Disadvantage Beyond Poverty:
Adverse Childhood Experiences, School Choice, and Educational Outcomes

Leah R. Clark

January 15, 2019

Updated version: https://www.leahrclark.com/research

Abstract: Adverse home and family circumstances play an important role in shaping children's educational outcomes, but typical administrative data cannot identify these disadvantages beyond a simple household income proxy (e.g., free or reduced lunch eligibility). By linking human service records to school data, this paper identifies a subset of more severely disadvantaged children: those linked to child welfare investigations prior to kindergarten (one-in-eight children in the city under study). Regardless of which elementary school they attend, these children miss 25% more school days than non-disadvantaged children and have suspension odds that are 57% higher. These gaps are significantly larger than those for low-income children not linked to investigations. Moreover, a multinomial choice model of kindergarten enrollment shows that children linked to child welfare investigations are systematically less likely to enroll in charter or magnet schools over traditional public schools, in contrast to other low-income children who are nearly as likely as their counterparts to enroll in charter or magnet schools. Thus, in aggregate at the school level, a handful of traditional public schools disproportionately enroll the most disadvantaged students. By carefully controlling for these sorting patterns and leveraging measures of early childhood disadvantage determined prior to school entry, I recover causal peer effect estimates, which show that having more disadvantaged peers significantly increases students' own suspension probabilities.

I. Introduction

Economic disadvantage has far-reaching implications for parental investments in children's human capital development and children's own educational outcomes (Cunha & Heckman 2007; Dahl & Lochner 2012).[1] Of course, not all low-income children and families are the same, and poverty may not be the direct cause of low parental investments or adverse outcomes. For some, poverty comes alongside physical, mental, and environmental health issues that worsen educational outcomes, even (or perhaps especially) when exposure comes very early in life (Case et al. 2004; Aizer et al. 2009; Aizer et al. 2018). Similarly, child maltreatment – another correlate of poverty and parental investments – is connected to worse educational outcomes via survey data (Pieterse 2014; Slade & Wissow 2007). While rich administrative school data are available to researchers, they typically contain few measures of children's family and early childhood experiences, limiting the ability of researchers to understand how the nuances of disadvantage factor into educational outcomes.

This paper leverages a unique dataset that links information on early childhood human services exposure to data on school enrollments and behavioral outcomes. These data enable analyses of the educational experiences of children linked to child welfare investigations early in life (12% of kindergartners in the city under study) relative to those who were low-income early in life but never linked to child welfare (52% of kindergartners). These analyses produce several new findings on the relationship between adverse early childhood experiences and educational outcomes. (For brevity, I will refer to students linked to child welfare investigations prior to kindergarten as Early Child Welfare students). First, Early Child Welfare students are less than one-half as likely to enroll in choice schools (in lieu of neighborhood public schools) relative to children without child welfare ties. This stands in contrast to low-income children, who are only slightly less likely to enroll in magnet and charter schools than their non-low-income counterparts. Second, while low-income students have significantly higher rates of absence and suspension than non-low-income students, these disparities are even larger for Early Child Welfare students. Regardless of which school they attend, Early Child Welfare students miss 25% more school days and are 57% more likely to be suspended than other students.

[1] I will use "parent" or "parents" to refer to the adult caretakers that make decisions regarding children's educational experiences. I will also use the singular pronoun "their/them."

Moreover, I find evidence of significant adverse peer effects on suspension. Students who have more disadvantaged peers – whether Early Child Welfare, low-income, or both – are themselves more likely to be suspended. Since Early Child Welfare students are disproportionately concentrated in a subset of traditional public schools (and not in magnet or charter schools), so, too, are these adverse spillover effects.

Why should early child welfare investigations provide insight into later educational outcomes? The most literal interpretation of being an Early Child Welfare student is that, at some point before kindergarten, someone made a phone call to a hotline to report concerns of abuse or neglect concerning the student, or another child who lives with them. Parental risk factors for child maltreatment include drug and alcohol abuse, intimate partner violence, and criminal justice involvement (Doyle & Aizer 2018). Potentially, the parents of Early Child Welfare students make lower investments, on average, in their child's wellbeing in early childhood, and this correlates with later disparities in educational investments. This interpretation is consistent with the findings of this paper: parents of children linked to child welfare may be less likely to invest the time and effort to get their children enrolled in choice schools and less likely to facilitate regular school attendance for their children. These disparities do not necessarily reflect mal-intent: they could also be attributable to differences in parental resources (financial resources but also time and health) or preferences and savvy with respect to education.

From the child's perspective, child welfare investigations may indicate exposure to abuse, neglect, or living in a resource-poor family, among other factors. Research on later-in-life outcomes for maltreated children find decreased earnings, more criminal activity, and a continuing cycle of violence (Currie & Spatz Widom 2010; Currie & Tekin 2012). Perhaps, then, it is unsurprising that Early Child Welfare students face significantly higher odds of suspension, regardless of which school they attend. Hopefully, child welfare investigations are effective in improving children's circumstances, but the experience of the investigation itself may also have an adverse impact. Approximately one-in-five investigations result in the child being removed from the home (potentially to a relative's house) for at least one day.

Critically for the causal framework employed in this paper, Early Child Welfare can be measured prior to starting school. Children are linked to child welfare investigations as early as birth. This is true of other human services variables, as well, enabling the use of early childhood Medicaid enrollments as a proxy for living in a low-income household prior to school entry. Most children who use Medicaid enroll in the year of their birth, and take-up among young children is high. Since enrollment is means-tested, this provides a reliable proxy for whether children live in low-income families in early childhood (Kenney et al. 2012). (For brevity, I will refer to students who were enrolled in Medicaid prior to kindergarten as Early Medicaid students). In contrast to Early Child Welfare and Early Medicaid, variables commonly available in school data – like indicators for whether a student receives free/reduced-price lunch or special education services – are not necessarily indicative of early childhood factors and may be influenced by policies and practices that vary across schools and over time (Domina et al. 2018; Michelmore & Dynarski 2017).

Whether disadvantaged children have equal access to choice schools is a major concern for education reform efforts that center around expanding school choice options in order to improve access to quality education for children living in low-performing school districts (Epple & Romano 1998). To investigate whether Early Child Welfare and Early Medicaid students are equally as likely to enroll in magnet, charter, or private schools (in lieu of traditional public schools) as their counterparts, I implement a multinomial choice model of the type of school students first enroll in for kindergarten as a function of Early Child Welfare and Early Medicaid (in addition to gender and race controls). To account for potential neighborhood effects that operate prior to school enrollment, the analysis includes zip code fixed effects, which offer reasonable approximations of neighborhoods in the city under study. Interestingly, zip code fixed effects have little-to-no impact on the estimates of interest, suggesting either Early Child Welfare and Early Medicaid operate orthogonally to zip code effects, or they explain a substantial portion of zip code effects, at least as they pertain to educational enrollments.

The findings of the selection model – that disadvantaged children are systematically less likely to enroll in alternatives to traditional public school – are new to the school selection literature, but do not refute current wisdom. Conditional on income, private school voucher decliners are more likely to be single mothers working full-time, suggesting that non-income family resources (like parental time) play an important role in accessing choice options (Cowen, 2010). Studies on the role of race in school choice selection suggest that choice schools may increase racial segregation by disproportionately enrolling students of a particular race relative to local district demographics (Bifulco & Ladd, 2006; Garcia, 2008). My model indicates that it is white students in this district who are disproportionately enrolling in charter schools. While the analysis does not consider special needs students directly, to the extent that Early Child Welfare students are more likely to receive special education services in school, their enrollment disparity is also in line with ongoing concerns that choice schools do not serve special needs students proportionally (Cowen & Winters, 2013).

These findings raise questions about whether the disparities in educational access of Early Child Welfare students extend to school outcomes. And, if their outcomes are worse, does their presence in the classroom cause other students to have worse outcomes, as well? In the second major analysis of this paper, I use data on elementary school attendance and suspension outcomes to address these questions. This paper is not the first to use administrative measures of child welfare investigations as a risk factor pertinent to educational outcomes, but those analyses use measures that are potentially endogenous to schools and are, thus, limited in their causal claims (Fantuzzo et al. 2014; Ryan et al. 2018).

Absences and suspensions are of particular interest in this context for several reasons. First, they are understudied in the economics literature, but play a critical role in student learning (Gershenson et al. 2017). Second, comprehensive human services histories are primarily available for younger students, many of whom are not yet taking comprehensive standardized exams (which start in third grade). Finally, these outcomes offer a useful contrast. For elementary school students, attendance is primarily a parent’s choice, while suspension reflects a student’s behavior in combination with how the school decided to respond to that behavior. These variables are not simply outcomes, but also measures of inputs to the education production function (Lazear, 2001). Thus, disparities in attendance provide insight into the challenges that schools face in effectively educating students.

Identifying the impact of Early Child Welfare and Early Medicaid on outcomes requires separating multiple pathways through which the effects could operate (directly and through peers) and multiple ways in which these pathways could be confounded (school policies and sorting). Fortunately, having pre-determined measures of children and their peers ensure that the independent variables of interest cannot be influenced by educational outcomes. This enables an approach like that of Carrell & Hoekstra (2010), who study the impact of domestic violence exposure (as proxied by protective orders) on educational outcomes. As with Early Child Welfare and Early Medicaid, students have no influence over whether their peers are exposed to domestic violence, so it offers a source of exogenous variation in peer group composition across schools and grades, after accounting for the challenge to identification stemming from nonrandom selection into schools.

Since this paper provides evidence of selection into schools along dimensions of early disadvantage, it is reasonable to assume there are other factors influencing school selection that are not captured in the data. Such unobserved variables create a challenge for identification: if students' outcomes at a certain school are correlated, does that represent the school's impact, or does that reflect the fact that students (or parents) at the school are similar? Thus, identifying the causal effect of early childhood human services exposure on absences and suspensions requires carefully controlling for sorting. This is done by employing a number of fixed effects. Ultimately, the identifying variation in the outcomes regressions occurs within school-grade cohorts over time, and separately from district- or school-level trends.

Analyses will repeatedly return to the question of whether a simple control for poverty can account for heterogeneity in the experiences of disadvantaged children. Throughout the analyses, I find evidence that the Early Child Welfare variable identifies children who experience worse outcomes, on average, than those who are "just" Early Medicaid. Moreover, Early Child Welfare students exert a unique impact on their peers. Section II describes the data. Section III explains the methodology and results of the school type selection model; section IV covers the educational outcomes model. Section V concludes.

II. Data

The data stem from a partnership between a midsize urban school district and the human services department that serves the city and the surrounding county. Child welfare and Medicaid histories, school enrollments, and attendance and suspension outcomes are observed for every student who enrolls in kindergarten in the district's traditional public schools or magnet schools from the 2012-13 school year through the 2016-17 school year.[2] From the 2013-14 school year forward, human service histories and enrollments are additionally observed for children who live in the city but enroll outside of the district – primarily in charter or private schools. However, attendance and suspension outcomes are never observed for non-district schools.

[Figure 1]

[2] The 2015-16 school year is excluded from the main analyses due to a censoring issue in the enrollment data for that year.

The selection and outcomes models (described in sections III and IV, respectively) rely on different but overlapping cuts of the dataset. Figure 1 offers a comparison of the two subsamples. The reasons for the data restrictions are explained in the methodology descriptions of each model, but, in short, the selection model utilizes kindergarten enrollments from 2013-14 on (when the universe of schooling options is observed), while the outcomes model makes use of data for all elementary grades from 2012-13 forward, but is limited to traditional public and magnet schools. Thus, to understand what demographics are typical of the city's children in recent years, one should focus on the selection model statistics.

[Figure 2]

Approximately 2,800 children in the city enroll in kindergarten for the first time each year (see Figure 2). 56% enroll in traditional public schools, 18% in the district's magnet schools, 9% in non-district-affiliated public charter schools, and 15% in private schools. A small number of children enroll in kindergarten in day cares or special needs schools (e.g., schools for the blind). Very few children are enrolled in home school for kindergarten, but it is possible that some home-schooled children are left out of the data system. On average, 45% of enrolling kindergartners are identified as black, 39% white, and 9% multi-racial, with the remaining 7% identified as Asian, Hispanic, or Native American. There are not many English Language Learners in the district.

Human service variables

Human service histories in the dataset are comprised of a set of indicators summarizing whether a child interacted with a given service area during a given school year – even if the child was not yet enrolled in school. Consider a child enrolling in kindergarten for the first time in Fall 2013 after their 5th birthday: the data indicate which services they were associated with in 2012-13 (approximately age 4), 2011-12 (approximately age 3), and so on. I will use "early childhood" to denote the years between birth and beginning kindergarten (typically between ages 5 and 6). Early childhood human service histories are missing only for children who did not live inside the county prior to enrolling in a city school. Unfortunately, the data do not differentiate between students who never used human services and students who lived outside of the county.

This paper focuses on two aspects of student's early childhood human service histories: child welfare investigations and Medicaid enrollments. "Early Child Welfare" or "ECW" and "Early Medicaid" or "EM" will be used to denote indicator variables for any involvement in the respective service in early childhood. Thus, an "ECW student" is a student who was linked to at least one child welfare investigation prior to starting kindergarten, and an "EM student" is a student who was enrolled in Medicaid at least once prior to starting kindergarten.

Medicaid is a means-tested program, with high take-up among children. Like free/reduced lunch eligibility, it is often used as a proxy for children living in low-income households. As with any voluntary program, concerns about selection into the program are reasonable. The population of Early Medicaid students, however, is fairly consistent as children age, with the vast majority enrolled in Medicaid every year between birth and kindergarten.

Most ECW students, on the other hand, are only associated with a child welfare investigation during one or two years (the investigations themselves are typically much shorter than a year). Involvement begins at any age, though more than half of ECW students are linked to their first investigation before their second birthday. Child welfare investigations typically arise after someone contacts a hotline to report suspicions of abuse or neglect concerning a particular child. Children who live with children referred for investigation are also referred. Many professionals who work with children are mandated reporters – that is, the law requires them to report suspicions of abuse or neglect within a strict time window. Mandated reporters include doctors and child care workers.

Child welfare investigations sometimes result in a child being removed from the home and placed in foster care, often with a relative. Approximately one-in-five ECW students are placed in foster care for some period of time in early childhood. Other ECW students receive supports or services to ensure their well-being while remaining in their homes.

Participation in child welfare is involuntarily, so there is no concern that ECW reflects selection on the part of parents or children. However, it is possible that some children who should have been referred to child welfare in early childhood were not. It is also possible that the referring party is biased in their choice to make a report. The latter concern has led to a number of studies investigating why black children are referred to child welfare at much higher rates than white children. These studies are not able to rule out racial bias, but they do present some potential alternative explanations. For example, Maloney et al. (2017) use birth records linked to administrative data on child welfare referrals and find that differences in marital status and maternal age fully explain the racial gap in referrals. Again, this cannot rule out bias, but it makes sense that resource limitations or limited parental experience might give rise to circumstances that prompt reports to child welfare. Conditional on referral to child welfare, black and white children are equally likely to be placed in foster care.

In sum, ECW students may have encountered a host of negative experiences before starting kindergarten, but there is no obvious way to model what this indicator should mean in terms of human capital formation. These children may also have been impacted – positively or negatively – by well-intentioned services and providers, including foster care placements. The simplest interpretation of the indicator is that, at some point early in a child's life, someone expressed concern about their well-being. Despite the variable's murky implications, ECW provides a useful signal about early childhood experiences.

School variables

School data is compiled at the student-enrollment-year level; that is, for each student and each school in which they enroll, I observe the school, grade, enrollment/withdrawal dates, and number of days absent or suspended. Additionally, I observe some demographic information, like race, gender, age at start of each year, and zip code of residence. Indicators of special need and free/reduced lunch eligibility are also in the data, but will not be used in subsequent analyses, since they may change over time and in response to school practices.

The dependent variables in the outcomes model are absence rates and suspensions. Absence rates are calculated as days missed – net of suspension days – as a share of total days enrolled.[3] Most children are not suspended, and those that are typically serve just one day of suspension, so the suspension outcome is simply an indicator for whether a child was suspended during a particular enrollment-year. Enrollments are excluded if they (a) last 50% of the school year or less, or (b) exhibit an absence rate of 75% or more. The latter exclusion reflects the concern that these records reflect administrative errors rather than real school enrollments.

[Figure 3]

Figure 3 shows average absence and suspension outcomes by grade for students with various human service histories. These data come from the outcomes model dataset, so 14% of students have both Early Child Welfare and Early Medicaid exposure; 2% of students have Early

[3] I do not differentiate between unexcused and excused absences out of a concern that the distinction is endogenous. For the same number of missed days, low-income children have fewer excused absences than do non-low-income children.

Child Welfare exposure without Early Medicaid exposure; 52% of students have Early Medicaid exposure without Early Child Welfare exposure; and 32% of students have neither. This latter group has the best average outcomes, while the "EM only" group has somewhat worse outcomes, on average, and the two ECW groups have the worst average outcomes of all.

The most striking observation from these figures is how closely the "ECW only" lines track the "ECW & EM" lines. This suggests that the Early Child Welfare measure – not the Early Medicaid measure – is more strongly associated with adverse outcomes. Of course, the outcomes model will explore this contention more rigorously.

Overall, absence rates decline slightly as students ascend grades. Parents may take school attendance more seriously as students get older. Suspension rates increase substantially as students ascend grades, though approximately 5% of kindergartners receive suspensions – enough to include kindergartners in the suspension outcomes model. The outcomes model will include district-level grade-by-year fixed effects to separate out these general trends from variation attributable to particular schools, cohorts, or students.

III. Selection into schools

School selection can occur any time during a child's education, but most children enroll in kindergarten in a particular elementary school and stay in that school through fifth grade. Since the human services data permit characterization of children prior to school entry, I will focus on modeling school selection at the point of first entry into elementary school. It is possible that parents will change their mind about their child's school and switch them into a different school type later on. However, later school selection could occur based on specific school experiences unobserved to the econometrician (e.g., a child being bullied), while initial school selection more likely reflects general expectations for a school, and those expectations may be shared across many parents.

What does a child's first kindergarten enrollment say about their parents' preferences? Choice schools (including private schools) likely require more advance planning – oversubscribed choice options typically require applications[4] more than six months prior to the start of school – but some will admit students through the start of the school year. Enrollment in a neighborhood

[4] To my knowledge, no schools administer aptitude tests for kindergarten entry, though private schools may consider attributes that charter/magnet schools are not allowed to consider.

public school could reflect a passive choice (parents did not consider alternatives), an active choice (parents searched for alternatives and concluded it was the best choice), or a temporary choice (parents are still searching for alternatives, but want their child to start school on time). Thus enrollment does not always reflect parents' preferences for schools, but for parents making the passive choice, enrollment in traditional public schools is practically guaranteed.

So if some degree of heterogeneity among parents is captured by which schools they select for their children in kindergarten, how does this heterogeneity extend to children? Is there any evidence that children enrolled in traditional public schools differ from those enrolled in choice schools? This section puts forth a multinomial logistic choice model of school type for first kindergarten enrollment on the basis of characteristics plausibly exogenous to realized school type.

III.A. Methodology

To sidestep the endogeneity issues facing a model of school switching, the kindergarten selection model will consider only the first observed kindergarten enrollment, regardless of whether the child subsequently remained at the school for six days or six years. For students who repeat kindergarten, only their first enrollment from their first year of kindergarten is included.

Schools included in the model fall into four broad categories: traditional public (reference category, $j = 0$), public magnet ($j = 1$), public charter ($j = 2$), and private (religious or secular, $j = 3$). Using a multinomial logistic regression, probability of selecting school type j is

$$\pi_{ij} = \frac{exp\{\eta_{ij}\}}{\sum_{k=1}^{J} exp\{\eta_{ik}\}}$$

where $\eta_{ij} = \beta_0 + \alpha_j ECW_i + \gamma_j EM_i + X_i'\beta_j + \lambda_{zip}$. The vector X_i includes indicators for gender and race.

Some alternative school types are excluded. First, I exclude schools specifically targeted to special needs populations (e.g., schools for the blind) on the assumption that parents who select these schools are driven by a different set of concerns than parents who enroll their children in regular schools. Special needs schools include public and private options. Second, cyber schools (public and charter) and home schooling were excluded. Enrollment in these options for kindergarten is quite low, and it is possible that these parents were never inclined to pick a regular schooling option to begin with.

Multinomial logistic choice models assume independence from irrelevant alternatives – that is, eliminating a choice parents did not select would not cause them to change their selection. As noted previously, it is possible that parents observed in the data are making a temporary choice, that is, they are waiting for a spot in their preferred school to open (Engberg et al. 2014). Suppose a child is temporarily enrolled in traditional public kindergarten while their parent tries to get them a spot in a charter school, but, if that charter school did not exist, the parent would have already enrolled them in a private school. The cautious distinction to make here is that a model of selection for first kindergarten enrollment is not necessarily a model of long-term school enrollment choices.

III.B. Results

[Figure 4]

Figure 4 shows the selection model results. Coefficients are risk ratios – that is, if they are larger than one then an increase in the corresponding independent variable predicts increased likelihood of choosing the relevant school type over traditional public school; if it is smaller than one it predicts decreased relative likelihood. Panels 2 and 3 include residential zip code fixed effects, while Panel 1 does not. Standard errors for all specifications are clustered at the zip code level. The fact that including zip code fixed effects has little impact on the estimates for ECW and EM suggests that location – while an important factor in school selection broadly – does not explain the relationship between early human services exposure and school choice.

Early Child Welfare is associated with substantially lower probability of enrollment in magnet, charter, or private schools. In fact, besides the “other minority” indicator, it is the only variable that predicts lower take-up of choice schools in general. Early Medicaid students are slightly less likely to enroll in magnet and charter schools, though the coefficients are not statistically significant in the preferred specification (Panel 2). Unsurprisingly, EM predicts substantially lower enrollment in private schools, where most students have to pay tuition.

Panel 3 estimates the same model as panel 2 except it excludes the Early Child Welfare variable. This exclusion causes the relative probability estimates of Early Medicaid students enrolling in choice schools to decline slightly, but their confidence intervals still overlap with the estimates in Panel 2. Thus, a simple low-income proxy cannot identify the subset of disadvantaged students unlikely to access choice schools.

[Figure 5]

Selection patterns of ECW students show up in the aggregate at the school level. Figure 5 plots the three-year weighted average kindergarten demographic profile for each traditional public, magnet, and charter school in the city. Magnet and charter schools have similar income compositions to traditional public schools, with the vast majority enrolling a disproportionate share (more than 63%) of Early Medicaid students. The ECW dimension, however, exposes substantial heterogeneity in the composition of magnet and charter schools compared to traditional public schools. Magnet and charter schools enroll far fewer ECW students in proportional terms than do traditional public schools. In fact, only one choice school enrolls a disproportionate share of ECW students, while more than half of traditional public schools do. Private schools (omitted from the figure for visual clarity) universally enroll kindergarten cohorts that are less disadvantaged (on both dimensions) than the city's overall demographics.

Figure 5 points to the same conclusion as the selection model: a simple low-income proxy variable cannot reliably identify the schools that serve the most disadvantaged students. To my knowledge, this analysis is the first to find evidence of a systematic disparity in choice school enrollment for disadvantaged children. This disparity may be driven by disadvantaged parents choosing not to choose, and accepting a default traditional public school enrollment without exploring alternatives. I turn next to the question of how ECW students fare in schools, and whether these enrollment disparities carry implications for outcomes via peer effects.

IV. Outcomes model

[Figure 6]

To consider the potential pathways through which ECW and EM might impact educational outcomes, consider Figure 6. The goal is to understand the causes of an educational outcome, Y_{isgt}, for student i enrolled in grade g in school s during period t. The subsequent analysis models elementary school absence and suspension outcomes. The vector X_i denotes all the attributes that contribute to a student's educational outcomes over which a school has no control – including attributes of their family and home environment. Researchers cannot hope to observe all the components of X_i, but examples of X_i observed in the subsequent analysis include ECW, EM, gender and race. For each school-grade cohort in each time period, average peer attributes are constructed by calculating the average attributes of student i's cohort, leaving out student i. These average peer attributes form the vector $\bar{X}_{-i,sgt}$, and peer average outcomes (constructed analogously) are denoted $\bar{Y}_{-i,sgt}$. Finally, S_{sgt} denotes a vector of school attributes over which students have no control. These attributes may vary by grade.

There are five potential effects on outcomes in this framework. The first three are illustrated in figure 6 under the assumption that X_i and S_{sgt} are exogenous. The last two effects will consider violations of those exogeneity assumptions.

1. Direct effect: student i's attributes and those of their family and home have a direct impact on their educational outcomes. Since X_i is exogenous (e.g., permanent characteristics or variables measured prior to school enrollment), there is no possibility for Y_{isgt} to, instead, cause X_i.
2. Peer effect: the average attributes of peers impact student i's outcomes. Since $\bar{X}_{-i,sgt}$ is exogenous, it cannot be influenced by Y_{isgt}. Notably, the peer effect can also operate in an endogenous fashion, in which student i's outcomes are affected by their peer's outcomes (not attributes). Hence, the graph shows pathways for both the exogenous peer effect (the direct pathway from $\bar{X}_{-i,sgt}$ to Y_{isgt}) and the endogenous peer effect (where $\bar{Y}_{-i,sgt}$ impacts and is impacted by Y_{isgt}). As indicated by the double-ended arrow between peer average outcomes and student i's outcomes, this relationship is subject to simultaneity. Thus, the subsequent estimates will identify an overall net peer effect as opposed to separating its endogenous and exogenous components.
3. School effect: school s impacts student i's outcomes directly through its policies and practices, and via the endogenous peer effect pathway by impacting peer outcomes. As long as X_i and S_{sgt} cannot impact each other, peer and direct effects can be separated from the school effect. Just as school attributes may vary by grade, so may school effects, so this can also be conceptualized as the school-grade effect.

The assumption of exogeneity of X_i and S_{sgt} is reasonable on its face – schools cannot change permanent or predetermined attributes of students, and individual students likely cannot change school policy – but comes into question when sorting and endogenous school responses come into the picture.

4. Sorting: parents select schools based, in part, on their children's attributes (which include parents' own attributes). Perhaps parents want their children to attend school with neighbors or family members, or school attributes appeal to certain types of

parents/students. This leads to a correlation between variables determined at the school-grade level (S_{sgt}, $\bar{X}_{-i,sgt}$, and $\bar{Y}_{-i,sgt}$) and X_i, thus conflating peer, school, and direct effects on outcomes by breaking exogeneity of X_i.

5. School response effect: schools change their policies in anticipation of or in response to the composition and/or outcomes of their student body. This breaks the exogeneity of S_{sgt} to both student attributes and outcomes. Ultimately, time-varying school response effects coincide with peer effects, and longer-lasting school response effects coincide with sorting patterns and school effects. This simultaneity problem prevents separate identification of these effects.

It is not reasonable to assume away any of these effects. In fact, one of the motivations for investigating the effect of Early Child Welfare on outcomes is the evidence that it predicts differential sorting into schools. Thus, the empirical strategy will control for sorting, while providing estimates of direct and peer effects, albeit with some portion of the school response effect folded in.

A note on teacher effects: the choice not to give teachers their own pathway in Figure 6 is more an anticipation of data limitations than a statement of theory. However, the data includes only elementary grades, so the school effect can be thought to include the average effect of a given school's teachers on student outcomes over time. It is likely that teachers comprise a substantial component of the school effect on student outcomes, but, ultimately, their impact will be lumped in with that of the school. It is also worth noting that teachers generally do not have direct control over the outcomes being considered: they cannot force parents to bring their children to school, and suspension decisions are typically made by school-level administrators using guidelines set at the district level.

Before moving into the estimation strategy, it is worthwhile to align the effects discussed above with the standard enumeration of challenges to identifying peer effects: simultaneity/the reflection problem, sorting, and correlated shocks (Manski 1993; Moffitt 2001). First, simultaneity is avoided by measuring the peer effect of characteristics that are permanent or determined prior to school enrollment (in reference to the Figure 6, peer effects will be estimated as the net impact of $\bar{X}_{-i,sgt}$ on Y_{isgt}, rather than the effect of $\bar{Y}_{-i,sgt}$ on Y_{isgt}). The sorting problem is clearly central to this analysis and will be addressed in the identification strategy. Finally, correlated shocks would impede analysis if changes in peer demographics correlated with other unobserved shocks.

While the method used will rule out a number of unobserved shocks, I will additionally run falsification tests to examine this possibility.

IV.A. Identification & Methodology

Conceptually, peer effects seem more complicated than direct effects, but their interpretation is quite literal, which makes explaining their identification easier. Thus, I will first describe how peer effects of Early Child Welfare and Early Medicaid are identified, and then explain the identification of their direct effects.

The outcomes model seeks to estimate the impact of peer percent ECW and EM on student outcomes. The approach taken here is similar to Carrell & Hoekstra (2010) in which an extensive set of fixed effects controls for sorting. First, consider a naïve linear specification:

$$Y_{isgt} = \beta_0 + \beta_1 ECW_i + \beta_2 \overline{ECW}_{-i,sgt} + \beta_3 EM_i + \beta_4 \overline{EM}_{-i,sgt} + X_i'\alpha + \bar{X}'_{-i,sgt}\gamma$$
$$+\delta_{zip} + S'_{sgt}\eta + \theta_{gt} + \epsilon_{isgt}$$

where X_i includes permanent or pre-determined student attributes (excluding ECW and EM) for which I can also construct peer averages. Residential zip code fixed effects (δ_{zip}) are included to control for the average effect of neighborhoods on outcomes, as well as underlying sorting attributable to neighborhood (and not school) attributes. School attributes (S_{sgt}) control for school characteristics that might be correlated with both outcomes and sorting: for example, suppose smaller average class size both improved attendance and attracted parents with certain underlying attributes. Grade-year fixed effects capture the fact that students tend to miss fewer days of school but receive more suspensions as they ascend grades (see Figure 3). They also subsume district- or state-level policy changes and annual changes to administrative systems or variable definitions that may systematically impact outcome measures.

Of course, it is difficult to make the case that school attributes (S_{sgt}) and neighborhood fixed effects (δ_{zip}) entirely capture school effects and sorting. For one thing, the vector of school attributes would have to be quite extensive in order to capture all the features parents value. Moreover, as the selection model showed, sorting into schools is not explained wholly by neighborhoods.

Thus, the naïve model needs to be amended to include school fixed effects, which can capture the impact schools have on outcomes due to factors unobserved to the econometrician. School fixed effects also capture the average impact on outcomes of unobserved variables that correlate with sorting. Due to the concern that sorting and school effects might change over time or across grades, school-by-year (α_{st}) and school-by-grade (λ_{sg}) fixed effects will be used in the preferred specification.

$$Y_{isgt} = \beta_0 + \beta_1 ECW_i + \beta_2 \overline{ECW}_{-i,sgt} + \beta_3 EM_i + \beta_4 \overline{EM}_{-i,sgt} + X_i'\alpha + \bar{X}'_{-i,sgt}\gamma$$
$$+\delta_{zip} + S'_{sgt}\eta + \theta_{gt} + \alpha_{st} + \lambda_{sg} + \epsilon_{isgt}$$

The identifying variation in peer ECW, therefore, comes from variation in peer ECW that is not attributable to trends for grades over time (e.g., this year the district has more ECW kindergartners than usual), trends for schools over time (e.g., this year this school is attracting fewer non-ECW students), of permanent differences within schools across grades (e.g., non-ECW students tend to switch out of this school before third grade). What does that leave? Variation within a school-grade over time, demeaned from district- and school-level trends. If this variation is as-good-as-random, then β_2 has a causal interpretation as the peer effect of ECW on Y_{isgt}, with one important caveat: the time-varying school-grade response effect.

Without a detailed accounting of a school's responses to its student body composition or an instrument, the time-varying school response effect cannot be separately identified from the peer effect. This is because the portion of the school response effect that varies within school-grades over time (e.g., an individual teacher's response to their classroom composition) is determined simultaneously with the peer effect. However, for our purposes, this is fine: our peer effect estimate will include the ways in which school-grades react to variation in peer group composition, which certainly belongs in our accounting of how concentrating at-risk students impacts schools.

Ideally, school response effects could also be separated from school direct effects. It would be nice to know if the concentration of various subgroups in schools causes schools to change their policies or practices. However, these shifts are simultaneous with sorting, so the fixed effects will subsume them as they do the school direct effects.

A potential threat to identification of peer effects would be a parental response to peer groups (or contemporaneous sorting), in which parents are more likely to remove their children from a cohort during the year once they realize that cohort is particularly disadvantaged. Appendix C estimates the probability that students exit during the year using the same framework employed for the outcomes analyses. No evidence of contemporaneous sorting is found.

The external validity of the peer effects estimate depends not only on a host of contextual features of the educational system (e.g., school district suspension policies), but also on the manner in which agencies responsible for handling child welfare investigations record referrals. But here is where the scope for concerned parties to take initiative offers a benefit: the set of such concerned individuals and the environments in which they may observe child maltreatment are broad. Plus, mandated reporter laws are similar across states. Similarly, while Medicaid programs vary across states, the population identified by early involvement in Medicaid is likely broadly consistent across states.

Turning now to estimation of the direct effect: with sorting controlled for, there is no possible feedback loop through which school outcomes could influence ECW or EM – they are determined prior to school enrollment. Thus, β_1 identifies the average disparity in outcomes experienced by Early Child Welfare students, irrespective of which school they attend, the demographics of their classmates, and their own demographic characteristics (including EM, gender, race, and residential zip code).

Interpreting β_1 as the causal effect of Early Child Welfare runs into two challenges: Early Child Welfare is difficult to conceive as a treatment variable, and, as with most investigations into the social determinants of outcomes, I cannot rule out omitted variable bias. On the former matter, recall that β_1 is not the causal effect of a child welfare investigation; it is the causal effect of being in the sort of family or environment that prompts someone to report a young child to child welfare. ECW implies, by definition, that someone observed something about a child's welfare that the econometrician does not. One must be careful not to assume that the unobserved component is necessarily abuse or neglect, just as it should not be assumed that the child welfare investigation resolved (or did not resolve) the underlying issue. Potentially, ECW is indicative of a lack of parental savvy or resources alongside its most disturbing potential explanations. Still, ECW captures student attributes that are outside of schools' control, making the estimation of its effect informative to educators, case workers, and policymakers.

Returning to the possibility of omitted variable bias, of particular concern are variables that might impact both ECW and outcomes. For example, if young parents are more likely to be referred to child welfare, but are also (and independently) less concerned with regular school attendance, then some component of β_1 might reflect the effect of having a young parent, conflated with the effect of ECW. Reassuringly, the effects shown in the subsequent section are quite stable to the inclusion of controls – like race and zip code – that are also likely correlated with omitted variables of interest, suggesting that omitted variables may not loom large.

Similarly, the estimate of the direct effect of EM (β_3) requires a nuanced interpretation. It should not be interpreted as the causal effect of Medicaid, but, rather, the causal effect on educational outcomes of being the sort of child who gets enrolled in Medicaid at young ages. Ideally, there would be better measures of family income and resources available to researchers, but the direct effect of EM offers a robust and useful benchmark for the direct effect of ECW.

High take-up rates limit concerns of selection into Medicaid, but it is worth noting that β_3 could be biased by differential selection into Medicaid. If higher-income eligible families are less likely to enroll then β_3 could be biased upward; if lower-income families are less likely to enroll then the bias would more likely be downward. The intent of this paper, however, is to use EM to benchmark the effect of ECW, which, as an involuntary program, does not suffer from selection concerns.

While the absence rate outcome is estimated using linear regression, since the suspension outcome is binary, I use logistic regression with the same set of covariates and fixed effects as in the absence model.[5] Logistic regression with a large number of fixed effects can suffer from the incidental parameters problem, in which the model's estimates are not reliable. The stability of the estimates as fixed effects are added (as will be shown in Figure 8) suggests the incidental parameters problem may not be an issue here, likely because school-grade cohorts are sufficiently large so the model has enough data to estimate the parameters. Since the absence and suspension outcomes can be reformulated as count variables, negative binomial regression offers some potential advantages, so both the absence and suspension models are run with negative binomial specifications in Appendix A, but the results do not appreciably change.

IV.B. Results

[Figure 7]

Figure 7 shows the absence model results, with each column adding more fine-grained fixed effects. Column 2 shows the results from the naïve specification, while Column 4 shows the specification that removes the effects of sorting. The direct effect of ECW on attendance is

[5] A linear probability model is not appropriate here as suspension probabilities are fairly low, so the linear model yields a large number of negative predicted probabilities.

strikingly robust across specifications – absence rates for ECW students are approximately 1.5 percentage points higher than non-ECW students. This translates to a 25% increase over the mean absence rate. The direct effect of EM is also significant and robust across specifications, but is roughly half the magnitude of the ECW effect.

The direct effects of ECW and EM are not additive. Specifically, when an interaction term between the two of them is included in the regression, I cannot reject the null hypothesis that the EM direct effect for ECW students is zero. When ECW is left out of the regression (Column 5), the Early Medicaid direct effect picks up a little of the Early Child Welfare effect. This makes sense, since nearly 90% of ECW students are also EM, but the estimate remains significantly smaller than the direct effect of ECW. This again highlights how adverse circumstances beyond poverty can drive educational disparities more than poverty itself.

No peer effects are significant for absences. This is as one would expect. For young children, parents play a central role in determining whether a child attends school. It seems unlikely that a child's peers would influence the decision of a child's parents, though it is not impossible. The importance of controlling for sorting is evident in the absence specifications. In the naïve specification in Column 2, both ECW and EM are linked with significant adverse peer effects in attendance, but, once sorting into school-grade cohorts is accounted for (see Column 3), the peer effects disappear.

[Figure 8]

Figure 8 shows the results for the suspension model. Coefficients are expressed as odds ratios so the null hypothesis is that they equal one. ECW and EM students have substantially higher odds of receiving suspensions, and, as expected, peer effects on suspension are significant. The direct and peer effect estimates are quite similar across specifications. All else equal, the odds of an ECW student getting suspended are 57% higher than a non-ECW student; and the odds of an EM student getting suspended are 22% higher than a non-EM student. A one percentage point increase in peer percent ECW increases students' own suspension odds by 1.9%. For a one percentage point increase in EM, the peer effect impact is 1.4%. The notable persistence of gender and racial gaps in suspension across specifications will be discussed further at the end of this section.

Once again, ECW captures heterogeneity that would otherwise be lost. When it is excluded, the EM direct effect climbs slightly but not significantly so. More surprisingly, the EM peer effect does not change at all. That is, the influence of ECW students on peers is entirely missed when accounting only for poverty.

When considering the impact of sorting along the Early Child Welfare dimension, it is useful to consider how changing peer percent ECW would impact the typical student. To investigate this, I evaluate predicted suspension probabilities at the means of all the right hand side variables with the exception of peer percent ECW. I find that a typical student in a school-grade cohort that is 10% ECW (approximately the 25th percentile of the peer percent ECW distribution) has a 5.3% probability of being suspended. The same student in a school-grade cohort that is 25% ECW (approximately the 75th percentile of peer percent ECW) has a 6.9% probability of being suspended. The 95% confidence intervals of these predictions do not overlap.

[Figure 9]

It is also worth investigating the possibility that different subgroups of students experience different peer effects. I run a fully interacted version of the model allowing direct and peer effects to differ based on whether students have both disadvantage indicators, just one, or neither. The predicted suspension probabilities of an average student in each of these subgroups versus peer percent ECW is plotted in Figure 9. The fact that each slope is positive and increasing at a similar rate shows that ECW students impact their peers regardless of their peers' own disadvantage. Moreover, the point estimates for non-low-income ECW students are right in line with those of low-income ECW students, suggesting, once again, that the challenges facing these students do not boil down simply to income.

While the previous specifications are robust to sorting across school-grades or school-years, those fixed effects subsume part of the story: parents who pick magnet schools may differ unobservably from parents who do not pick magnet schools in a manner that impacts school outcomes. Is it additionally the case that ECW or EM students in magnet schools differ unobservably from non-magnet ECW or EM students? Or, might they experience differential outcomes in magnet schools?

To inspect this, I ran the preferred specification but interacted an indicator for magnet enrollment with every covariate (including every peer measure). There is no difference in absence rates for ECW or EM students at magnet schools versus traditional public schools. Apparently, in attendance terms, ECW parents do not respond to magnet schools differently, aside from the differences attributable to school sorting. Nor is there any difference in effect of peer percent ECW on absences or suspensions in magnet schools versus traditional public schools. However, ECW students at magnet schools have significantly higher relative odds of suspension, while EM students do not. In a magnet school, the odds of an ECW student being suspended are 86% higher than a non-ECW student, while at a traditional public school, the ECW student's suspension odds are 51% higher than a non-ECW student. This suggests either that magnet schools suspend ECW students at higher rates, and/or ECW students are more likely to engage in behavior warranting suspension at magnet schools. This latter story becomes less plausible when one considers that absence rates are no different for ECW students at magnet schools, and baseline suspension rates across magnet schools are considerably more variable than those at traditional public schools. Thus, it is certainly plausible that some magnet schools suspend ECW students more aggressively, and this school response effect is being captured in the direct effect. It is important to note here that both traditional public and magnet schools fall under the same school district administration and operate under the same suspension guidelines, so this gap likely reflects implementation differences, not policy differences.

In Appendix B, the regressions are run with ECW split by whether children are linked to child welfare investigations in just one year or in multiple years prior to kindergarten. Children with multiple years of ECW fare worse than children with just one year in terms of suspension, but for absence the outcomes are approximately equivalent.

Results by race

Conditioning on ECW and EM, black and other minority students have absence rates that are 0.3 percentage points lower, on average, than white students (see Column 4 in Figure 7). The gap is reversed for suspension, in which minority students – and black students, in particular – have considerably higher suspension odds. As shown in Column 4 of Figure 8, suspension odds of black students are 197% larger than those of white students, while those for other minority students are 50% higher than for white students. Female students are considerably less likely to be suspended than male students. Taken together for the average student in the average cohort, this translates to baseline suspension probabilities of 2% for white girls, 6% for black girls and white boys, and more than 13% for black boys.

[Figure 10]

The ECW measure offers no explanation for this large and persistent racial disparity in suspension probability (compare Column 5 to Column 4 in Figure 8). Thus, it is worthwhile to explore whether direct and peer effects of Early Child Welfare persist across racial groups. In Figure 10 I plot suspension probabilities for black students compared to non-black students for absence and suspension. This plot is drawn from a specification of the suspension model where an indicator for black students is interacted with every individual and peer measure. As with Figure 9, the upward-sloping lines show that peer effects persist for ECW and non-ECW students alike, for both black and non-black students. However, a non-black student linked to a child welfare investigation in early childhood has slightly lower suspension probabilities than a black student not linked to child welfare.

As is somewhat apparent in Figure 10, peer effects appear proportionally larger for non-black ECW students than other subgroups. There are slight differences in the coefficient estimates of direct and peer effects of ECW by race, but all are positive and significant. Given the substantial differences in baseline suspension probabilities across subgroups, these differences can be made to seem big or small depending on how they are framed, but looking at the predicted suspension probabilities in Figure 10 offers a clear bottom line. ECW students have significantly higher suspension probabilities regardless of race, but these disparities operate on top of – not in place of – enormous racial disparities in suspension.

V. Discussion

This paper discovers a way to use administrative data to gain insight into a dimension of childhood disadvantage that has long concerned economists: disparities in parental investment in children's well-being and education that may not be attributable to disparities in family income. While more prevalent among low-income families, the effect of Early Child Welfare on educational outcomes applies whether or not the student is also low-income. By using administrative data that tracks virtually all elementary school-aged children living in a midsized city, I uncover evidence that Early Child Welfare students are less likely to enroll in charter, magnet, and private schools. In contrast, other low-income children are nearly as likely as higher-income children to enroll in these schools. Finally, Early Child Welfare students generate robust adverse peer effects in suspension that operate independently from the peer effects of low-income children.

This paper joins others in highlighting the need for better measures of childhood disadvantage than what typical school data permit. Variables more nuanced than simple poverty proxies are needed to understand the boundaries that prevent some children from accessing the best educational opportunities available to them. After all, if the costs to parents of school choice enrollment and attendance come through information and effort, not tuition dollars or transportation costs, then a proxy indicator of poverty will, at best, only partially capture the heterogeneity in family resources that may contribute to school selection and outcomes. I demonstrate this throughout my analyses by showing that when the Early Child Welfare variable is excluded, the estimates of the effect of Early Medicaid are relatively unchanged.

The findings of this paper are relevant to both child welfare and education policymakers and practitioners, including architects of school assignment systems. However, developing interventions and concrete policy recommendations on the basis of the patterns I identify in this paper requires a more detailed decomposition of the early childhood and subsequent classroom experiences of Early Child Welfare and Early Medicaid students. For example, if it can be shown that ECW is a good proxy for exposure to trauma, a case could be made for increasing school resources to support ECW students, potentially by hiring more school counselors (Reback 2010; Carrell & Hoekstra 2014). This idea aligns with evidence in the broader literature that disadvantaged students particularly benefit from increased school resources in the form of smaller class sizes (Aizer 2008; Krueger & Whitmore 1999).

However, there is also the potential for additional empirical analyses centered around policy changes and expanded data. Policy changes that are of particular interest concern school district changes in suspension practices, particularly for young children. The district under study is increasing the use of restorative practices and taking steps to reduce the number of out-of-school suspensions for young students. Future analyses can evaluate how Early Child Welfare students fare with these changes. On the expanded data front, linking information on parents' human services experiences to their children might enable an analysis of the cycles of abuse, poverty, and systemic involvement.

Appendix A: Negative binomial specification

This appendix provides an updated version of the educational outcomes models using negative binomial regressions, in lieu of linear and logistic regressions. Instead of an absence rate, the attendance outcome is the count of absent days (excused or unexcused). Instead of an indicator for whether a student is suspended during a given year, the suspension outcome is the count of suspension incidents. To account for varying lengths of enrollment – which prompt variation in the number of opportunities to be absent or suspended – the exposure variable in both models is set equal to the number of days a student is enrolled in a given school.

There are two key advantages of using the negative binomial approach. First is the ability to adjust for exposure, enabling inclusion of students enrolled for short periods of time in the model. The second advantage pertains specifically to the suspension model: including fixed effects as indicator variables in a negative binomial regression may avoid the incidental parameters problem (Allison & Waterman 2002).[6] A drawback of this approach, also noted by Allison & Waterman, is that the standard errors may be too small.

[Figures A1 & A2]

The results of the negative binomial specification (Figures A1 and A2) do not qualitatively alter the results from section IV. Specifically, an ECW student has absence incidence rates that are 23% higher than a non-ECW student, while an EM student has rates that are 18% higher than a non-EM student. In terms of suspension, an ECW student has rates that are 68% higher than a non-ECW student, while the disparity for EM students is 24%. The increased disparity for ECW students relative to the main specification indicates that ECW students are more likely to receive *multiple* suspensions relative to other students.

One slight difference from the main specification is that Figure A1 shows significant – but very small – peer effects for Early Medicaid students in attendance (at $p<.1$). This is not surprising since the EM peer effects in the main attendance specification are of similar magnitude, and nearly as precise.

[6] Since the dependent variables are over-dispersed – but not zero-inflated – I opt for a negative binomial specification, but Poisson regression is also consistent and does not suffer from the incidental parameters problem. The Poisson estimates are quite similar to those of the negative binomial.

Appendix B: Heterogeneity in exposure to disadvantage

Early Child Welfare and Early Medicaid encompass a range of adverse early life experiences. The data, however, show only whether a child was linked to an investigation or enrolled in Medicaid within a given year (defined to correspond with the school calendar as August of one year through July of the next). This enables two potential variants on the disadvantage variables: a measure of the number of years of involvement a child has prior to kindergarten, and a measure of approximate age at first or last involvement prior to kindergarten. Given the limited number of years a child receives services prior to kindergarten (in which most children enroll at age 5), these potential variants are largely redundant to each other. For the sake of simplifying the following analysis, I will focus on the number of years prior to kindergarten a child is linked to child welfare or Medicaid.

For Early Child Welfare, roughly equal proportions of children are linked to child welfare for one, two, or three or more years prior to kindergarten. Without overextending the interpretation of the data, there is some more nuance about early childhood experiences that can be gleaned by comparing children linked to an investigation during one year to those linked for multiple years. Suppose a referral was made because the reporter was biased or misunderstood a child's circumstances, not because of clear evidence of a threat to a child's well-being. When the child welfare agency investigates, if they do not find any evidence of such a threat, the case can be resolved quickly. The likelihood of another investigation is low. For children linked in multiple years, on the other hand, they may be repeatedly or continually exposed to dangerous circumstances. Thus, even though this is not a precise measure of severity, children linked to an investigation during just one year may have substantively better early childhood experiences than those linked in multiple years.

In contrast, there is little variation in the number of years children are enrolled in Medicaid. Among kindergartners in the outcomes dataset, 30% are never enrolled in Medicaid, and 55% are enrolled for 5 or 6 years prior to kindergarten (i.e., essentially enrolled from birth). Only 16% show some sort of heterogeneity, enrolling in Medicaid between one and four years. Children may have inconsistent Medicaid enrollment due to changes in household income, migration into or out of the county, or parents who choose to stop participating in Medicaid.

In the analysis that follows, in place of ECW, I use indicators for one year or multiple years of Early Child Welfare,[7] maintaining the reference category as children with no involvement. Additionally, in place of peer percent ECW, I use the percentage of peers with one year ECW and the percentage with multiple years ECW. Additionally, I use two "doses" of Early Medicaid in place of the EM indicator: 1-4 years, or 5+ years, along with peer measures for both of these subgroups.

[Figure B1]

The results in Figure B1 indicate that the effect attributed to ECW is stronger for students with multiple years of involvement than for students linked to an investigation during just one year. This finding is consistent with the intuition laid out above that postulates that children facing more difficult circumstances are more likely to have multiple years of ECW. However, students linked for just one year still have significantly worse outcomes than non-ECW students. This is consistent with findings of Ryan et al. (2018) who show than unsubstantiated child welfare cases are still predictive of adverse outcomes for children. The negative effect of unsubstantiated or brief interactions with child welfare could be attributable to the negative effect of system involvement, proper referrals that could not be verified, or referrals resulting from bias, which impacts the family beyond the child welfare interaction.

The Early Medicaid estimates also yield an interesting nuance: children enrolled in Medicaid for 1-4 years in early childhood have outcomes that are no worse than children never enrolled in Medicaid. This result in line with the findings of Michelmore & Dynarski (2017) who find that students continually eligible for free lunch fare significantly worse than those who are occasionally eligible.

In terms of peer effects, the subgroups with more involvement show stronger peer effects – including significant adverse peer effects in attendance from students with 5 or more years of Early Medicaid enrollment. However, peers with limited involvement in ECW and EM still show adverse – if less precise – peer effect estimates.

[Figure B2]

[7] This analysis was initially run with ECW doses of 2 or 3+ separated, but the results are statistically equivalent for those two subgroups. Combining them permits more precise estimates.

Unlike the main specification where the effects of disadvantage were not additive, interaction terms tell a different story in this model. This is best understood by investigating the predicted absence and suspension rates for students by disadvantage subgroup, as shown in Figures B2. These predictions reflect a version of the model where the effects of all covariates are allowed to vary for each subgroup. While the results in Figure B1 imply that students with two or more years of ECW have higher absence rates than students with just one year of involvement, the difference between the predicted attendance rates is quite small. Even one year of ECW is predictive of high absence rates. The same story is not true for suspension: predicted suspension rates for students with multiple years of ECW are nearly double those of students with 0-1 years of ECW involvement.

With respect with Early Medicaid, the predicted outcomes are consistent with the regression results in B1, with the notable observation that children with 0-4 years EM actually have lower predicted absence rates than those with no Medicaid involvement. This could reflect selection into Medicaid among households near the income cut-off, or increased efforts by schools to encourage attendance among EM students (to the extent they can identify these students, possibly through free/reduced lunch eligibility). While breaking the ECW variable out by dosage does not yield a clear story across these outcomes, the patterns for EM make the case that continual Medicaid involvement is a more useful predictor than any Medicaid involvement.

Appendix C: Probability of mid-year school exit

One potential threat to the identification of peer effects would be selective exit, where non-disadvantaged students are more likely to withdraw from a given school-grade during the school year when their peer group is more disadvantaged. This would lead to a change in peer group composition over time that might be conflated with a peer effect, since disadvantaged students with worse outcomes would be more likely to remain in the cohort.

The data enable a direct test of this possibility. I designate a student as exiting mid-year if they (a) are enrolled within the first 10 days of the school year, and (b) exit before the last month of the school year. Recall, the outcomes analysis is limited to traditional public and magnet schools. Early exits may include expulsions, which should be rare in elementary grades, but cannot be reliably separated from voluntary exits in the data.

[Figure C1]

Using the same logistic regression approach used in the suspension outcomes model in the main text, I find that, within a given school-grade, having more disadvantaged peers does not make a student more likely to withdraw mid-year. Thus, on average, idiosyncratic[8] changes in the percentage of peers that are Early Child Welfare and Early Medicaid do not result in the exit of non-disadvantaged students from the classroom.

However, this analysis does highlight an interesting difference among these two categories of disadvantage: Early Child Welfare students have exit probabilities 45% larger than those of other students, while Early Medicaid students are *less* likely to exit mid-year than other students. This gets to the heart of the potential difference between these two measures of disadvantage, in which Early Child Welfare captures a subgroup of children confronting systematically different treatment by parents and schools (in the case of expulsion) that is not exclusively attributable to poverty.

[8] Idiosyncratic in the sense that this variation is not attributable to systematic sorting into schools over time, or sorting in and out of the district over time.

References

Aizer, Anna. “Home Alone: Supervision after School and Child Behavior.” *Journal of Public Economics* 88, no. 9–10 (August 2004): 1835–48.

———. “Peer Effects and Human Capital Accumulation: The Externalities of ADD.” Cambridge, MA, September 2008.

Aizer, Anna, Janet Currie, Peter Simon, and Patrick Vivier. “Do Low Levels of Blood Lead Reduce Children’s Future Test Scores?” *American Economic Journal: Applied Economics* 10, no. 1 (January 2018): 307–41.

Aizer, Anna, Laura Stroud, and Stephen Buka. “Maternal Stress and Child Well-Being: Evidence from Siblings,” April 2009.

Allison, Paul D., and Richard P. Waterman. “Fixed Effects Negative Binomial Regression Models.” *Sociological Methodology* 32, no. 1 (August 2002): 247-265.

Bifulco, Robert, and Helen F. Ladd. “School Choice, Racial Segregation, and Test-Score Gaps: Evidence from North Carolina’s Charter School Program.” *Journal of Policy Analysis and Management* 26, no. 1 (2007): 31–56.

Carrell, Scott E., and Mark Hoekstra. “Are School Counselors an Effective Education Input?” *Economics Letters* 125, no. 1 (October 2014): 66–69.

Carrell, Scott E, and Mark L Hoekstra. “Externalities in the Classroom: How Children Exposed to Domestic Violence Affect Everyone’s Kids.” *American Economic Journal: Applied Economics*, January 2010, 19.

Case, Anne, Angela Fertig, and Christina Paxson. “The Lasting Impact of Childhood Health and Circumstance.” *Journal of Health Economics* 24, no. 2 (March 2005): 365–89.

Cowen, Joshua. “Who Chooses, Who Refuses? Learning More from Students Who Decline Private School Vouchers.” *American Journal of Education* 117, no. 1 (November 2010): 1-24.

Cowen, Joshua, and Marcus A. Winters. “Choosing Charters: Who Leaves Public School as an Alternative Sector Expands?” *Journal of Education Finance* 38, no. 3 (2013): 210–29.

Cunha, Flavio, and James Heckman. “The Technology of Skill Formation.” *AEA Papers and Proceedings* 97, no. 2 (2007): 17.

Currie, Janet, and Cathy Spatz Widom. “Long-Term Consequences of Child Abuse and Neglect on Adult Economic Well-Being.” *Child Maltreatment* 15, no. 2 (May 2010): 111–20.

Currie, Janet, and Erdal Tekin. “Understanding the Cycle: Childhood Maltreatment and Future Crime.” *The Journal of Human Resources*, 2012, 42.

Dahl, Gordon B, and Lance Lochner. "The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit." *American Economic Review* 102, no. 5 (August 2012): 1927–56.

Domina, Thurston, Nikolas Pharris-Ciurej, Andrew M. Penner, Emily K. Penner, Quentin Brummet, Sonya R. Porter, and Tanya Sanabria. "Is Free and Reduced-Price Lunch a Valid Measure of Educational Disadvantage?" *Educational Researcher*, September 6, 2018, 0013189X1879760.

Doyle, Joseph J., and Anna Aizer. "Economics of Child Protection: Maltreatment, Foster Care, and Intimate Partner Violence." *Annual Review of Economics* 10, no. 1 (August 2, 2018): 87–108.

Engberg, John, Dennis Epple, Jason Imbrogno, Holger Sieg, and Ron Zimmer. "Evaluating Education Programs That Have Lotteried Admission and Selective Attrition." *Journal of Labor Economics* 32, no. 1 (January 2014): 27–63.

Epple, Dennis, and Richard E. Romano. "Competition between Private and Public Schools, Vouchers, and Peer-Group Effects." *American Economic Review* 88, no. 1 (1998): 33–62.

Fantuzzo, John W., Whitney A. LeBoeuf, and Heather L. Rouse. "An Investigation of the Relations Between School Concentrations of Student Risk Factors and Student Educational Well-Being." *Educational Researcher* 43, no.1 (January 2014): 25-36.

Finkelhor, David, Heather A. Turner, Anne Shattuck, and Sherry L. Hamby. "Prevalence of Childhood Exposure to Violence, Crime, and Abuse: Results From the National Survey of Children's Exposure to Violence." *JAMA Pediatrics* 169, no. 8 (August 1, 2015): 746.

Garcia, David R. "The Impact of School Choice on Racial Segregation in Charter Schools." *Educational Policy* 22, no. 6 (November 2008): 805–29.

Gershenson, Seth, Alison Jacknowitz, and Andrew Brannegan. "Are Student Absences Worth the Worry in U.S. Primary Schools?" *Education Finance and Policy* 12, no. 2 (April 2017): 137–65.

Kenney, Genevieve M, Victoria Lynch, Michael Huntress, Jennifer M Haley, and Nathaniel Anderson. "Medicaid/CHIP Participation Among Children and Parents." *Timely Analysis of Immediate Health Policy Issues*, December 2012, 15.

Lazear, Edward P. "Educational Production." *The Quarterly Journal of Economics*, 2001.

Maloney, Tim, Nan Jiang, Emily Putnam-Hornstein, Erin Dalton, and Rhema Vaithianathan. "Black–White Differences in Child Maltreatment Reports and Foster Care Placements: A Statistical Decomposition Using Linked Administrative Data." *Maternal and Child Health Journal* 21, no. 3 (March 2017): 414–20.

Manski, Charles F. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies* 60, no. 3 (July 1993): 531.

Michelmore, Katherine, and Susan Dynarski. “The Gap Within the Gap: Using Longitudinal Data to Understand Income Differences in Educational Outcomes.” *AERA Open* 3, no. 1 (February 2017): 233285841769295.

Moffitt, Robert A. “Policy Interventions, Low-Level Equilibria and Social Interactions.” In *Social Dynamics*. MIT Press, 2001.

Pieterse, Duncan. “Childhood Maltreatment and Educational Outcomes: Evidence from South Africa.” *Health Economics* 24, no. 7 (July 2015): 876–94.

Reback, Randall. “Schools’ Mental Health Services and Young Children’s Emotions, Behavior, and Learning.” *Journal of Policy Analysis and Management* 29, no. 4 (September 2010): 698–725.

Ryan, Joseph P., Brian A. Jacob, Max Gross, Brian E. Perron, Moore, Andrew Moore, and Sharlyn Ferguson. *Child Maltreatment* 23, no. 4 (2018): 365-375.

Slade, Eric P., and Lawrence S. Wissow. “The Influence of Childhood Maltreatment on Adolescents’ Academic Performance.” *Economics of Education Review* 26, no. 5 (October 2007): 604–14.

Figure 1: Summary statistics.

	Selection model	Outcomes model
Students	8,502	20,020
School years	3	4
Grades	K	K-5
School types	All major	Trad. & magnet
Avg. obs. per student	1	2.2 (1.1)
Female	48.1%	48.7%
White	38.9%	31.4%
Black	45.1%	52.7%
Other minority	16.0%	15.9%
Early child welfare	12.4%	16.6%
Early medicaid	63.3%	66.2%
ECW & EM	11.1%	14.5%

Standard deviation in parentheses.

Figure 2: Kindergarten enrollment by school type.

School type	2013-14	2014-15	2016-17
Traditional public	1,794	1,603	1,373
Magnet	527	540	486
Charter	241	268	272
Private	460	481	340
Other*	40	36	41
Total	3,062	2,928	2,512

*Includes day cares, special needs schools, and home school.

Figure 3: District average outcomes by grade and Early Child Welfare and Early Medicaid exposure.

Absence rate (% of days)

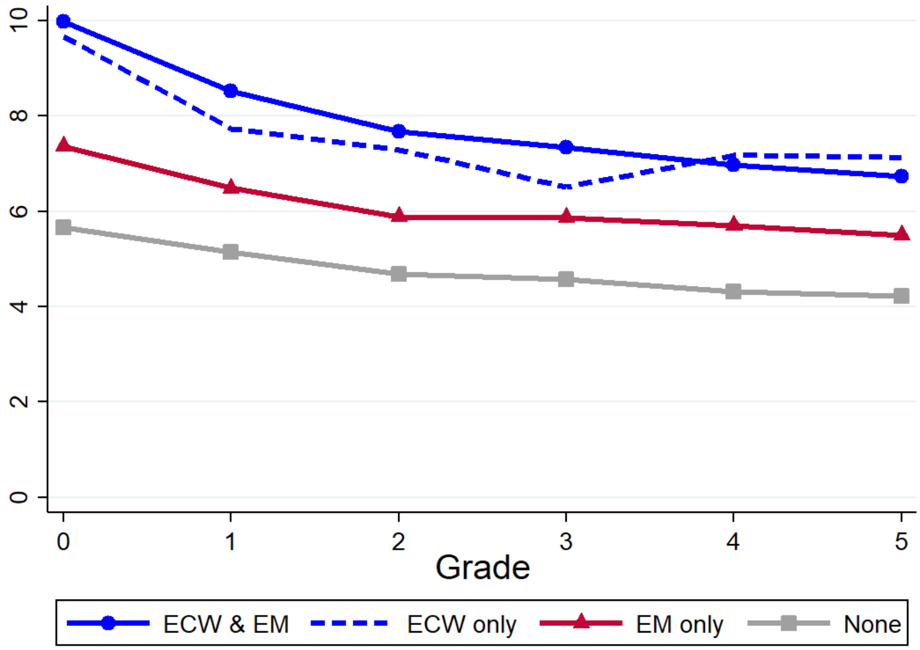

Percent suspended

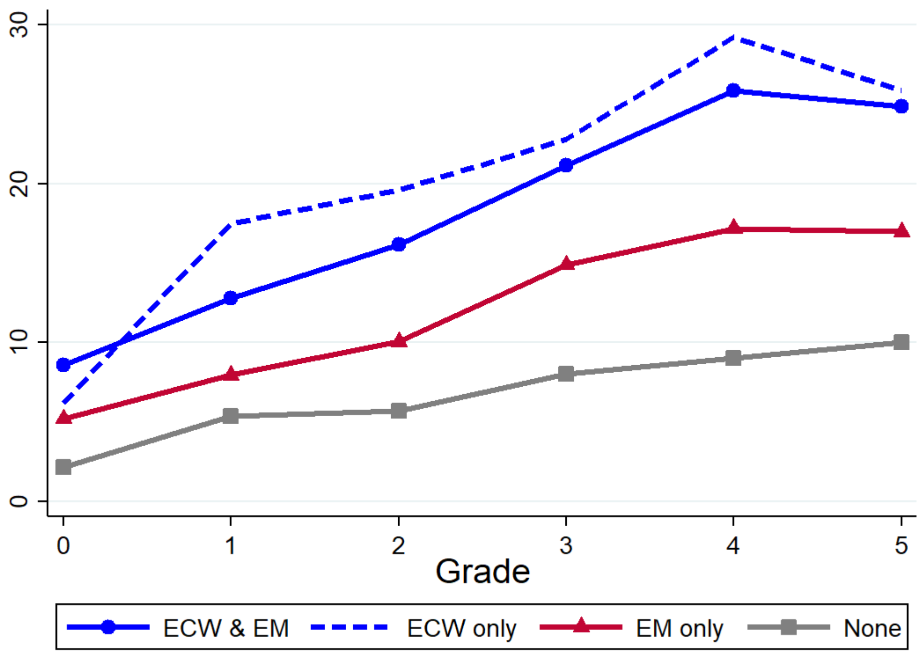

Figure 4: Multinomial logistic regression, Probability of school type enrollment relative to traditional public schools.

	(1)			(2)			(3)		
	Magnet	Charter	Private	Magnet	Charter	Private	Magnet	Charter	Private
Early Child Welfare	0.45***	0.50***	0.26***	0.42***	0.48***	0.27***	-	-	-
	(0.04)	(0.06)	(0.05)	(0.04)	(0.06)	(0.05)			
Early Medicaid	0.82*	0.92	0.26***	0.89	0.94	0.34***	0.81**	0.88	0.30***
	(0.09)	(0.09)	(0.0)	(0.07)	(0.08)	(0.03)	(0.07)	(0.08)	(0.03)
Female	0.98	1.07	1.04	0.98	1.04	1.04	0.98	1.04	1.05
	(0.05)	(0.06)	(0.07)	(0.06)	(0.06)	(0.07)	(0.06)	(0.06)	(0.07)
Black	1.47	0.89	0.15***	0.78	0.57***	0.09***	0.77	0.56***	0.09***
	(0.36)	(0.18)	(0.04)	(0.17)	(0.13)	(0.03)	(0.16)	(0.12)	(0.03)
Other minority	0.82	0.51***	0.31***	0.68**	0.43***	0.24***	0.68**	0.44***	0.24***
	(0.14)	(0.11)	(0.04)	(0.13)	(0.09)	(0.02)	(0.13)	(0.09)	(0.02)
Observations	8,379			8,379			8,379		
Year FEs	Y			Y			Y		
Zip code FEs	N			Y			Y		

Standard errors clustered at the zip code level. *** p<0.01, ** p<0.05, * p<0.1

Coefficients are displayed as relative risk ratios: the null hypothesis is coeff.=1.

Figure 5: Percent of school's kindergarten cohort linked to child welfare investigations in early childhood versus percent enrolled in Medicaid in early childhood.

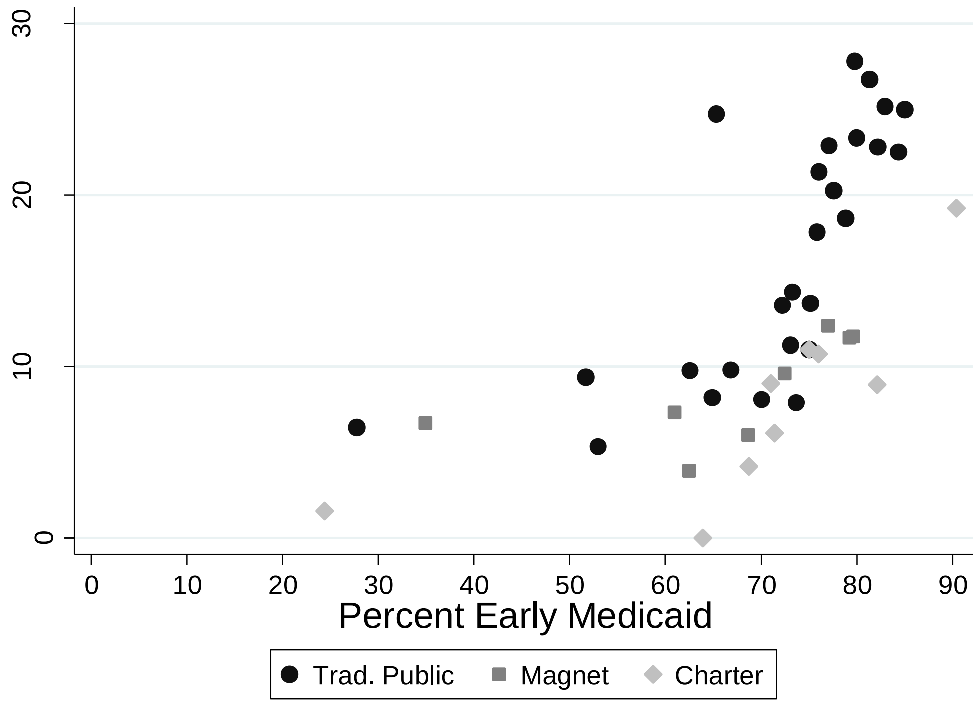

Note: Each point represents the weighted three-year average demographic composition of kindergarten cohorts in one of the city's schools. Schools enrolling fewer than 30 kindergartners over three years were excluded from the figure. Private schools (n=16) are excluded for visual clarity, but all except one have less than 60% Early Medicaid, and all have less than 10% Early Child Welfare.

Figure 6: Causal pathways under exogeneity of X_i and S_{sgt}.

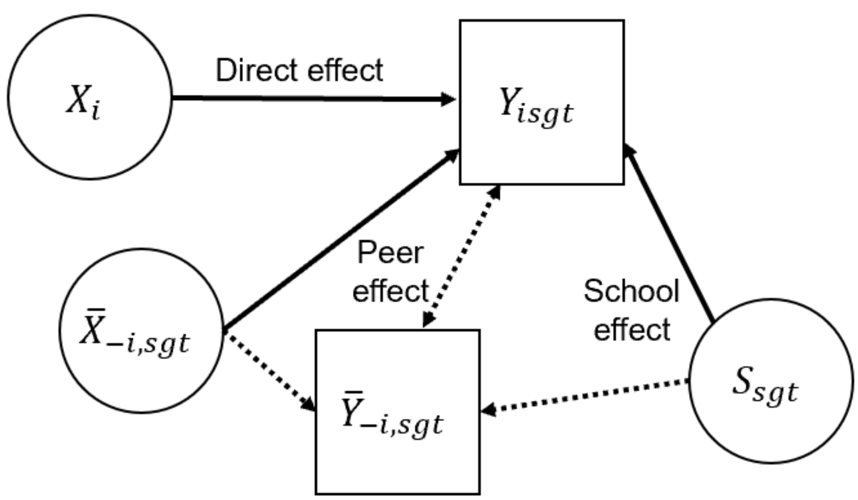

Figure 7: Linear regression, Absence rate.

	(1)	(2)	(3)	(4)	(5)
Early Child Welfare	1.618***	1.624***	1.495***	1.499***	-
	(0.101)	(0.101)	(0.100)	(0.101)	
Peer	0.049***	0.066***	0.004	0.005	-
	(0.007)	(0.007)	(0.007)	(0.007)	
Early Medicaid	0.962***	0.885***	0.858***	0.863***	1.034***
	(0.061)	(0.063)	(0.064)	(0.064)	(0.065)
Peer	0.031***	0.018***	0.006	0.009	0.009
	(0.004)	(0.005)	(0.006)	(0.005)	(0.005)
Female		0.027	0.030	0.025	0.018
		(0.059)	(0.058)	(0.058)	(0.059)
Black		-0.338***	-0.331***	-0.333***	-0.318***
		(0.084)	(0.086)	(0.085)	(0.086)
Other minority		-0.288***	-0.291***	-0.291***	-0.322***
		(0.086)	(0.086)	(0.086)	(0.087)
Cohort size		0.004	-0.005	0.003	0.004
		(0.003)	(0.004)	(0.004)	(0.004)
N	43406	43406	43406	43406	43406
R-sq	0.048	0.072	0.098	0.104	0.095
Other peer controls+	N	Y	Y	Y	Y
Zip code FEs	N	Y	Y	Y	Y
Grade-by-year FEs	N	Y	Y	Y	Y
School-by-grade FEs	N	N	Y	Y	Y
School-by-year FEs	N	N	N	Y	Y

Standard errors clustered at the school-grade-year level. *** p<0.01, ** p<0.05, * p<0.1

The units of the dependent variable (Absence Rate) are percentage points.

+ "Peer" variables are calculated as the share of student i's peers - excluding student i - for whom the relevant variable equals one. Controls for gender and race peer groups are also included in specifications 2-5. None are statistically significant.

Figure 8: Logistic regression, Probability of suspension.

	(1)	(2)	(3)	(4)	(5)
Early Child Welfare	1.613***	1.526***	1.540***	1.567***	-
	(0.059)	(0.057)	(0.060)	(0.061)	
Peer	1.035***	1.012***	1.012**	1.019***	-
	(0.005)	(0.004)	(0.005)	(0.005)	
Early Medicaid	1.385***	1.217***	1.205***	1.218***	1.280***
	(0.058)	(0.050)	(0.051)	(0.053)	(0.055)
Peer	1.011***	1.010***	1.010**	1.014***	1.014***
	(0.003)	(0.004)	(0.005)	(0.004)	(0.004)
Female		0.401***	0.388***	0.378***	0.378***
		(0.016)	(0.016)	(0.015)	(0.015)
Black		3.026***	2.940***	2.965***	2.974***
		(0.175)	(0.174)	(0.174)	(0.176)
Other minority		1.430***	1.490***	1.495***	1.483***
		(0.098)	(0.104)	(0.105)	(0.104)
Cohort size		0.999	1.007**	1.006**	1.006**
		(0.002)	(0.003)	(0.003)	(0.003)
N	43406	43379	43379	42901	42901
Other peer controls+	N	Y	Y	Y	Y
Zip code FEs	N	Y	Y	Y	Y
Grade-by-year FEs	N	Y	Y	Y	Y
School-by-grade FEs	N	N	Y	Y	Y
School-by-year FEs	N	N	N	Y	Y

Standard errors clustered at the school-grade-year level. *** p<0.01, ** p<0.05, * p<0.1
Coefficients expressed as odds ratios: the null hypothesis is coeff.=1.
+ "Peer" variables are calculated as the share of student i's peers - excluding student i - for whom the relevant variable equals one. Controls for gender and race peer groups are also included in specifications 2-5. None are statistically significant.

Figure 9: Predicted probability of suspension by Early Child Welfare and Medicaid exposure versus peer percent Early Child Welfare.

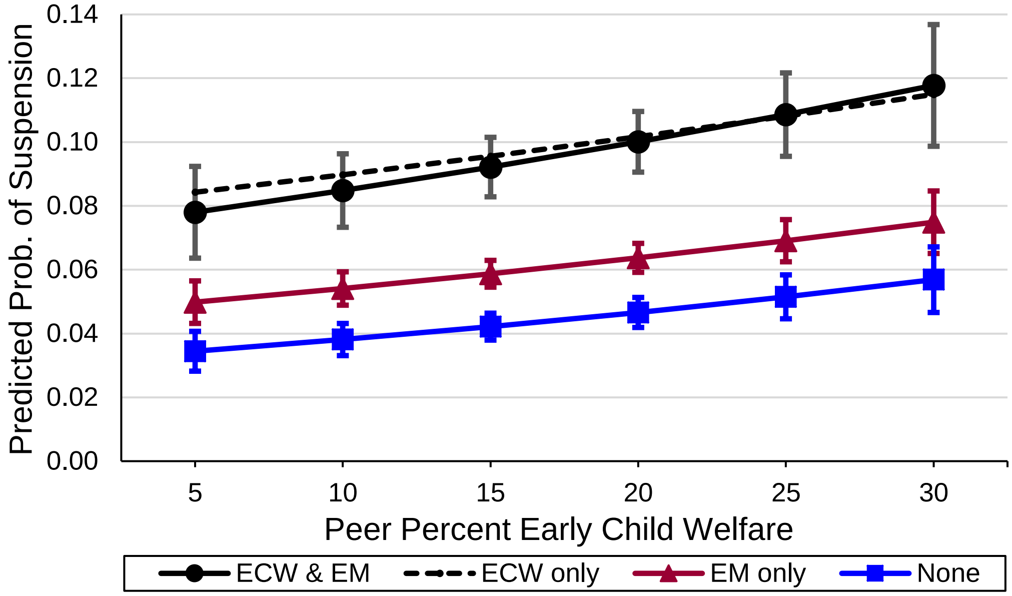

Note: Results from fully interacted model, where direct and peer effects can vary for each displayed subgroup. All covariates and fixed effects held to their means. Vertical bars show 95% confidence interval. Confidence intervals for the ECW only subgroup are omitted for visual clarity. They are much wider than those of the ECW & EM subgroup, but lie above those of the EM only subgroup at all points except for x=5 and x=30.

Figure 10: Predicted probability of suspension by Early Child Welfare exposure for black and non-black students versus peer percent Early Child Welfare.

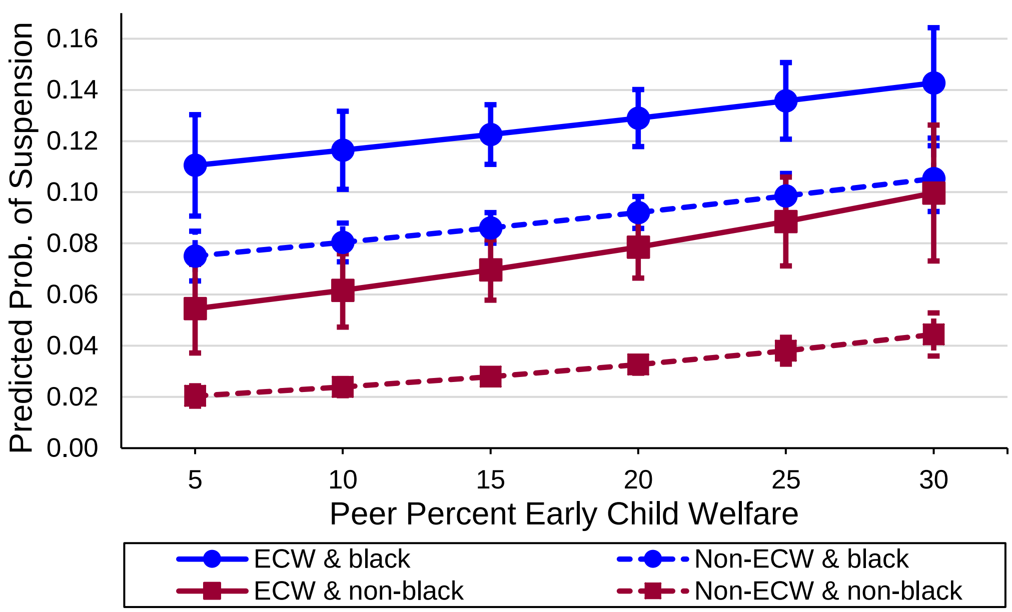

Note: Results from fully interacted model, where direct and peer effects can vary for each displayed subgroup. All covariates and fixed effects held to their means. Vertical bars show 95% confidence interval.

Figure A1: Negative binomial regression, Count of absences.

	(1)	(2)	(3)	(4)	(5)
Early Child Welfare	1.264***	1.258***	1.228***	1.227***	-
	(0.017)	(0.017)	(0.016)	(0.016)	
Peer	1.008***	1.011***	1.000	1.000	-
	(0.001)	(0.001)	(0.001)	(0.001)	
Early Medicaid	1.202***	1.181***	1.178***	1.179***	1.207***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Peer	1.006***	1.003***	1.001	1.002*	1.002*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Female		1.007	1.006	1.005	1.004
		(0.010)	(0.010)	(0.009)	(0.010)
Black		0.953***	0.954***	0.954***	0.958***
		(0.013)	(0.013)	(0.013)	(0.013)
Other minority		0.950***	0.947***	0.948***	0.945***
		(0.014)	(0.013)	(0.013)	(0.013)
Cohort size		1.000	0.999	1.001	1.001
		(0.000)	(0.001)	(0.001)	(0.001)
N	43416	43416	43416	43416	43416
Log pseudo-likelihood	-145133.2	-144419.7	-143730.1	-143585.8	-143764.7
Other peer controls+	N	Y	Y	Y	Y
Zip code FEs	N	Y	Y	Y	Y
Grade-by-year FEs	N	Y	Y	Y	Y
School-by-grade FEs	N	N	Y	Y	Y
School-by-year FEs	N	N	N	Y	Y

Standard errors clustered at the school-grade-year level. *** p<0.01, ** p<0.05, * p<0.1
+ "Peer" variables are calculated as the share of student i's peers - excluding student i - for whom the relevant variable equals one. Controls for gender and race peer groups are also included in specifications 2-5. None are statistically significant.

Figure A2: Negative binomial regression, Count of suspension incidents.

	(1)	(2)	(3)	(4)	(5)
Early Child Welfare	1.869***	1.710***	1.672***	1.680***	-
	(0.081)	(0.074)	(0.069)	(0.067)	
Peer	1.042***	1.020***	1.014**	1.017***	-
	(0.005)	(0.005)	(0.006)	(0.005)	
Early Medicaid	1.440***	1.257***	1.236***	1.238***	1.306***
	(0.066)	(0.058)	(0.055)	(0.056)	(0.059)
Peer	1.012***	1.011***	1.014***	1.013***	1.013***
	(0.003)	(0.004)	(0.005)	(0.004)	(0.004)
Female		0.400***	0.384***	0.378***	0.375***
		(0.017)	(0.016)	(0.016)	(0.016)
Black		3.118***	3.052***	3.065***	3.103***
		(0.207)	(0.194)	(0.192)	(0.197)
Other minority		1.499***	1.548***	1.566***	1.581***
		(0.123)	(0.120)	(0.122)	(0.125)
Cohort size		0.999	1.007**	1.005	1.005*
		(0.002)	(0.003)	(0.003)	(0.003)
N	43416	43416	43416	43416	43416
Log pseudo-likelihood	-21312.6	-20099.9	-19507.4	-19217.0	-19299.2
Other peer controls+	N	Y	Y	Y	Y
Zip code FEs	N	Y	Y	Y	Y
Grade-by-year FEs	N	Y	Y	Y	Y
School-by-grade FEs	N	N	Y	Y	Y
School-by-year FEs	N	N	N	Y	Y

Standard errors clustered at the school-grade-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

+ "Peer" variables are calculated as the share of student i's peers - excluding student i - for whom the relevant variable equals one. Controls for gender and race peer groups are also included in specifications 2-5. None are statistically significant.

Figure B1: Outcomes regressions with heterogeneous disadvantage.

	(1) Absence rate *(percentage pts.)*	(2) Suspension *(odds ratio)*
1 year ECW	1.017***	1.251***
	(0.148)	(0.069)
Peer	-0.010	1.012
	(0.010)	(0.008)
2+ years ECW	1.541***	1.694***
	(0.120)	(0.044)
Peer	0.011	1.021***
	(0.008)	(0.005)
1-4 year EM	-0.029	0.993
	(0.080)	(0.062)
Peer	0.004	1.010*
	(0.007)	(0.005)
5+ years EM	1.184***	1.279***
	(0.073)	(0.045)
Peer	0.012**	1.016***
	(0.006)	(0.004)
Female	0.019	0.377***
	(0.060)	(0.041)
Black	-0.457***	2.904***
	(0.088)	(0.058)
Other minority	-0.295***	1.495***
	(0.087)	(0.070)
Cohort size	0.004	1.005**
	(0.004)	(0.003)
N	43416	42911
R-sq	0.106	
Other peer controls[+]	Y	Y
Zip code FEs	Y	Y
Grade-by-year FEs	Y	Y
School-by-grade FEs	Y	Y
School-by-year FEs	Y	Y

Standard errors clustered at the school-grade-year level. *** p<0.01, ** p<0.05, * p<0.1
[+] "Peer" variables are calculated as the share of student i's peers - excluding student i - for whom the relevant variable equals one. Controls for gender and race peer groups are also included in both specifications. None are statistically significant.

Figure B2: Predicted outcomes by disadvantage subgroups.

I. Predicted absence rate

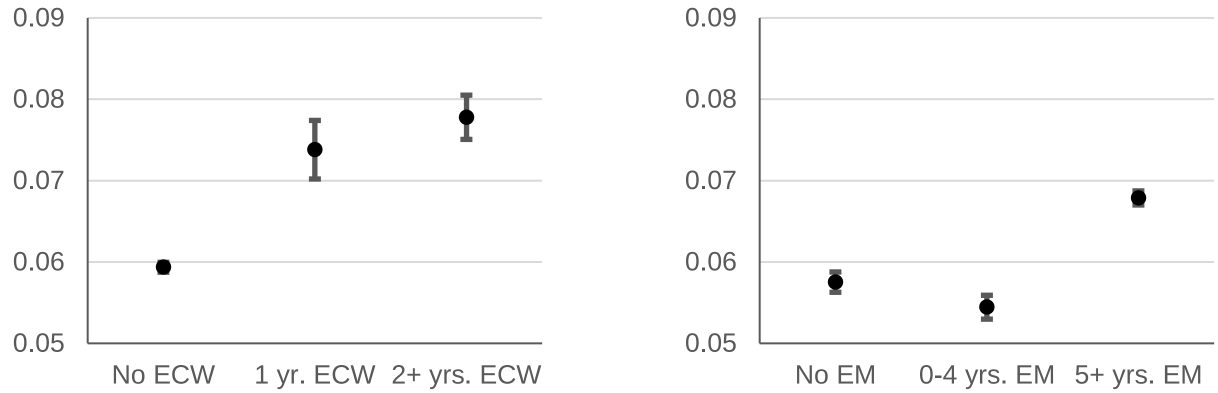

II. Predicted probability of suspension

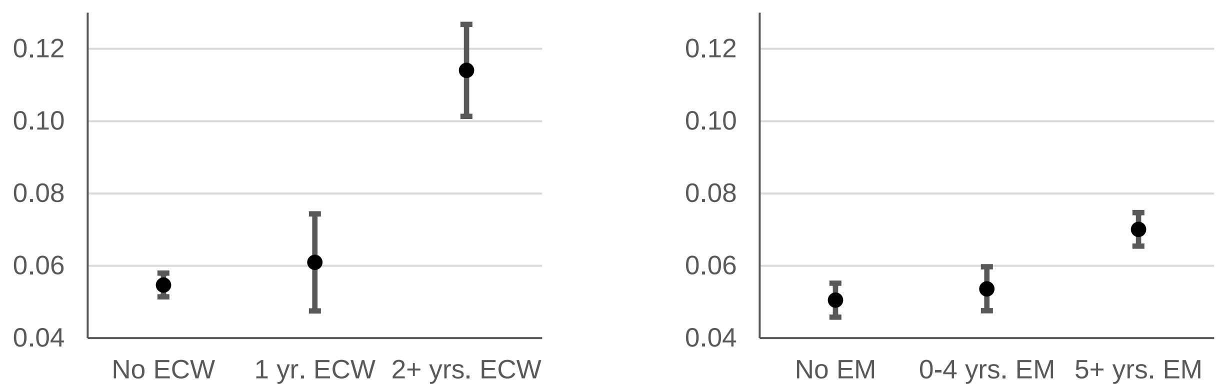

Note: Results from fully interacted model, where direct and peer effects across all covariates can vary for each displayed subgroup. All other covariates and fixed effects are held to their means. Vertical bars show 95% confidence interval.

Figure C1: Logistic regression, Probability of mid-year exit from school.

	(1)	(2)	(3)	(4)
Early Child Welfare	1.536***	1.536***	1.441***	1.447***
	(0.073)	(0.074)	(0.071)	(0.071)
Peer	1.043***	1.042***	1.003	1.003
	(0.003)	(0.004)	(0.004)	(0.003)
Early Medicaid	0.826***	0.796***	0.792***	0.792***
	(0.035)	(0.034)	(0.034)	(0.035)
Peer	0.995**	0.994**	0.998	0.998
	(0.003)	(0.003)	(0.003)	(0.003)
Female		0.920**	0.923**	0.924**
		(0.035)	(0.035)	(0.035)
Black		1.528***	1.517***	1.518***
		(0.092)	(0.095)	(0.096)
Other minority		1.603***	1.581***	1.572***
		(0.112)	(0.112)	(0.112)
Cohort size		0.998	1.002	1.000
		(0.001)	(0.002)	(0.002)
N	42468	42468	42468	42468
Other peer controls+	N	Y	Y	Y
Zip code FEs	N	Y	Y	Y
Grade-by-year FEs	N	Y	Y	Y
School-by-grade FEs	N	N	Y	Y
School-by-year FEs	N	N	N	Y

Standard errors clustered at the school-grade-year level. *** p<0.01, ** p<0.05, * p<0.1

+ "Peer" variables are calculated as the share of student i's peers - excluding student i - for whom the relevant variable equals one. Controls for gender and race peer groups are also included in specifications 2-4. None are statistically significant.