Demography (2017) 54:1653–1676
DOI 10.1007/s13524-017-0603-1

Neighborhoods, Schools, and Academic Achievement: A Formal Mediation Analysis of Contextual Effects on Reading and Mathematics Abilities

Geoffrey T. Wodtke[1] · **Matthew Parbst**[1]

Published online: 28 July 2017

Abstract Although evidence indicates that neighborhoods affect educational outcomes, relatively little research has explored the mechanisms thought to mediate these effects. This study investigates whether school poverty mediates the effect of neighborhood context on academic achievement. Specifically, it uses longitudinal data from the Panel Study of Income Dynamics, counterfactual methods, and a value-added modeling strategy to estimate the total, natural direct, and natural indirect effects of exposure to an advantaged rather than disadvantaged neighborhood on reading and mathematics abilities during childhood and adolescence. Contrary to expectations, results indicate that school poverty is not a significant mediator of neighborhood effects during either developmental period. Although moving from a disadvantaged neighborhood to an advantaged neighborhood is estimated to substantially reduce subsequent exposure to school poverty and improve academic achievement, school poverty does not play an important mediating role because even the large differences in school composition linked to differences in neighborhood context appear to have no appreciable effect on achievement. An extensive battery of sensitivity analyses indicates that these results are highly robust to unobserved confounding, alternative model specifications, alternative measures of school context, and measurement error, which suggests that neighborhood effects on academic achievement are largely due to mediating factors unrelated to school poverty.

Keywords Neighborhoods · Schools · Academic achievement · Poverty · Mediation

Electronic supplementary material The online version of this article (doi:10.1007/s13524-017-0603-1) contains supplementary material, which is available to authorized users.

✉ Geoffrey T. Wodtke
geoffrey.wodtke@utoronto.ca

[1] Department of Sociology, University of Toronto, 725 Spadina Avenue, Toronto, ON M5S 2J4, Canada

Introduction

Why does living in an advantaged rather than disadvantaged neighborhood improve academic achievement? Although evidence from a variety of different study designs indicates that neighborhood context affects educational outcomes (Aaronson 1998; Chetty et al. 2016; Harding 2003; Rosenbaum 1995; Wodtke et al. 2011), few studies have investigated the mechanisms thought to mediate these effects. *Neighborhood effect mediation* refers to the causal process whereby changes in neighborhood context lead to changes in an intermediate variable—known as a *mediator*—which in turn lead to changes in an outcome of interest. A frequent criticism of research on neighborhood effects is that the contextual mediators remain obscured in a so-called black box (Galster 2012; Jencks and Mayer 1990; Sampson et al. 2002)—that is, "research findings . . . are too scant to draw any firm conclusions about the potential pathways through which neighborhood effects may be transmitted" (Leventhal and Brooks-Gunn 2000:322).

Researchers have commonly hypothesized that neighborhood effects are mediated by the school environment to which children are exposed by virtue of their residential location (Arum 2000; Ferryman et al. 2008; Jencks and Mayer 1990; Johnson 2012; Leventhal and Brooks-Gunn 2000; Sanbonmatsu et al. 2006; Wilson 1987). Neighborhood context directly affects the socioeconomic composition of schools because school assignment rules are based at least partly on residential location, and the socioeconomic composition of schools is thought to affect student achievement in turn because schools with predominantly low-income students tend to have less-effective instructors, more disorderly classrooms, and fewer academic resources (Harris 2010; Willms 2010). Thus, differences in neighborhood context engender differences in school composition, which are in turn expected to engender differences in student achievement.

Although differences in the school environment are widely thought to mediate neighborhood effects on academic achievement, no prior study has provided a formal mediation analysis that evaluates the extent to which the total effect of neighborhood context is in fact explained by this particular mechanism. Several prior studies have investigated the joint effects of neighborhood and school contexts on educational outcomes, with some finding mainly neighborhood effects (Ainsworth 2002; Card and Rothstein 2007), some finding mainly school effects (Carlson and Cowen 2015; Cook et al. 2002; Goldsmith 2009), and others finding both (Owens 2010; Rendón 2014). However, none have appropriately decomposed the total effect of neighborhood context into direct and indirect components, which is essential for evaluating hypotheses that posit an important mediating role for schools. Moreover, prior studies are limited by their reliance on measurements of neighborhood and school contexts taken simultaneously rather than sequentially over time. As Cook et al. (2002:1303–1034) astutely noted, it is exceedingly difficult to evaluate whether "neighborhoods exercise their influence through their effects on schools" without the appropriate sequential measurements because any assumption about "the simultaneity of multiple causal relations is surely an oversimplification of reality."

In this study, we investigate whether school poverty mediates the effects of neighborhood context on reading and mathematics achievement, using longitudinal data that provide the requisite sequential measurements of these variables during both childhood and adolescence. We focus on measures of reading and mathematics achievement because these outcomes are closely linked with other dimensions of social stratification among adults, such as educational attainment, income, and health (Auld and Sidhu 2005; Murnane and Levy 2006). We focus on school poverty because it is widely analyzed in prior research on school effects (e.g., Ainsworth 2002; Halpern-Manners 2016; Lauen and Gaddis 2013), because school poverty has a direct causal connection with neighborhood composition, and because the socioeconomic composition of students is more closely related to educational outcomes than most other characteristics of the school environment (Coleman et al. 1966). In ancillary analyses, however, we also consider many alternative measures of school context, including the racial composition of the student body, the teacher-pupil ratio, per pupil expenditures, and a variety of teacher-reported classroom characteristics. Finally, we conduct separate analyses during childhood and adolescence to account for the possibility that contextual effects operate differently across developmental periods (Chetty et al. 2016; Cunha and Heckman 2007; Wodtke 2013; Wodtke et al. 2016).

To investigate whether school poverty mediates neighborhood effects, we use novel counterfactual methods (Pearl 2000; VanderWeele 2015) that allow for the decomposition of total effects into direct and indirect components under a weaker set of modeling restrictions than is required with more conventional approaches to mediation analysis (e.g., Baron and Kenny 1986), which are appropriate only if the effects of interest are linear and additive. Specifically, we focus on decomposing the total effect of neighborhood context into (1) a natural direct effect that measures differences in achievement due to residence in an advantaged versus a disadvantaged neighborhood if subjects are subsequently exposed to the level of school poverty that they would experience living in a disadvantaged neighborhood; and (2) a natural indirect effect that measures differences in achievement resulting from exposure to the level of school poverty that subjects would experience living in an advantaged neighborhood rather than the level of school poverty that they would experience living in a disadvantaged neighborhood.

We estimate these effects by fitting value-added models to longitudinal data from the Panel Study of Income Dynamics (PSID) that control for lagged measures of achievement, neighborhood context, and (where appropriate) also school poverty, among a variety of other individual- and family-level characteristics. This approach to estimation, which involves conditioning on lagged measures of the treatment, mediator, and outcome in an effort to identify the effects of future levels of the treatment and mediator on future levels of the outcome, provides some of the strongest protection against confounding bias in observational research (Pearl 2000; VanderWeele 2015). Nevertheless, it still requires strong assumptions about accurate measurement, correct model specification, and unobserved confounding that may not be satisfied in practice. Thus, we also conduct an extensive sensitivity analysis that evaluates whether our findings are robust to hypothetical violations of these assumptions.

Contrary to expectations, results indicate that school poverty is not a significant mediator of neighborhood effects on academic achievement during childhood or adolescence. Although total effect estimates indicate that moving from a disadvantaged neighborhood to an advantaged neighborhood significantly improves reading and mathematics achievement, natural direct and indirect effect estimates indicate that school poverty is not an important mediator of these effects because the differences in school composition induced by differences in neighborhood context do not have an appreciable impact on achievement. These findings are highly robust to hypothetical violations of the measurement, modeling, and confounding assumptions on which they are based, indicating that neighborhood effects are likely due to mediating factors unrelated to school poverty.

Neighborhood Effect Mediation by School Poverty

Institutional resource theory highlights the mediating role of schools in transmitting neighborhood effects on educational outcomes (Arum 2000; Jencks and Mayer 1990; Johnson 2012; Leventhal and Brooks-Gunn 2000; Wilson 1987). According to this perspective, differences in neighborhood context generate differences in the school environment to which children are exposed, which in turn lead to differences in academic achievement.

Neighborhood context directly affects the socioeconomic composition of schools because the public schooling options available to residents are, with some important exceptions, geographically determined. In most U.S. districts, public schools have designated attendance areas that restrict enrollment to residents from a set of local neighborhoods (National Center for Education Statistics (NCES) 2014a). These assignment rules engender an important connection between neighborhood and school contexts: changes in neighborhood composition due to family mobility or residential turnover lead to changes in the pool of eligible students from which local schools draw their enrollment. As a result, children living in disadvantaged neighborhoods typically attend schools with a greater number of low-income students, while children living in advantaged neighborhoods typically attend schools with fewer low-income students (Saporito and Sohoni 2007).

Despite the close link between neighborhood and school composition, a nontrivial number of children attend schools and live in neighborhoods with starkly different socioeconomic profiles. Many public schools serve catchment areas comprising multiple neighborhoods, and some of these neighborhoods may differ in their socioeconomic composition. In addition, due to the proliferation of magnet schools, charter schools, and intradistrict open enrollment policies, approximately 25 % of urban residents currently enroll their children in schools outside the local neighborhood (NCES 2014a). Similarly, private schools, which tend to enroll substantially fewer poor students than public schools because of the additional tuition costs, are also an option for high-income families or for low-income families with access to school vouchers or targeted scholarships. Consequently, it is not uncommon for children to attend a school with a socioeconomic composition that differs from their neighborhood (Saporito and Sohoni 2007).

The socioeconomic composition of schools is thought to be closely linked with school quality because schools with many low-income students tend to suffer from a variety of educational deficiencies (Battistich et al. 1995; Choi et al. 2008; Hedges et al. 1994; Kahlenberg 2001; Steinberg 1997; Willms 1986, 2010). For example, schools with a large proportion of low-income students often have less-effective teachers, a slower pace of instruction, a less-rigorous curriculum, and a greater number of in-class disruptions than schools with fewer low-income students (Kahlenberg 2001; Raudenbush et al. 2011; Willms 2010). In addition, because school funding comes in part from property taxes levied by local governments, the socioeconomic composition of schools may be linked to the financial resources at their disposal, although compensatory disbursements from state and federal governments tend to offset financial disparities that emerge between schools at the local level (Heuer and Stullich 2011; NCES 2015). Nevertheless, the socioeconomic composition of schools may be linked to other more intangible resources that are also important for student learning. For example, high-income parents tend to be more closely involved with their children's school and thus may provide a variety of social, cultural, and in-kind resources with spillover benefits for all students (Kahlenberg 2001; Steinberg 1997). Similarly, schools with a large proportion of high-income students tend to provide greater exposure to peers with expansive vocabularies and advanced subject knowledge, which may diffuse through student networks (Kahlenberg 2001; Willms 2010).

Consistent with arguments about the link between school poverty and school quality, a large volume of research has suggested that children who attend schools with a greater proportion of low-income students tend to have significantly lower academic achievement than children who attend schools with fewer low-income students but who are otherwise comparable on observed individual- and family-level characteristics (e.g., Battistich et al. 1995; Halpern-Manners 2016; Kahlenberg 2001; Rumberger and Palardy 2005; Schellenberg 1999; Willms 1986, 2010). Moreover, a number of other studies have reported significant effects of classroom and teacher characteristics that are thought to vary systematically with school poverty (Duncan and Murnane 2011). For example, evidence indicates that students perform better on reading and mathematics assessments when they learn in smaller classes with peers who achieve at higher levels (Burke and Sass 2011; Konstantopoulos and Chung 2009), when they attend classes with less absenteeism and student mobility (Raudenbush et al. 2011), and when they receive instruction from high-quality teachers (Hanushek 2011; Rivkin et al. 2005).

Findings from research on school effects, however, can be somewhat inconsistent, and the degree to which schools contribute to achievement disparities remains the subject of considerable debate (e.g., Downey and Condron 2016a, b). For example, some studies reported that attending a high- versus low-poverty school actually improves educational outcomes, at least for certain types of students (Attewell 2001; Crosnoe 2009). Several other studies raised serious questions about whether the associations between school context and student achievement documented in prior research warrant a causal interpretation, and they suggested that the true effects of school poverty may be small in practical terms (Coleman et al. 1966; Lauen and Gaddis 2013). Consistent with this view, seasonal learning comparisons show that socioeconomic differences in achievement grow much more slowly during the months when school is in session than during the summer when students are not in school (Downey et al. 2004; Heyns 1978), suggesting that these disparities are primarily a product of nonschool factors related to the family or neighborhood.

In sum, theory and prior research suggest a strong, albeit imperfect, causal link between the socioeconomic composition of neighborhoods and schools. Moving from a disadvantaged neighborhood to an advantaged neighborhood is therefore expected to substantially reduce subsequent exposure to school poverty for most students. Research on the causal link between school poverty and student achievement, on the other hand, is more ambiguous, but the overall weight of the evidence suggests that exposure to schools with many low-income students leads to at least somewhat lower achievement, net of other factors. Thus, the reductions in exposure to school poverty induced by moving from a disadvantaged neighborhood to an advantaged neighborhood are expected to improve academic achievement.

Neighborhood Effects Via Alternative Pathways

In addition to the school environment, theories of neighborhood effects also highlight a number of other mechanisms through which neighborhoods may influence educational outcomes. Social and cultural isolation theories of neighborhood effects posit that living in a disadvantaged neighborhood isolates resident children from influential adults who value education and discourage risky behaviors (Jencks and Mayer 1990; Wilson 1987), which is thought to curb educational aspirations and ultimately lead to disengagement from school. Social disorganization theories contend that disadvantaged neighborhoods engender lower levels of collective efficacy, which in turn may hinder academic progress (Sampson 2001; Sampson et al. 1997). For example, a breakdown of collective efficacy in disadvantaged neighborhoods is associated with high levels of violent crime (Sampson et al. 1997), and exposure to violent crime is a risk factor for many different cognitive, emotional, and behavioral problems in children (Sharkey 2010; Sharkey et al. 2012). Environmental theories of neighborhood effects focus on the disparate health hazards encountered in different neighborhood contexts. Because of the poor physical condition of disadvantaged neighborhoods, together with their proximity to major industrial centers, residents of these neighborhoods are disproportionately exposed to pollutants, toxins, and allergens (Crowder and Downey 2010; Rosenfeld et al. 2010), which may lead to place-based educational disparities (Sharkey and Faber 2014). Finally, although institutional resource theory focuses largely on the mediating role of schools, it also suggests that several other local institutions are important for explaining neighborhood effects. For example, in addition to high-quality schools, advantaged neighborhoods are more likely than disadvantaged neighborhoods to have high-quality childcare centers, grocery stores with healthy food options, and safe recreational facilities, all of which may promote positive educational outcomes for children (Bader et al. 2010; Johnson 2012; Weiss et al. 2011). In sum, although the school environment is thought to be a particularly important mediator of neighborhood effects, these effects may also be transmitted through several other potentially powerful pathways. Thus, theory and prior research additionally suggest a significant direct effect of neighborhood context on academic achievement that does not operate through differential exposures to school poverty.

Methods

Data

We use data from the PSID linked to information from the U.S. Census on the composition of neighborhoods and to information from the U.S. National Center for Education Statistics (NCES) on the characteristics of schools. The PSID began in 1968 with a probability sample of approximately 4,800 households (PSID 2014), conducting annual surveys until 1997 and biannual surveys thereafter. Data on academic achievement were collected as part of the Child Development Supplement (CDS), a component of the broader survey designed to assess early human capital formation. The CDS began in 1997 with a sample of 3,563 children aged 0 to 12, and it reinterviewed these children again in 2002 and 2007. The analytic sample for this study includes the 2,208 children who were interviewed at the 1997 wave of the CDS when they were between ages 3 and 12. We focus on this subset of children because it is the group for which repeated measures of academic achievement are available during childhood and/or adolescence.

We match sample members to their neighborhoods—here defined as *census tracts*—and to their schools using the PSID restricted-use geocode and school identification files, respectively.[1] Data on the composition of census tracts come from the GeoLytics Neighborhood Change Database, which contains tract-level data from the 1970–2010 U.S. Censuses and from the 2006–2010 American Community Surveys (GeoLytics 2013). Tract characteristics are imputed using linear interpolation for intercensal years. Data on the characteristics of schools come from the NCES Common Core of Data and Private School Universe Survey, which contain aggregate measures of student and staff characteristics from all public and private schools in the United States (NCES 2014b, c).

Measures

Treatment in this study is denoted by A_{it} and represents the socioeconomic composition of a sample member's neighborhood. We use principal components analysis to generate a composite measure of neighborhood context based on seven tract characteristics: the poverty rate, the unemployment rate, median household income, the proportion of households that are female-headed, aggregate levels of education (the proportion of residents age 25 or older without a high school diploma, and the proportion of residents aged 25 or older with a college degree), and the occupational structure (the proportion of residents aged 25 or older in managerial or professional occupations). This composite measure is scaled so that higher values represent more advantaged neighborhoods, and lower values represent more disadvantaged neighborhoods. Part A of Online Resource 1 describes the construction of this variable in detail. In all multivariate analyses, treatment is standardized to have 0 mean and unit variance.

The mediator of interest in this study is denoted by M_{it} and represents the percentage of students in a sample member's school who are eligible for a free lunch through the U.S. National School Lunch Program. To qualify for a free lunch, a student's family must have an income at or below 130 % of the federal poverty threshold. This measure is therefore an approximate school-level poverty rate. In a set of ancillary analyses, we also consider many alternative measures of school context, including the racial composition of students, the teacher-pupil ratio, per pupil expenditures, and the district dropout rate. In addition, for the subset of sample members whose elementary school

[1] Some of the data used in this analysis are derived from Restricted Data Files of the PSID, obtained under special contractual arrangements designed to protect the anonymity of respondents. These data are not available from the authors. Persons interested in obtained PSID Restricted Data Files should contact PSIDhelp@umich.edu.

teacher completed a phone interview as part of the 2002 wave of the CDS, we conduct ancillary analyses based on characteristics of the classroom environment as reported by these teachers, including the ability levels and socioeconomic status of classroom peers, tardiness and absenteeism, and disorderly behavior. Results from these analyses, presented in Part B of Online Resource 1, are highly consistent with those based on the measure of school poverty. In all multivariate analyses, the mediator is rescaled by its standard deviation but is not mean-centered.

The outcome in this study, academic achievement, is denoted by Y_{it} and measured using the letter-word and applied problem tests from the Woodcock-Johnson Psycho-Educational Battery-Revised (Woodcock and Johnson 1989). These tests assess reading and mathematics abilities, respectively, and they have excellent psychometric properties, including high test-retest reliability and high criterion validity (LaForte et al. 2014). Specifically, we use “W scores” from these tests, which provide an equal-interval measure designed to capture differences in learning over time. A W score of 500 is the reference score. This score represents achievement equivalent to the average fifth grader in the United States, and by extension, a test question with 500-level difficulty is answered correctly by fifth graders approximately 50 % of the time. A 10-point increase anywhere on the scale represents a sample member’s ability to perform, with 75 % success, academic tasks that she could previously perform with only 50 % success (Jaffe 2009).

To control for potential confounding, we measure and adjust for an extensive set of individual- and family-level covariates, including the race, gender, and age of the sample member; the age and education level of the sample member’s primary caregiver; the marital and employment status of the family head; and the net worth, income, homeowner status, size, and regional location of the sample member’s family. Race and gender are both dummy variables coded, respectively, as 1 for black and 0 for nonblack, and as 1 for female and 0 for male. Measures of race and gender are time-invariant and denoted by the vector $\mathbf{V}_i$. The marital and employment status of the family head are also both dummy variables coded, respectively, as 1 for married and 0 for unmarried, and 1 for employed and 0 for not employed. Age is measured in years, as is the education level of the primary caregiver. Family size is equal to the total number of people present in the household. Homeownership is a dummy variable coded as 1 for families who own their homes and 0 for families who do not. Regional location is also a dummy variable coded as 1 for residence in a southern census division and 0 otherwise. A family’s net worth is equal to the value of all assets minus all debts. This measure is adjusted for inflation and expressed in cube-root real dollars to correct for its extreme positive skew. Family income is measured using an income-to-needs ratio equal to a family’s annual real income divided by the official poverty threshold. All these measures are time-varying and are denoted by the vector $\mathbf{C}_{it}$. In multivariate analyses, measures of both the time-invariant and time-varying controls are centered on their sample means.

Table 1 depicts the longitudinal measurement strategy used in this analysis. The time index, t, is used to distinguish between baseline (t = 0), childhood (t = 1), and adolescent (t = 2) measurements of the variables outlined previously. In general, baseline measures are taken when sample members are aged 7 or younger; childhood measures are taken several years later, when sample members are between ages 8 and 12; and adolescent measures are taken several more years later, when sample members

Table 1 Longitudinal measurement strategy

	Time					
	1995	1997	1999	2002–2003	2005	2007
Main Survey	PSID95	PSID97	PSID99	PSID03	PSID05	PSID07
CDS Survey	—	CDS97	—	CDS02	—	CDS07
Analytic Sample						
3- to 7-year-olds at CDS97	A_0	Y_0, $\mathbf{C}_0$	A_1	M_1, Y_1, $\mathbf{C}_1$	A_2	M_2, Y_2
Age	1–5	3–7	5–9	8–12	11–15	13–17
8- to 12-year-olds at CDS97	A_1	M_1, Y_1, $\mathbf{C}_1$	A_2	M_2, Y_2	—	—
Age	6–10	8–12	10–14	13–17	—	—

Note: A_t = neighborhood advantage, M_t = school poverty, Y_t = academic achievement, and $\mathbf{C}_t$ = covariates.

are between ages 13 and 17. For notational simplicity, we use the same time subscript for measures of the treatment, mediator, and outcome taken within the same developmental period; however, these measures are in fact sequentially ordered because in each developmental period, neighborhood context is measured two years prior to the outcome, and because the measure of school poverty refers to the academic year immediately preceding measurement of the outcome. Thus, these data have the following temporal structure: $\{\mathbf{V}_i, A_{i0}, Y_{i0}, \mathbf{C}_{i0}, A_{i1}, M_{i1}, Y_{i1}, \mathbf{C}_{i1}, A_{i2}, M_{i2}, Y_{i2}\}$.

Our identification strategy in this study is to control for lagged measures of the treatment, outcome, covariates, and, where appropriate, also the mediator in an effort to estimate the effects of future levels of the treatment and mediator on future levels of the outcome. Prior research on this approach to identification suggests that it significantly mitigates confounding bias (Pearl 2000; VanderWeele 2015), although it is still premised on a number of strong assumptions (detailed later) that we submit to an extensive sensitivity analysis. More specifically, in analyses focused on childhood, we implement this identification strategy by controlling for baseline measures of neighborhood context (A_{i0}), achievement (Y_{i0}), and covariates ($\mathbf{C}_{i0}$) in order to estimate the effects of childhood neighborhood context (A_{i1}) and primary school poverty (M_{i1}) on childhood measures of achievement (Y_{i1}).[2] Because baseline measures are available only for sample members who were interviewed at the 1997 wave of the CDS when they were between ages 3 and 7, analyses of contextual effects during childhood are limited to this smaller group (n = 1,135). In analyses focused on adolescence, we implement this identification strategy by controlling for childhood measures of neighborhood context (A_{i1}), primary school poverty (M_{i1}), achievement (Y_{i1}), and covariates ($\mathbf{C}_{i1}$) in order to estimate the effects of adolescent neighborhood context (A_{i2}) and secondary school poverty (M_{i2}) on adolescent measures of achievement (Y_{i2}). Because these measures are available for all respondents, analyses of contextual effects during adolescence are based on the full analytic sample (n = 2,208). Descriptive statistics for all variables are presented in Tables 2 and 3.

[2] We do not attempt to control for a lagged measure of the mediator in analyses focused on childhood because many sample members were not yet in school at the time baseline measures were taken.

Table 2 Time-invariant sample characteristics

Variables	Mean	SD
Black	0.43	0.49
Female	0.49	0.50

Notes: The sample includes respondents who were interviewed at the 1997 wave of the CDS between ages 3 and 12 (n = 2,208). Results are combined estimates from 100 imputations.

Estimands

Total, natural direct, and natural indirect effects are defined using potential outcomes and the counterfactual framework (Rubin 1974; VanderWeele 2015). First, let a_t indicate exposure to a specific level of neighborhood advantage at time t. Next, let the potential outcome $Y_{it}(a_t)$ denote the achievement level of respondent i at time t had she previously been exposed to these neighborhood conditions, possibly contrary to fact. Similarly, let $M_{it}(a_t)$ represent the level of school poverty to which respondent i would subsequently be exposed under prior exposure to the level of neighborhood advantage given by a_t. The mediator is also defined as a potential outcome because it may be affected by treatment. Finally, note that $Y_{it}(a_t) = Y_{it}(a_t, M_{it}(a_t))$. This indicates that the potential outcomes, which are conventionally defined only in terms of treatment, can also be defined as a function of both treatment and the value of the mediator

Table 3 Time-varying sample characteristics

	Baseline (t = 0)		Childhood (t = 1)		Adolescence (t = 2)	
Variables	Mean	SD	Mean	SD	Mean	SD
Letter-Word Test Score	398.62	46.08	501.96	25.84	525.60	23.74
Applied Problem Test Score	437.11	35.07	504.08	18.30	524.25	19.78
Neighborhood Advantage Index	–1.01	2.32	–0.80	2.30	–0.53	2.35
School Poverty Rate	—	—	40.30	29.96	32.27	26.04
Respondent Age	4.96	1.41	10.00	1.41	—	—
Primary Caregiver's Age	32.54	7.30	37.65	7.08	—	—
Primary Caregiver's Education	12.92	2.27	12.84	2.45	—	—
Wealth (cube-root)	23.23	27.34	29.29	29.35	—	—
Income-to-Needs Ratio	2.68	2.36	3.01	2.48	—	—
Family Size	4.14	1.27	4.27	1.33	—	—
Head Is Married	0.62	0.49	0.64	0.48	—	—
Head Is Employed	0.82	0.39	0.83	0.38	—	—
Family Owns Home	0.55	0.50	0.65	0.48	—	—
Southern Residence	0.48	0.50	0.46	0.50	—	—

Notes: Baseline measures are reported only for respondents who were interviewed at the 1997 wave of the CDS between ages 3 and 7 (n = 1,135). Childhood and adolescence measures are reported for these respondents and additionally for respondents who were interviewed at the 1997 wave between ages 8 and 12 (n = 2,208). Results are combined estimates from 100 imputations.

under treatment. In other words, $Y_{it}(a_t, M_{it}(a_t))$ is the academic achievement level for respondent i at time t under exposure to the level of neighborhood advantage given by a_t and, by extension, under subsequent exposure to the level of school poverty the respondent would experience as a result of residence in these neighborhood conditions.

In the counterfactual framework, each subject is conceived to have a set of potential outcomes corresponding to all possible values of treatment, and contrasts between these potential outcomes define the causal effects of interest. Specifically, the average total effect at time t is defined as

$$ATE_t = E\left(Y_{it}\left(a_t^*\right) - Y_{it}(a_t)\right).$$

In words, ATE_t is the expected difference in academic achievement had respondents previously been exposed to the level of neighborhood advantage given by a_t^* rather than a_t. The average total effect can be additively decomposed into direct and indirect components as follows:

$$\begin{aligned} &E\left(Y_{it}\left(a_t^*\right) - Y_{it}(a_t)\right) = E\left(Y_{it}\left(a_t^*, M_{it}\left(a_t^*\right)\right) - Y_{it}(a_t, M_{it}(a_t))\right) \\ &= E\left(Y_{it}\left(a_t^*, M_{it}(a_t)\right) - Y_{it}(a_t, M_{it}(a_t))\right) + E\left(Y_{it}\left(a_t^*, M_{it}\left(a_t^*\right)\right) - Y_{it}\left(a_t^*, M_{it}(a_t)\right)\right). \end{aligned}$$

The first term in this decomposition is the average natural direct effect at time t,

$$NDE_t = E\left(Y_{it}\left(a_t^*, M_{it}(a_t)\right) - Y_{it}(a_t, M_{it}(a_t))\right).$$

NDE_t represents the expected difference in achievement under exposure to the level of neighborhood advantage given by a_t^*, rather than a_t, if each subject were subsequently exposed to the level of school poverty they would experience under the neighborhood conditions given by a_t. The second term in this decomposition is the average natural indirect effect at time t,

$$NIE_t = E\left(Y_{it}\left(a_t^*, M_{it}\left(a_t^*\right)\right) - Y_{it}\left(a_t^*, M_{it}(a_t)\right)\right).$$

NIE_t represents the expected difference in academic achievement under exposure to the level of neighborhood advantage given by a_t^* if each subject were subsequently exposed to the level of school poverty they would experience as a result of exposure to neighborhood conditions given by a_t^* rather than a_t.

For example, with $a_1^* = 1$ and $a_1 = -1$, NDE_1 represents the expected difference in academic achievement during childhood linked to residence in an advantaged neighborhood that is 1 standard deviation above the national mean of the composite socioeconomic index, rather than a disadvantaged neighborhood 1 standard deviation below the national mean, if each subject were subsequently exposed to the level of primary school poverty that they would experience by virtue of living in the more disadvantaged neighborhood. Similarly, NIE_1 represents the expected difference in academic achievement during childhood if, after initially being exposed to an advantaged neighborhood 1 standard deviation above the national mean of the composite socioeconomic index, respondents were then exposed to the level of primary school poverty that they would experience living in this advantaged neighborhood rather than the level of primary school poverty that they would experience living in a disadvantaged neighborhood 1 standard deviation below the national mean.

Substantively, ATE_t measures the effect of neighborhood context operating through all mediating pathways, including the pathway operating through subsequent exposure to school poverty. NDE_t, by contrast, measures the effect of neighborhood context on academic achievement operating through all pathways other than school poverty by fixing the mediator to the level it would have "naturally" been for each respondent under the reference level of treatment and then comparing achievement levels across differences in neighborhood context. This deactivates the pathway operating through school poverty but leaves all other pathways intact. NIE_t measures the effect of neighborhood context operating specifically through subsequent exposure to school poverty by fixing the level of treatment for each subject and then comparing achievement levels across the differences in school poverty that children would have experienced under prior exposure to different neighborhood conditions. This deactivates all pathways except for that operating through the socioeconomic composition of schools.

Identification

ATE_t, NDE_t, and NIE_t can be identified from the observed data under a set of so-called ignorability assumptions (Pearl 2000; VanderWeele 2015). Formally, these assumptions can be expressed as

$$Y_{it}(a_t, m_t) \perp A_{it} | \mathbf{C}_{it-1}; \quad Y_{it}(a_t, m_t) \perp M_{it} | \mathbf{C}_{it-1}, A_{it}; \quad M_{it}(a_t) \perp A_{it} | \mathbf{C}_{it-1};$$

and

$$Y_{it}(a_t, m_t) \perp M_{it}\left(a_t^*\right) | \mathbf{C}_{it-1} \text{ for } t = \{1, 2\},$$

where $\perp$ denotes statistical independence. For notational simplicity, lagged measures of the treatment, outcome, and—during adolescence—also the mediator are here and henceforth subsumed into the vector of prior covariates, $\mathbf{C}_{it-1}$, as are measures of the time-invariant controls. Informally, these assumptions respectively state that there must not be any unobserved treatment-outcome confounding; any unobserved mediator-outcome confounding; any unobserved treatment-mediator confounding; or any treatment-induced mediator-outcome confounding, whether observed or unobserved, during both childhood and adolescence. Figure 1 presents a directed acyclic graph (DAG) (Pearl 2000) that describes a set of hypothesized causal relationships in which all these assumptions are satisfied. The graph shows that (1) exposure to different neighborhood contexts directly affects subsequent exposure to school poverty and also academic achievement, (2) exposure to school poverty directly affects academic achievement, and (3) each of these effects is confounded only by a set of observed covariates. In this situation, NDE_t can be expressed in terms of the observed data as

$$\begin{aligned} NDE_t = &\textstyle\sum_{\mathbf{c}_{t-1}} \left[\sum_{m_t} \left(E\left(Y_{it} | \mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t^*, M_{it} = m_t\right)\right.\right. \\ &\left.\left. - E(Y_{it} | \mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t, M_{it} = m_t)\right) P(M_{it} = m_t | \mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t)\right] \\ &P(\mathbf{C}_{it-1} = \mathbf{c}_{t-1}) \text{ for } t = \{1, 2\}; \end{aligned}$$

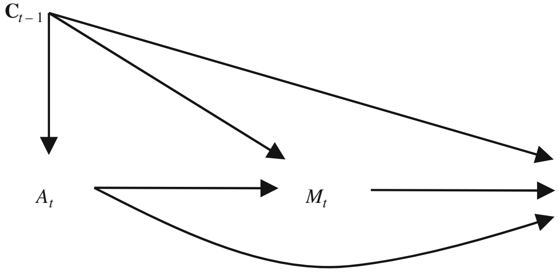

Fig. 1 Directed acyclic graph depicting the hypothesized causal relationships between neighborhood context, school poverty, and academic achievement. A_t = neighborhood advantage, M_t = school poverty, Y_t = academic achievement, and $\mathbf{C}_t$ = covariates

NIE_t can be expressed as

$$NIE_t = \sum_{\mathbf{c}_{t-1}} \Big[\sum_{m_t} \Big(P\big(M_{it} = m_t \big| \mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t^*\big) \\ -P(M_{it} = m_t | \mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t)\Big) E(Y_{it} | \mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t^*, M_{it} = m_{it})\Big] \\ P(\mathbf{C}_{it-1} = \mathbf{c}_{t-1}) \text{ for } t = \{1, 2\};$$

and ATE_t can be expressed as the sum of NDE_t and NIE_t.

Figure 2, by contrast, presents a series of DAGs that describe different hypothetical scenarios in which each of these confounding assumptions is violated. In any of these situations, decomposition of the total effect into natural direct and indirect components cannot be achieved with the observed data.

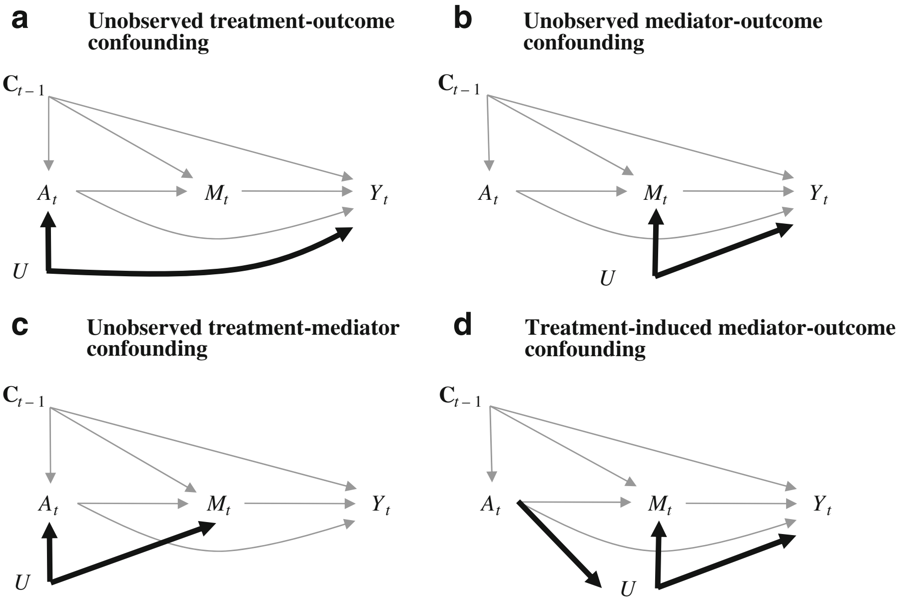

Fig. 2 Directed acyclic graphs depicting patterns of confounding that would lead to bias in mediation analyses of neighborhood effects. A_t = neighborhood advantage, M_t = school poverty, Y_t = academic achievement, $\mathbf{C}_{t-1}$ = prior covariates, and U = a hypothetical unobserved covariate

Estimation

Natural direct and indirect effects can be estimated from a set of two regression models: the first for the conditional mean of the mediator given treatment and prior covariates, and the second for the conditional mean of the outcome given treatment, the mediator, and prior covariates. These models are fit separately during childhood (t = 1) and adolescence (t = 2) and can be expressed as

$$E(M_{it}|\mathbf{C}_{it-1}, A_{it}) = \theta_{0t} + \theta'_{1t}\mathbf{C}_{it-1} + \theta_{2t}A_{it} + \theta_{3t}A_{it}^2 + \theta_{4t}A_{it}^3 \quad (1)$$

and

$$E(Y_{it}|\mathbf{C}_{it-1}, A_{it}, M_{it}) = \lambda_{0t} + \lambda'_{1t}\mathbf{C}_{it-1} + \lambda_{2t}A_{it} + \lambda_{3t}M_{it} + \lambda_{4t}M_{it}A_{it}. \quad (2)$$

Equation 1 includes linear, quadratic, and cubic terms for treatment to accommodate evidence of nonlinearity in the effects of neighborhood context on the mediator, and Eq. (2) includes linear terms for treatment and the mediator as well as a treatment-mediator interaction to account for the possibility that contextual effects on achievement may be nonadditive. Experimentation with a variety of more complex specifications failed to improve model fit. Under the assumption that Eqs. (1) and (2) are both correctly specified, in addition to the set of ignorability assumptions outlined previously, NDE_t is equal to

$$\begin{aligned}&NDE_t = \textstyle\sum_{\mathbf{c}_{t-1}}\big[\sum_{m_t}\big(E\big(Y_{it}|\mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t^*, M_{it} = m_t\big)\\ &-E(Y_{it}|\mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t, M_{it} = m_t)\big)P(M_{it} = m_t|\mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t)\big]P(\mathbf{C}_{it-1} = \mathbf{c}_{t-1})\\ &= \big(\lambda_{2t} + \lambda_{4t}\big(\theta_{0t} + \theta_{2t}a_t + \theta_{3t}a_t^2 + \theta_{4t}a_t^3\big)\big)\big(a_t^* - a_t\big) \text{ for } t = \{1, 2\};\end{aligned}$$

and NIE_t is equal to

$$\begin{aligned}&NIE_t = \textstyle\sum_{\mathbf{c}_{t-1}}\big[\sum_{m_t}\big(P\big(M_{it} = m_t|\mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t^*\big)\\ &-P(M_{it} = m_t|\mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t)\big)E\big(Y_{it}|\mathbf{C}_{it-1} = \mathbf{c}_{t-1}, A_{it} = a_t^*, M_{it} = m_t\big)\big]P(\mathbf{C}_{it-1} = \mathbf{c}_{t-1})\\ &= \big(\lambda_{3t} + \lambda_{4t}a_t^*\big)\big(\big(\theta_{2t}a_t^* + \theta_{3t}a_t^{*2} + \theta_{4t}a_t^{*3}\big) - \big(\theta_{2t}a_t + \theta_{3t}a_t^2 + \theta_{4t}a_t^3\big)\big) \text{ for } t = \{1, 2\}.\end{aligned}$$

By extension, ATE_t is equal to the sum of these two expressions.

In the upcoming Results section, we focus on total, natural direct, and natural indirect effects that contrast exposure to a neighborhood at the 80th percentile of the national treatment distribution with exposure to a neighborhood at the 20th percentile, which is roughly equivalent to a 1.5 standard deviation difference on the composite measure of neighborhood advantage. The contrast between the 80th versus the 20th percentile returns the effect of living in an advantaged neighborhood with low poverty, few female-headed households, and many highly educated adults versus living in a disadvantaged neighborhood with high poverty, many female-headed households, and few well-educated adults.

Because the vector of prior covariates, $\mathbf{C}_{it-1}$, includes lagged measures of the outcome, this analysis is based on a set of value-added models, which are commonly used in research on school effects and student learning (e.g., Chetty et al. 2014; Rowan et al. 2002). Estimates of the parameters in these models are computed by ordinary least squares and then used to construct the effects of interest with the formulas outlined previously. Standard errors are computed using the delta method after adjusting the appropriate variance-covariance matrices for the clustering of sample members within families. This analysis is then replicated across 100 complete data sets with missing values for all variables simulated via multiple imputation, and estimates are combined across replications following Rubin (1987). Overall, the proportion of missing information is roughly 10 % in analyses of both childhood and adolescence, most of which is due to sample attrition over time.[3] Finally, although the CDS is based on a complex sample design, we focus on unweighted estimates because they are very similar to weighted estimates that adjust for unequal probabilities of selection. In this situation, unweighted estimates are preferred because they are more efficient (Pfeffermann 1993; Winship and Radbill 1994). For reference, we report weighted estimates in Part C of Online Resource 1.

Results

The Joint Distribution of Neighborhood and School Composition

Table 4 describes the joint distribution of neighborhood context and school poverty during childhood and adolescence, where measures of these variables have been grouped by national quintile and then cross-tabulated. Several patterns are evident in these data. First, neighborhood context and school poverty are highly correlated. For example, among sample members in disadvantaged first-quintile neighborhoods, approximately 53 % attend high-poverty fifth-quintile schools, and only approximately 10 % attend schools with lower poverty levels in the first and second quintiles during childhood. By contrast, among sample members in advantaged fifth-quintile neighborhoods, approximately 73 % attend low-poverty first-quintile schools, and only approximately 4 % attend schools with higher poverty levels in the fourth and fifth quintiles during childhood.

Second, even though neighborhood context and school poverty are tightly coupled, a nontrivial number of sample members still live in neighborhoods and attend schools that differ in their socioeconomic composition by a considerable margin. For example, during adolescence, approximately 6 % of sample members attend a school that is at least two quintiles poorer than their neighborhood, and approximately 23 % attend a school that is at least two quintiles wealthier.

Total, Direct, and Indirect Effects of Neighborhood Context

Table 5 presents estimates of total, natural direct, and natural indirect effects, separately by developmental period. During childhood, total effect estimates suggest that exposure to different neighborhood contexts has a large effect on academic achievement.

[3] In addition, because the Private School Universe Survey does not include information on free lunch eligibility, sample members attending private schools, who compose between 6 % and 9 % of the analytic sample, are missing data on the mediator. For this group, we use measures of school racial composition, which are included in the survey, along with all the other variables outlined previously, to impute school poverty rates.

Table 4 Joint treatment-mediator distribution during childhood and adolescence

	Childhood School Poverty Quintile						Adolescence School Poverty Quintile					
Neighborhood Advantage Quintile	1	2	3	4	5	Total	1	2	3	4	5	Total
1												
N	12	26	47	90	201	376	46	69	130	226	234	705
Row	.03	.07	.13	.24	.53		.06	.19	.18	.32	.33	
Cell	.01	.02	.04	.08	.18		.02	.03	.06	.10	.11	
2												
N	22	39	48	70	59	238	65	106	109	91	67	438
Row	.09	.16	.20	.29	.25		.15	.24	.25	.21	.15	
Cell	.02	.03	.04	.06	.05		.03	.05	.05	.04	.03	
3												
N	30	56	43	28	18	176	97	120	73	36	28	353
Row	.17	.32	.25	.16	.10		.28	.34	.21	.10	.08	
Cell	.03	.05	.04	.03	.02		.04	.05	.03	.02	.01	
4												
N	63	51	28	21	9	172	147	81	49	31	12	320
Row	.37	.30	.16	.12	.06		.46	.25	.15	.10	.04	
Cell	.06	.05	.02	.02	.01		.07	.04	.02	.01	.01	
5												
N	126	23	16	5	2	172	265	78	28	16	4	392
Row	.73	.14	.09	.03	.01		.68	.20	.07	.04	.01	
Cell	.11	.02	.01	.00	.00		.12	.04	.01	.01	.00	
Total	253	196	183	213	290	1,135	621	454	389	400	345	2,208

Notes: The adolescence sample includes respondents who were interviewed at the 1997 wave of the CDS between ages 3 and 12; the childhood sample includes only the subset of these respondents who were interviewed at this wave between ages 3 and 7. Results are combined estimates from 100 imputations.

Specifically, estimates of ATE_1 indicate that childhood exposure to an advantaged neighborhood at the 80th percentile of the national treatment distribution, rather than a disadvantaged neighborhood at the 20th percentile, increases performance on the letter-word and applied problem tests by 4.72 and 4.42 points, respectively. These effects are substantively large and statistically significant at the $\alpha = .01$ level. To put them in perspective, note that a 5-point increase on these tests represents the ability to perform, with 62.5 % success, academic tasks that could previously be performed with only 50 % success, which is roughly equivalent to the learning typically achieved over three to five months of schooling during childhood.

Contrary to expectations, however, estimates of natural direct and indirect effects during childhood provide no evidence that neighborhood effects are mediated by subsequent exposure to primary school poverty. For example, estimates of NDE_1 indicate that if children were exposed to an advantaged neighborhood at the 80th percentile of the national treatment distribution, rather than a disadvantaged

Table 5 Total, natural direct, and natural indirect effects of neighborhood poverty on academic achievement during childhood and adolescence

	Letter-Word Scores			Applied Problem Scores		
Estimand	Est.	SE	p Value	Est.	SE	p Value
Childhood (n = 1,135)						
ATE_1	4.72	1.72	.006	4.42	1.19	<.001
NDE_1	5.16	1.73	.003	4.68	1.19	<.001
NIE_1	–0.44	0.89	.618	–0.27	0.62	.667
Adolescence (n = 2,208)						
ATE_2	1.24	1.11	.264	3.04	0.92	.001
NDE_2	1.04	1.11	.348	2.72	0.93	.003
NIE_2	0.20	0.28	.469	0.32	0.27	.241

Notes: The adolescence sample includes respondents who were interviewed at the 1997 wave of the CDS between ages 3 and 12; the childhood sample includes only the subset of these respondents who were interviewed at this wave between ages 3 and 7. Effect estimates are based on models that control for lagged measures of the treatment; outcome; covariates; and, in adolescence, also the mediator. Results are combined estimates from 100 imputations. Delta-method standard errors are reported in parentheses. p values are from two-sided Wald tests of no effect.

neighborhood at the 20th percentile, and then were subsequently exposed to the level of primary school poverty they would have experienced in the disadvantaged neighborhood, their performance on the letter-word and applied problem tests would still increase by 5.16 and 4.68 points, respectively. These effects are substantively large, statistically significant at stringent thresholds, and essentially indistinguishable from the corresponding total effect estimates. By extension, estimates of NIE_1 indicate that if children were to live in an advantaged neighborhood and then were exposed to the level of primary school poverty they would have experienced living in this neighborhood, rather than the level of primary school poverty they would have experienced living in a disadvantaged neighborhood, their performance on both the letter-word and applied problem tests would barely change at all.

The results for adolescence are generally consistent with those for childhood. Estimates of ATE_2 indicate that adolescent exposure to an advantaged neighborhood at the 80th percentile of the national treatment distribution, rather than a disadvantaged neighborhood at the 20th percentile, increases performance on the letter-word and applied problem tests by 1.24 and 3.04 points, respectively. The estimated total effect on letter-word scores during this developmental period is small and not statistically significant, but the estimated total effect on applied problem scores is substantively large and statistically significant at the α = .01 level. Specifically, it is roughly equivalent to the typical performance gains achieved over nine months of schooling during adolescence.[4]

[4] Although smaller in absolute terms when compared with the corresponding total effect during childhood, the total effect on applied problem scores during adolescence is larger when expressed relative to the performance gains typically achieved over the same developmental period because test score growth slows substantially at older ages.

Despite evidence of a strong total effect on mathematics achievement during adolescence, estimates of natural direct and indirect effects provide no indication that the total effect is mediated by subsequent exposure to secondary school poverty. For example, the estimate of NDE_2 on applied problem scores indicates that if adolescents were exposed to an advantaged neighborhood at the 80th percentile of the national treatment distribution, rather than a disadvantaged neighborhood at the 20th percentile, and then were subsequently exposed to the level of secondary school poverty they would have experienced in the disadvantaged neighborhood, their performance on this test would still increase by 2.72 points. The corresponding estimate of NIE_2, by contrast, indicates that if adolescents were to live in an advantaged neighborhood and then were exposed to the level of secondary school poverty they would have experienced living in this neighborhood, rather than the level of secondary school poverty they would have experienced living in a disadvantaged neighborhood, their performance on the applied problem test would increase by only 0.32 points.

Table 6 presents estimates of the parameters associated with treatment and the mediator from Eqs. (1) and (2), which were used to construct estimates of the total, natural direct, and natural indirect effects discussed previously. These results illuminate *why* school poverty does not appear to explain the effects of neighborhood context on academic achievement. Specifically, estimates from Eq. (1) indicate that the effect of neighborhood context on subsequent exposure to school poverty is substantively large and statistically significant, as expected. For example, during childhood, moving from a

Table 6 Selected parameter estimates from models of achievement test scores and exposure to school poverty during childhood and adolescence

	Eq. (1)		Eq. (2)			
	School Poverty		LW Scores		AP Scores	
Variable	Est.	p Value	Est.	p Value	Est.	p Value
Childhood (n = 1,135)						
A_1 (neighborhood advantage)	–0.475	<.001	2.252	.147	2.449	.020
A_1^2	0.050	.004	—	—	—	—
A_1^3	0.040	.001	—	—	—	—
M_2 (school poverty)	—	—	–0.064	.952	–0.013	.987
A_1M_1	—	—	0.766	.283	0.433	.341
Adolescence (n = 2,208)						
A_2 (neighborhood advantage)	–0.255	<.001	0.470	.602	2.125	.005
A_2^2	0.022	.137	—	—	—	—
A_2^3	0.021	.057	—	—	—	—
M_2 (school poverty)	—	—	–0.745	.208	–0.685	.193
A_2M_2	—	—	0.165	.678	–0.232	.508

Notes: The adolescence sample includes respondents who were interviewed at the 1997 wave of the CDS between ages 3 and 12; the childhood sample includes only the subset of these respondents who were interviewed at this wave between ages 3 and 7. All models include controls for lagged measures of the treatment; outcome; covariates; and, in adolescence, also the mediator. Results are combined estimates from 100 imputations. p values are from two-sided t tests of no effect.

disadvantaged neighborhood at the 20th percentile of the national treatment distribution to an advantaged neighborhood at the 80th percentile is estimated to reduce subsequent exposure to primary school poverty by nearly two-thirds of a standard deviation (i.e., $(1.0(-0.475) + (1.0)^2(0.050) + (1.0)^3(0.040)) - (-0.5(-0.475) + (-0.5)^2(0.050) + (-0.5)^3(0.040)) = -0.630$), which is roughly equivalent to 20 percentage points. Estimates from Eq. (2), however, indicate that even the large reductions in exposure to school poverty associated with moving from a disadvantaged neighborhood to an advantaged neighborhood fail to explain the total effects of neighborhood context because school poverty does not have an appreciable effect on achievement, net of other factors. For example, during childhood, even an extreme reduction in exposure to primary school poverty of 2 full standard deviations is estimated to increase performance on the letter-word and applied problem tests by only 0.13 and 0.03 points, respectively (i.e., $-2(-0.064) = 0.128$, and $-2(-0.013) = 0.026$).

Sensitivity Analyses

Effect estimates from this analysis only have a causal interpretation under a number of strong assumptions about correct model specification, the absence of confounding, and accurate measurement. First, if Eqs. (1) or (2) are incorrectly specified in any way, then effect estimates may be biased. Part D of Online Resource 1 presents results from an ancillary analysis that explores a variety of alternative specifications for these models. Experimentation with many different and more flexible specifications suggests that the reported estimates are robust. Second, if any of the confounding assumptions outlined previously are violated, then effect estimates may also be biased. Part E of Online Resource 1 presents a formal sensitivity analysis that investigates whether any of our inferences would change if certain of these assumptions are violated in different ways. Results indicate that our central conclusions about neighborhood effect mediation remain valid even under extreme violations of the confounding assumptions on which they are based. Finally, measurement error in the mediator would also lead to bias in estimates of the degree to which differences in school poverty can explain neighborhood effects on academic achievement. This is concerning because free lunch eligibility is known to be an imperfect proxy for school poverty (Cruse and Powers 2006). Nevertheless, estimates of natural direct and indirect effects that adjust for measurement error in the mediator are very similar to those discussed previously. These estimates are presented in Part F of Online Resource 1.

Discussion

Although the educational effects of neighborhood context have been extensively studied, relatively little research has investigated the mechanisms commonly hypothesized to mediate these effects. In this study, we investigate whether school poverty mediates the effects of neighborhood context on academic achievement during both childhood and adolescence. Using counterfactual methods and a value-added modeling strategy, which permit a decomposition of total effects into direct and indirect components under a defensible set of assumptions, we find that exposure to an advantaged rather than disadvantaged neighborhood significantly improves academic achievement.

Contrary to expectations, however, we find no evidence that school poverty mediates these effects because the differences in school poverty linked to differences in neighborhood context appear to have no appreciable impact on achievement. Moreover, we find that these results are highly robust to the use of alternative school-level measures, to alternative model specifications, to hypothetical patterns of unobserved confounding, and to measurement error in the mediator.

These findings are difficult to reconcile with institutional resource theory, which posits that the school environment is a particularly important mediator of neighborhood effects on educational outcomes. Rather, consistent with recent arguments that "socioeconomic achievement gaps . . . are more a product of factors outside of schools than pernicious school processes" (Downey and Condron 2016b:207), our results suggest that neighborhood effects are most likely explained by alternative mediators that are not directly linked with schools, such as neighborhood subcultures, local violent crime, or environmental health hazards.

These findings are also difficult to reconcile with policy prescriptions that advocate for school-based interventions—and in particular, for the socioeconomic desegregation of schools—as a means to mitigate the disparities in academic performance engendered by differences in neighborhood environments (e.g., Kahlenberg 2001; Oreopoulos 2012). In other words, it does not appear that reducing socioeconomic segregation across schools would remedy the harms of persistent socioeconomic segregation across neighborhoods. This type of school-based intervention can, of course, be motivated on alternative grounds and implemented to achieve alternative ends. The results presented in this study merely suggest that it would be ineffectual specifically with regard to the goal of attenuating neighborhood effects on academic achievement, even if it may benefit students in other ways.

An important methodological implication of this study is that effects of school poverty estimated from research designs that do not control for neighborhood context are likely to be inflated. This is because neighborhood context strongly affects both subsequent exposure to school poverty and academic achievement, which makes it an important confounder of school effects that, if left uncontrolled, would lead to bias. Indeed, in a set of ancillary analyses not reported here, we find that estimated school effects become much larger and statistically significant in models that omit measures of neighborhood context. Thus, given that so few studies of school effects control for the residential environment, these results suggest that a reconsideration of the large literature reporting significant effects of school poverty on achievement may be in order.

Although this study has important implications for theory, policy, and methods, it also suffers from several limitations. The first is our reliance on an imperfect and nonexhaustive set of school-level measures (i.e., the student poverty rate; and in ancillary analyses, the racial composition of the student body, the teacher-pupil ratio, per pupil expenditures, the district dropout rate, the ability level of classroom peers, tardiness and absenteeism, and disorderly behavior) when other dimensions of the school environment, such as teacher quality (Rivkin et al. 2005), may be more important mediators of neighborhood effects on achievement. The second limitation is our narrow focus on test scores, when major transitions such as high school graduation or college attendance may be more closely linked to later life chances and also to differences in the school environment earlier during childhood (Jackson 2012). Finally, a third limitation is that we focus only on population average effects, when it remains possible that neighborhood effect mediation via schools may be more pronounced among certain subgroups of children or in certain states, cities, and school districts.

Despite these limitations, we provide considerable evidence that neighborhood effects on academic achievement are not due to compositional differences in the schools attended by resident children. Future research should build on these findings while also addressing the limitations mentioned previously—for example, by investigating other educational outcomes, other characteristics of the school environment, or different subgroups of children. Although we find little evidence that school differences mediate neighborhood effects on academic achievement, this study directs the focus of future research toward alternative outcomes, mediators, and subpopulations, which may shed new light into the black box through which neighborhood effects are transmitted.

Acknowledgments The authors thank Ken Wodtke, David Harding, Felix Elwert, Alvin Pearman, Doug Lauen, and participants in both the Toronto Inequality Workshop and Contexts of Social Inequality Conference at the Berlin Social Science Center for helpful comments and discussions. This research was supported by a Connaught New Researcher Award from the University of Toronto.

References

Aaronson, D. (1998). Using sibling data to estimate the impact of neighborhoods on children's educational outcomes. *Journal of Human Resources, 33,* 915–946.

Ainsworth, J. W. (2002). Why does it take a village? The mediation of neighborhood effects on educational achievement. *Social Forces, 81,* 117–152.

Arum, R. (2000). Schools and communities: Ecological and institutional dimensions. *Annual Review of Sociology, 26,* 395–418.

Attewell, P. (2001). The winner-take-all high school: Organizational adaptations to educational stratification. *Sociology of Education, 74,* 267–295.

Auld, M. C., & Sidhu, N. (2005). Schooling, cognitive ability and health. *Health Economics, 14,* 1019–1034.

Bader, M. D., Purciel, M., Yuosefzadeh, P., & Neckerman, K. M. (2010). Disparities in neighborhood food environments: Implications of measurement strategies. *Economic Geography, 86,* 409–430.

Baron, R. M., & Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology, 51,* 1173–1182.

Battistich, V., Solomon, D., Kim, D.-i., Watson, M., & Schaps, E. (1995). Schools as communities, poverty levels of student populations, and students' attitudes, motives, and performance: A multilevel analysis. *American Educational Research Journal, 32,* 627–658.

Burke, M. A., & Sass, T. R. (2011). *Classroom peer effects and student achievement* (Public Policy Discussion Paper No. 11-5). Boston, MA: Federal Reserve Bank of Boston.

Card, D., & Rothstein, J. (2007). Racial segregation and the black-white test score gap. *Journal of Public Economics, 91,* 2158–2184.

Carlson, D., & Cowen, J. M. (2015). Student neighborhoods, schools, and test score growth: Evidence from Milwaukee, Wisconsin. *Sociology of Education, 88,* 38–55.

Chetty, R., Friedman, J. N., & Rockoff, J. E. (2014). Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates. *American Economic Review, 104,* 2593–2632.

Chetty, R., Hendren, N., & Katz, L. F. (2016). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity project. *American Economic Review, 106,* 855–902.

Choi, K. H., Raley, R. K., Muller, C., & Riegle-Crumb, C. (2008). Class composition: Socioeconomic characteristics of coursemates and college enrollment. *Social Science Quarterly, 89,* 846–866.

Coleman, J. S., Campbell, E. Q., Hobson, C. J., McPartland, J., Mood, A. M., Weinfeld, F. D., & York, R. L. (1966). *Equality of educational opportunity* (National Center for Educational Statistics, Department of Health, Education, and Welfare, Report No. OE-38001). Washington, DC: U.S. Government Printing Office.

Cook, T. D., Herman, M. R., Phillips, M., & Settersten, R. A. (2002). Some ways in which neighborhoods, nuclear families, friendship groups, and schools jointly affect changes in early adolescent development. *Child Development, 73,* 1283–1309.

Crosnoe, R. (2009). Low-income students and the socioeconomic composition of public high schools. *American Sociological Review, 74,* 709–730.

Crowder, K., & Downey, L. (2010). Interneighborhood migration, race, and environmental hazards: Modeling microlevel processes of environmental inequality. *American Journal of Sociology, 115,* 1110–1149.

Cruse, C., & Powers, D. (2006). *Estimating school district poverty with free and reduced-price lunch data* (Research report of the Small Area Estimates Branch of the U.S. Census Bureau). Washington, DC: U.S. Census Bureau. Retrieved from https://www.census.gov/did/www/saipe/publications/files/CrusePowers2006asa.pdf

Cunha, F., & Heckman, J. J. (2007). The technology of skill formation. *American Economic Review, 97,* 31–47.

Downey, D. B., & Condron, D. J. (2016a). Two questions for sociologists: A rejoinder. *Sociology of Education, 89,* 234–235.

Downey, D. B., & Condron, D. J. (2016b). Fifty years since the Coleman report: Rethinking the relationship between schools and inequality. *Sociology of Education, 89,* 207–220.

Downey, D. B., von Hippel, P. T., & Broh, B. (2004). Are schools the great equalizer? Cognitive inequality during summer months and the school year. *American Sociological Review, 69,* 613–635.

Duncan, G. J., & Murnane, R. J. (2011). *Whither opportunity? Rising inequality, schools, and children's life chances*. New York, NY: Russell Sage.

Ferryman, K. S., de Souza Briggs, X., Popkin, S. J., & Rendón, M. (2008). *Do better neighborhoods for MTO families mean better schools?* (Brief No. 3 of the Metropolitan and Housing Communities Center). Washington, DC: Urban Institute.

Galster, G. C. (2012). The mechanism(s) of neighbourhood effects: Theory, evidence, and policy implications. In M. van Ham, D. Manley, N. Bailey, L. Simpson, & D. Maclenna (Eds.), *Neighbourhood effects research: New perspectives* (pp. 23–56). New York, NY: Springer.

GeoLytics, Inc. (2013). *Census CD neighborhood change database tract data from 1970–2010*. Retrieved from http://www.geolytics.com/USCensus,Neighborhood-Change-Database-1970-2000,Products.asp

Goldsmith, P. R. (2009). Schools or neighborhoods or both? Race and ethnic segregation and educational attainment. *Social Forces, 87,* 1913–1942.

Halpern-Manners, A. (2016). Measuring students' school context exposures: A trajectory-based approach. *Social Science Research, 58,* 135–149.

Hanushek, E. A. (2011). The economic value of higher teacher quality. *Economics of Education Review, 30,* 466–479.

Harding, D. J. (2003). Counterfactual models of neighborhood effects: The effect of neighborhood poverty on dropping out and teenage pregnancy. *American Journal of Sociology, 109,* 676–719.

Harris, D. N. (2010). How do school peers influence student educational outcomes? Theory and evidence from economics and other social sciences. *Teachers College Record, 112,* 1163–1197.

Hedges, L. V., Laine, R. D., & Greenwald, R. (1994). Does money matter? A meta-analysis of studies of the effects of differential school inputs on student outcomes. *Educational Researcher, 23,* 5–14.

Heuer, R., & Stullich, S. (2011). *Comparability of state and local expenditures among schools within districts: A report from the study of school-level expenditures*. Washington, DC: U.S. Department of Education.

Heyns, B. L. (1978). *Summer learning and the effects of schooling*. New York, NY: Academic Press.

Jackson, C. K. (2012). *Non-cognitive ability, test scores, and teacher quality: Evidence from 9th grade teachers in North Carolina* (NBER Working Paper No. 18624). Cambridge, MA: National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w18624

Jaffe, L. E. (2009). *Development, interpretation, and application of the W score and the relative proficiency index* (Woodcock-Johnson III Assessment Service Bulletin No. 11). Rolling Meadows, IL: Riverside.

Jencks, C., & Mayer, S. E. (1990). The social consequences of growing up in a poor neighborhood. In L. E. Lynn & M. G. H. McGreary (Eds.), *Inner-city poverty in the United States* (pp. 111–186). Washington, DC: National Academies Press.

Johnson, O., Jr. (2012). A systematic review of neighborhood and institutional relationships related to education. *Education and Urban Society, 44,* 477–511.

Kahlenberg, R. D. (2001). *All together now: Creating middle-class schools through public school choice*. Washington, DC: Brookings Institution Press.

Konstantopoulos, S., & Chung, V. (2009). What are the long-term effects of small classes on the achievement gap? Evidence from the lasting benefits study. *American Journal of Education, 116,* 125–154.

LaForte, E. M., McGrew, K. S., & Schrank, F. A. (2014). *WJ IV technical abstract* (Woodcock-Johnson IV Assessment Service Bulletin No. 2). Rolling Meadows, IL: Riverside.

Lauen, D. L., & Gaddis, S. M. (2013). Exposure to classroom poverty and test score achievement: Contextual effects of selection? *American Journal of Sociology, 118,* 943–979.
Leventhal, T., & Brooks-Gunn, J. (2000). The neighborhoods they live in: The effects of neighborhood residence on child and adolescent outcomes. *Psychological Bulletin, 126,* 309–337.
Levanthal, T., & Brooks-Gunn, J. (2004). A randomized study of neighborhood effects on low-income children's educational outcomes. *Developmental Psychology, 40,* 488–507.
Murnane, R. J., & Levy, F. (2006). *Teaching the new basic skills: Principles for educating children to thrive in a changing economy.* New York, NY: Free Press.
National Center for Education Statistics. (2014a). *State support for school choice and other options* (Tables 4.2–4.4). Washington, DC: U.S. Department of Education. Retrieved from http://nces.ed.gov/programs/statereform/sss.asp
National Center for Education Statistics. (2014b). *Common core of data: School and district-level datasets*. Washington, DC: U.S. Department of Education. Retrieved from https://nces.ed.gov/ccd/ccddata.asp
National Center for Education Statistics. (2014c). *Private school universe survey datasets*. Washington, DC: U.S. Department of Education. Retrieved from https://nces.ed.gov/surveys/pss/pssdata.asp
National Center for Education Statistics. (2015). *Revenues and expenditures for public elementary and secondary school districts: School year 2011–12*. Washington, DC: U.S. Department of Education. Retrieved from https://nces.ed.gov/pubsearch/pubsinfo.asp?pubid=2014301
Oreopoulos, P. (2012). Moving neighbourhoods versus reforming schools. *Cityscape, 14*(2), 207–212.
Owens, A. (2010). Neighborhoods and schools as competing and reinforcing contexts for educational attainment. *Sociology of Education, 83,* 287–311.
Panel Study of Income Dynamics (PSID). (2014). *Public- and restricted-use datasets.* Ann Arbor: University of Michigan, Institute for Social Research, Survey Research Center.
Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, MA: Cambridge University Press.
Pfeffermann, D. (1993). The role of sampling weights when modeling survey data. *International Statistical Review, 61,* 317–337.
Raudenbush, S. W., Jean, M., & Art, E. (2011). Year-by-year and cumulative impacts of attending a high-mobility elementary school on children's mathematics achievement in Chicago, 1995 to 2005. In G. J. Duncan & R. J. Murnane (Eds.), *Whither opportunity? Rising inequality, schools, and children's life chances* (pp. 359–376). New York, NY: Russell Sage.
Rendón, M. G. (2014). Drop out and "disconnected" young adults: Examining the impact of neighborhood and school contexts. *Urban Review, 46,* 169–196.
Rivkin, S. G., Hanushek, E. A., & Kain, J. F. (2005). Teachers, schools, and academic achievement. *Econometrica, 73,* 417–458.
Rosenbaum, J. E. (1995). Changing the geography of opportunity by expanding residential choice: Lessons from the Gautreaux program. *Housing Policy Debate, 6,* 231–269.
Rosenfeld, L., Rudd, R., Chew, G. L., Emmons, K., & Acevedo-Garcia, D. (2010). Are neighborhood-level characteristics associated with indoor allergens in the household? *Journal of Asthma, 47,* 66–75.
Rowan, B., Correnti, R., & Miller, R. (2002). What large-scale survey research tells us about teacher effects on student achievement: Insights from the Prospects Study of Elementary Schools. *Teachers College Record, 104,* 1525–1567.
Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology, 66,* 688–701.
Rubin, D. B. (1987). *Multiple imputation for nonresponse in surveys*. New York, NY: John Wiley & Sons.
Rumberger, R., & Palardy, G. (2005). Does resegregation matter? The impact of social composition on academic achievement in southern high schools. In J. C. Boger & G. Orfield (Eds.), *School resegregation: Must the south turn back?* (pp. 127–147). Chapel Hill: University of North Carolina Press.
Sampson, R. J. (2001). How do communities undergird or undermine human development? Relevant contexts and social mechanisms. In A. Booth & A. C. Crouter (Eds.), *Does it take a village? Community effects on children, adolescents, and families* (pp. 3–30). Mahwah, NJ: Lawrence Erlbaum.
Sampson, R. J., Morenoff, J. D., & Gannon-Rowley, T. (2002). Assessing "neighborhood effects": Social processes and new directions in research. *Annual Review of Sociology, 28,* 443–478.
Sampson, R. J., Raudenbush, S. W., & Earls, F. (1997). Neighbourhoods and violent crime: A multilevel study of collective efficacy. *Science, 277,* 918–924.
Sanbonmatsu, L., Kling, J. R., Duncan, G. J., & Brooks-Gunn, J. (2006). Neighborhoods and academic achievement: Results from the Moving to Opportunity Experiment. *Journal of Human Resources, 41,* 649–691.
Saporito, S., & Sohoni, D. (2007). Mapping educational inequality: Concentrations of poverty among poor and minority students in public schools. *Social Forces, 85,* 1227–1253.

Schellenberg, S. J. (1999). Concentration of poverty and the ongoing need for Title 1. In G. Orfield & E. H. DeBray (Eds.), *Hard work for good schools: Facts not fads in Title I reform* (pp. 130–146). Cambridge, MA: Harvard University Civil Rights Project.

Sharkey, P. T. (2010). The acute effect of local homicides on children's cognitive performance. *Proceedings of the National Academy of Sciences, 107,* 11733–11738.

Sharkey, P. T., & Faber, J. W. (2014). Where, when, why, and for whom do residential contexts matter? Moving away from the dichotomous understanding of neighborhood effects. *Annual Review of Sociology, 40,* 559–579.

Sharkey, P. T., Tirado-Strayer, N., Papachristos, A. V., & Raver, C. (2012). The effect of local violence of children's attention and impulse control. *American Journal of Public Health, 102,* 2287–2293.

Steinberg, L. (1997). *Beyond the classroom: Why school reform has failed and what parents need to do*. New York, NY: Simon & Schuster.

VanderWeele, T. J. (2015). *Explanation in causal inference: Methods for mediation and interaction*. New York, NY: Oxford University Press.

Weiss, C. C., Purciel, M., Bader, M., Quinn, J. W., Lovasi, G., Neckerman, K. M., & Rundle, A. G. (2011). Reconsidering access: Park facilities and neighborhood disamenities in New York City. *Journal of Urban Health, 88,* 297–310.

Willms, J. D. (1986). Social class segregation and its relationship to pupils' examination results in Scotland. *American Sociological Review, 51,* 224–241.

Willms, J. D. (2010). School composition and contextual effects on student outcomes. *Teachers College Record, 112,* 1008–1038.

Wilson, W. J. (1987). *The truly disadvantaged: The inner city, the underclass, and public policy*. Chicago, IL: University of Chicago Press.

Winship, C., & Radbill, L. (1994). Sampling weights and regression analysis. *Sociological Methods & Research, 23,* 230–257.

Wodtke, G. T. (2013). Duration and timing of exposure to neighborhood poverty and the risk of adolescent parenthood. *Demography, 50,* 1765–1788.

Wodtke, G. T., Elwert, F., & Harding, D. J. (2016). Neighborhood effect heterogeneity by family income and developmental period. *American Journal of Sociology, 121,* 1168–1222.

Wodtke, G. T., Harding, D. J., & Elwert, F. (2011). Neighborhood effects in temporal perspective: The impact of long-term exposure to concentrated disadvantage on high school graduation. *American Sociological Review, 76,* 713–736.

Woodcock, R. W., & Johnson, M. E. B. (1989). *Tests of achievement, standard battery* (Form B). Chicago, IL: Riverside.