Are Schools Protective of Health, Even During a Pandemic? The Effects of School Openings During the COVID-19 Pandemic on Children's Health Care Usage

Maria D. Fitzpatrick

Cornell University and NBER

Katharine Sadowski

Cornell University

October 15, 2021

Abstract: Most of the research on the effects of school environments on children, particularly during the global COVID-19 pandemic, has focused on children's academic outcomes. However, schools provide a range of social and health supports for children and families. In the proposed research, we examine how school closures and openings affected health outcomes of children, with a particular focus on mental and developmental health outcomes (e.g., depression and ADHD-related care). To determine the effects of school environments during the pandemic, we utilize a national sample of electronic health records across the country. We pair this with newly collected district-level data on the closure, opening, and operational modes (in-person or online) of school districts across the country. We examine disparities in effects across children from different socioeconomic background and across communities with different geography, density, and economic conditions. Preliminary results suggest that, although the pandemic itself led to increases in depression-related health care visits, access to in-person school environments decreases the use of services for depression. The results of this research will be useful for understanding the myriad of ways that school closure policies affect child wellbeing during this pandemic and better prepare responses to future public health emergencies.

Corresponding author: Maria Fitzpatrick, Cornell University, maria.d.fitzpatrick@cornell.edu

I. Introduction

The global COVID-19 pandemic significantly altered many dimensions of children's lives. Increased unemployment rates and financial insecurity likely increased parental and child stress. At the same time, school closures and stay-at-home orders significantly decreased children's access to schools and other environments that provide them with social, cognitive, and physical health-supporting opportunities. Although researchers, public health experts, and the policy community have recognized the potential for the pandemic to have impacted children's health and development, much of the existing work to date has focused on academic outcomes (e.g., Kogan and Lavertu 2020, Maldondo and De Witte 2020, Kuhfeld, Soland, and Tarasawa 2020), or has involved survey research with a limited sample size (e.g., Hawrilenko et al. 2020, Weaver et al. 2021) or a specific population (e.g., Gassman-Pines, et al. 2020). Very few have used methods able to distinguish between the role of schools and the effects of the pandemic more broadly (Baron et al. 2020 is an exception).

We provide an important contribution to this fast-growing literature by examining child health during the pandemic using an innovative database of electronic health records for a large portion of the U.S. population. These data cover the period from January 2020 until the present and provide a rich set of information about the exact timing and nature of visits to health care providers. These features allow us to measure changes in children's health care usage at a daily level for a range of diagnoses and treatments not possible with previously utilized sources. We provide evidence about the relationship between children's health care usage the onset of the pandemic, COVID-19 prevalence in an area, local economic conditions, and, most centrally, local access to public school environments. We pay particular attention to the effects of the

pandemic on children's mental health, including depression, anxiety, and ADHD, as well as to indicators of its effects on child maltreatment.

To determine the effects of exposure to school environments on child health during the COVID-10 pandemic, we make use of the staggered timing of school openings in the fall of 2020 coupled with the variation in modality of instruction. We use a differences-in-differences framework to compare child health care usage before and after schools opened with varying degrees of in-person instruction. We also control for underlying prevalence of COVID-19 infection in an area to separately identify the effects of school exposure from the broader effects of the pandemic.

[EXPLAIN RESULTS]

II. Data

In this study, we use school openings and modality information compiled from multiple sources and electronic health records from HealthJump. In this section, we describe each of these data sources in turn.

School Openings and Modality Data

In the Fall of 2020, decisions about when to start the school year and what mode of instruction to use were mostly left to local districts. These decisions were often based on information about the local spread of COVID-19, the physical and staffing capacity for social distancing and other health and safety requirements, and the risk preferences of staff and

families. For the purposes of this study, these decisions can largely be characterized as having two dimensions: the opening of the school year and the modality of instruction.[1]

Our analysis hinges on having accurate school opening data and learning modalities. Our final dataset uses a limited set of school opening dates collected by MCH with the Covid19 School Data Hub's state reports on learning modality broken down by grade level. Collecting accurate school opening data has proven difficult and both datasets comes with their limitations. While the MCH data is the most comprehensive publicly available survey data representing the 2020-21 school opening dates and learning modalities for all 50 states and DC, the opening dates range is unrealistic extending from June 8, 2020 to December 18, 2020. This leads to concerns that school districts might have misinterpreted the survey questions and were reporting the district modality and date as of the day they took the survey not the first day of school. Additionally, we know districts often implemented different modality plans for different grade levels; however, MCH only reports one overall learning modality by district.

The Covid19 School Data Hub allows us to improve on a few of these concerns, while presenting drawbacks of its own. The Covid19 School Data Hub collects data on learning modality based on a states' desire to report at the school or school district level with records given with varying frequency from daily to annually. This data is more nuanced providing a breakdown of modality for some of the districts and schools in the dataset. However, this database only covers 29 states and doesn't include an indicator for the school opening date.

To exploit the best features of each dataset and limit the possibility of error, we use the following protocol. First, we limit the MCH data to the school openings between August 1 and

[1] To more precisely characterize children's exposure to school environments, it also would have been useful to understand the frequency with which schools 'shut down' and reopened during the fall of 2020. Unfortunately, data on closings and reopenings is even harder to collect than data on the initial openings and modalities upon opening. In future work, we will explore other ways of characterizing shifts in exposure to schools during the fall of 2020.

September 31. Second, we rely on the learning modality reported by the Covid19 School Data Hub for grades K-5, 6-8, and 9-12. If modality by grade is missing, we assign the overall learning modality to each grade level. If the overall learning modality is missing, then we rely on the number of students enrolled by modality to determine the learning modality by grade. Enrollment by modality is not broken down by grade level, so we use the following rules. If enrollment is exclusive to one modality, then we assign that modality to all grade levels. If, however there are any students enrolled in any combination of in-person, virtual, and hybrid, we assign all grade levels as hybrid. Finally, because the Covid19 School Data Hub does not contain school openings, we connect the two datasets using the NCES ID and then keep the record with the closest matching period start date to the MCH opening date.

We weight the exposure to school as follows. Full in-person schools receive a weight of 1. Hybrid schools receive a weight of 0.8 for elementary and 0.4 for middle and high schools. Fully virtual receive a weight of 0. Our choice of weights for the hybrid options are based on the notion that while schools defined their hybrid models differently due to practical constraints, overall reports suggest school districts prioritized moving elementary schools back to normal before middle and high schools.

Electronic Health Records

We also use data from Healthjump, an EHR system which has partnered with the COVID-19 Research Database to make data available to researchers interested in understanding more about COVID-19 and the pandemic. Data come from participating providers across the country that use the EHR vendors that in turn participate in the HealthJump platform for homogenizing and sharing records across vendors. Though the data cover more than 40 million

people in the U.S., they include only the visits and other relevant interactions within this system, a point we return to later.

The HealthJump data on the COVID-19 Research Database system includes data from the beginning of 2019 to late 2021, since the data on the system are updated daily. It includes information on appointments, demographics, diagnoses, immunizations, lab orders and results, medications, procedures, social history, vital statistics, and allergies. We use ICD9 and ICD10 codes to determine the reason for each visit. We create indicators for routine visits, depression, ADHD, confirmed maltreatment, and suspected maltreatment based on reports from the American Academy of Pediatrics of the 25 most common diagnoses (AAP, 2013) as well an extensive literature review of work focused on classifying child maltreatment in Electronic Health Records (EHR).[2]

We limit our sample to unique, non-cancelled encounter records for children aged 0 to 18 between January 1, 2019 to June 30, 2021. We limit our sample to only children living in the 50 United States or DC and for who gender is reported. This represents a total of 1,306,468 unique users and 7,430,464 records. Of these, 23% are elementary aged (5-10 years old), 12% are middle school aged (11-13 years old), and 34% are high school aged (14-18 years old).

Unfortunately, race is unknown for half of the sample. This is important to keep in mind as we discuss differences in the pattern of results across children of different race and ethnicity, so we also present results separate for those with missing race information.

[2] We also looked at visits related to asthma, vaccination, and other ICD-9 and ICD-10 codes.

III. Empirical Strategy

Our goal is to understand how access to in-person school environments during the COVID-19 pandemic affected children's health care usage. In the ideal experiment, identifying this effect would involve randomly assigning children to be in school settings or not, as well as perhaps randomly assigning them to different levels of exposure to the school setting (full week, 1 day a week, etc.). Of course, it is not possible to conduct such an experiment. We therefore make use of the variation in school district decisions about when to open and the modality of instruction to create a measure of children's exposure to schools during the fall of 2020.

This measure makes use of two sources of variation in district operations: (i) the timing of when the 2020-2021 school year commenced and (ii) the modality of instruction (in-person, online, or hybrid). For a given district, we characterize children's exposure to school as the interaction between two variables, *open* and *proportion*. $Open_{dt}$ is a dummy variable for whether district *d*'s school year had commenced by time *t*. $Proportion_{dt}$ is the number of days per week the average child was in school divided by 5. This is defined using the modality reported by the district and the age of the children, as described in the data section.

We do not have information on children's location at the district-level, rather at the zip3 level. So, we aggregate our measure of exposure to school environments by calculating the percent of the student population in a zip3 area that is enrolled in each district and multiplying that percentage by the measure of school exposure created using the school start date and the modality. In practice, we define the exposure of children in a given zip3, *z*, in week *t* as: $Exposure_{zt} = \sum_{d \in z} \frac{Students_d}{Students_z} \times Open_{dt} \times Proportion_{dt}$. Exposure can therefore be

interpreted as the share of students in an area in a given week that are exposed to a full week of school.

We use this measure of exposure to estimate the effects of exposure to school environments on child health care usage. Our estimation equation is the following:

$$Visits_{zt} = \alpha + \beta Exposure_{zt} + \gamma COVID_{zt} + \delta Unemp_{zt} + \sigma_z + \tau_t \,. \qquad (1)$$

We include fixed effects for each zip3 and week of our sample. The main coefficient of interest is β, which is identified off of changes in health care usage over time across zip3s, which is in turn driven by district decisions. The intuition behind the estimate is that it is a comparison of health outcomes for children before and after schools start the school year weighted by the varying degrees of in-person attendance and the proportion of students with access to open schools. In our main estimates, we use data from the fall of 2020 (specifically we limit our time horizon to the month before the first reported school opening and the month after the last school opening, i.e., July 1, 2020 to August 31, 2020) and only include controls for school exposure. Here, the assumption is that there is nothing else varying across zip3 areas over time in the fall of 2020 other than the children's exposure to school. We use robust standard errors at the state-level. To control for the fact that states are represented in the data to different degrees, we weight all of our models by the inverse of the proportion of 2019 users per state population.

Because school decisions about when and how to open may be affected by the local pandemic-related conditions, in other specifications, we control for the prevalence of COVID-19 and the unemployment rate. This allows us to identify the effects of children's exposure to school environments separately from the impacts of other pandemic-related conditions. Our estimates in this specification will only be biased if there are confounding factors affecting both

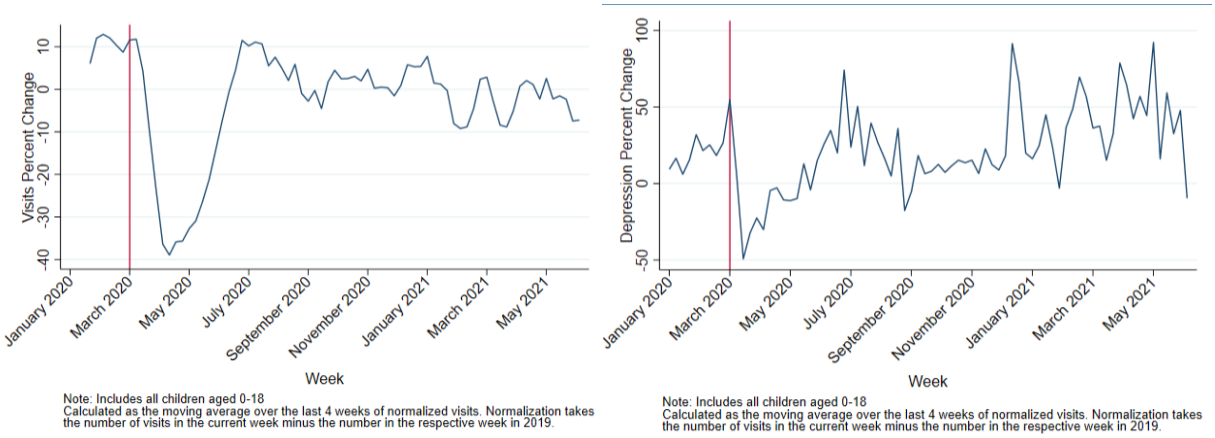
children's health and the decision about whether and how to open schools that are orthogonal to the prevalence of COVID-19 and to the unemployment rate.

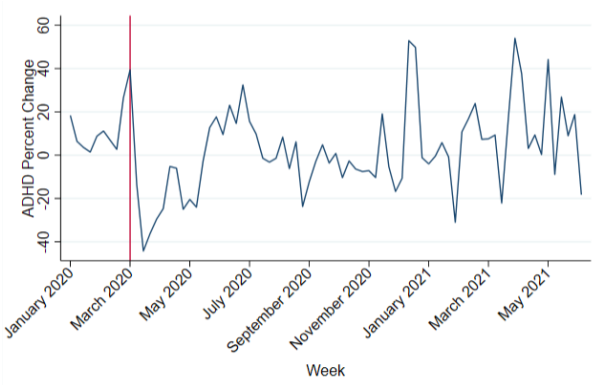
Under these assumptions, we interpret β as the effect of increasing the share of students in an area attending school for a full week by 1 percentage point. If other organizations, like child care centers, after school programs, and other educational, service, or child related business are tying their opening decisions to the school district decision making, to the extent that they are similarly timed, that will be incorporated into our estimates.

IV. Results

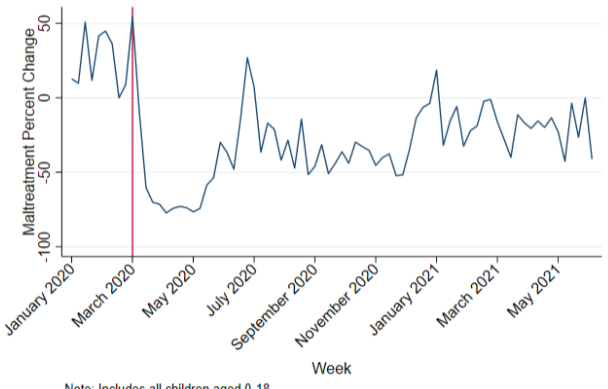
Changes in Child Health Care Use During the COVID-19 Pandemic

Overall

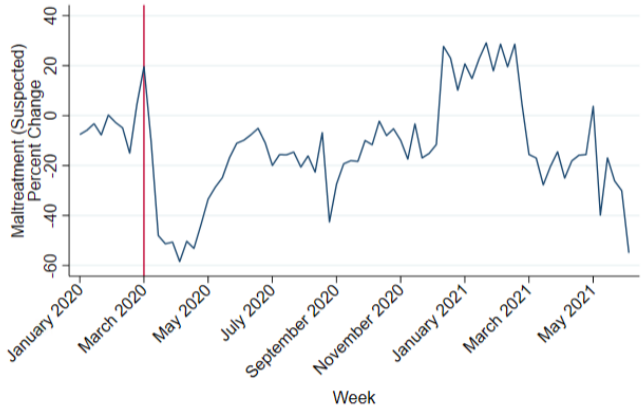




Note: Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

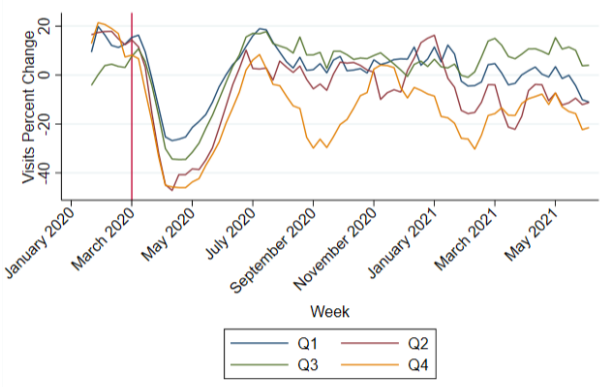


Note: Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

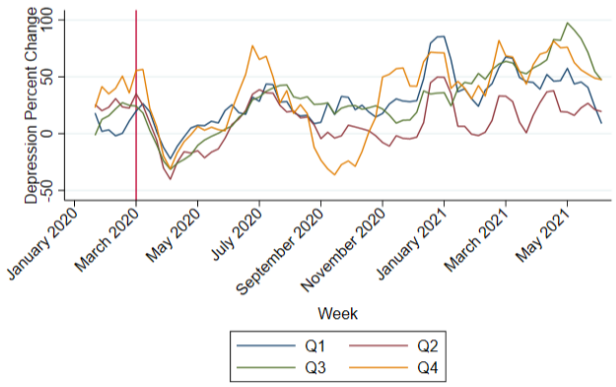


Note: Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

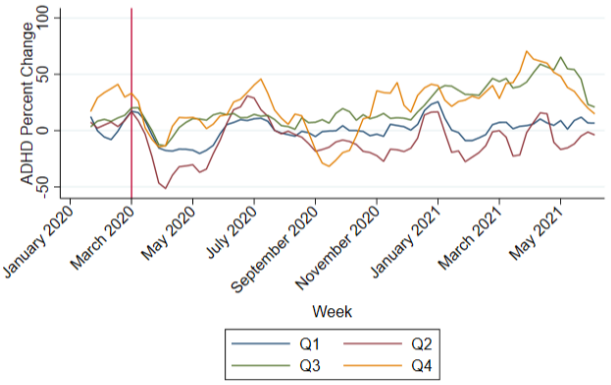
White



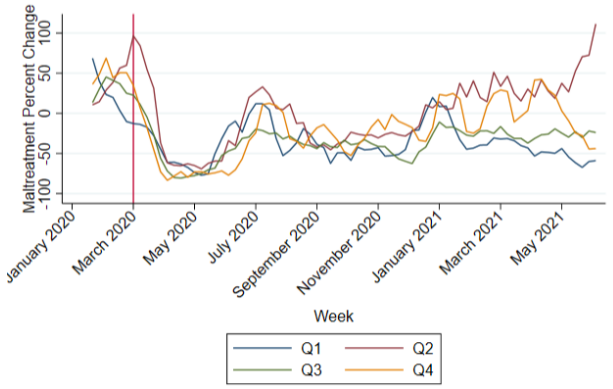
Note: Figures show zip3 percent white quartiles Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



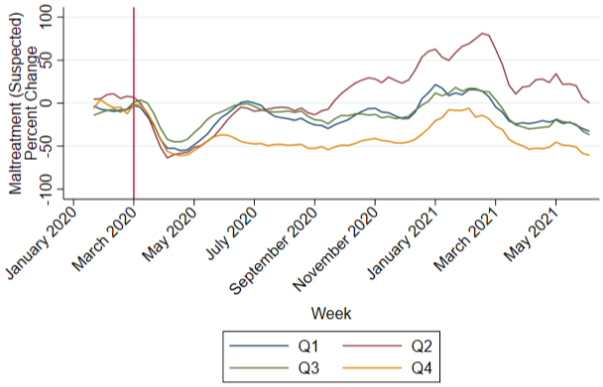
Note: Figures show zip3 percent white quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



Note: Figures show zip3 percent white quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

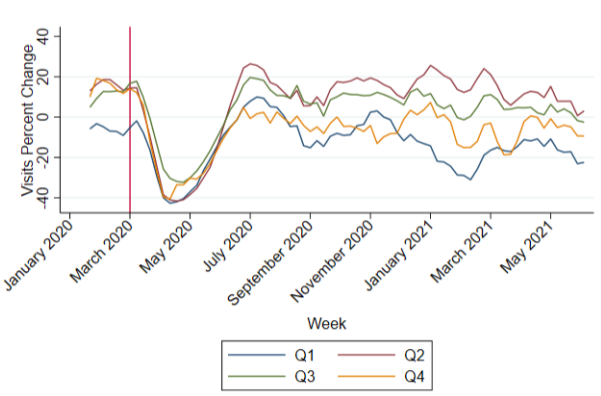


Note: Figures show zip3 percent white quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

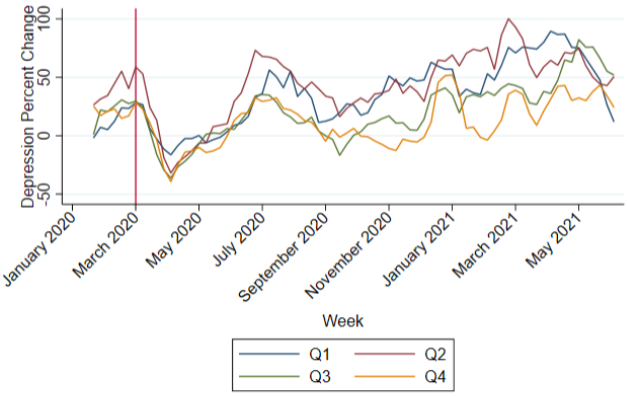


Note: Figures show zip3 percent white quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

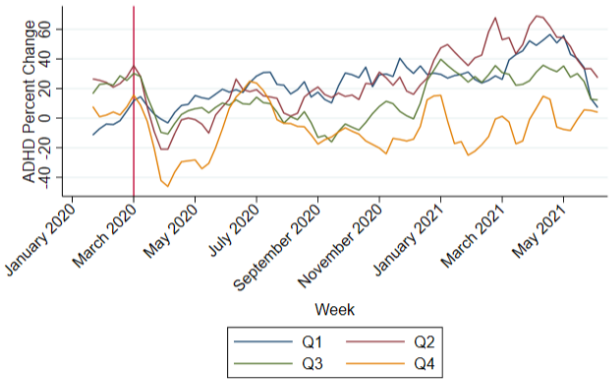
Black



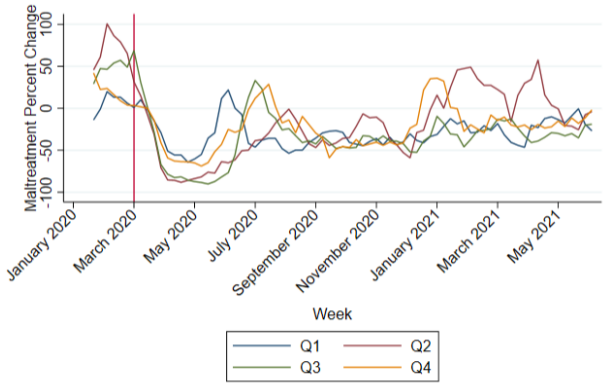
Note: Figures show zip3 percent Black quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



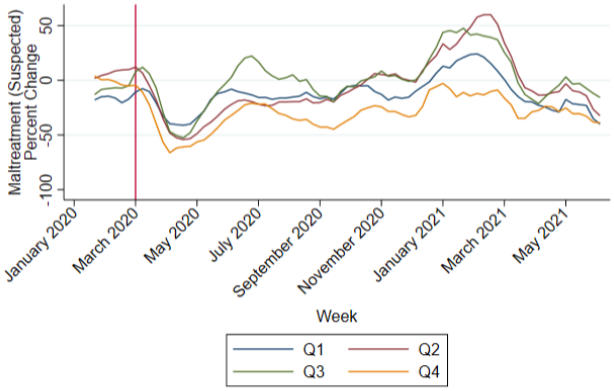
Note: Figures show zip3 percent Black quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



Note: Figures show zip3 percent Black quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

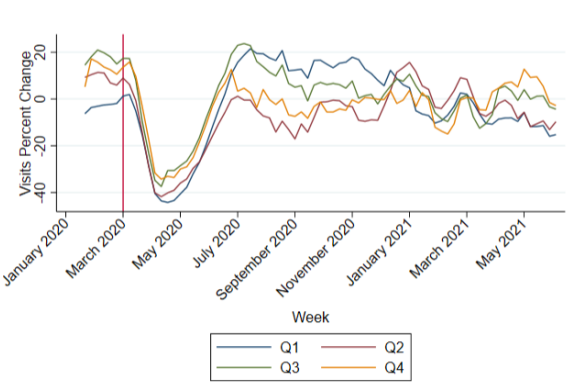


Note: Figures show zip3 percent Black quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

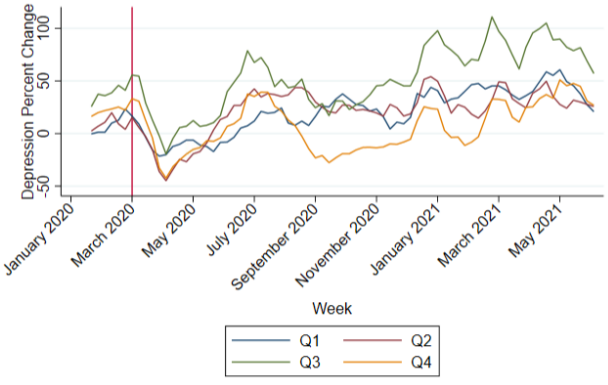


Note: Figures show zip3 percent Black quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

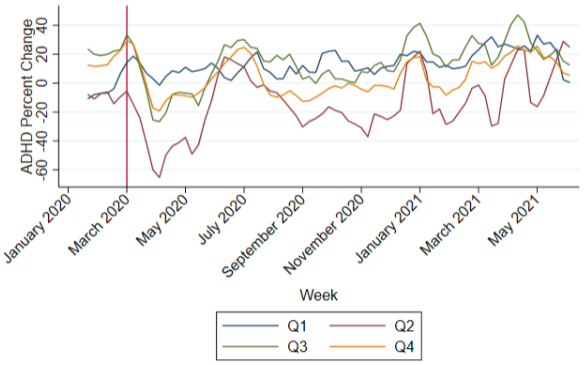
Poverty



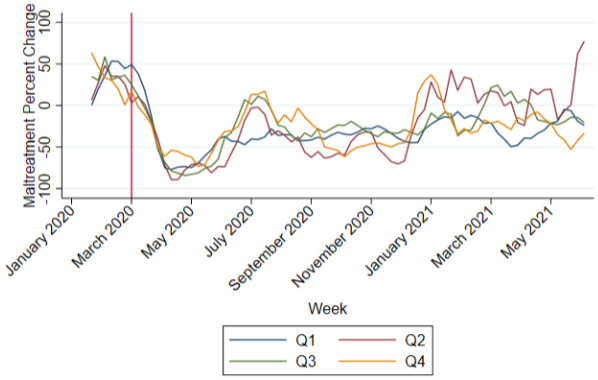
Note: Figures show zip3 percent poverty quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



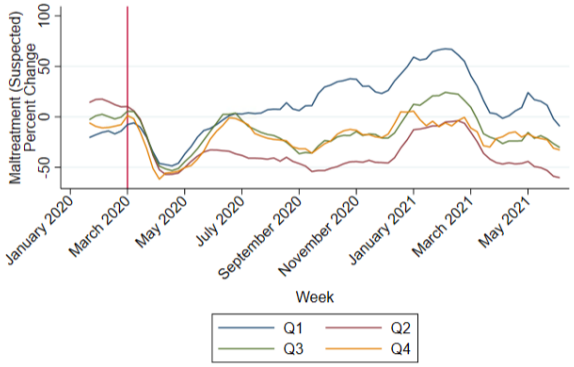
Note: Figures show zip3 percent poverty quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



Note: Figures show zip3 percent poverty quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

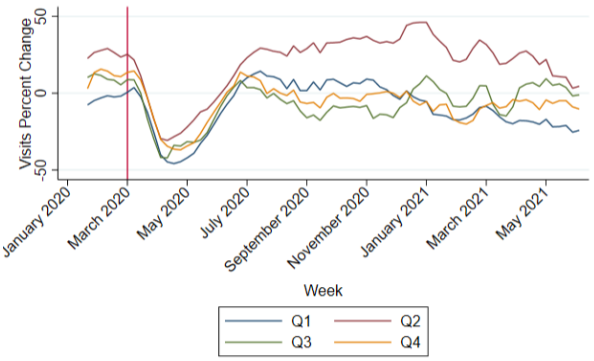


Note: Figures show zip3 percent poverty quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

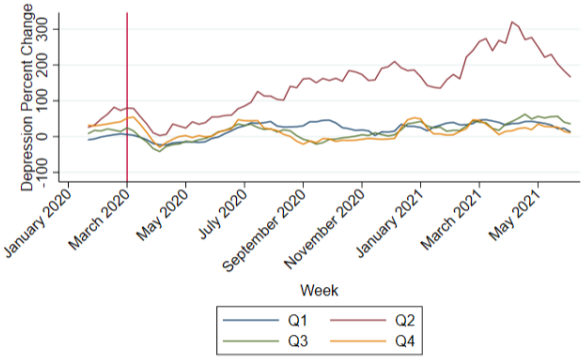


Note: Figures show zip3 percent poverty quartiles. Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

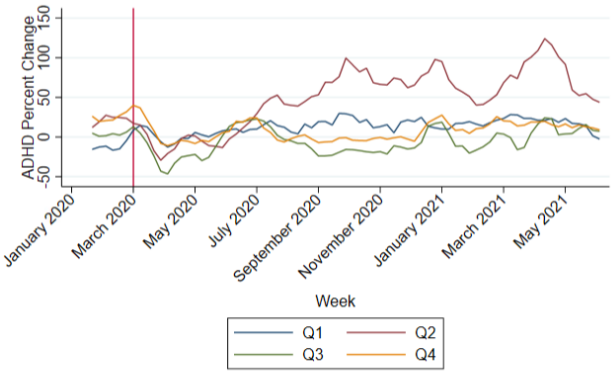
High School or Less



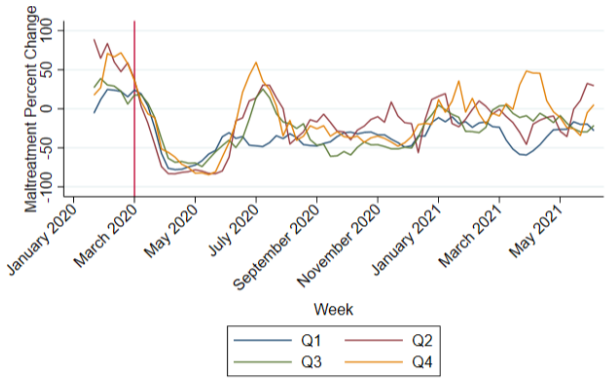
Note: Figures show zip3 percent with no high school degree quartiles.
Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



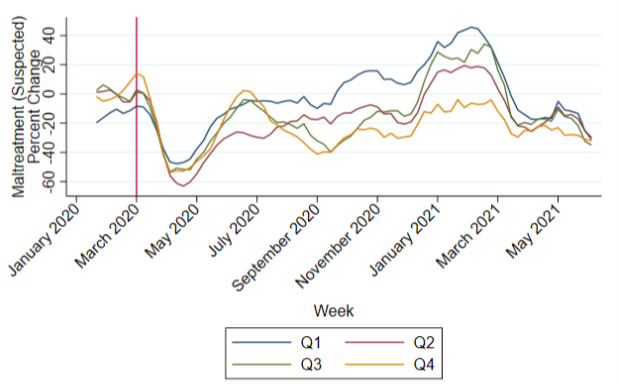
Note: Figures show zip3 percent with no high school degree quartiles.
Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



Note: Figures show zip3 percent with no high school degree quartiles.
Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



Note: Figures show zip3 percent with no high school degree quartiles.
Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.



Note: Figures show zip3 percent with no high school degree quartiles.
Includes all children aged 0-18
Calculated as the moving average over the last 4 weeks of normalized visits. Normalization takes the number of visits in the current week minus the number in the respective week in 2019.

Table 1. Total Visits per Unique Users *1000

	Visits	Depression	ADHD	Maltreatment	Maltreatment (Suspected)
Exposure	-2.072*	-0.166***	0.122*	-0.0007+	-0.317
	(0.901)	(0.043)	(0.053)	(0.004)	(0.563)
Constant	1125.813***	1.596	3.353	-0.280	203.946**
	(66.740)	(2.006)	(3.241)	(0.305)	(71.420)
Controls^{+}	X	X	X	X	X
N	9768	9768	9768	9768	9768
R^2	0.50	0.61	0.23	0.28	0.56

V. Conclusion

References

Appendix

We focus only on the tables that contain ICD9 or ICD10 codes. We begin with the complete list of encounters for our sample of children. We remove the encounters marked as patient communications, chart updates, void, audited, legacy, or migrated. We next merge appointment data to each record removing appointments that were cancelled. Finally we connect the diagnosis and procedure codes to each record. We search through all codes associated with a record to determine whether a record is associated with one of our outcomes of interest.

ICD9/10 codes are used for medical billing. Each code has a general parent category (e.g., 890), with associated specific sub-categories (e.g., 890.1). For our analysis, we compiled a list of codes for depression, ADHD, confirmed maltreatment, suspected maltreatment: burns, fracture/injury level 1 (head to pelvis), fracture/injury level 2 (other extremities), sexual, other. We determined our depression, ADHD, and confirmed maltreatment codes using the online ICD10 code handbook. Due to the fact that maltreatment codes are rarely used by physicians researchers often use codes and rules to determine "suspected" maltreatment. We compiled a list of these papers and synthesized their codes. Most of these papers were written prior to 2015, therefore they used ICD9codes. We used the online ICD9/10 converter tool to convert the ICD9 codes to ICD10. ICD10 codes can change over time, therefore we include all relevant codes for our study period. Importantly, we only convert specific codes to general if the general (or parent code) reflects the appropriate categorization. For instance, because we consider all burns as "suspected" maltreatment, we use the parent code E924. However, E910.4is accidental drowning in a bathtubwhich is likely associated with maltreatment. We do not include the parent category E910 (accidental drowning and submersion) because this includes codes for water-skiing and

recreational activities that are less likely due to maltreatment. Additionally, dueto the fact that codes can be general or specific, we match our data as follows. If a code in our data matches a code in our list exactly, we count the instance. We also take our data and truncate the codes with 2 or more decimal points to 1 decimal point. If the truncated code matches a code in our master list, then we count it. We do not count instances in our data that are general if the general category in our master list is not included due to issues of ambiguity.