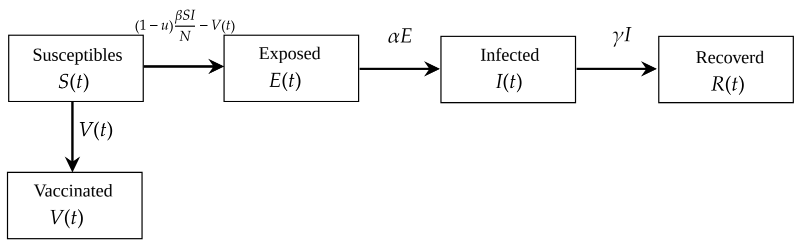

Figure 1. Dynamics of the SEIR Model with Social Distancing and Vaccinations

all agents in the economy are susceptible to the virus, then agents either get exposed and thereafter infected and eventually recover from the virus, or the agent becomes vaccinated and cannot become infected. Thus, in this simple model, the vaccination plays a huge role in limiting the rate a which agents become exposed to the virus. Moreover, when all agents in the economy are vaccinated, the model predicts no more spread of the virus, explicitly assuming that vaccinated agents cannot procure the virus and therefore not spread the virus.

For simplicity, we assume that recovered agents cannot be infected again and

[3]See for instance Collins et al. (2021) and Mwalili et al. (2020)

that all infected cases predicted by the model are the actual number of cases; there are no unreported symptomatic individuals. Demographic processes in the model are also neglected, and the death rate of those infected by Covid-19 is assumed to be negligible.

Then, the dynamics of the model can be described as a system of ordinary differential equations (ODEs)

$$\frac{ds}{dt} = -(1-u)\beta si - v'(t) \tag{3.1}$$
$$\frac{de}{dt} = (1-u)\beta si - \alpha e \tag{3.2}$$
$$\frac{di}{dt} = \alpha e - \gamma i \tag{3.3}$$
$$\frac{dr}{dt} = \gamma i \tag{3.4}$$
$$\frac{dv}{dt} = v'(t) \tag{3.5}$$

where each lower case letter corresponds to the fraction of the population in each state i/N for each $i = S, E, I, R, V$. The terms β, α and γ are country-specific[4] parameters, and $0 \leq u \leq 1$ is the intensity of social distancing measurements. If society is completely closed, $u = 1$ and no transmission will occur since there will be no contact between susceptible agents.

Since school closures are a type of social distancing interventions, closing schools will imply that $u > 0$, meaning that the number of susceptible agents in the economy decreases. For an illustration of this, see Figure 9 in the appendix for simulation results of equations (3.1)-(3.5).

4 Data and Methodology

4.1 Data

This paper utilizes a novel dataset on school closures among lower secondary schools in Sweden during the Covid-19 pandemic. We have collected data on 40 out of Sweden's 290 municipalities. The data includes information during which periods each school in a specific municipality had distance teaching.[5]

To measure the intensity of school closures, we create a variable that takes on values between 0 and 3, where a value of zero implies that all classes in that particular school have on-site lectures, whereas a value between one and three indicates the number of lower secondary grades at home; if one, then one grade is

[4]See appendix for explanation of parameters

[5]In Swedish, distance teaching is defined as either "distansundervisning" or "fjärrundervisning."

at home; if two, then two grades are at home; and if three, all three grades are at home. The school-level data is aggregated to the municipality level by taking the average intensity for each week per municipality. The treatment variable used in this paper is derived from the intensity variable. If the school intensity is zero, the treatment variable will be zero. In the baseline model, treatment occurs if the intensity is equal to three; that is, when the average weekly municipal intensity is equal to three, the treatment variable will be one. By deriving treatment from the intensity variable, we are able to vary the intensity measurement to investigate the sensitivity of the estimates derived in the baseline model. In addition, this specification of treatment is also required since some school data does not provide exact measurements of which grades had on-site learning, only that a particular share of the grades were home or in school.

The period studied in this paper ranges from week 45 of 2020 until week 24, 2021. Using the selected period, it is possible to capture initial school closures for a few schools in 2020[6] and all school closures up until the summer holiday 2021. In the sample, we identify two distinct periods where school closures were particularly prevalent: the second week of January 2020, which is after the winter holiday; the week after the "sportlov," which is a one week holiday that occurs in late February or beginning of March depending on the municipality. In the week after the Easter holiday in 2021, there was also a slight increase in school closures. Figure 2 shows the mean weekly evolution of school closure intensity and the mean number of new Covid-19 cases under the same period. During the winter holiday, the societal transmission was high, so schools close and then gradually opened up until the week after the sportlov when the number of cases increased again. Likely, the increase in cases was due to the increased mobility and in-country traveling following the sportlov. Moreover, the trends in Figure 2 suggest that the effect of school closures — given that they have an effect — comes with a lag; therefore, closing schools is hypothesized not to have an immediate effect on virus transmission; instead, the effects are seen weeks later.

4.1.1 Covariates and Raw Relationships

To isolate the variation of school closures, we include 12 different covariates. The criteria for including covariates is that they both could affect the number of Covid-19 cases and the likelihood of closing schools. Firstly, the weekly number of Covid-19 cases before treatment is a strong predictor of the number of weekly cases the week after. Therefore, we include a lagged variable for the number of Covid-19 cases one week before the municipality close. We also

[6] In our sample no schools closed earlier than week 45 2020

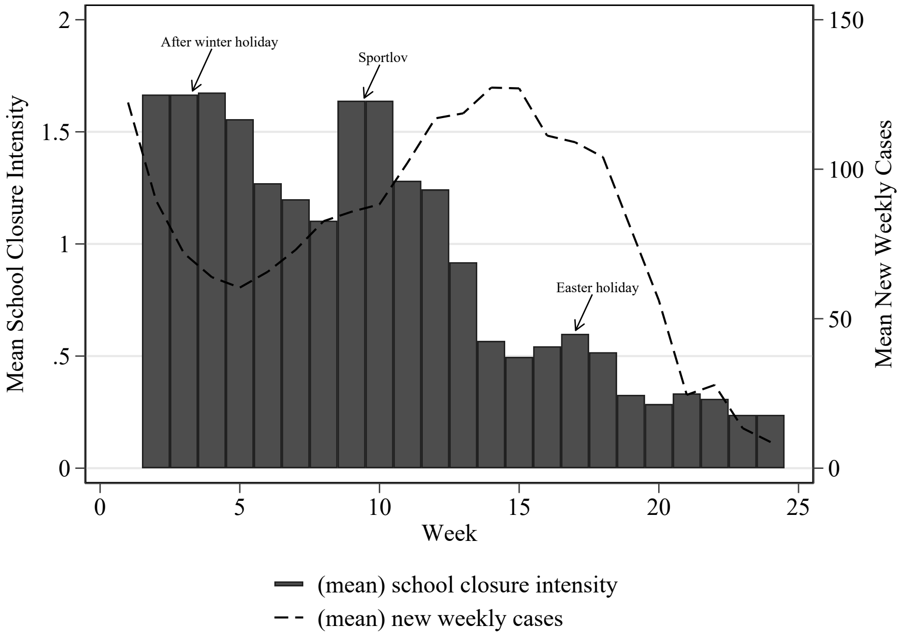

Figure 2. Weekly average school closure intensity and number of new Covid-19 cases

include regional dummies to control for unobserved characteristics across regions. Although regions and municipalities most often followed state recommendations, by including region dummies, it is possible to control for any unobserved heterogeneity that differed between regions, including other non-pharmaceutical interventions. Moreover, we include a range of predictors: demographic variables (the population size, population density, and mean age to control for spread around schools); geographic variables (municipality size and average net flow of commuters across municipalities); labor and educational variables (employment rate and student density); income variable (disposable household income); and lastly vaccination variables indicating the share of individuals with at least one and two doses of the Covid-19 vaccine.[7]

Table 1 shows descriptive statistics of each covariate for the population and the treated and untreated samples. The treated sample, on average, is significantly smaller in size compared to the population and the untreated sample. However, the treated municipalities tend to have a higher population density than the population and untreated ones. Municipalities with higher intensity of school closures were, on average, larger in population size but lower in area. Accordingly, the student density for the treated municipalities is also higher than in the population and the untreated sample. In terms of the other

[7]For an extensive motivation of covariates see Appendix.

covariates, the treated and untreated samples seem to be relatively similar to the population.

Table 1. Descriptive Statistics

	Population			Treated Sample			Untreated Sample		
	Obs	Mean	Std. dev.	Obs	Mean	Std. dev.	Obs	Mean	Std. dev.
Weekly new cases	10,560	89	121	224	86	103	1,078	87	120
Population size	10,560	95276	214027	224	112871	144692	1,078	107141	137571
Mean age	10,560	43	3	224	42	3	1,078	43	3
Size (km^2)	10,560	1297	2341	224	568	452	1,078	910	935
Population density	10,560	439	1160	224	645	961	1,078	555	919
Income (100tsek)	10,560	219	24	224	213	15	1,078	216	17
Employment rate (change this)	10,560	0.008	0.0004	224	0.008	0.0006	1,078	0.008	0.0006
Commuters (in)	10,560	22672	65005	224	21931	31698	1,078	19823	30112
Commuters (out)	10,560	13117	28380	224	12638	14906	1,078	11558	14226
Student density	10,560	14	38	224	20	29	1,078	17	28
Share of, at least. 1 dose of vaccine	10,560	0.014	0.016	224	0.015	0.016	1,078	0.019	0.018
Share of, at least, 2 doses of vaccine	10,560	0.014	0.016	224	0.009	0.011	1,078	0.015	0.019

Note: The Population includes all 290 Swedish municipalities. The treated sample is defined as the municipalities in which the intensity of school closures is equal to 3; the untreated sample is the rest. Table shows weekly averages of each variable.

The outcome variable of interest is new weekly cases for each municipality during the chosen sample period. The collected data is from the Public Health Agency of Sweden (Folkhälsomyndigheten). Figure 3 shows the raw relationship between the mean number of weekly Covid-19 cases and the mean weekly intensity of school closures. Each point in the figure represents the mean number of Covid-19 cases against the weekly municipal average intensity of school closures. We also draw the linear and quadratic regression lines on the underlying data. Firstly, we observe a positive linear relationship between the weekly Covid-19 cases and school closures. The quadratic fit has an inverse U-shape suggesting that the intensity of school closures does not mitigate the number of Covid-19 cases until an intensity of one. However, running these regressions, both the quadratic and linear fit are highly insignificant with $p = 0.9161$ and $p = 0.8658$, respectively.[8]

4.2 Identification

To identify the causal effect of school closures on the number of Covid-19 cases, we employ a synthetic control method (SCM) to estimate the treatment effect of school closures. To combat potential heterogeneity problems by comparing treatment effects across municipalities, the SCM approach allows for the construction

[8]Running a region and time fixed-effect model using the covariates gives a positive and significant ($p = 0.025$) estimate on the treatment dummy, suggesting that municipalities close their schools in times of high transmission.

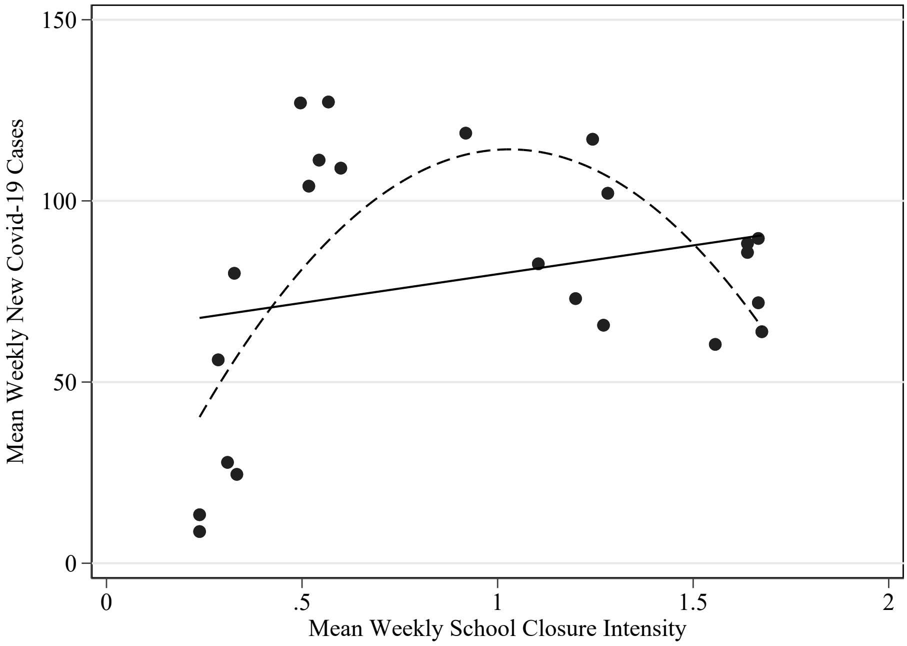

Figure 3. Raw relationship between weekly Covid-19 cases and school closures

of a weighted average of the untreated units that closely match the treated units in covariates.

The first part of this section will define some notation following the traditional SCM framework introduced by Abadie & Gardeazabal (2003) and Abadie et al. (2010). Secondly, we will expand the traditional framework following Abadie (2021) by allowing for more than one treated unit. Doing this renders the possibility to turn the aggregate municipality-specific effects into an average treatment effect across all treated municipalities.

Suppose that there are $t = 1, ..., T$ time periods, where $T_0 \in t$ denotes the number of periods before school closures with $1 \leq T_0 < T$. Furthermore, suppose that we have data on $j = 1, ..., J + 1$ municipalities where, without any loss of generality, we make the imposing restriction that municipality $j = 1$ is the only treated unit, that is, the only municipality that closed their schools, and the other $j = 2, ..., J$ are untreated municipalities referred to as the "donor pool."

Following, Abadie et al. (2010), let $\mathbf{Y}^I_{jt}$ be a $(T \times J + 1)$ outcome matrix describing the number of new weekly Covid-19 cases for municipality j at time t. Decomposing the outcome matrix $\mathbf{Y}^I_{jt}$ into treated and untreated municipality outcomes, the treated outcome becomes a $(T \times 1)$ vector denoted $\mathbf{Y}^I_{1t}$ and the untreated outcomes become a $(T \times j)$ matrix denoted $\mathbf{Y}^N_{jt}$. By considering school closures as an exogenous variation, outcomes before treatment should be unaffected for both treated and untreated units; that is, for time periods $t \in \{1, ..., T_0\}$ and for all municipalities $j \in \{1, ..., J+1\}$ it should be true that $\mathbf{Y}_{1t}^{I} = \mathbf{Y}_{jt}^{N}$.[9] This assumption follows immediately in a theoretical setting where $\mathbf{Y}_{jt}^{N}$ denotes the unobserved counterfactual outcome for the treated units.

For time periods $t > T_0$, we are interested in the treatment effect between the treated municipality and the untreated ones. Hence, for any time $t = T_0 + 1, T_0 + 2, ..., T$, the treatment effect of school closures is determined by

$$\tau_{1t} = Y_{1t}^{I} - Y_{jt}^{N}. \tag{4.1}$$

Rewriting equation (4.1) as $Y_{1t}^{I} = Y_{jt}^{N} + \tau_{it}$ and defining an indicator variable D_{jt} that takes the value 1 if municipality j is exposed to the treatment and 0 otherwise, the observed number of Covid-19 cases for any municipality j at time t can be written as

$$Y_{jt} = Y_{jt}^{N} + \tau_{jt} D_{jt}. \tag{4.2}$$

Since only municipality $j = 1$ is exposed to treatment after period T_0, the indicator variable is

$$D_{jt} = \begin{cases} 1 \text{ if } j = 1 \text{ and } t > T_0 \\ 0 \quad \text{otherwise.} \end{cases} \tag{4.3}$$

Since we are interested in the treatment effect of school closures, this implies that we want to estimate the parameters

$$\boldsymbol{\tau}_{1t} = \{\tau_{1,T_0+1}, ..., \tau_{1,T}\}.$$

The counterfactual outcome Y_{jt}^{N} is unobserved for each treated unit j and has to be estimated. Equation (4.1) therefore implies that the estimated treatment effect is a function of the estimated counterfactual outcome of Y_{jt}^{N}:

$$\widehat{\tau}_{1t} = Y_{1t} - \widehat{Y}_{jt}^{N}. \tag{4.4}$$

That is, the estimated treatment effect of school closures is just the difference between the observed outcome for the treated municipality and the estimated unobserved outcome.

To estimate the unobserved outcome, Abadie (2021) suggest using a weighted average of the outcomes for the donor pool municipalities. That is, let $\mathbf{W} = (w_2, ..., w_{J+1})'$ be a $(j \times 1)$ vector of weights assigned to each untreated municipality,

[9] This assumption is essential for the determination of causality since it explicitly assumes that before treatment outcomes are (virtually) identical implying that potential dispersion after treatment is possibly due to treatment itself.

with restrictions $\sum_{j=2}^{J+1} w_j = 1$ and $w_j \geq 0 \ \forall j \in \{2, ..., J+1\}$. Then, the synthetic control estimator is given by

$$\widehat{Y}_{1t}^{N} = \sum_{j=2}^{J+1} w_j Y_{jt} = \mathbf{Y}_0 \mathbf{W}, \tag{4.5}$$

thus the treatment effect of school closures is just the difference between the observed number of Covid-19 cases for the treated municipalities and the weighted average outcome of Covid-19 cases generated by the synthetic control group[10]:

$$\widehat{\tau}_{1t} = Y_{it} - \sum_{j=2}^{J+1} w_j Y_{jt}. \tag{4.6}$$

4.3 The Average Treatment Effect

In the framework presented by Abadie et al. (2010), we only observed one treated unit ($j = 1$), which for the purposes of this paper is not enough. Since school closures were not concentrated to only one municipality, there are more than one treated unit. Following, Abadie (2021) the extension is fairly straightforward. This extension allows us to examine the treatment effect when there is more than one treated unit and allow the treatment to occur continuously during the examined period. If treated units are indexed with $g \in \{1, ..., G\}$ and J denotes the donor pool that does not undergo treatment, the treatment effect for a single municipality g is still given by $\widehat{\tau}_{gt}$. Taking the average across all observed municipalities treatment effects yields the average treatment effect in the sample:

$$\overline{\tau} = (\overline{\tau}_{T_0+1} + ... + \overline{\tau}_T) = \frac{1}{G} \sum_{g=1}^{G} (\widehat{\tau}_{g,T_0+1}, ... + \widehat{\tau}_{g,T}). \tag{4.7}$$

That is, the average treatment effect is just the average of all individual treatment effects for each treated municipality g. Notice here also that $t > T_0$ is allowed to vary for each g. This allows for continuous and heterogeneous treatment effects, in our case, weekly treatment effects that get averaged out when computing $\overline{\tau}$.

4.4 Dealing with the Endogeneity of School Closures

The novelty of the synthetic control method in our setting is that it automatically deals with the reverse causality problem that schools tend to close under high transmission. By constructing a synthetic control group that fulfills the assumption that the outcome between the average treatment unit and the

[10]See appendix for estimation of the donor pool weights.

control group are identical before treatment, the number of Covid-19 cases will be identical for both the treated and untreated municipalities. Therefore, we are able to isolate the variation of Covid-19 cases and let treatment, that is, school closures, be the only variation affecting the outcome.

5 Result

5.1 Average Treatment Effect

As discussed in section 4, the average treatment effect is calculated by computing the average treatment effect for each municipality that closed its schools. To decide which week should constitute treatment, the average weekly intensity is aggregated over each week, and the week that received the highest intensity was chosen as the treatment week. In line with what was shown in Figure 3, the highest intensity was present during weeks 1-3 and 6-10. To extend the pre-treatment period, week eight is chosen as the week where treatment occurs.

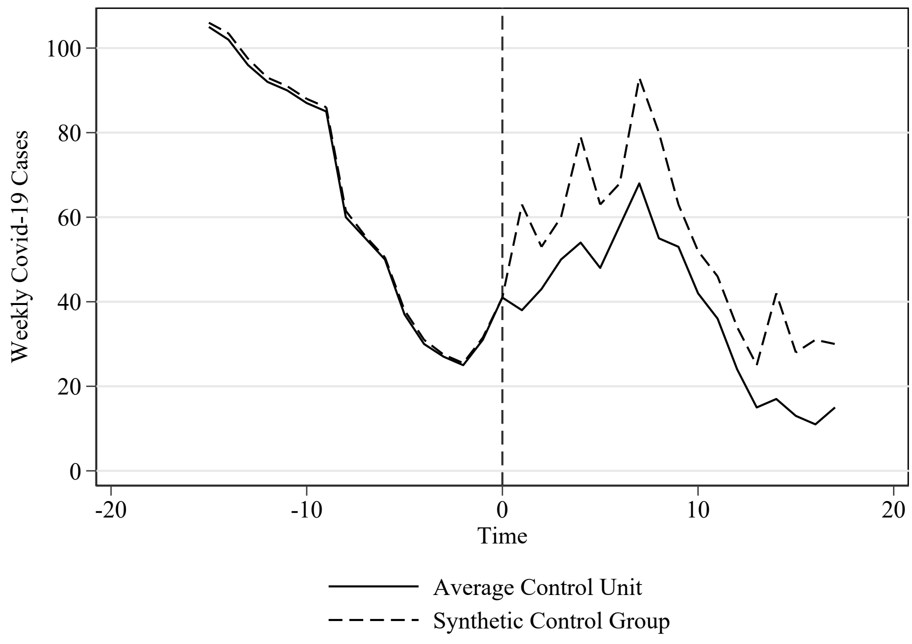

Figure 4. Average treatment effect of school closures

Figure 4 displays the new weekly Covid-19 cases from week 45, 2020, to week 24, 2021. Treatment here is indicated by the value zero. As we can see, the synthetic control group tracks the average treatment unit relatively well before treatment, which suggests that the synthetic control group is a good approximation of how the number of Covid-19 cases would evolve in the absence of school closures. Furthermore, the similarities in pre-treatment trends also suggest that the endogeneity problems with school closures are dealt with since the nature of SCM controls the data so that the number of Covid-19 cases was equal for the treated and untreated municipalities in the pre-treatment period.

To quantify the effect of school closures, we observe the difference between the average treatment unit and the synthetic control group after treatment. As we can see, Figure 4 suggests that the synthetic control group, that is, the municipalities that did not close their schools had more weekly new cases of Covid-19 after the other municipalities closed their schools. The average treatment unit follows the synthetic control group relatively close in the number of cases but remains lower during the examined period. The estimated magnitude indicates that closing schools seemingly affected the difference in the number of new weekly cases varying between 15 and 25 during the examined period.

5.2 Inference

The gap between the synthetic control group and the average treatment unit indicates that school closures may have affected the transmission of Covid-19. Although there is an observed difference, this observation is not sufficient to prove that the difference is statistically significant.

To test the statistical significance of the estimates, we construct multiple placebo tests where treatment is assigned to the untreated municipalities. By creating a permutation distribution, we are able to extract the statistical significance of the estimates obtained from the synthetic control (Abadie 2021). By artificially assigning treatment to untreated units, the variation between the true estimated effect and the synthetic controls should occur randomly. Therefore, if the estimated magnitude of the placebo units is similar to the estimated effect in the actual synthetic control, the estimated effect will be statistically insignificant.

Figure 6 shows the estimated gap between the actual outcome for each of the untreated units and the synthetic control group. The black line represents the difference between the average treatment unit and the synthetic average treatment unit, while the gray lines represent the difference between the true outcome and synthetic control outcome for the placebo units. First we observe that some of the untreated placebo units are not suitable matches to their synthetic controls due to the large dispersion in the pre-treatment period. Second, for the average treatment unit (black line), the deviation in the pre-treatment period is close to zero. After the treatment, the line diverges from most of the other lines suggesting that the average treated unit had fewer cases in relation to

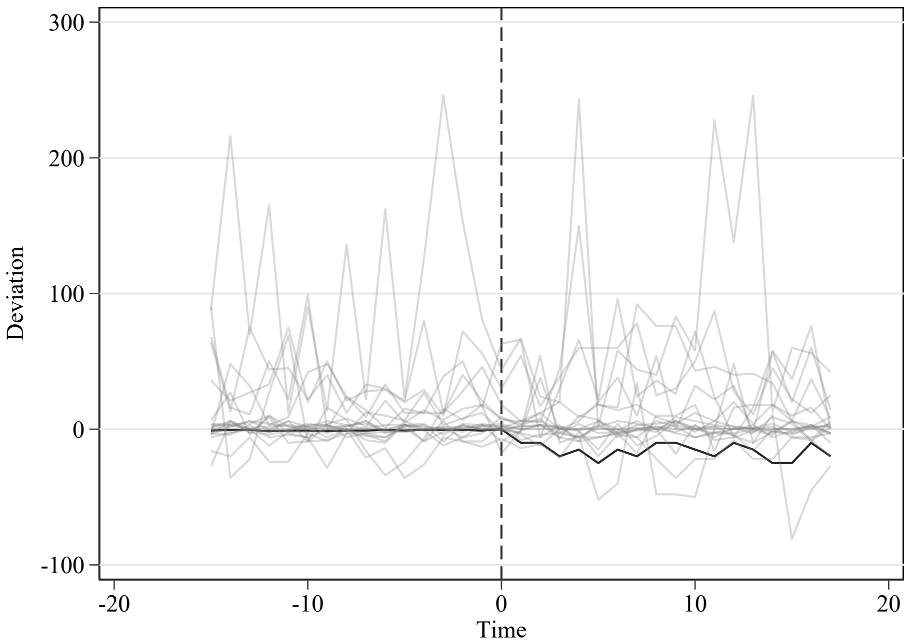

Figure 5. Placebo permutation distribution

its corresponding synthetic control group and other placebo units. However, there is no substantial difference between the placebo units and the average treatment unit, implying that the estimated difference displayed in Figure 4 is not statistically significant.

To check the sensitivity of the placebo distribution, we remove municipalities with a pre-treatment mean squared prediction error larger than two compared to the synthetic average control unit. As can be observed in 4, this includes the gray lines with substantial pre-treatment deviations from their synthetic control. Removing these from the placebo distribution makes it easier to disentangle the post-treatment effect of the average treatment unit since the potential large dispersion across other placebos after treatment may be a result of poor fit before treatment rather than a large effect after. Again, the difference is smaller than most other placebo units, and now the difference is large enough to be classified as statistical significant. That is, after removing municipalities with a high MSPE, we find evidence that school closures among lower secondary schools had a significant, albeit small, effect on the transmission of Covid-19 across Swedish municipalities.

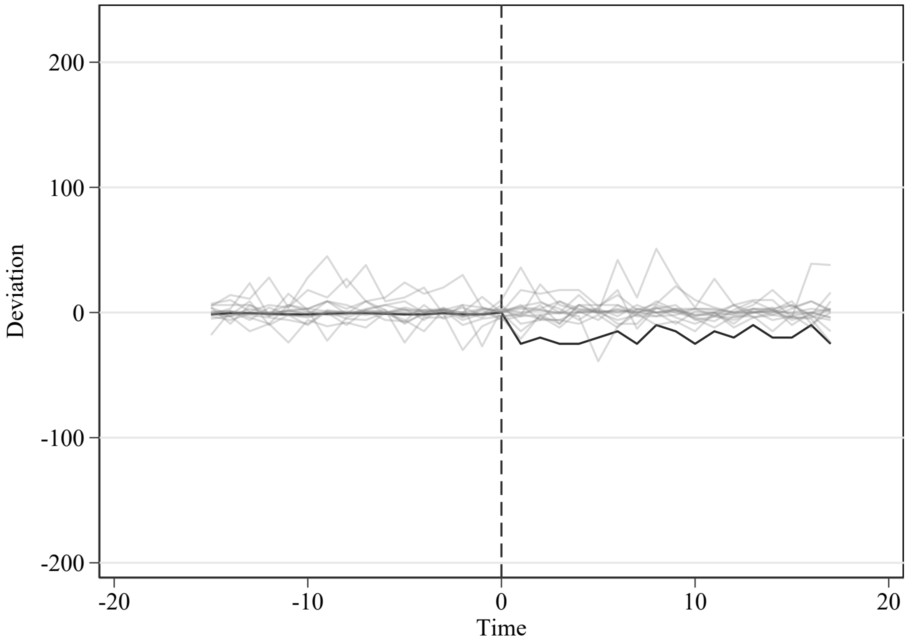

Figure 6. Placebo permutation distribution excluding municipalities with MSPE larger than two to their pre-treatment fit

5.3 Robustness Checks

5.3.1 Time Placebo

First, we run six time-placebo tests where we artificially shift the treatment date backward in time across the pre-treatment period prior to the selected treatment date. This yields, in effect, fictitious school closure treatments across all weeks at the beginning of 2021, excluding the two first where all schools were off due to the winter holiday. The construction of the synthetic controls follows immediately from an equivalent derivation to the one detailed in section 4. Since the pre-treatment periods differ in length for the six placebos, the synthetic control units will differ from the original; however, for the estimates to be robust, there should be no substantial divergence between the actual control and the placebos before treatment.

Figure 7 shows the time-placebo permutation distribution backdating the intervention to the weeks before treatment. The black line shows the baseline number of cases, and the gray lines represent the placebo controls. As can be seen, the baseline synthetic closely follows the placebos before and after treatment. This implies that there are no random variations before treatment that are directly affected by selecting a particular treatment date. This result provides credibility to the robustness of the estimated effect.

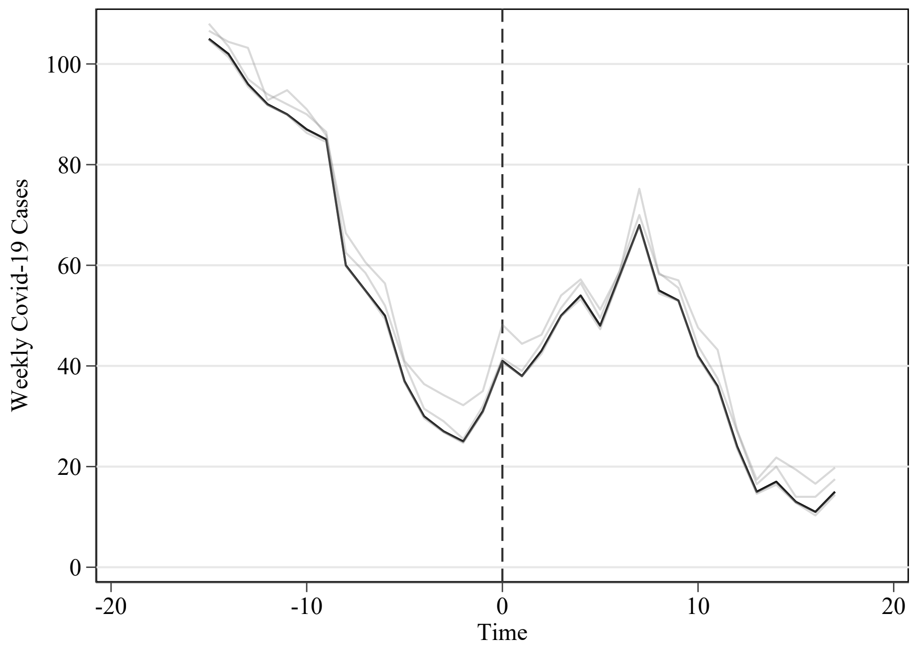

Figure 7. Time-placebo permutation distribution

5.3.2 Restricting the Donor Pool

Second, we also check the robustness of the SCM estimates by restricting the available control units when constructing the synthetic control group. To do this, we iteratively omit municipalities from the donor pool that contribute with more than one percent to the baseline model. Doing this tests the sensitivity of changes in the donor pool.

In Figure 8 the black line shows the baseline synthetic control compared to the subsequent treatment effects (gray lines) produced with the restricted donor pool. We observe some variation in the estimated effect; however, the restricted donor pools still center around the baseline model, suggesting that the baseline estimates are seemingly robust to changes in the donor pool.

5.4 School Closure Heterogeneity

By defining treatment from the intensity variable, we are able to vary the treatment intensity to investigate potential heterogeneous treatment effects across municipalities with all two and one grade(s) at home. To do this, we use the following difference-in-difference model

$$Y_{it} = \beta_0 + \beta_1 T_{it} + \beta_2 D_{it} + \beta_3 (T_{it} \times D_{it}) + \mathbf{X}'\boldsymbol{\gamma} + \epsilon_{it} \tag{5.1}$$

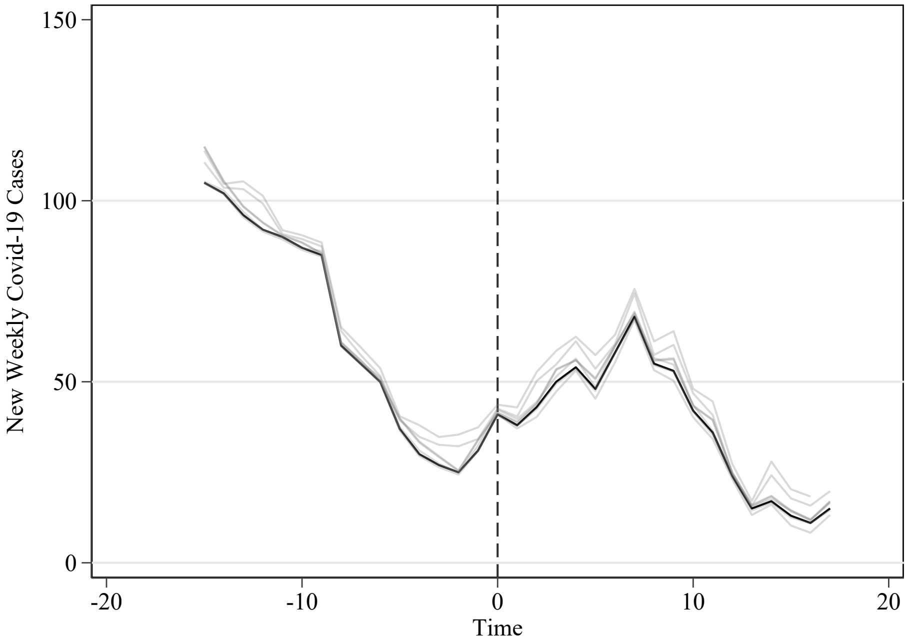

Figure 8. Placebo distribution with restricted donor pool

where T is a time indicator, D is the treatment dummy, and $\mathbf{X}$ is a vector of controls including the covariates presented in Table 1 for each time period t and municipality i. Here, the outcome variable, Y, is the number of weekly Covid-19 cases for the baseline, medium, and low-intensity scenarios. The outcome of the control group is derived from the synthetic control outcomes.

The estimates in Table 2 show that the number of Covid-19 cases in our sample is increasing for each model; the insignificant estimates of the treatment variable suggest that the difference in outcomes before treatment between the actual value and the synthetic control group for each model was small. The interaction estimate between the time indicator and treatment dummy gives us the treatment effect of school closures. We can observe that it is only in the baseline model where school closures significantly affected the number of Covid-19 cases. The negative value of the estimate implies that treated municipalities that enforced all lower secondary students to work from home experienced lower societal transmission. Furthermore, the estimated magnitude of the treatment effect is consistent with what was observed in Figure 4.

Table 2. Heterogeneous effects in intensity of school closures

	Dependent variable		
	Baseline	**Medium**	**Low**
Time	0.930^{***}	0.997^{***}	1.009^{***}
	(0.229)	(0.243)	(0.247)
Treatment	1.036	0.981	1.67
	(3.966)	(4.657)	(7.857)
Time × Treatment	-21.417^{***}	-13.332	-18.347
	(5.926)	(12.248)	(16.252)
Observations	1302	18022	18022

Note: Standard errors in parentheses.*, ** and *** denotes significance level at 10, 5 and 1 percent, respectively. The dependent variable is new weekly Covid-19 cases. The baseline model estimates the treatment effect when school intensity is equal to three. The medium model and low model estimates the treatment effect when school intensity is equal to two and one, respectively. The control variables are the seen in 1.

6 Concluding Discussion

This paper investigates the effect of lower secondary school closures in Sweden on the municipal transmission of Covid-19. Identification is reached by utilizing school closure variations across 40 different municipalities. By regarding school closures as an exogenous shock and using a generalized synthetic control method, we are able to estimate the treatment effect of school closures, bridging both empirical evidence on the effectiveness of school closures as a non-pharmaceutical intervention and methodological guidance in how to deal with endogeneity and confounding with other NPIs. In sum, we find that lower secondary school closures had a minuscule effect on the municipal transmission of Covid-19 in Sweden. We estimate that closing schools are associated with a reduction of approximately 15-25 cases on a municipal level. The results are robust to both time shifts and restrictions of the donor pool.

In addition to the baseline results, several insightful results are presented. This paper investigates the effectiveness of different compositions of on-site contra distance teaching. We find, interestingly, that only completely restricting student interactions by enforcing all lower secondary school students to have distance teaching simultaneously had a mitigating effect on municipal transmission. This is in line with what is expected from theory; more restrictions imply less transmission between individuals. On the contrary, our results also suggest that only having one or two classes at home did not have any mitigating effect on the societal transmission of Covid-19. We interpret these results to be in line with what is previously found by Vlachos et al. (2021) in that closing schools only had a minuscule effect on societal transmission and only hold under specific circumstances. Moreover, we find no significant effects of municipal heterogeneity when comparing the effects on school closures for large and small municipalities in terms of area, population, student density, and the number of Covid-19 cases.[11] Suggesting that the treatment effect for the municipalities that closed their schools was relatively homogeneous.

These results provide important policy implications for the evaluation of school closures, both in terms of information for the Covid-19 pandemic but also for future pandemics. The results in this paper indicate that closing schools may not have been as useful as initially thought, at least to the extent that it by itself would relieve societal transmission. This is especially true when considering the documented adverse effects of closing schools. In the case of Sweden, there is still a lack of available studies examining the effect on children, such as learning outcomes and mental health issues, as a result of school closures; thus, it is still too early to conclude that closing schools were worth it in Sweden, despite our results in this paper.

In terms of generalizability and external validity, our results should be regarded with care since Sweden adopted a relatively unique strategy to mitigate Covid-19. By taking a more lenient approach to imposing other NPIs, the effect of closing schools could very much vary quite substantially compared to countries that adopted a more restrictive policy on other NPIs. Likewise, we investigate the effect during the second wave of the pandemic, and it would be reasonable to believe that endogenous behavioral responses would have changed between the first, second, and third waves, which could affect the effectiveness of school closures.

Lastly, this paper also contributes to the literature by paying particular attention to causal inference. In terms of internal validity, our results are robust even in relatively small samples dealing well with both the endogeneity of school closures and identification of a true school closure effect independent of other NPIs.

[11] See appendix for result of municipal heterogeneity analysis.

References

Abadie, A. (2021), 'Using synthetic controls: Feasibility, data requirements, and methodological aspects', *Journal of Economic Literature* **59**(2), 391–425. 13, 14, 15, 17

Abadie, A., Diamond, A. & Hainmueller, J. (2010), 'Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program', *Journal of the American statistical Association* **105**(490), 493–505. 13, 15, 27

Abadie, A. & Gardeazabal, J. (2003), 'The economic costs of conflict: A case study of the basque country', *American economic review* **93**(1), 113–132. 13

Andrew, A., Cattan, S., Costa Dias, M., Farquharson, C., Kraftman, L., Krutikova, S., Phimister, A. & Sevilla, A. (2020), 'Inequalities in children's experiences of home learning during the covid-19 lockdown in england', *Fiscal Studies* **41**(3), 653–683. 4

Askitas, N., Tatsiramos, K. & Verheyden, B. (2021), 'Estimating worldwide effects of non-pharmaceutical interventions on covid-19 incidence and population mobility patterns using a multiple-event study', *Scientific reports* **11**(1), 1–13. 6

Bravata, D., Cantor, J. H., Sood, N. & Whaley, C. M. (2021), Back to school: The effect of school visits during covid-19 on covid-19 transmission, Technical report, National Bureau of Economic Research. 7

Castex, G., Dechter, E. & Lorca, M. (2020), 'Covid-19: cross-country heterogeneity in effectiveness of non-pharmaceutical interventions', *Covid Economics* **14**, 175–99. 4

Chernozhukov, V., Kasahara, H. & Schrimpf, P. (2021), 'Causal impact of masks, policies, behavior on early covid-19 pandemic in the us', *Journal of econometrics* **220**(1), 23–62. 6

Collins, C., Ruppanner, L., Christin Landivar, L. & Scarborough, W. J. (2021), 'The gendered consequences of a weak infrastructure of care: School reopening plans and parents' employment during the covid-19 pandemic', *Gender & Society* **35**(2), 180–193. 4, 8

Courtemanche, C., Garuccio, J., Le, A., Pinkston, J. & Yelowitz, A. (2020), 'Strong social distancing measures in the united states reduced the covid-19 growth rate: Study evaluates the impact of social distancing measures on the growth rate of confirmed covid-19 cases across the united states.', *Health affairs* **39**(7), 1237–1246. 6

Engzell, P., Frey, A. & Verhagen, M. D. (2021), 'Learning loss due to school closures during the covid-19 pandemic', *Proceedings of the National Academy of Sciences* **118**(17). 4

Fukumoto, K., McClean, C. T. & Nakagawa, K. (2021), 'No causal effect of school closures in japan on the spread of covid-19 in spring 2020', *Nature Medicine* pp. 1–9. 7

Goldhaber, D., Imberman, S. A., Strunk, K. O., Hopkins, B. G., Brown, N., Harbatkin, E. & Kilbride, T. (2020), 'To what extent does in-person schooling contribute to the spread of covid-19? evidence from michigan and washington', *Journal of Policy Analysis and Management* . 7

Hanushek, E. A. & Woessmann, L. (2020), 'The economic impacts of learning losses', *OECD Education Working Papers* (225).
URL: *https://www.oecd-ilibrary.org/content/paper/21908d74-en* 4

Harris, D. N., Ziedan, E. & Hassig, S. (2021), 'The effects of school reopenings on covid-19 hospitalizations', *National Center for Research on Education Access and Choice* . 7

Liu, Y., Morgenstern, C., Kelly, J., Lowe, R. & Jit, M. (2021), 'The impact of non-pharmaceutical interventions on sars-cov-2 transmission across 130 countries and territories', *BMC medicine* **19**(1), 1–12. 6

Mwalili, S., Kimathi, M., Ojiambo, V., Gathungu, D. & Mbogo, R. (2020), 'Seir model for covid-19 dynamics incorporating the environment and social distancing', *BMC Research Notes* **13**(1), 1–5. 8

Rauscher, E. & Burns, A. (2021), 'Unequal opportunity spreaders: Higher covid-19 deaths with later school closure in the united states', *Sociological Perspectives* **64**(5), 831–856. 7

Reinbold, G. W. (2021), 'Effect of fall 2020 k-12 instruction types on covid-19 cases, hospital admissions, and deaths in illinois counties', *American journal of infection control* **49**(9), 1146–1151. 7

Svaleryd, H. & Vlachos, J. (2022), 'Covid-19 and school closures'. 6, 7

Takaku, R. & Yokoyama, I. (2021), 'What the covid-19 school closure left in its wake: evidence from a regression discontinuity analysis in japan', *Journal of public economics* **195**, 104364. 4

Vlachos, J., Hertegård, E. & Svaleryd, H. B. (2021), 'The effects of school closures on sars-cov-2 among parents and teachers', *Proceedings of the National Academy of Sciences* **118**(9). 5, 6, 7, 22

Yamamura, E. & Tsustsui, Y. (2021), 'School closures and mental health during the covid-19 pandemic in japan', *Journal of Population Economics* pp. 1–38. 4

Appendix

Estimating the Synthetic Control Group

By partitioning the outcome vectors into two different vectors indicating pre- and post-treatment $\mathbf{Y}_j = (\mathbf{Y}^{pre}/\mathbf{Y}^{post})$, we can define a set of k pre-treatment predictors $\mathbf{X}$ that includes a set of observable covariates $\mathbf{Z}$ and a set of M linear combinations of pre-treatment outcomes $\mathbf{Y}^{pre}$. Likewise, $\mathbf{X}_0$ defines a $(k \times j)$ matrix of predictors for the donor pool.

To estimate the weights assigned to the untreated units, Abadie et al. (2010) suggest to assign the weights in such a way that minimizes the distance:

$$\mathbf{W}^* = \min_{\mathbf{W}} \|\mathbf{X}_1 - \mathbf{X}_0\mathbf{W}\| = \sqrt{(\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})'\mathbf{V}(\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})} \tag{6.1}$$

subject to the restrictions put on the weights (sum to one and greater or equal to zero), given a $(k \times k)$ symmetric positive and semidefinite matrix $\mathbf{V}$.

It is clear from equation (6.1) that the choice of $\mathbf{V}$ affects the minimized values of the weights $\mathbf{W}^*$. Although, the inferential strategy is unaffected by the choice of $\mathbf{V}$, the optimal $\mathbf{V}$ should be chosen in such a way that it assigns the optimal weights as a linear combination of the predictors and thus minimizing the mean squared error of the synthetic control group. In this paper, $\mathbf{V}$ will be chosen in such a way that the synthetic control matches the pre-treatment outcomes for all M linear combinations of the treated units outcome,[12]

$$\|\mathbf{Y}_1^{pre} - \mathbf{Y}_0^{pre}\mathbf{W}\| = 0 = \|\mathbf{X}_1 - \mathbf{X}_0\mathbf{W}\|. \tag{6.2}$$

Calibration of SEIR Simulation

The simulation of the ordinary differential equations in (3.1)-(3.5) is done in Python using the inbuilt ODE solver. For purposes of replication, let us consider the calibration of the model.

Let t_0 denote the start of the pandemic. Then, at time t_0 there are no exposed, infected nor vaccinated individuals. That is, $E(t_0) = 0$, $I(t_0) = 0$ and $V(t_0) = 0$. Naturally it also follows that $R(t_0) = 0$. Furthermore, let τ^* and $\tau^\star$ denote the incubation time and the number of days that an individual is infected,

[12] Alternatively one could define a training period $P \subseteq M$ that spans only half of the pre-treatment period, and then validate the fit of the model trough out-of-sample validation on the second half of the pre-treatment period. This method is also used in the paper to validate the model fit and to reduce the chance of overfitting.

respectively. The parameters α, γ and β in (3.1)-(3.5) are given by

$$\alpha = \frac{1}{\tau^*}, \tag{A.1}$$

$$\gamma = \frac{1}{\tau^\star}, \tag{A.2}$$

$$\beta = R_0\gamma, \tag{A.3}$$

where α is the rate of change of going from exposed to infected; γ is the rate of change of going from infected to recovered; and β is the rate of change of going from susceptible to exposed. Here R_0, in equation (A.3), tell us how many individuals one person transmits the virus to. Table 3 shows the parameter calibrations used in the model.

Table 3. SEIR parameter calibration

Parameter	Value	Description
τ^*	5	Incubation time
$\tau^\star$	7	Infection time
R_0	2.4	Number of people a given person transmit the virus to
N	10^6	Number of individuals in simulation
α	0.2	ROC of going from exposed to infected
γ	0.143	ROC of going from infected to recovered
β	0.343	ROC of going from susceptible to exposed

Note: ROC is short for rate-of-change.

Using these parameter estimates, we are able to simulate different scenarios in equation (3.1)-(3.5) by varying the intensity of school closures trough the parameter u. Figure 9 shows this simulation in a low, medium and high intensity scenario.

Municipal Heterogeneity

Table 4 show the estimated heterogeneous treatment effects across different covariate samples. Here, S denotes a dummy variable that takes the value one for a particular condition (as seen in the parenthesis in the table). If the condition is fulfilled, then $S = 1$ and otherwise zero. It is reasonable to assume

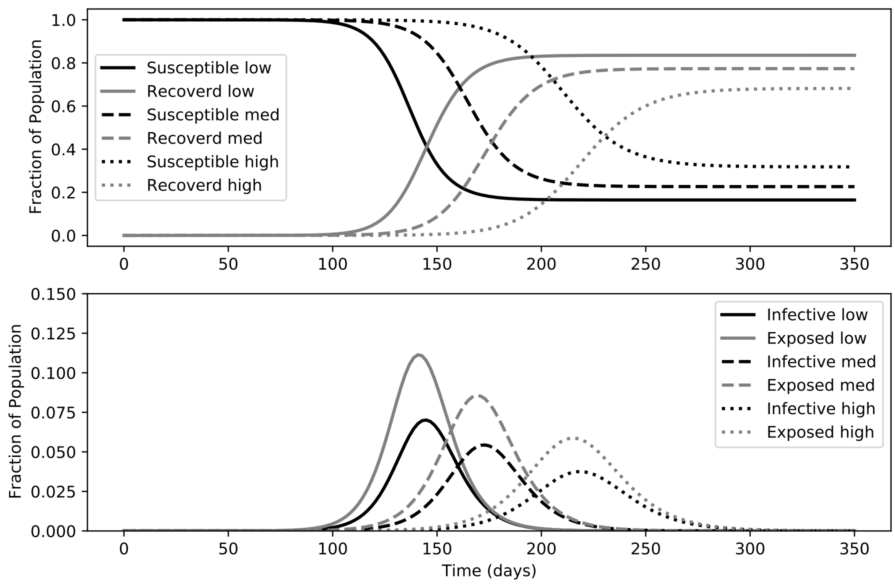

Figure 9. SEIR simulation

that the data generating process differs between the municipalities that fulfill the condition and the ones that do not; therefore, the estimated coefficients in Table 4 are the estimates obtained from the split sample difference-in-difference regression in (5.1) conditional on the population mean as a threshold for each condition.

As can be seen, we find no significant differences between small and large municipalities in the municipal area, population size, student density, and no differences between the municipalities that had high and low transmission. Partly, our results are driven by low efficiency due to the small sample size when splitting the data; indeed, for the split samples where we find significant effects, the sample size exceeds 800 observations.[13] Nevertheless, our results suggest that there are no heterogeneous treatment effects across municipalities in the variables that substantially differ between the treated and untreated samples.

[13]1085 and 868 observations, respectively.

Table 4. Heterogeneous municipal effects

	New weekly Covid-19 cases		
	$S = 1$	$S = 0$	p-value
Municipal area	-74.36	-16.952***	0.626
($S = 1$: Area> 1297)	(114.6)	(3.144)	
Population size	-25.05	-15.25	0.883
($S = 1$: Population size> $95,276$)	(37.18)	(29.3)	
Student density	-14.70***	-16.84	0.920
($S = 1$: Student density> 14)	(1.935)	(19.4)	
Covid-19 cases	-14.72	-24.1	0.826
($S = 1$: Weekly Covid-19 cases> 89)	(15,03)	(27.8)	

Note: Standard errors in parentheses.*, ** and *** denotes significance level at 10, 5 and 1 percent, respectively. The dependent variable is new weekly Covid-19 cases. p-value refers to statistical significance between the difference of coefficients. The estimates shows the heterogeneous treatment effects of municipal size, population size, student density and Covid-19 cases

Motivation of Covariates

- **Previous outcomes:** The weekly number of Covid-19 cases before treatment increases the likelihood of closing schools if the spread is high. This data is collected from Folkhälsomyndigheten
- **Region dummies:** Although regions and municipalities most often followed state recommendations, including region dummies, we can account for any characteristics that differed between regions, including other non-pharmaceutical interventions.
- **Demographic variables:** Population size, population density in a municipality increase the potential number of Covid-19 cases, which increases the likelihood of closing schools. This is be collected from Statistics Sweden (Statistiska Centralbyrån)
- **Geographic variables:** The larger the inhabitable area and number of bordering municipalities, the higher the number of Covid-19 cases is, which will affect the likelihood of closing schools. This is be collected from Statistics Sweden
- **Municipality income and financial variables:** Wealthier municipalities tend to be more urban, which reduces the likelihood of "natural" social distancing, which increases the number of Covid-19 cases, and thus the likelihood of closing schools. This is be collected from Statistics Sweden
- **Educational Variables:** More students in schools will likely increase the number of Covid-19 cases, especially in schools, which increases the

likelihood of closing schools. This is be collected from the Swedish National Agency for Education (Skolverket)

- **Labor variables:** The higher the employment rate, the more interaction between workers, which will increase the number of Covid-19 cases and thus the likelihood of closing schools. This is be collected from Statistics Sweden
- **Medical Variables:** Less vaccinated inhabitants will increase the outcome. Expecting that, municipalities will close schools. This is be collected from Folkhälsomyndigheten
- **Movement across municipalities:** The more people commute across municipalities, the more Covid-19 cases there will be. Expecting that, municipalities will close their schools. This is collected from Statistics Sweden