

NBER WORKING PAPER SERIES

BEYOND READING, WRITING, AND ARITHMETIC:
THE ROLE OF TEACHERS AND SCHOOLS IN REPORTING CHILD MALTREATMENT

Maria D. Fitzpatrick
Cassandra Benson
Samuel R. Bondurant

Working Paper 27033
<http://www.nber.org/papers/w27033>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2020

Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau or the National Bureau of Economic Research. We appreciate the helpful comments from seminar participants at the NBER Summer Institute Children's Program Meeting, New York University's Wagner Graduate School of Public Service and Steinhardt School, Columbia University, the University of Virginia's Batten School of Public Policy and Curry School of Education. We thank Carina Chien and Martha Johnson for excellent research assistance. Data used in this article are a version of the National Child Abuse and Neglect Data System (NCANDS, HHS ACF Children's Bureau 2015) Child File, which was shared with us as part of a unique pilot secure micro-data program. All errors and omissions are our own.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Maria D. Fitzpatrick, Cassandra Benson, and Samuel R. Bondurant. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Beyond Reading, Writing, and Arithmetic: The Role of Teachers and Schools in Reporting
Child Maltreatment

Maria D. Fitzpatrick, Cassandra Benson, and Samuel R. Bondurant

NBER Working Paper No. 27033

April 2020

JEL No. I29,I31,J12

ABSTRACT

Nearly 4 in 10 children report experiencing maltreatment by adulthood. Early detection mitigates maltreatment's negative effects. Yet factors that drive early detection remain understudied. We examine one possible source of early detection: educators in school settings. Administrative data on reports of child maltreatment across the U.S. over a 14-year period allows us to use two different regression discontinuity methods, one based on school-entry laws and one on school calendars. Both methods show education professionals are reporting cases that would have been missed otherwise. These findings suggest that improved training and support of educators may improve society's ability to help children and families.

Maria D. Fitzpatrick
Department of Policy and Management
Cornell University
103 Martha Van Rensselaer Hall
Ithaca, NY 14853
and NBER
maria.d.fitzpatrick@cornell.edu

Samuel R. Bondurant
Dallas-Fort Worth Research Data Center
Federal Reserve Bank of Dallas
2200 N Pearl Street
Dallas, TX 75201
samuel.r.bondurant@census.gov

Cassandra Benson
Cornell University
cmb465@cornell.edu

I. Introduction

Child maltreatment is a vexing problem in the United States. In 2018, there were approximately 678,000 victims of confirmed child abuse and neglect.¹ While, nearly 13 percent of all children have a confirmed case of maltreatment by age 18 (Wildeman et al. 2014), survey data suggest that rates of actual child maltreatment are even higher than officially reported rates. For example, in 2011, researchers found that 4 in 10 children report experiencing maltreatment by the time they are ages 14 to 17 (Finkelhor et al. 2013). Additionally, a growing body of evidence suggests that experiencing maltreatment is detrimental to children's health (Bruce et al. 2009; Felitti et al. 1998) and that maltreatment has significant costs for society (Currie and Widom 2010; Fang et al. 2012; Currie and Tekin 2012).

Preventing child maltreatment may be the ideal goal of policymakers, but it is likely infeasible to prevent it entirely. This makes early detection essential for at least two reasons. First, earlier detection of maltreatment leads to quicker intervention. Research has shown that children's development is most fluid early in life (Fox, Levitt, and Nelson 2010; Shonkoff and Phillips 2000), thus limiting negative shocks early in life will benefit child development and wellbeing. Second, most interventions share a common goal: to provide the child with a safe, permanent home. To achieve this goal, Child Protective Services (CPS) may provide additional social services to the family or may recommend removal of the child to a foster care setting. In the former scenario, interventions aimed at altering abusive behavior are most likely to be successful when abusive patterns are less ingrained (Dozier et al. 2006; Fisher et al. 2007). In the latter scenario, evidence has suggested that parent-child relationships are stronger for children

¹ As determined by child protective service agencies. See <https://www.cdc.gov/violenceprevention/childabuseandneglect/index.html>

adopted at earlier ages (HHS 2011). Therefore, early detection of maltreatment is crucial for improving child outcomes.

In practice, identifying child maltreatment as early as possible depends on early, consistent observation of the child by individuals likely to report the maltreatment, something previously understudied by the literature.² Mandatory reporting laws were passed in all fifty states in some form by 1967 and are in place to compel most individuals who have regular contact with children (physicians, police, social workers, caregivers, teachers, etc.) to report maltreatment (Brown and Gallagher, 2015). Despite these laws, training and support of individuals mandated to report is lacking, leaving many mandatory reporters unaware of their responsibilities and obligations (Hawkins & McCallum, 2001; Kenny, 2004; Payne, 1991; Dinehart and Kenny 2015). Because of this, it is unclear how large of a role mandatory reporters play in reporting child maltreatment.

In this study, we focus on the effect of time with teachers and other education professionals on child maltreatment reporting.³ Given the significant proportion of a child's day spent with educators, it is logical to expect they will heavily contribute to reports of child maltreatment. However, there are two potential limiting factors: (1) teachers may not be fully equipped in identifying and reporting maltreatment, and (2) teachers may report children who would have otherwise been reported independent of the teacher's report. We discuss each of these two factors briefly. First, as mentioned above, training for teachers is nearly nonexistent.

² Most of the literature, particularly in economics, has focused on identifying the causes (Lindo, Schaller, and Hansen 2018; Raissan and Bullinger 2016; Berger et al. 2017; Zhai, Waldfogel and Brooks-Gunn 2013) and consequences (e.g. Currie and Widom 2010; Currie and Tekin 2012) of child maltreatment. To a lesser extent, there has been some research on interventions aimed at helping children (e.g. Doyle 2007a, 2007b, 2008; Doyle and Peters 2007; Aizer and Doyle 2013).

³ In what follows, we will use teachers as shorthand for all education professionals in traditional school settings.

For teachers who are provided formal training, either through their district or professional education setting, its quality varies (Crosson-Tower 2002, Child Welfare Information Gateway, 2003). As such, teachers may not be particularly adept at identifying abuse in practice or they may be reluctant to report it when they do observe it. Second, even if educators are effective at identifying and reporting child maltreatment, they may be reporting maltreatment that would have been seen and reported by someone else in the child's life.

To identify the role of teachers in reporting child maltreatment, we use two forms of exogenous variation in children's exposure to settings involving teachers. First, we compare reports of child maltreatment at age five for children who are age-eligible to attend kindergarten at age five to the investigated reports of child maltreatment for those too young to attend kindergarten until the following year. Second, we use public school calendar start and end dates, which vary across districts and across years, to examine how the number of investigated reports differs between the academic year and summer break. In both sets of analyses, we have access to the universe of child maltreatment reports across almost all states over a 14 year period, though our estimation sample varies depending on the setting. Importantly for interpretation of our estimates, these reports stem from referrals sent to CPS agencies that met criteria for screening into the system for an investigation or alternative response. (Nationally, about 56 percent of referrals to CPS agencies became reports in 2018. For more information on this process and the data, see Sections II and III.) In both analyses, we use regression discontinuity methods, which allow us to control flexibly for differences either across children who are born at different times of the year (in the first setting) or in seasonal patterns (in the second setting).

We find that additional time in school leads to marked increases in reports of child maltreatment. The number of reports for children who are age five is 5 to 10 percent higher for

children who are eligible to enroll in kindergarten at age five compared to those who are not. Moreover, the number of investigated child maltreatment reports is 30 to 65 percent higher at the beginning and end of the school year compared to the beginning and end of summer when children are not regularly interacting with teachers.

Having identified that more time in school increases investigated reports of maltreatment, we address the concern that reports by teachers replace reports by others. For example, teachers in schools may report maltreatment that was identified previously and reported by a child's physician. If this is the case, the reporting by teachers is not useful in identifying new maltreatment cases, although it could be useful for proving or substantiating a suspected case of child maltreatment. We test this potential issue by examining the subset of first reports made for each child. Our results indicate that teachers are making reports both for children who have not had prior cases in the system and those that have been reported before. Both might be helpful for child well-being since follow-on reports can be useful for confirming a case of suspected maltreatment.

Undoubtedly, not all these new reports represent cases where maltreatment actually occurred. The substantiation rate for teachers is about 1 in 5, which is lower than the rate for police officers or physicians, but higher than other types of reporters. However, since almost all reports by teachers take place during the school year, this means that at least 20 percent of the school-exposure-related reporting involves instances where child maltreatment occurred and would have gone unnoticed otherwise. (The fraction is higher if we assume that substantiation is a lower-bound indicator of true abuse because of the high legal bar for substantiation.) However, this indicates that school exposure results in many new reports that may not be true instances of maltreatment. Many families become involved in a difficult investigation and surveillance

process because the symptoms of poverty are sometimes similar to those of maltreatment. As a society, we need to determine how to balance over-reporting versus under-reporting, which play out differently for different types of families. In doing so, one main goal should be to maximize accurate or high-quality reporting (Waldfogel 2001). In speaking to the importance of teachers and other education professionals in early identification of child maltreatment, our findings suggest that expanded and improved teacher training would lead to higher-quality detection of child maltreatment.

II. Reporting of Child maltreatment

Federal law defines maltreatment as, at a minimum, “Any recent act or failure to act on the part of a parent or caretaker which results in death, serious physical or emotional harm, sexual abuse or exploitation; or an act or failure to act, which presents an imminent risk of serious harm.”⁴ In general, CPS agencies are responsible for receiving, managing, investigating, and responding to allegations of child maltreatment. In official statistics, allegations of maltreatment made to CPS agencies are called *referrals*. A referral is either screened-in (for investigation or response, 56 percent of referrals) or screened-out (44 percent of referrals). The most common reasons referrals are screened out include there not being enough information in the referral for a CPS response (e.g. lacking name or other information), the referral not involving child abuse or neglect, the referral involved a case where another agency was more appropriate for responding (e.g. the report did not involve children, the report was on tribal lands), and the referral did not meet a state’s standards for intake into the CPS system (HHS

⁴ The Child Abuse Prevention and Treatment Act (CAPTA), (P.L. 100-294), as amended by the CAPTA Reauthorization Act of 2010 (P.L. 111-320).

2018). Some states also “screen out” referrals that receive an alternative response, which we discuss below.

Screened-in referrals are called *reports*. Reports are the focus of this study. In 2018, there were 2.4 million reports of child maltreatment for a national rate of 32.5 per 1,000 children.⁵ Reports are either *investigated* or designated to an *alternative response*.

Investigations are the process through which caseworkers determine if a child was maltreated or is at-risk of maltreatment (if so, the case is termed *substantiated* or *indicated*, respectively; 17 percent of reports are substantiated or indicated) and whether an intervention is warranted.

Alternative response (which make up 14 percent of reports) is a process that usually sidesteps determination of whether maltreatment occurred and instead focuses on helping provide support and services to the child and family involved in the report. About 38 percent of children who are the subject of reports receive some sort of post-response service, including both children who are and are not defined as victims of substantiated maltreatment.

Some states screen-out referrals that receive an alternative (also called differential) response. In 2018, there were 27 states that had a differential response system in place. Of these, 6 did not include referrals that received an alternative response as reports. Not having data on the referrals that receive alternative responses in these states means that we are underestimating the true effect of school exposure on child maltreatment reporting that results in investigation or alternative response. That is particularly likely to be true if teachers are more likely to report in an alternative response environment because they expect the process to be more helpful or less intrusive for the child and family.

⁵ Information on reports for 2018 comes from the HHS Child Maltreatment Report for 2018. (<https://www.acf.hhs.gov/cb/resource/child-maltreatment-2018>) Summary statistics for our data are reported below.

Most states recognize four major types of maltreatment: neglect, physical abuse, psychological maltreatment, and sexual abuse.⁶ Neglect is the most common form of abuse, occurring in over half of all reports. Physical abuse is the second most common type (and makes up about 18 percent of all reports). Most reports are made by people who have contact with children as part of their professional responsibilities, including police officers (18 percent), social services personnel (11 percent), and education personnel (19 percent). Another 18 percent come from friends, neighbors and relatives, while 17 percent were from anonymous or unknown sources.

In addition to the fact that educators make up nearly one-fifth of all reporters, there are reasons to think that educators may play a key role in identifying and reporting child maltreatment. Nearly all states legally mandate that educators report suspected child maltreatment and neglect (Crosson-Tower, 2003). Typically, the law requires educators to report suspected abuse and neglect, provides protection for those educators who become involved in the case, and may penalize those who fail to meet their obligations (Crosson-Tower, 2003).

However, while many districts may provide employees with training on how to identify and report child maltreatment, these trainings are heterogeneous across districts and/or schools. Further, the district/school may decide which employees will receive the training, thereby potentially excluding some employees of the school or district. Despite limitations to training, mandatory report laws coupled with the amount of time teachers spend with children are expected to generate an increase in reports to CPS from education professionals when children are exposed to schools.

⁶ Many cases involve more than one type of maltreatment. If so, the official statistics count each separate type of maltreatment. This means reporting rates across types may sum to a number greater than 100.

III. Data Description

Our main source of information on child maltreatment is a version of the National Child Abuse and Neglect Data System's Child File, which was shared with us as part of a unique pilot secure micro-data program (NCANDS, HHS ACF Children's Bureau 2015). Since 1988, the NCANDS has been collecting data from states on all investigated or assessed reports of maltreatment to state CPS agencies. As discussed earlier, these reports are the result of a process where referrals (usually phone calls) to CPS agencies are assessed to determine whether they contain enough information and are relevant for the CPS agencies to respond. If so, they are screened-in, become reports, and are part of the NCANDS data. Information collected on reports includes demographic characteristics of the children and their perpetrators, the type of maltreatment, outcomes of the report (including investigation or alternative response), and types of services provided to children, if any.

In the first years of data collection, only a few states participated. Over time, more and more states began reporting. In 2003, 44 states participated.⁷ By 2005, the number had climbed to 48 states, plus the District of Columbia; the non-participating states are relatively small.⁸ We use the data from 2003 onward.⁹

The data are recorded at the child-by-report level. There are three key features of the data used to identify the role of education professionals in reporting child maltreatment. First, the data have information on a child's date of birth, which we use to define a child's age relative

⁷ The nonparticipating states in 2003 include Alabama, Alaska, Georgia, North Dakota, Oregon, and Wisconsin.

⁸ By 2005, the only states not participating regularly were North Dakota and Oregon.

⁹ Our results are robust to using the larger set of states from 2005 onwards.

to the school entry cutoff date in his or her state of residence.¹⁰ Second, the data also contain information on the report date of the incident. In our first analyses, we use this information, coupled with the child's date of birth, to determine a child's age at the time of the report.¹¹ Report date is also used to determine whether cases are more likely to be reported during the school year. Finally, the data include information on the reporter of the maltreatment. This allows us to clearly pinpoint the role of educators in reporting child maltreatment.

Our sample selection choices will be different across the two identification strategies. In our first set of analyses, we use cutoff dates for determining school entry eligibility. These generally vary at the state-level, and some changed over time. In some states, the determination of cutoff dates defining school entry eligibility is left up to local school districts, rather than set at the state level.¹² Therefore, in our first set of analyses, we limit our sample to a balanced panel of 35 states that both report valid data and have state-level cutoff dates determining school

¹⁰ In some states in some years, there is an overabundance of January 1 or January 15 birthdays. This is likely due to states assigning January 1 or January 15 as the date of birth when the true information is not recorded. Because this heaping could bias our estimates (Barreca, Lindo and Waddell, 2016), we drop all children with January 1 or January 15 birthdays in the years where there is an excess of January 1 or January 15 birthdays in that state. This leads us to drop 21,662 incidents of child maltreatment from the sample, or 0.99 percent of the reported incidents over this period. It is unlikely that the measurement error caused by this missing information is systematically related to a child's age in relation to the cutoff for school enrollment. Moreover, since few states have cutoff dates within two months of January (and our optimal bandwidth procedures suggest about two months is optimal), the loss of children for whom January 1 or January 15 was the true date of birth likely has very small effects on our estimates.

¹¹ About 30 percent of reports also include information about the date of the incident. Information on incident dates is missing in 12 states and for many observations in other states. Among those that have a valid incident date, 92 percent have the same incident and report date. Another 6 percent have a report date within one week of the report date.

¹² States that leave the determination of cutoff dates defining school entry eligibility up to local school districts include Colorado, Massachusetts, New Hampshire (after 2005), New Jersey, New York (after 2001), Ohio (after 2002), Pennsylvania (after 2004), Vermont, and Washington (until 2006).

entry eligibility in each year between 2003 and 2015.¹³ In the second methodology we make use of school district calendars that vary at the county level. The sample for this analysis consists of a panel of 30 counties for which we have valid school start and end dates, generally from 2007 to 2015. We elaborate more on the samples and identification strategies in Sections IV and V.

In Table 1, we present information about reports of child maltreatment for all children between 2003 and 2015. Over the period, there were nearly 48 million reports of child maltreatment. About half of the children involved in these reports were male; 25 percent of children involved in a report were African American. The average age of children in these reports is 7.5. Information on the demographic and other characteristics of the alleged perpetrator is only logged when the case is substantiated, and is therefore available for about 20 percent of reports. When the information is available, a child's parent is the perpetrator of the abuse in over 90 percent of cases. The most common form of abuse is neglect (51 percent), followed by physical abuse (18 percent). Education professionals are responsible for 16 percent of the reports in our sample. Most cases of child maltreatment are unsubstantiated (62 percent). About one quarter are substantiated and the remainder have some other resolution in the system.¹⁴

¹³ Five states (Maryland, Michigan, Nevada, South Dakota, and Tennessee) are missing one year of data over this period. We include these in the panel. Our results are not sensitive to excluding those states.

¹⁴ "A finding of substantiated (sometimes referred to as founded) typically means that the child protective services (CPS) agency believes that an incident of child abuse or neglect, as defined by State law, has happened." <https://training.cfsrportal.acf.hhs.gov/section-2-understanding-child-welfare-system/3013>

IV. The Regression Discontinuity Comparison Using School Entry Laws

The ideal experiment aimed at identifying the role of teachers and other school professionals in identifying and reporting child maltreatment would involve randomly assigning children either to school settings or to their standard environments when not enrolled in school. Such an experiment is difficult, if not impossible, to conduct. Instead, to identify the role of teachers and other school professionals in identifying and reporting child maltreatment, we make use of exogenous variation in the timing of a child's exposure to a school setting. In our first regression discontinuity method, exogenous variation in school exposure stems from the statewide policies regarding the earliest age at which a child is allowed to enter school.

IV.a. Background on School Entry Laws

Our research design rests on institutional policies within each state that determine the age a child is eligible to enter kindergarten. Most states require a child to turn five on or before a statewide cutoff date in order to enroll in kindergarten, be it voluntary or mandatory kindergarten, in a particular year.¹⁵ For example, Texas requires students to turn five on or before September 1 of the school year they enter kindergarten. Thus, the kindergarten class in the fall of 2010 is made up (largely) of children born between September 2, 2004 and September 1, 2005. Children born on September 2, 2005 wait to enroll in kindergarten in the fall of 2011.

These rules are not strictly binding. Many children, particularly those who would be young for their grade, wait a year to enroll in kindergarten. Others enter school before they are technically eligible according to the law in their state of residence. Figure 1, which is taken from

¹⁵ In states where kindergarten is not mandatory, the requirement is that any child must be six before the statewide cutoff date in order to enroll in first grade in that year.

Dobkin and Ferreira (2010) illustrates these enrollment facts. In the figure, Panel A depicts kindergarten enrollment of children in Texas, and Panel B shows enrollment of children in California. The samples include children in the 2000 Decennial Census Long Form Census Data who became age five within 180 days of the school entry eligibility cutoff in that particular state, September 1 in Texas and December 2 in California. Those with negative values for relative age were born before the cutoff for kindergarten entry, which was in time to enroll in public school at age five. As can be seen in both panels, compliance with the law is not perfect, however, there is an approximately 60 percentage point increase in kindergarten enrollment for those born just before the cutoff date in their state (relative to those born just after the cutoff). Although the information is from just two states, it is representative of the enrollment rates of five year olds in the U.S. more broadly.

States vary their entry cutoff dates and these dates have changed over time. In general states moved their cutoff date earlier in the school year such that many states now adhere to an August 31 or September 1 cutoff date. For children who turned 5 in 2003, 5 states used a July or August cutoff date, 27 used a September cutoff date, and 11 states used a cutoff date in October, December or January. By 2015, 9 states used an August cutoff date, 29 used a September cutoff date, and only 3 used an October or January cutoff date.¹⁶ Table 2 summarizes the 2015 cutoff dates and previous changes to the statewide cutoff date by state.

¹⁶ During transition periods, states tend to phase-in a new cutoff date over several years. For instance, California moved from a December 2 cutoff date in 2011 to a September 1 cutoff date in 2014 by pushing the cutoff date forward 30 days per year for three consecutive years.

IV.b. Regression Discontinuity Framework

Using this information on school entry laws across states and over time, we test whether the number of investigated reports of child maltreatment for five year olds to CPS is greater for children born just before the statewide school entry eligibility cutoff date (those who enroll in kindergarten in the year they turn five) relative to those born just after the entry eligibility cutoff date (those who enroll in kindergarten in the year they turn six).

We start with the balanced panel of 35 states with state-level cutoff dates determining school entry eligibility for the years 2003 to 2015 (Table 2). We then keep only the reports for five-year-olds. Define d to be a child's age in days relative to the cutoff date for kindergarten in his/her state of residence. We define d such that positive values indicate children who were born before their state's eligibility cutoff and are therefore eligible to enter school earlier (at age five). We then aggregate reports according to this relative age measure such that a single observation in our model represents the total number of reports at age 5 for children born on a particular day relative with respect to their state's eligibility cutoff for kindergarten enrollment. Our estimation equation is therefore the following:

$$Y_d = \delta_0 + \delta_1 I_d + f(d \cdot I_d) + \varepsilon_d . \quad (1)$$

where Y is an outcome of interest (e.g., the number of CPS reports) measured at age five for children born on relative date d . I is an indicator for $d > 0$. In other words, it is an indicator for children who were eligible to enter school at age five, given the statewide eligibility cutoff in place at the time they turned five in their state of residence at age five. The function $f(d \cdot I_d)$ represents the polynomial used to control for the age-child maltreatment relationship. We allow the relationship between d and child maltreatment reporting to vary on either side of the discontinuity. The error term, ε_d , is clustered on the running variable.

There are two types of approaches to estimating regression discontinuity models: flexible global parametric models and local regression with a triangular kernel that places more weight on observations closest to the cutoff point. In our main tables and analyses, we present results of the effect of school exposure on child maltreatment reports using the local regression techniques, which place additional weight on observations nearest to the cutoff date. Results using the global parametric models are very similar (see Appendices). Similarly, there are different methods for choosing the bandwidth in local regressions and for choosing the order of the polynomial. As we will show in our main tables and Appendices, our conclusions are not sensitive to these choices.

The coefficient of interest, δ_1 , measures the difference in reports to state CPS agencies for five-year-olds who are eligible to enroll in school at age five versus those not eligible until a year later. Our outcome measure is the total number of investigated reports of five-year-olds. Thus, in this specification we are comparing the total number of reports of maltreatment for children aged 5 at the time of the report between our treated and comparison children. Treated children are those who were age-eligible to enter kindergarten in the year they turn 5, whereas the comparison children are those who were age-ineligible to enter kindergarten in the year they turn 5. Since most of our sample consists of states where the cutoff date for kindergarten eligibility is in August or September, the “treated” children were eligible to be enrolled in kindergarten during the time they were age 5, while the comparison children will not have enrolled until the following academic year. Thus a child’s exposure to educational professionals varies during the child’s fifth year of life, where treated children had greater exposure to education professionals during the year they are age 5.¹⁷ Note that this definition of the outcome

¹⁷ Consider an example, Illinois, where the cutoff date is September 1. The students in our sample in the “treated” group are those born within about 45 days before the cutoff date. Therefore, it is children born between July 17 and September 1. These children are aged 5 for the period from

measure will not cover some of the time spent in kindergarten for states with cutoffs later in the year (like California, which had a December cutoff for part of the period studied). This means that our estimated difference in reports across all states will be an underestimate of the effect of exposure to school settings at age 5.

The underlying assumption with this identification strategy is that other factors related to child maltreatment reporting do not systematically change around the entry eligibility cutoff dates in ways that are not captured by our flexible polynomial in relative age, i.e. $E[\varepsilon_a I_a] = 0$. In turn, this assumption implies three things. First, it implies that, in the neighborhood around the school eligibility cutoff, date of birth does not vary systematically across families of different types in ways that are not captured by our flexible polynomial in *Relative Date*. This assumption is supported by the evidence in Dickert-Conlin and Elder (2010), who show that there are no discontinuities in the density of births or in maternal characteristics around state eligibility cutoffs. Second, since we are measuring outcomes at age a for children relative to the eligibility cutoff in the state in which they reside at age five, it must be the case that migration does not differ systematically for children born before and after the eligibility cutoff in ways that are not captured by our flexible polynomial in *Relative Date*. This assumption is supported by previous research showing that observable characteristics of children age five do not vary with the side of the eligibility cutoff their birthday is on (e.g., Fitzpatrick 2012). Third, in order for δ_1 to represent an increase in reports, it must be the case that actual rates of abuse and neglect do not increase with exposure to school. Since parents are by far the main perpetrators of maltreatment

July 17 (to September 1) to July 16 (to August 31), which for most of them will include all of the days they are actually in kindergarten. The “comparison group” is born between September 2 and October 16. The time they are age 5 will span from September 2 (to October 16) to September 1 (to October 15). Since they are not eligible for kindergarten, most will not be enrolled in kindergarten during this period.

(over 95 percent of confirmed perpetrators are parents or their partners) and time with parents often decreases significantly when children are in school, we interpret our estimates as an increase in reporting. However, this assumption is difficult to support empirically because true maltreatment is a latent variable; for further discussion of this assumption, see the Section VI.

Our data lack information about whether children are enrolled in school. Therefore, in the above analyses, we are comparing outcomes of children born on adjacent days regardless of whether they are enrolled in school. As such, this is a sharp regression discontinuity design using eligibility for school enrollment as the treatment. Not all children enroll in kindergarten on-time. Many are held back a year and some receive exceptions to the rules and enroll earlier (see Figure 1). Similarly, some children, like many of those in universal preschool states, are exposed to school settings before they reach kindergarten eligibility. Therefore, the effects we estimate are net intent-to-treat estimates of the exposure to school settings on child maltreatment reporting. As we interpret the effects, it is useful to keep in mind that 80 to 90 percent of five year olds born just before the cutoff are enrolled in kindergarten at age five, while just about 20 percent of those born after the cutoff date are enrolled in kindergarten at age five. Also, with considerations towards measuring the entirety of exposure to school settings, our estimates do not account for universal preschool where the differential exposure to school settings occurs at age 4 rather than age 5. Our estimates only represent the net additional reporting by kindergarten teachers (and adjacent educational personnel) above and beyond the reporting by other reporting entities at age 5. Moreover, many children are exposed to formal care settings earlier than age 5. Caregivers in those settings may also be responsible for reporting child maltreatment. Our estimates represent the net additional reporting by kindergarten teachers above and beyond the reporting by other caregivers at age 5.

IV.c. Estimates of the Effects of Contact with Teachers on Child Maltreatment Reporting Using the Regression Discontinuity Design

First, we present graphical information about the data in the neighborhood of the cutoff for eligibility. In Figure 2, we plot the number of reports of maltreatment of five year olds to CPS agencies by *Relative Date*. As a reminder, positive values of *Relative Date* indicate children who were eligible for school in the year they turned five and negative values indicate the group of children who had to wait another year before enrolling in school. In Panel A of the figure, each dot measures the total number of reports that were made by education professionals. In Panel B, each dot measures the total number of reports across all children born on that day that were made by reporters other than education professionals. Across the panels, there is a clear increase in reports by education professionals that is not accompanied by a change in reporting by others.

To measure the size of the increase in reports by education professionals, we turn to our estimation results. In Table 3, we present the estimated change in reports at age five for those eligible to attend kindergarten at age five relative to those who are not eligible until the following year. The local regression estimates range from 339 to 582. All are statistically significant at the one percent level. Since the average number of reported instances of child maltreatment while age 5 among the children ineligible for kindergarten until age six is about 6,300 per relative day of birth, the estimates represent an increase in reported instances of child maltreatment of between 5.4 and 9.2 percent.¹⁸ Given relevant populations sizes, if this is all new reporting for a

¹⁸ 6,300 is the number of reports of five-year-olds born on the day after the cutoff date for kindergarten eligibility in their state of residence.

different child on each report, 0.3 percent of 5 year-olds are being newly reported to the CPS agencies because of their kindergarten eligibility.¹⁹

In Appendix Table 1, we confirm that these results hold when using global polynomial techniques instead of the local regression specifications. Across the specifications, and in accord with the visual evidence, we find a statistically significant increase in the number of child maltreatment reports at age five for children who are eligible to enroll in school at age five relative to those that have to wait a year to enroll.

To ensure the increased reporting is driven by school contact, we disaggregate the reports of child maltreatment by the type of reporter and estimate the differences in reports at age five by reporter type. In columns (2) and (3) of Table 3, we report the results using the sample of reports by education professionals and all other reporters, respectively. All of the estimates of the increase in reporting by education professionals are large in magnitude (as compared to those for other types of reporters) and statistically significant. The number of reports by education professionals goes up by between 357 and 401. All are statistically significant at the one percent level. The estimated effect on child maltreatment reporting by other types of reporters is between -8 and 169, but only the local cubic specification estimate of 169 is statistically significant at conventional levels. Importantly, the only coefficient in column (3) of Table 3 that

¹⁹ About 4 million children were born each year from 1997 to 2010. This implies 4 million 5 year-olds in each year of our sample, or about 11,000 kids born each day (if kids are evenly distributed across the year). That means, in the entire sample, there are 143,000 5-year olds with each relative day of birth. (13 years x 11,000 kids per day.) Take the estimate from column 2 of table 2 and round it to 400. So, 5 year-olds born on “Sept 1” have 400 less reports while age 5 than 5 year-olds born on Sept 2. If each of the reports is a new report, the 400 report increase means 0.3 percent more 5 year-olds are being reported because of their kindergarten eligibility. Since Figure 2 shows the estimate is steady out to large bandwidths, it applies to a broader population than just those born within one day of the eligibility cutoff, but is still an increase of 0.3 percent for this population.

is negative is very small, suggesting that the reporting by education professionals consists of reports that would not have occurred without education professionals, showing there is little to no crowd-out of reports by others.

To get a better sense of the nature of the increase in child maltreatment reporting, in Table 4, we present estimates of Equation (1) using the local linear specification segmenting the sample by the different types of maltreatment. We report the estimates of δ_1 for the specific categories of neglect and physical abuse; we report our estimates for the whole sample in the top row for the ease of comparison. The omitted category is all other forms of maltreatment. The increases in daily reports at age five related to school eligibility are statistically significant and are comprised of 39 percent cases of neglect, 47 percent cases of physical abuse, and 14 percent other types of maltreatment. Note that this is quite different from the broader composition of maltreatment at age five (51 percent neglect, 18 percent physical, and 31 percent other types). The increased reporting by teachers across abuse types closely mirrors the overall increases (35 percent neglect, 49 percent physical abuse, and 16 percent other types).

Of interest is whether the increase in reports of child maltreatment at age five that occurs due to school contact is reporting of child maltreatment that would not have been identified without the school contact. To investigate the extent of this claim, we examine whether there exists increases in the number of child maltreatment reports that are a child's first report in his or her life. The results in Table 5 suggest that school contact increases the number of "first reports" by 143 reports. In other words, of the 339 reports at age five that are due to school contact, over 40 percent are the first reported experience of child maltreatment for a particular child. Among those reports that have to do with physical abuse, 55 percent are the first reported experience of child maltreatment for a particular child. The results in column 2 and 3 of Table 5 show that the

increase in first reports is driven entirely by education professionals reporting. Therefore, teachers are identifying maltreatment for some children who otherwise would not have been identified and making additional reports for children who have already been identified as victims of child maltreatment. Both types of reporting are important to reduce the prevalence of child maltreatment.

To summarize the results so far, the combination of results in Table 3 and Table 4 suggests that eligibility for school enrollment at age five, resulting in increased contact with education professionals, increases reports of child maltreatment at age five. Notably, education professionals are much more likely to identify physical abuse at this stage than reporters are more generally. Forty percent of the new reports are the first-reported instance of child maltreatment for a given child. Education professionals are responsible for the new reporting, and little of the child maltreatment reported by education professionals is maltreatment that would have been reported by other people had the children not been enrolled in school. The combination of these estimates suggests that teachers and other education professionals play a key role in the early detection and reporting of child maltreatment.

We expect that the increase in reported child maltreatment at age five for those who are eligible for school at age five is reporting that occurs earlier than it would have had the children delayed their entry to school. If this is the case, we should see a pattern where our estimate is different from zero when the two groups of children (those eligible for public kindergarten at five versus those eligible at age six) have differential exposure to school and zero otherwise. We rerun Equation (1) changing the age sample cutoff from 5 to all of the different ages reported on child maltreatment reports. In Figure 3, we present coefficient estimates (and confidence intervals) for our intent-to-treat estimates of the effect of school contact on child maltreatment

reporting at all ages from zero to seventeen. In the figure, one can see that there is elevated reporting at age five for children eligible for school enrollment at age five. Many of the estimates at other ages are quite close to zero in magnitude and are statistically indistinguishable from zero at conventional levels. Exceptions include the estimates at ages six and seventeen. At age six, children who did not enroll in kindergarten in the previous year (since in many places it is not mandatory to enroll in kindergarten or to enroll at age five) will also experience school contact for the first time, which would lead to a positive estimate at that age. At age seventeen, the number of child abuse reports for children eligible for entry to school at age five is less than that of their counterparts who were ineligible until age six. This is likely driven by the fact that at age seventeen some of the cohort that was eligible to enter school at age five will have already left school and therefore no longer have contact with education professionals. Regardless, none of the estimates at other ages are as large as those at age five. This is not surprising since none of the other shifts in exposure to school contact are as large in magnitude as the one that occurs at age five.

Therefore, we see these results across different aged children as consistent with a story in which (i) education professionals identify abuse and neglect that other people do not and (ii) because some children are eligible to enter school at a younger age the maltreatment they suffer gets reported earlier. We now turn to our results using information on the timing of reports across the calendar year to offer additional evidence about the role of education professionals in reporting maltreatment.

V. *The Regression Discontinuity Comparison Using School Calendars*

In this section, we use exogenous variation in school calendars to identify the effects of exposure to school settings on child maltreatment reporting. The intuition for the use of the variation in school calendars is straightforward. Consider two identical children, one of whom is abused in the week before school starts, the other is abused the following week. On the one hand, we might expect the latter child's abuse to have a better chance of being reported because the child is more likely to be observed by a set of adults with some training in maltreatment identification and a responsibility for reporting it. On the other hand, if there are plenty of qualified observers in the child's life outside of school and/or if educators are not good at identifying or reporting child maltreatment, there may be no difference in the likelihood that the maltreatment of each child is reported. To determine whether this is the case, we examine reporting patterns for children in the 25 largest public school systems (30 largest counties) in the US (as of 2014, reported in US Department of Education 2017).

V.a. *Background on School Calendars*

This regression discontinuity design stems from the fact that school districts set their school start and end dates, and these calendar dates vary across districts. Traditionally, schools have based their operation dates on the agrarian schedule. While the U.S. economy is no longer as reliant on agriculture as it once was, school calendars are still largely based on this type of seasonality. For example, Michigan, Minnesota, and Virginia have laws on the books restricting local districts from starting before specific dates in August (ECS 2014).

Despite state restrictions on calendars, for the most part, the decisions about which day children start classes, which day they end classes, and which days they attend school are left up

to local districts. These decisions are made by local school boards within the confines of state regulations and are made based on a number of factors including resource management and the timing of holidays.²⁰ The calendars, the resulting school start dates, and the school end dates vary across districts and vary from year-to-year within districts.

For this study, we coded the school start and end dates for 25 of the largest districts in the U.S. (U.S. Department of Education, 2017). This included New York City, Los Angeles, Chicago, Miami-Dade, Houston, the state of Hawaii, and many others. These districts cover 30 counties.²¹ Where possible, we included school start and end dates from the 2005-2006 school year to the 2015-2016 school-year.²² For a given calendar year, the relevant start and end dates come from adjacent school-years. For example, in 2015, the relevant end date in the spring is from the 2014-2015 school year, while the relevant start date in the fall is from the 2015-2016 school year.

V.b. Regression Discontinuity Framework

Using this information on school calendars across districts and over time, we test whether the number of reports of child maltreatment is greater at calendar dates when children are attending school compared to dates when they are not. In these analyses, we include all children ages 6 to 17, not just those age five as we did in the previous analyses.

²⁰ One common set of issues is the timing of Labor Day in relation to the start of the school year and the timing of Memorial Day in relation to the end of the school year. Another is the timing of winter break with relation to end-of-semester exam periods.

²¹ Because the finest level of geography in our NCANDS data is county, we only use districts that are contiguous with counties.

²² Inclusion in the sample is based solely on our ability to find calendar information from a given district in a given year. Appendix Table 2 contains information on which districts are included in which years.

Define *Relative Start* to be the number of days between a particular day of the year, d , and the school start date in county c in the fall of that year, t .²³ Similarly, define *Relative End* to be the number of days between a given day of the year and the last day of school in the spring of that year in county c . We define both such that positive values indicate dates that occur later in the year. Then we also define two variables of key interest: *After Start* and *After End*. The first is defined as one if $Relative\ Start_{cdt} > 0$, and zero otherwise, and the second is defined as $Relative\ End_{cdt} > 0$, and zero otherwise. In other words, *After Start* is defined such that positive values for $Relative\ Start_{cdt}$ indicate times during which children are attending the local public schools, while *After End* is defined such that negative values for $Relative\ End_{cdt}$ indicate times when children are attending the local public schools. Given these definitions for relative start and end dates, we expect a positive coefficient at the start-of-school date and a negative coefficient at the end-of-school date. Both estimates would capture the effect of being in school relative to being out of school.

When we examine how child maltreatment reporting changes at the beginning of the school-year, our estimation equation is the following:

$$Y_{cdt} = \alpha + \beta AfterStart_{cdt} + f(Relative\ Start_{cdt} \cdot AfterStart_{cdt}) + \varepsilon_{cdt} . \quad (2)$$

Y is an outcome of interest (e.g., the number of CPS reports) in county c on a given date dt . The function $f(Relative\ Start_{cdt} \cdot AfterStart_{cdt})$ represents the flexible polynomial used to control for

²³ We use county as the level of analysis for two reasons. First, the finest level of geographic information on reports in NCANDS data is the county of the incident. Second, most of the largest districts cover entire counties. Some, like the New York City Public School District and the Hawaii Public School District, cover multiple counties. In other counties, there are more than one district. For example, Baltimore County contains the Baltimore County School District and the Baltimore City School District. When this is the case, we include the county as long as relevant districts have the same start or end date.

the relationship between school start dates and child maltreatment reporting.²⁴ The error term, ε_{cct} , is clustered on the running variable. An analogous equation defines our analyses examining the change in child abuse reports at the end of the school-year.

The coefficient of interest, β , measures the difference in reports to state CPS agencies on days in the fall (or spring) when school is in session relative to days when it is not. The assumption underlying our use of this specification is that there are no other discontinuous changes in reporting when the school-year starts or ends in a particular county that are unrelated to the school start or end date itself. This assumption would be violated if school enrollment itself increases the prevalence of abuse and neglect. For example, if school enrollment changes child-perpetrator interactions in a way that leads to more maltreatment, our interpretation of our estimates as an increase in reporting would be misplaced. Given that parents are the most likely source of maltreatment, longer exposure to school reduces that amount of time in which parents can inflict physical abuse, sexual abuse, and psychological maltreatment. All else equal, a greater number of hours in school should reduce the prevalence of maltreatment perpetrated against children. Thus, we expect that less abuse occurs when children are enrolled in school.

V.c. Estimates of the Effects of Contact with Teachers on Child Maltreatment Reporting Using the Regression Discontinuity Design Stemming from School District Calendars

Before turning to our estimates, we present two figures using raw data that demonstrate the role of school start and end dates in the reporting of child maltreatment. In Figure 4, we plot the number of reports involving children of any age across the entire country over the whole

²⁴ As in the previous section, we present results using local linear estimation techniques in the main text and estimates using global polynomial specifications in the Appendix.

period of our sample (2003 to 2015) by the calendar day of the year on which they are reported. Several things about the information in the figure are worth discussing. First, there are dramatic decreases in reporting on holidays, most notably Christmas, New Year's (and the week between it and Christmas), Thanksgiving, and July 4th. Second, there is a seven-day cyclical pattern in reporting that is driven by the fact that there is much less reporting on weekend days than on weekdays. Although the data come from multiple years and particular calendar days do not fall on the same day of the week in each year, the calendar days in the figure have different proportions of weekdays and weekends, which results in the pattern seen in the figure. Third, more directly related to our setting, there is a distinct drop-off in reports during the summer months. The decline begins in mid-May and reverses course beginning in mid-August. Not coincidentally, school start dates largely range from mid-August to mid-September (the range of the increase) and the school end dates range largely from mid-May to mid-June (the range of the decrease).

Next, we present data around the neighborhood of the start and end date of the school year graphically for the 25 school districts. In Panel A and B of Figure 5, we present reports per day relative to the start and end date of school in the relevant county and year, respectively. In the figure, there is a clear increase in reports when school starts (in Panel A) and decrease in reports when school ends (Panel B). In Panel A, after the initial jump in reports at the start of the school year, there is a continued upward trend in reports. This may be the result of the fact that it takes teachers time to observe or confirm instances of maltreatment. Also, we wondered if we might see an increase in reports leading up to the end of the school year as teachers make reports in anticipation of children being out of school and in less well-monitored situations. Consistent with this theory, reports in Panel B are at a peak one month before the school year ends.

However, it is a local peak and there is a similarly sized peak in reports about 3 months before school ends as well.

In Table 6, we present regression discontinuity estimates of the effects of school starting (columns 1 to 3) and school ending (columns 4 to 6) on child maltreatment reporting. At the start of the school-year, reports go up between 48 and 64 percent (estimates of 7.7 and 10.3). Depending on which specification we focus on (local linear, quadratic or cubic), education professionals are responsible for between 33 and 41 percent of the increase in child maltreatment reports at the beginning of the school-year. Unlike in the analyses in the previous section, reporting by others sources increases quite a bit at the beginning of the school-year. About half of the increased reporting by others is driven by increases in reporting from social services, mental health professionals and others.²⁵ At the end of the school-year, reports decrease by anywhere from 20 to 24 percent, depending on the specification. Similar to what happens at the start of the school-year, this is a decrease in reporting by both educators and others.

The increase in reporting by non-educational professionals is different than in our first set of analyses where teachers were responsible for almost all of the increased reporting. There are differences across the analyses samples that may lead to these differences. First, the children are of different ages across the two analyses. Schools often have staff in these health and social service areas on campus, like school counselors or school nurses. The prevalence of those staff may be higher in schools for older children. Also, in many places, older children need a physical at the beginning of the school year, particularly if they play sports. It may also be the case that teachers are more likely to refer older students to social services and mental health professionals who then report the maltreatment. Second, the sample for the school start and end date analysis

²⁵ References available from the authors upon request.

includes a subset of children in the largest districts in the country. It may be the case that these districts operate differently or have different reporting processes in place relative to the rest of the country, which is the coverage of our sample in the first set of analyses.

In Table 7, we present the estimated effects of exposure to school disaggregated by report source and type of abuse. When we disaggregate by type of report, we find that the distribution of reports instigated by the start or end of school dates remains similar to the distribution in overall reports. For example, at the beginning of the school-year, 34 percent of the increase in reports are cases of neglect and 22 percent are cases involving physical abuse. This is probably a better match to overall reporting patterns than we saw in the school eligibility setting because in this setting, as we saw in Table 7, the school contact associated with the start or end of the school-year leads to increased reporting by a wider set of reporters than the exposure to school for children aged five. If we look closely at teachers, in column (2), the reports are equally spread across neglect, physical abuse, and other types of abuse (nearly one-third of the increase in reports by teachers is in each category). This, combined with the information on how educators are more likely to identify physical abuse for children who start school at age five, suggests that teachers are more likely to identify and report physical abuse than other reporters.

VI. Discussion

We have shown that an increase in time spent in school leads to an increase in reporting of child maltreatment using two different identification strategies. Our interpretation is that this increase represents a change in reporting. An alternative explanation is that there is something about school attendance that leads to an increase in the underlying rate of abuse. We think this is unlikely for a few reasons. First, we note that parents are the primary perpetrator of maltreatment

(over 95 percent of confirmed perpetrators are parents or their partners) and that physical abuse, sexual abuse, and psychological maltreatment all require the perpetrator and victim to be in the same place. For many families, removing the child from the perpetrators' home during school hours reduces the child's exposure to that perpetrator. For these families, less abuse occurs when children are enrolled in school. Some families may experience the opposite change. With a shift from full-time daycare to school hours (if after-school programming is not used) it may mean some children spend more time with their parents. If so, the rates of underlying maltreatment may increase for these children. However, this pattern is unlikely to be driving our estimates using school calendar dates.

Second, many theories for child abuse suggest that abuse is more likely to occur when a family is under financial stress (Lindo and Shaller 2014). When children are in school, more resources are available to the household. By getting children out of the home environment, school may provide parents with more time flexibility at a low cost. Also, many children get access to school lunches and breakfasts, which should ease the strain on families.

Third, many rational-actor theories for child abuse also posit that abuse is less likely to occur when the chances of being caught increase (Berger 2004, 2005). Since children spending time in school increases monitoring of children for potential abuse, perpetrators should be less likely to commit abuse when children are in school.

Finally, the discontinuity in the number of reports occurs immediately when children start or end school in the fall or spring. Most explanations for why parental abuse would increase likely necessitate enough time elapsing for parents to observe poor in-school behavior (low grades, misbehavior) by their children. Therefore, these changes in maltreatment would not appear immediately upon school entry.

For all of these reasons, we argue that the increases in reporting of child maltreatment with time in school are changes in maltreatment reporting rather than changes in underlying rates of maltreatment. However, we acknowledge that some may consider that to be based on a strong assumption and that, because true maltreatment is a latent variable, it is difficult to find empirical support for the assumption.

VII. Conclusion

In this study, we have shown that time spent in school, and the resulting contact with education professionals leads to increases in the number of reports of child maltreatment. The results indicate that the increased reporting by education professionals is high-quality new reporting. It is not over-reporting, nor is it reporting of child maltreatment that would have been identified and reported by someone else in contact with the child. As such, we conclude that teachers are playing a key role in the early detection and reporting of child maltreatment.

These findings have several potential implications. First, since training of education professionals in the identification and reporting of child maltreatment is uneven across and within districts, our estimates are likely a lower-bound on how effective teachers can be at identifying and reporting child maltreatment. There could be returns to more consistent, higher-quality training of teachers and other education professionals. Second, any discussion of the benefits of time spent in school should include estimates of the improvement in child wellbeing that stems from the resulting early detection and reporting of child maltreatment. Such benefits should be a part of the discussions of measures of teacher and school quality, as well as of particular policies that extend the amount of time children spend in school settings (e.g. extending the school year, public preschool provision).

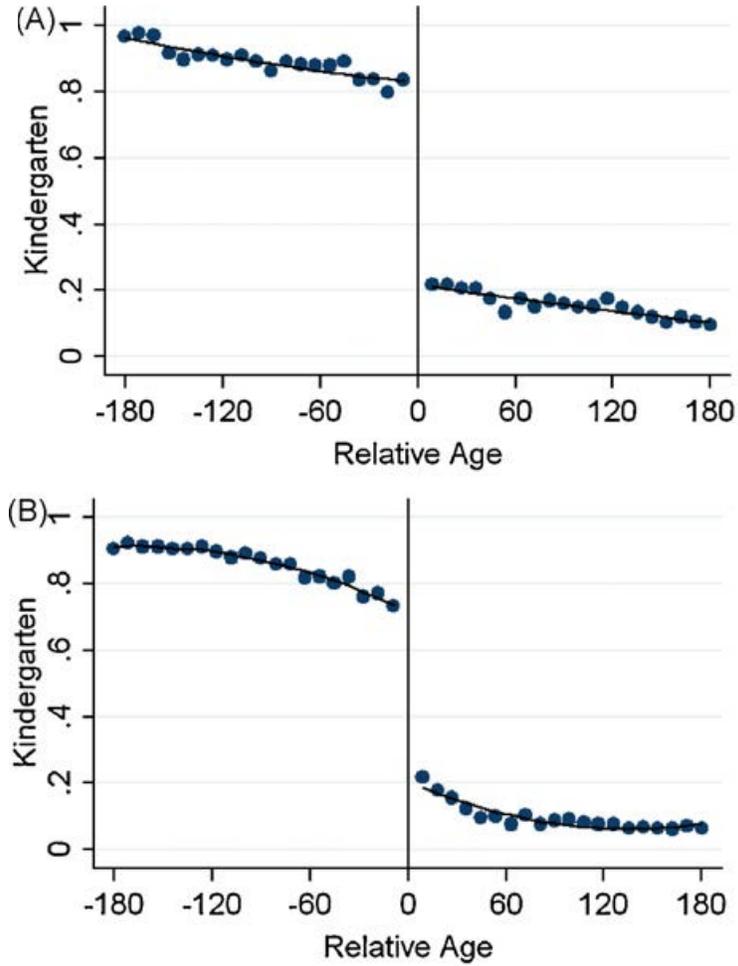
References

- Aizer, Anna and Joseph J Doyle, Jr. 2013. "Economics of Child Wellbeing" In Ben-Arieh, Asher, Casas, Ferran, Fronese, Ivar. and Korbin, Jill E. (Eds.) 2013 Handbook of Child Well-Being. Theories, Methods and Policies in Global Perspective . Dordrecht: Springer.
- Barreca, Alan I., Jason Lindo, and Glen R. Waddell. 2016. Heaping-induced Bias in Regression-Discontinuity Designs. *Economic Inquiry*, 54(1): 268-293.
- Berger, Lawrence, Sarah A Font, Kristen S Black, Jane Waldfogel. 2017. "Income and child maltreatment in unmarried families: evidence from the earned income tax credit" *Rev Econ Household* (2017) 15:1345–1372
- Berger, Lawrence M. 2004. "Income, Family Structure, and Child Maltreatment Risk," *Children and Youth Services Review*, Vol. 26, No. 8, pp. 725–748.
- Berger, Lawrence M. 2005. "Income, Family Characteristics, and Physical Violence Toward Children," *Child Abuse & Neglect*, Vol. 29, No. 2, pp. 107–133.
- Brown, L.B., III and Gallagher, K. (2015). Mandatory reporting of abuse: A historical perspective on the evolution of states' current mandatory reporting laws with a review of the laws in the Commonwealth of Pennsylvania. *Villanova Law Review Tolle Lege*, 59(6), 1-44.
- Bruce, J., Fisher, P. A., Pears, K. C., & Levine, S. (2009). Morning cortisol levels in preschool-aged foster children: Differential effects of maltreatment type. *Developmental Psychobiology*, 51(1), 14-23.
- Calonico, Sebastian, Mattias Cattaneo, and Rocio Titiunik. 2014a. Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6):2295-2326.
- Calonico, Sebastian, Mattias Cattaneo, and Rocio Titiunik. 2014b. Robust Data-Driven Inference in the Regression-Discontinuity Design. *Stata Journal*, 14(4):909-946.
- Child Welfare Information Gateway. 2003. The role of educators in preventing and responding to child abuse and neglect. Washington, DC: U.S. Department of Health and Human Services, Children's Bureau.
- Crosson-Tower, Cynthia. 2002. *When Children Are Abused: An Educator's Guide to Intervention*. Boston, MA: Allyn and Bacon.
- Currie, Janet and Erdal Tekin. 2012. "Understanding the Cycle: Childhood Maltreatment and Future Crime" *Journal of Hum Resources* 47(2): 509–549.
- Currie, J., & Widom, C. S. (2010). Long-term consequences of child abuse and neglect on adult economic well-being. *Child Maltreatment*, 15(2), 111–120.
- Dickert-Conlin, Stacy, and Todd Elder. 2010. "Suburban Legend: School Cutoff Dates and the Timing of Births." *Economics of Education Review*. 29: 826-841.
- Dinehart, Laura and Maureen C. Kenny (2015) Knowledge of Child Abuse and Reporting Practices Among Early Care and Education Providers, *Journal of Research in Childhood Education*, 29:4, 429-443, DOI: 10.1080/02568543.2015.1073818
- Dozier, M., Peloso, E., Lindhiem, O., Gordon, M. K., Manni, M., Sepulveda, S., & Levine, S. (2006). Developing evidence-based interventions for foster children: An example of a randomized clinical trial with infants and toddlers. *Journal of Social Issues*, 62(4), 767-785.
- Doyle, Joseph J Jr. 2007a. "Can't buy me love? Subsidizing the care of related children." *Journal of Public Economics* 91: 281–304.

- Doyle, Joseph J Jr. 2007b. "Child Protection and Child Outcomes: Measuring the Effects of Foster Care" *American Economic Review*. 97(5): 1583-1610.
- Doyle, Joseph J Jr. 2008. "Child Protection and Adult Crime: Using Investigator Assignment to Estimate Causal Effects of Foster Care" *Journal of Political Economy*, vol. 116, no. 4
- Doyle, Joseph J Jr, and H. Elizabeth Peters. 2007. "The market for foster care: an empirical study of the impact of foster care subsidies" *Rev Econ Household* (2007) 5:329–351.
- Education Commission on the States. 2014. "Number of Instructional Days/Hours in the School Year" Report downloaded from http://www.ecs.org/wp-content/uploads/Number-of-Instructional-Days-Hours-in-a-School-Year_Revised.pdf on February 14, 2018.
- Felitti VJ, Anda RF, Nordenberg D, et al. Relationship of childhood abuse and household dysfunction to many of the leading causes of death in adults: the Adverse Childhood Experiences (ACE) Study. *Am J Prev Med*. 1998;14(4):245-258.
- Fisher, P. A., Stoolmiller, M., Gunnar, M. R., & Burraston, B. O. (2007). "Effects of a therapeutic intervention for foster preschoolers on diurnal cortisol activity." *Psychoneuroendocrinology*, 32(8-10), 892-905.
- Fitzpatrick, Maria D. 2012. "Revising Our Thinking about the Relationship between Maternal Labor Supply and Preschool." *Journal of Human Resources*. 47: 583-612.
- Fox SE, Levitt P, Nelson CA III. "How the timing and quality of early experiences influence the development of brain architecture." *Child Dev*. 2010; 81(1):28-40.
- Gelman, Andrew and Guido Imbens, 2014. Why High-order Polynomials Should Not Be Used in Regression Discontinuity Designs. *NBER Working Paper 20405*.
- Hawkins, R., & McCallum, C. (2001). "Effects of mandatory notification training on the tendency to report hypothetical cases of child abuse and neglect." *Child Abuse Review*, 10(5), 301–322. doi:10.1002/(ISSN)1099-0852
- Health and Human Services. 2011. Children Adopted from Foster Care: Child and Family Characteristics, Adoption Motivation and Well-Being. ASPE Research Brief. <https://aspe.hhs.gov/system/files/pdf/76246/rb.pdf>
- Kenny, M. C. (2004). "Teachers' attitudes toward and knowledge of child maltreatment." *Child Abuse & Neglect*, 28(12), 1311–1319. doi:10.1016/j.chiabu.2004.06.010
- Lee, David S. and Thomas Lemieux. 2010. Regression Discontinuity Designs in Economics. *Journal of Economic Literature* 48: 281-355.
- Lindo, Jason, Jessamyn Schaller, and Benjamin Hansen. 2018. "Caution! Men Not at Work: Gender Specific Labor Market Conditions and Child Maltreatment." *Journal of Public Economics* 163: 77-98.
- Lindo, Jason and Jessamyn Shaller. 2014. "Economic Determinants of Child Maltreatment." *Encyclopedia of Law and Economics*.
- Payne, B. (1991). The principal's role in reporting child abuse. Alexandria, VA: National Association of Elementary School Principals. (ERIC Document Reproduction Service No. ED 333 594)
- Rassian, Kerri M, and Lindsey Rose Bullinger. 2017. "Money matters: Does the minimum wage affect child maltreatment rates?" *Children and Youth Services Review* 72: 60–70
- Shonkoff JP, Phillips DA, eds. *From Neurons to Neighborhoods: The Science of Early Childhood Development*. Washington, DC: National Academy Press; 2000.
- U.S. Department of Education, National Center for Education Statistics, Common Core of Data (CCD), "Public Elementary/Secondary School Universe Survey," 2014-15; "Local Education Agency Universe Survey," 1990-91, 2000-01, 2010-11, and 2014-15; and

- Regulatory Adjusted Cohort Graduation Rates (ACGR), 2010-11 through 2014-15, retrieved June 26, 2017, from <http://www2.ed.gov/about/inits/ed/edfacts/data-files/index.html>. (Table was prepared June 2017.)
- U.S. Department of Health and Human Services, Administration for Children and Families, Administration on Children, Youth and Families, Children's Bureau (2015). National Child Abuse and Neglect Data System (NCANDS) Child File, FFY 2015 [Dataset]. Available from the National Data Archive on Child Abuse and Neglect Web site, <http://www.ndacan.cornell.edu>
- U.S. Department of Health & Human Services, Administration for Children and Families, Administration on Children, Youth and Families, Children's Bureau. (2020). Child maltreatment 2018. Available from <https://www.acf.hhs.gov/cb/resource/child-maltreatment-2018>
- Waldfoegel, Jane. 2001. *The Future of Child Protection. How to Break the Cycle of Abuse and Neglect*. Harvard University Press. Cambridge, MA.
- Wildeman, Christopher, Natalia Emanuel, John M. Leventhal, Emily Putnam-Hornstein, Jane Waldfoegel, and Hedwig Lee. 2014. "The Prevalence of Confirmed Maltreatment Among US Children, 2004-2011." *JAMA Pediatrics* 168:706-713.
- Zhai, Fuhua, Jane Waldfoegel, and Jeanne Brooks-Gunn. 2013. Estimating the effects of Head Start on parenting and child maltreatment." *Children and Youth Services Review* 35 (2013) 1119–1129

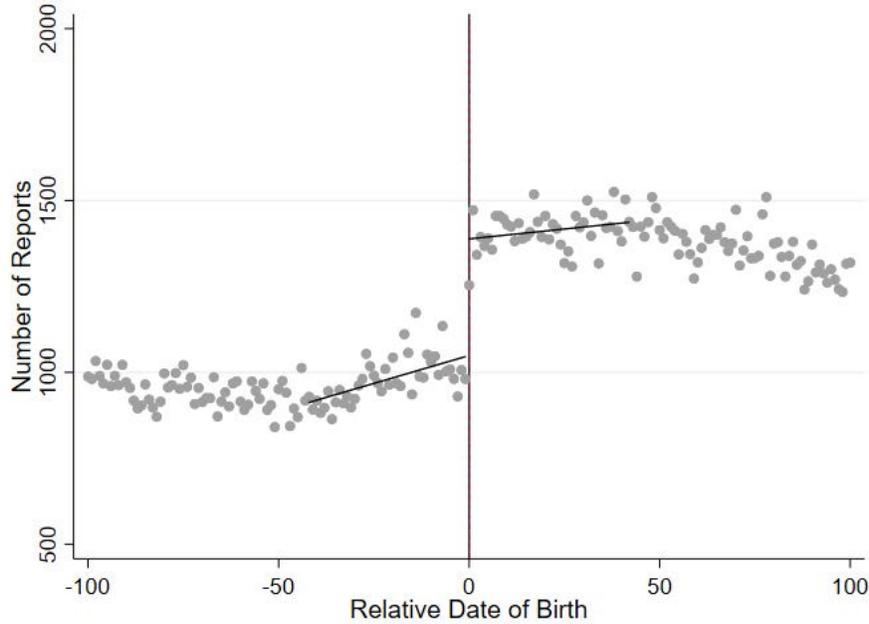
Figure 1. Enrollment in Kindergarten, by Date of Birth Relative to School Entry Eligibility Cutoff



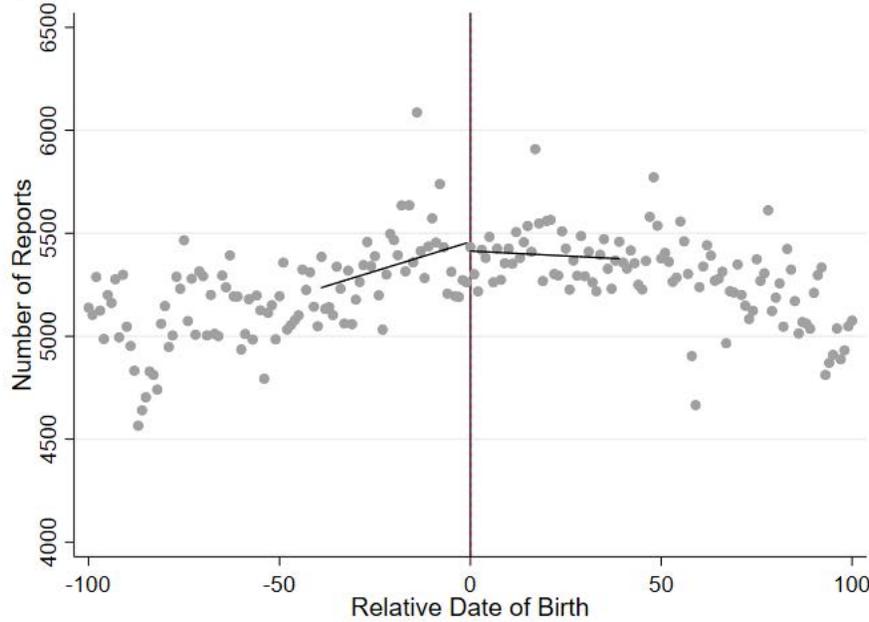
Note: From Dobkin and Ferreira (2010). (A) shows kindergarten enrollment of children in Texas, and (B) shows enrollment of children in California. The samples include children in the 2000 Decennial Census Long Form Census Data who became age five within 180 days of the school entry eligibility cutoff in that particular state, September 1 in Texas and December 2 in California. Here, those with negative values for relative age were born before the cutoff for kindergarten entry, which was in time to enroll in public school at age five.

Figure 2. Reports to CPS Agencies at Age 5 by Date of Birth Relative to School Entry Eligibility Cutoff

Panel A. All Reported Occurrences by Education Professionals

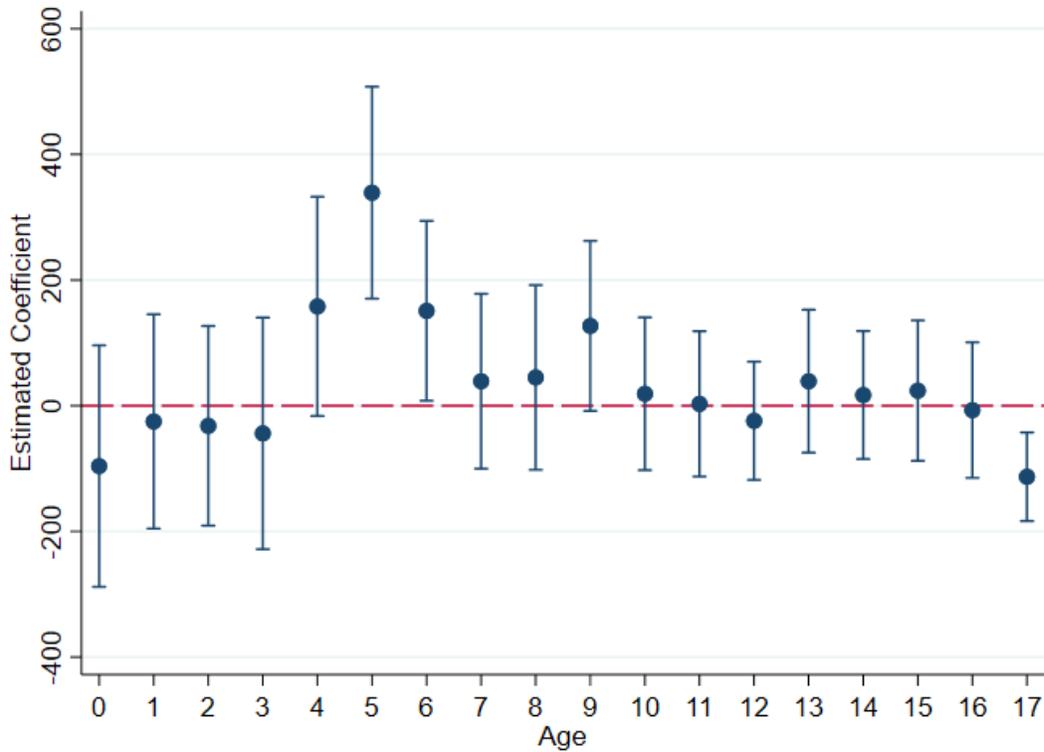


Panel B. All Reported Occurrences by Others



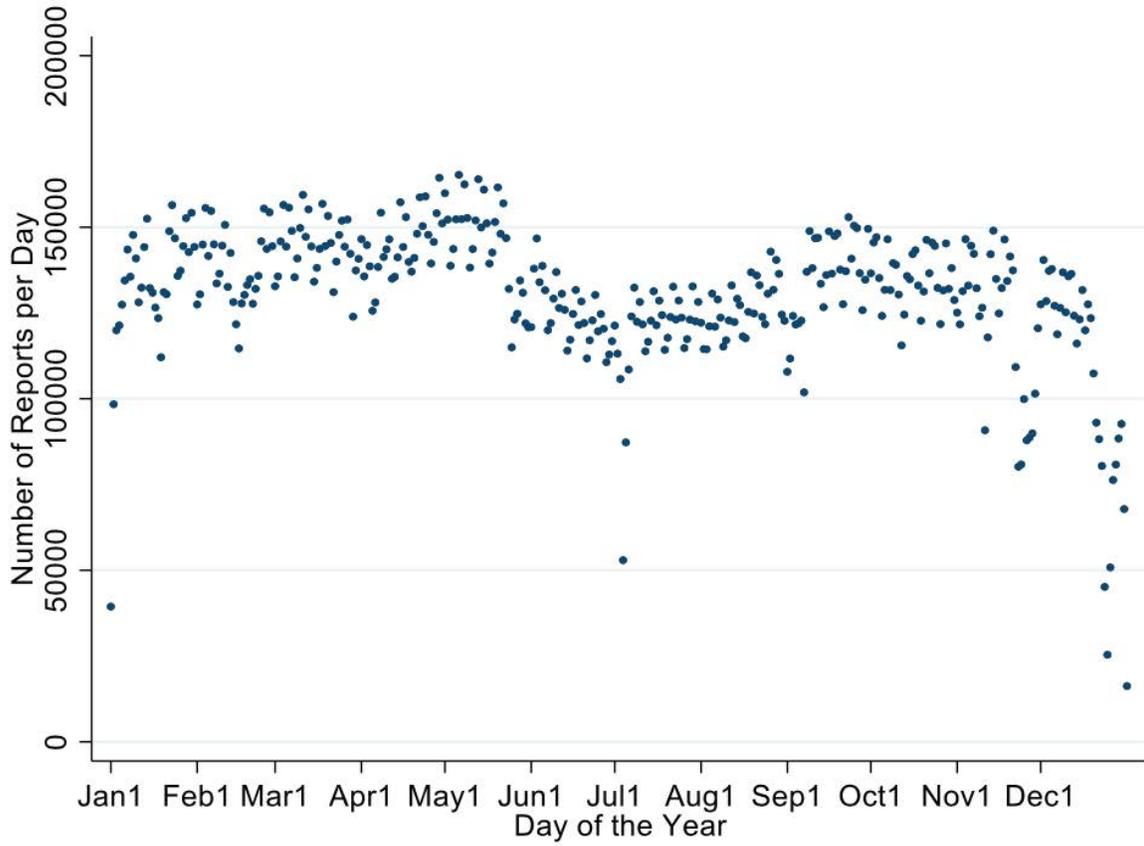
Note: Authors' calculations using the restricted-use versions of the National Child Abuse and Neglect Data System and include information reported between 2003 and 2015. Each dot in the figure represents the number of reports for children born on a given day relative to the cutoff for school entry in their state of residence. Here, those with positive values for relative age were born before the cutoff for kindergarten entry, which was in time to enroll in public school at age five.

Figure 3. Estimates of the Increase in Reporting to CPS Agencies at Age 0 through 17 for Children Eligible for School at Age 5 (Relative to Those Eligible at Age 6)



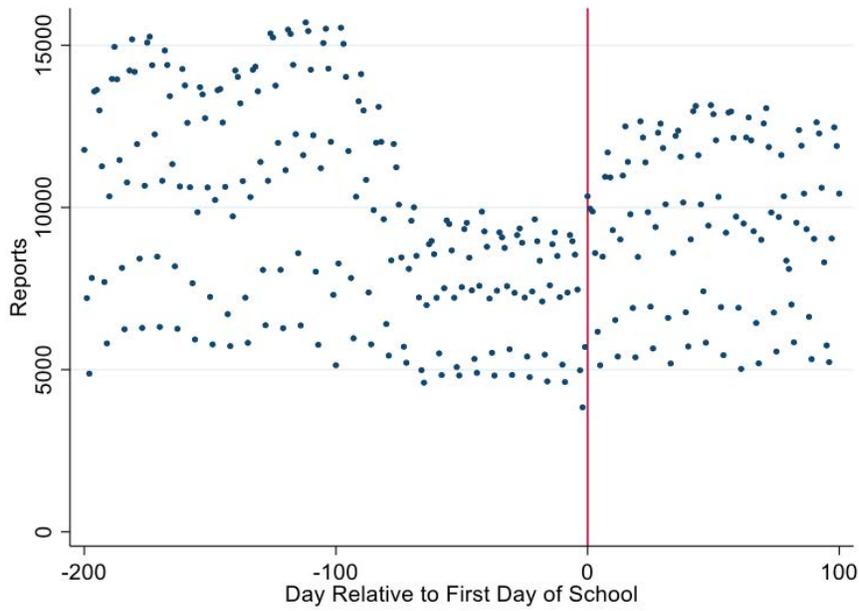
Note: Data is from restricted-use versions of the National Child Abuse and Neglect Data System and include information reported between 2003 and 2015. Figures present estimated coefficients (dots) and 95% confidence intervals (bars) of the difference in reports to CPS agencies for children of a give age (indicated on the horizontal axis) between children eligible for school at age five relative to those not eligible until age six. The estimates are from local linear regression discontinuity specification.

Figure 4. Reports of Child Maltreatment, 2003 to 2015, by Day of the Year

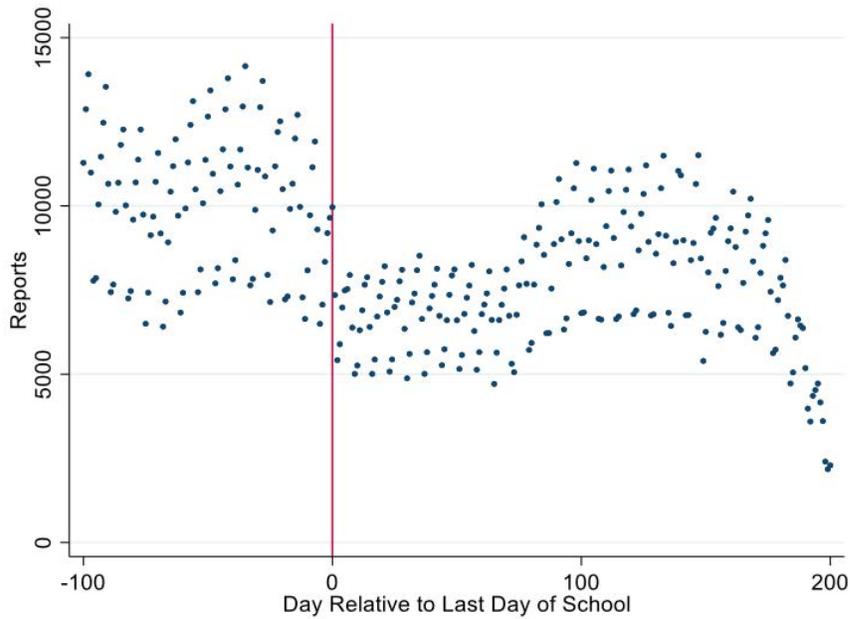


Note: Data is from restricted-use versions of the National Child Abuse and Neglect Data System and include information reported between 2003 and 2015.

Figure 5. Reports of Child Maltreatment, 2005 to 2015, by Day of the Year Relative to the Beginning or End of School, for 25 Districts
Panel A. School Start Date



Panel B. School End Date



Note: Data is from restricted-use versions of the National Child Abuse and Neglect Data System and include information reported between 2005 and 2015. Only information on reports from 25 districts (30 counties) are included in the figures.

Table 1. Summary Statistics, National Child Abuse and Neglect Data System, All States

Variable	Mean	Std. Dev.
Fraction Male	0.49	0.50
Fraction Black	0.25	0.43
Age at Report	7.56	5.07
Perpetrator is Own Parent	0.91	0.28
Physical Abuse	0.18	0.39
Neglect	0.51	0.50
Reporter is an Education Professional	0.16	0.37
Report is Substantiated	0.25	0.43
Report is Unsubstantiated	0.62	0.48
Child is Removed from the Home	0.06	0.24
Number of Observations	47,877,529	
Number of Observations with Information about Perpetrator	9,260,128	

Notes: Data is from restricted-use versions of the National Data Archive on Child Abuse and Neglect data and include information reported between 2003 and 2015.

Table 2. School Entry Eligibility Cutoff Dates Across States and Over Time, 2002 to 2015

State	2015 Policy		Previous Policy	
	Cutoff Date	State Legislation Code	Cutoff Date	Year Changed
Alabama	September 1	AL Code §16-28-4(b)		
Alaska	September 1	AK Stat §14.03.080(d)	August 15	2010
Arizona	September 1	AZ Rev Stat §15-821(c)		
Arkansas	August 1	AR Code §6-18-207(a)	September 15 ²⁶	2009
California	September 1	CA Educ Code §48000(a)	December 2 ²⁷	2011
Colorado	LEA	CO Rev Stat §22-1-115		
Connecticut	January 1	CT General Stat Sec §10-15c(a)		
Delaware	August 31	DE Code §14-27-02		
District of Columbia	September 30		December 31	2011
Florida	September 1	FL Stat §1003.21		
Georgia	September 1	GA Code §20-2-150		
Hawaii	July 31	HI Stat §302A-411	August 31	2009
Idaho	September 1	ID Code §33-201		
Illinois	September 1	IL Compiled Stat §105-5-26		
Indiana	August 1	IN Code §20-33-2-7		
Iowa	September 15	IA Code §282.3 (b)		
Kansas	August 31	KS Stat §72-1107(c)		
Kentucky	October 1	KY Stat §158.030		
Louisiana	September 30	LA Rev Stat §17:222(a)		
Maine	October 15	ME Rev. Stat Title 20-A §5201		
Maryland	September 1	MD Reg 13A.08.01.02 (b)	December 31 ²⁸	2002
Massachusetts	LEA	M.G.L. 603 CMR §8.02		
Michigan	September 1	M.C.L. §380.1147	December 1 ²⁹	2012
Minnesota	September 1	MN Stat §124D.02		
Mississippi	September 1	MS Code §37-15-9		
Missouri	August 1	MO Rev Stat §160.053.1		
Montana	September 10	MT Code §20-7-117		
Nebraska	July 31	NE Rev Stat §79-214		
Nevada	September 30	NV Rev Stat §392.040	October 15	2011
New Hampshire	LEA	Not specified in statute	September 30	2004
New Jersey	LEA	NJ Rev Stat §18A:44-2		
New Mexico	September 1	NM Stat §22-13-3 (d)		
New York	LEA	NY Educ L §1712	December 1	2000
North Carolina	August 31	NC Gen Stat §115C-364	October 16	2008
North Dakota	August 1	ND Cent Code §15.1-06-01	September 1	2010
Ohio	LEA	OH Rev Code §3321.01	September 30	2001
Oklahoma	September 1	OK Stat §70-18-108		
Oregon	September 1	ORS §336.092		
Pennsylvania	LEA		February 1	2003
Rhode Island	September 1	RI Gen Laws §16-2-27	December 31	2003
South Carolina	September 1	SC Code §59-63-20		
South Dakota	September 1	SD Code §13-28-2		

²⁶ Arkansas phased in the August 1 cutoff date using September 1 in 2010 and August 15 in 2011.

²⁷ California phased in the September 1 cutoff date using November 1 in 2012 and October 1 in 2013.

²⁸ Maryland phased in the September 1 cutoff date using November 30 in 2003, October 30 in 2004, and September 30 in 2005.

²⁹ Michigan phased in the September 1 cutoff date using November 1 in 2013 and October 1 in 2014.

Tennessee	August 15	TN Code §49-6-201	September 30 ³⁰	2012
Texas	September 1	TX Educ Code §29.151		
Utah	September 2	UT Code §53A-3-402(6)		
Vermont	LEA	16 VSA §1073		
Virginia	September 30	VA Code §22.1-199		
Washington	August 31		LEA	2005
West Virginia	September 1	WV Code §18-5-18		
Wisconsin	September 1	WI Stat §118.14		
Wyoming	September 15	WY Stat §21-4-302		

Note: Note: School entry cutoff legislation dates were collected from published reports by the *Education Commission of the States* in 2010, 2011, and 2014. Using specific legislative codes reported in the 2010 publication, we corroborated each state's cutoff date and documented more recent legislative changes, many of which were reported in the 2014 publication. The 2011 publication provided historical cutoff dates for each state in 1990 and 2005. These dates were verified using the state statutes and compared to the dates reported in Appendix 1 of Bedard and Dhuey (2007). When dates conflicted among sources, we reported the date recorded by state statute.

³⁰ Tennessee phased in the August 15 cutoff date by using August 31 in 2013.

Table 3. Estimates of the Increase in the Number of Reports to CPS Agencies at Age 5 for Children Eligible for School at Age 5 (Relative to Those Eligible at Age 6)

	(1)	(2)	(3)
	Reports by All Sources	Reports by Educators	Reports by Other Sources
<i>Local nonparametric regressions</i>			
Local linear using data-driven bandwidth	339***	357***	-8
	(86)	(30)	(70)
Data-driven bandwidth	43	43	39
Local quadratic using data-driven bandwidth	489***	380***	77
	(79)	(40)	(68)
Data-driven bandwidth	43	49	47
Local cubic using data-driven bandwidth	582***	401***	169*
	(105)	(47)	(88)
Data-driven bandwidth	52	62	55
Average Number of Reports per Relative Day	6,311	979	5,345

Notes: * denotes $p < 0.10$, ** denotes $p < 0.05$, and *** denotes $p < 0.01$. Data are from restricted-use versions of the NCANDS and include information reported between 2003 and 2015. The nonparametric regressions are estimated using the “robust data-driven” procedures of Calonico, Cattaneo and Titiunik (2014a; 2014b). We use a triangular kernel, robust standard errors clustered on the running variable, and their bandwidth selection and bias correction procedures.

Table 4. Estimates of the Change in Reporting to CPS Agencies at Age 5 for Children Eligible for School at Age 5 (Relative to those Eligible at Age 6), by Type of Abuse and Type of Reporter

	(1) Reports by All Sources	(2) Reports by Educators	(3) Reports by Other Sources
All Reports	339*** (86) 43 6311	357*** (30) 43 979	-8 (70) 39 5345
Reports of Neglect	132*** (49) 43 3245	124*** (10) 38 330	18 (42) 38 2925
Reports of Physical Abuse	161*** (19) 48 1095	176*** (11) 48 312	-15 (14) 47 783

Notes: * denotes $p < 0.10$, ** denotes $p < 0.05$, and *** denotes $p < 0.01$. Data are from restricted-use versions of the NCANDS and include information reported between 2003 and 2015. The nonparametric regressions are estimated using the “robust data-driven” procedures of Calonico, Cattaneo and Titiunik (2014a; 2014b). We use a linear polynomial, triangular kernel, robust standard errors clustered on the running variable, and their bandwidth selection and bias correction procedures.

Table 5. Estimates of the Increase in the Number of Reports to CPS Agencies at Age 5 for Children Eligible for School at Age 5 (Relative to Those Eligible at Age 6) that Were the Child's First Reported Case of Maltreatment

	(1)	(2)	(3)
	Reports by All Sources	Reports by Educators	Reports by Other Sources
All Reports	143*** (41) 52 3,233	199*** (13) 42 559	-51 (39) 49 2,674
Reports of Neglect	28 (26) 52 1572	52*** (7) 32 163	-29 (24) 52 1409
Reports of Physical Abuse	89*** (13) 45 608	106*** (7) 46 200	-17* (10) 47 408

Notes: * denotes $p < 0.10$, ** denotes $p < 0.05$, and *** denotes $p < 0.01$. Data are from restricted-use versions of the NCANDS and include information reported between 2003 and 2015. The nonparametric regressions are estimated using the “robust data-driven” procedures of Calonico, Cattaneo and Titiunik (2014a; 2014b). We use a triangular kernel, robust standard errors clustered on the running variable, and their bandwidth selection and bias correction procedures.

Table 6. Estimates of the Increase in the Number of Reports to CPS Agencies for Children ages 6 to 17 at the Beginning and End of the School Year in 25 Districts

	(1)	(2)	(3)	(4)	(5)	(6)
	Start of School Year			End of School Year		
	Reports by All Sources	Reports by Educators	Reports by Other Sources	Reports by All Sources	Reports by Educators	Reports by Other Sources
<i>Local nonparametric regressions</i>						
Linear using data-driven bandwidth	7.759*** (1.380)	3.155*** (0.314)	4.885*** (1.413)	-5.144*** (1.981)	-2.481*** (0.609)	-1.307 (1.163)
Data-driven bandwidth	44	18	31	29	15	55
Quadratic using data-driven bandwidth	8.797*** (1.865)	2.927*** (0.331)	5.988*** (1.656)	-4.341** (2.119)	-1.709** (0.691)	-1.026 (1.429)
Data-driven bandwidth	49	37	48	55	26	78
Cubic using data-driven bandwidth	10.29*** (2.228)	3.519*** (0.396)	7.477*** (1.974)	-3.986* (2.249)	-2.536*** (0.603)	-0.942 (1.617)
Data-driven bandwidth	59	38	59	85	62	107
Average Number of Reports per Relative Day in the Summer	16	0.51	15	16	0.53	15

Notes: * denotes $p < 0.10$, ** denotes $p < 0.05$, and *** denotes $p < 0.01$. Data are from restricted-use versions of the NCANDS and include information reported between 2003 and 2015. The nonparametric regressions are estimated using the “robust data-driven” procedures of Calonico, Cattaneo and Titiunik (2014a; 2014b). We use a triangular kernel, robust standard errors clustered on the running variable, and their bandwidth selection and bias correction procedures.

Table 7. Estimates of the Increase in the Number of Reports to CPS Agencies at the Beginning and End of the School Year in 25 Districts, by Type of Maltreatment

	(1)	(2)	(3)	(4)	(5)	(6)
	Start of School Year			End of School Year		
	Reports by All Sources	Reports by Educators	Reports by Other Sources	Reports by All Sources	Reports by Educators	Reports by Other Sources
All Reports	7.759*** (1.380) 44 16	3.155*** (0.314) 18 0.51	4.885*** (1.413) 39 15	-5.144*** (1.981) 29 16	-2.481*** (0.609) 15 0.53	-1.307 (1.163) 55 15
Reports of Neglect	2.607*** (0.490) 54 8	0.863*** (0.0867) 19 0.14	2.012*** (0.551) 35 7	-1.753** (0.701) 34 8	-0.834*** (0.186) 18 0.15	-0.440 (0.466) 56 8
Reports of Physical Abuse	1.670*** (0.235) 29 2	0.914*** (0.0925) 17 0.12	0.689*** (0.176) 30 2	-1.025*** (0.343) 46 2	-0.581*** (0.149) 45 0.12	-0.443** (0.187) 47 2

Notes: * denotes $p < 0.10$, ** denotes $p < 0.05$, and *** denotes $p < 0.01$. Data are from restricted-use versions of the NCANDS and include information reported between 2003 and 2015. The nonparametric regressions are estimated using the “robust data-driven” procedures of Calonico, Cattaneo and Titiunik (2014a; 2014b). We use a linear polynomial, triangular kernel, robust standard errors clustered on the running variable, and their bandwidth selection and bias correction procedures.