Discussion Paper Series – CRC TR 224

Discussion Paper No. 618
Project A 04

Neighborhood Exposure Effects in Cognitive Skills and the Role of Primary Schools

Xi Lin[1]

January 2025

[1]University of Mannheim, Email: xi.lin@uni-mannheim.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

Neighborhood Exposure Effects in Cognitive Skills and the Role of Primary Schools

Xi Lin*

October 2024

Abstract

This study examines how childhood residential location affects cognitive skills, focusing on the roles of neighborhood and primary school quality in shaping children's school performance. Using administrative data from the Netherlands, I estimate the causal effect of neighborhood exposure—defined as the impact of time spent in a neighborhood—on children's test scores at the end of their primary education. By comparing children who move at different ages, I separate the effects of exposure from those of sorting into neighborhoods. The results show that for each additional year a child spends in a neighborhood with higher expected test scores, their test scores improve by approximately 2.5% relative to the total gap between the lower- and higher-performing neighborhoods. As families can choose primary schools without geographical restrictions in the Netherlands, I can further isolate improvements attributable to school quality. Approximately 40% of the observed improvements in test scores can be explained by differences in primary school quality. These findings highlight the critical roles of neighborhood environments and school quality in reducing spatial educational inequalities.

*Lin: University of Mannheim; xi.lin@uni-mannheim.de. I thank Antonio Ciccone, Katja Maria Kaufmann, and participants of the CRC TR 224 Summer School on the Econometrics of Peer Effects and Social Interactions, the Summer School on the Development of Cognitive and Non-Cognitive Skills in Childhood and Adolescence, the Young Researchers Workshops of the CRC TR 224, and seminars at the University of Mannheim for helpful comments and feedback. Financial support by the German Research Foundation (DFG) through CRC TR 224 (project A04) is gratefully acknowledged.

1 Introduction

Recent experimental and quasi-experimental studies provide evidence that the neighborhood in which a child grows up is associated with long-term differences in educational and labor market outcomes (Chetty and Hendren, 2018a; Deutscher, 2020; Chetty et al., 2020; Laliberté, 2021). However, the mechanisms through which neighborhoods influence these outcomes remain less well understood. Emerging research suggests that neighborhoods exert influence through both contemporaneous (situational) effects shaped by the current environment and developmental (exposure) effects that accumulate over time[1]. Further investigation into these mechanisms is essential for informing policies aimed at improving childhood environments and expanding opportunities for disadvantaged populations.

In this study, I investigate whether neighborhood exposure—defined as the effect of time spent in a specific neighborhood—affects children's cognitive skills development, particularly their school performance, as measured by standardized test scores at the end of primary education. Additionally, I explore how much of the observed neighborhood exposure effects can be attributed to differences in school quality. By isolating the role of school quality, I aim to determine whether improvements in test scores are primarily driven by the neighborhood environment itself or by the quality of the schools children attend within those neighborhoods.

My empirical analysis draws on detailed administrative data from the Netherlands that offer several advantages for studying these issues. First, the data include standardized measures of academic performance at the end of primary school and long-term outcomes, enabling consistent comparisons across regions. Second, the dataset tracks children's residential histories from birth, allowing for precise measurement of neighborhood exposure over time. Third, the Dutch system of free school choice decouples residential location from school attendance, facilitating a clear separation of neighborhood and school quality effects.

In my main analysis, I define neighborhood at the municipality level, the lowest tier of government in the Netherlands. I begin by documenting the variation in school performance across municipalities. The results show substantial differences in standardized test

[1]See recent review by Chyn and Katz (2021).

scores between municipalities, even after accounting for family background. For instance, among children with parental income at the 25th percentile of the national income distribution, the difference in expected test scores between the highest- and lowest-performing municipalities can be as large as 15 percentage points in the national test score ranking.

To further explore the relationship between neighborhood exposure and school performance, I apply a *mover design*, following the framework developed by Chetty and Hendren (2018a). This approach compares children who move between neighborhoods at different ages. It allows me to estimate how much their academic outcomes converge toward those of children who have always lived in the destination neighborhood (i.e., permanent residents). The results suggest that moving to a neighborhood with higher expected school performance is associated with gradual improvements in own school performance over time: for each additional year of exposure, children close the gap in end-of-primary school performance by approximately 2.5%, relative to the difference between lower- and higher-performing neighborhoods. This identification strategy assumes that selection effects related to moving into neighborhoods with different school performance levels do not vary systematically with the child's age at the time of the move. I control for family fixed effects and examine subject-specific convergence patterns to test this assumption. The results of these robustness checks are consistent with my main findings.

I then examine the granularity of neighborhood dynamics. I replicate the baseline analysis, restricting the sample to moves at a finer geographic level—specifically, *buurten*, the smallest geographical unit used by the Statistical Bureau Netherlands with populations ranging from 1,000 to 5,000 residents. This allows for a more granular analysis of neighborhood effects. The results indicate that neighborhood effects are primarily localized. I also find that only the closest *buurten* have a significant impact on children's outcomes, with effects weakening as distance from the child's home increases. My findings suggest that policies aiming to improve educational outcomes may need to target very local areas to be most effective.

Finally, the institutional context of the Dutch education system, which has a free school choice policy(Patrinos, 2011), allows me to decompose the total neighborhood exposure effects into components attributable to school quality, neighborhood quality, and family quality. The results suggest that school quaöolity accounts for around 40% of the observed variation in test scores due to neighborhood exposure. While these findings highlight the potential importance of school quality, other neighborhood characteristics also appear to play a role in shaping children's academic outcomes.

This paper contributes to two strands of literature: the impact of neighborhoods on child development and the role of schools in mediating this impact. First, this study adds to the literature on neighborhood effects on children's outcomes. Chetty and Hendren (2018a) and Chetty and Hendren (2018b) document the long-term impacts of neighborhood environments on children's economic and educational trajectories, showing that the duration of exposure to better neighborhoods-measured using outcomes of the permanent residents- is an important determinant of future success. Chyn and Katz (2021) review this literature, emphasizing the importance of understanding how neighborhood effects translate into different outcomes. Moreover, studies such as those by Guryan et al. (2021) and Wodtke (2018) show that neighborhood characteristics, such as access to resources and exposure to violence, influence cognitive development. At the same time, Jackson (2020) and Rossin-Slater (2018) highlight the role of early childhood environments in shaping long-term achievement. Building on this literature, my paper provides new evidence from the Netherlands by analyzing how neighborhood exposure affects school performance, using data from a national standardized test administered to nearly all pupils. This earlier measure of academic outcomes helps bridge the gap between childhood exposure and longer-term educational outcomes, while the use of large-scale administrative data allows for a comprehensive analysis. Parallel to my study, Webbink et al. (2023) examines neighborhood exposure effects on income at age 30 in the same Dutch context, noting that defining movers based on parental addresses can introduce measurement errors, especially for older children (ages 16-24). In my study, which focuses on children up to age 16, this issue is less pronounced, as most children still live and move with their parents in the Netherlands. However, I exclude cases where children do not move with their parents to minimize potential bias.

Second, this paper contributes to the literature on the role of schools in mediating neighborhood effects. While extensive literature has explored the effects of school and neighborhood separately, studies combining these two contexts are still underdeveloped. Sykes and Musterd (2011) uses Dutch data to examine the correlation between the characteristics of neighborhoods and schools and educational outcomes. Unlike their approach, I use outcome-based measures of neighborhood and school quality, which allows me to avoid the issue of selecting specific observable characteristics to proxy for quality. Card et al. (2018) uses state- and county-level data from the early 20th century, showing that variations in school quality were key drivers of regional differences in upward mobility. Similarly, Rothstein (2019) examines the impact of K-12 school quality on intergenerational mobility using aggregate data at the commuting zone (CZ) level across U.S. cities. His findings suggest that school quality explains only a small portion of the variation in intergenerational income mobility across regions, with broader factors like neighborhood characteristics and economic conditions playing a more substantial role. In comparison, my use of micro-level data allows for a more detailed analysis by linking childhood environments directly to later educational outcomes, offering more precision in identifying the specific effects of school and neighborhood quality. Consistent with the findings of Gibbons and Silva (2018) and Aizer and Currie (2019), I find that school quality is an important determinant of later education attainment. My study also provides evidence that variation in test scores due to neighborhood differences can be attributed to differences in school quality.

A directly related study is Laliberté (2021), which developed an innovative approach to decomposing neighborhood exposure effects into the portion attributable to school quality and the portion due to non-school neighborhood factors. Set in Montreal, Canada, the study leverages the special institutional context, where the primary school catchment area boundaries for French and English schools do not perfectly overlap, to disentangle the long-term effects of school quality and neighborhood characteristics on educational outcomes. The findings reveal that 50 to 70 percent of the total effect of living in a better neighborhood on educational attainment is attributable to school quality, with the remaining portion tied to non-school neighborhood characteristics.

While my study build on the decomposition approach developed in Laliberté (2021), my study features several differences. First, the primary school system in the Netherlands has no formal school catchment areas, meaning parents are free to send their children to the school of their choice (Patrinos, 2011). This unique feature disentangles school choice from residential location choice, allowing for a clearer separation of school and neighborhood effects. Second, unlike in Laliberté (2021), where children with different mother tongues might face cultural barriers, children attending different schools in the Netherlands typically encounter no language-based barriers with their neighbors. This situation is more common globally and may be more representative. This difference could explain why, in my study, the neighborhood share of the effect is higher than in Laliberté (2021)'s findings. Third, given that the government equally funds both private and public primary schools in the Netherlands, and all school enrollments are well documented, I can examine the entire population of the Netherlands, capturing a more diverse cross-section of the population compared to the sample of public schools from Montreal Island used in Laliberté (2021). Lastly, although I could not employ a boundary discontinuity design, observing family income and schooling from the administrative data allows me to better proxy for family environment and separate those effects from school and neighborhood quality.

The remainder of this paper proceeds as follows. Section 2 provides an overview of the Dutch education system and details the data used in the study. Section 3 outlines the empirical strategy and presents the main results. Section 4 conducts a mediation analysis to disentangle the effects of school quality from other neighborhood factors. Section 5 concludes with policy implications.

2 Data and Institutional Background

2.1 Data Source and Sample Selection

I use a comprehensive administrative database from the Netherlands, managed by Statistics Netherlands (Centraal Bureau voor de Statistiek, CBS).[2] This database covers the country's entire population and provides detailed information that enables tracking individuals both geographically and over time. Different datasets in the database can be linked through individual random identification numbers.

The selection of birth cohorts is based on the availability of three primary datasets: the municipal registry (1995–2022), the CITO test scores at the end of primary education (2006–2019), and education enrollment records (2003–2022). I restrict the sample to individuals whose moving history is fully observable from age 1, including birth cohorts from 1994 onward. Due to the cancellation of the CITO test in 2019 during the pandemic,

[2]The database is accessible via a remote-access computer after a confidentiality statement has been signed.

children born after 2007 are excluded, as some of them did not take the test as they otherwise would have.

I assess educational attainment at age 24, by which time nearly all individuals in the Netherlands have typically completed or are concluding their education. This age selection aligns with previous studies that evaluate educational outcomes at entry into the labor market. Given that 2022 is the last available year for education enrollment records, the latest birth cohort for which I can observe educational attainment at age 24 is 1998.

The main sample comprises all children who meet the following criteria: (i) inclusion in the municipal population register from birth, (ii) birth between 1994 and 2007, (iii) at least one parent is identifiable, and (iv) their CITO test scores are available for the years 2006 to 2019. Every individual who has resided in the Netherlands since 1995 is included in the municipal population register. Undocumented immigrants and asylum seekers, who are typically not registered, are the primary groups excluded from the sample. Children are linked to parents using data on legal parent(s) from the municipal population registers, and only those with at least one identifiable legal parent are included in the sample.

2.2 Neighborhoods and Movers

I determine the neighborhoods of parents and children using their home addresses registered in the municipal population register, which are continuously updated based on administrative data from schools and social security agencies. The municipal population register is the government's primary means for communicating with citizens on various issues, including taxes, income, and social security matters. The Dutch data uniquely enable identifying the addresses of children and those of their parents separately, allowing for precise analysis. I code relocations based on the child's address in the baseline analysis. When individuals relocate, they must notify the municipal administration of their old and new addresses, along with the exact date of the move. Using birth date data, I can calculate the age at the time of the move with day-level precision.

The geographic information available for this study is exceptionally detailed, allowing observations of an individual's address down to the specific building and enabling precise calculations of moving distances. To define relevant neighborhoods, I use two levels of granularity. My baseline analysis focuses on municipalities. As of 2011, the Netherlands has 419 municipalities with an average population of approximately 40,000. For comparison, the study of Chetty and Hendren (2018a) is at the level of 741 commuting zones, averaging approximately 380,000 inhabitants. Municipalities are crucial in administering various local services, including education, as they implement national education laws and regulations at the local level. This administration includes overseeing the establishment and maintenance of schools, ensuring compliance with educational standards, and supporting special education needs.

For a more detailed analysis of local interactions, I define neighborhoods at a smaller scale, specifically, *buurten* in Dutch. A *buurt* is the smallest geographical unit used by the Statistical Bureau Netherlands, typically corresponding to a well-defined area within a city or town. There are approximately 11,000 *buurten* across the country with populations ranging from 1,000 to 5,000 residents. The number of schools per *buurt* varies, with urban areas having more schools, while many rural *buurten* possibly having none. The number of primary school-aged children per *buurt* ranges from 100 to 300. This level of granularity allows for a more precise examination of local dynamics and social interactions, which is crucial for studying the impact of localized factors on educational outcomes.

I divide the sample into permanent residents (or stayers) and movers. Permanent residents in each neighborhood are defined as the subset of children born and residing in a single neighborhood up to age 15. Movers are non-permanent residents, with one-time movers defined as those who move exactly once with their parents across neighborhoods before the age of 15[3]. The main sample has approximately 1.7 million children for whom I observe CITO scores, including about 1.3 million permanent residents and 295,000 one-time movers.

2.3 Institutional Setting

Education in the Netherlands is compulsory from age 5 to 16, with many children beginning at age 4. The system is split into primary, secondary, and tertiary education.

[3] In a recent study, Webbink et al. (2023) argue that including parental moves without children could introduce measurement error in this type of design. Therefore, I only consider cases where the addresses of both children and parents change simultaneously when defining movers.

Primary education in the Netherlands is characterized by a unique system of free school choice: families are not limited by residential catchment areas when selecting schools. Parents can enroll their children in any primary school—public, private, religious, or special education—regardless of their geographic location. All schools must meet the same national educational standards and receive government subsidies to ensure equitable funding across different types of schools. The central government provides the majority of financial resources for schools, including funding for teacher salaries, instructional materials, and student support services. This centralized funding model is designed to ensure that all schools, irrespective of their type or location, can deliver high-quality education and adhere to the national curriculum guidelines. It highlights the freedom of choice, the financial equality across school types, and the role of government subsidies in ensuring consistent educational standards across all schools in the Netherlands.

This structure contrasts with the U.S. context studied by Chetty and Hendren (2018a), where educational opportunities, especially those provided by the states, are more directly tied to the neighborhood in which a child resides. In the U.S., school catchment areas typically restrict families to public schools within their residential zone (e.g. Epple and Romano, 2003), making it difficult to disentangle place effects from school effects. In contrast, the Dutch system's freedom of choice allows parents to bypass geographic limitations, offering a unique opportunity to identify the separate effects of neighborhoods and schools on educational outcomes.

A structured and diverse secondary education system complements the freedom of primary school choice in the Dutch primary system. After primary school, students are tracked into one of three main paths—VMBO (pre-vocational), HAVO (senior general), or VWO (pre-university)—based on their performance. These tracks cater to students' varying abilities and aspirations, preparing them for tertiary education in vocational training (MBO) or at universities of applied sciences (HBO) or research universities (WO).

MBO offers vocational education and training at four levels, preparing students for the labor market or further studies. Students are typically directed into one of these four pathways after completing secondary education, following the same track as in their secondary schooling. HBO institutions focus on professionally oriented programs, while WO institutions emphasize academic and research-based education. Both HBO and WO offer bachelor's and master's degrees, with the possibility of advancing to doctoral programs at research universities.

The entry into tertiary education is highly structured, with most students beginning their bachelor's programs immediately after secondary school. Most students complete their education efficiently, with many reaching the final stage of their studies or having already finished their degrees by 24 (OECD, 2023; European-Commission, 2019).

2.4 Variable Definitions

In this section, I define the key variables used in the analysis: *parental income*, *CITO test scores*, and *educational attainment*.

Parental Income. Parental income is defined as the sum of the disposable incomes of the father and mother when the child is between 9 and 12. If parents are separated, parental income is calculated as the average of the mother's disposable income and that of her spouse. If this information is unavailable, the income is based on the mother's disposable income alone. The income variable is top-coded at €1 million. Following Chetty and Hendren (2018a), I convert incomes into percentile ranks relative to the national distribution for the child's birth cohort. This method improves comparability across individuals and reduces the influence of outliers and lifecycle income variability.

CITO Test Scores. Children's academic performance at the end of their primary education is measured using the so-called CITO test. In the Netherlands, primary education consists of six grade levels, and children typically complete their primary education at age 12. Schools can choose the provider of the end-of-primary education test they administer. Approximately 85% schools opt for the CITO test, renowned for its comprehensive assessment of key subjects such as mathematics, Dutch language, and study skills. Participation is mandatory for all enrolled students once a school decides to administer the CITO test. End-of-primary education test scores are high-stakes as teachers consider them in making recommendations for each student's secondary school track. The CITO scores range from 500 to 550 and are standardized across years to ensure consistency. As with parental income, I convert the raw test scores into percentile ranks within the child's national birth cohort. This approach allows for a more nuanced comparison of relative academic performance across regions and periods. Percentile ranks are calculated both overall and by subject area.

Educational Attainment. Educational attainment is measured use the highest degree or academic qualification obtained by the age of 24. Degrees are then converted into years of schooling, following the International Standard Classification of Education (ISCED) system. This conversion ensures a consistent and comparable measure of educational achievement across individuals, facilitating a robust analysis of how early-life factors affect long-term educational outcomes.

2.5 Summary Statistics

Table 1 lists the variables included in the study and provides summary statistics for the primary analysis samples—which consist of birth cohorts from 1994 to 2007—and presents these statistics for the entire population and by moving status.

Consistent with previous studies, I find that movers and permanent residents have similar characteristics. For instance, the median family disposable income is nearly €43,000 for both groups. The same pattern applies to children's educational attainment and test scores, which are comparable between both groups.

3 Identifying Childhood Exposure Effects

I use the empirical framework of Chetty and Hendren (2018a) to estimate how childhood exposure affects school performance. This method involves two key steps. First, I use the outcomes of permanent residents to predict the expected outcomes for children growing up in various neighborhoods. In the second step, I focus on children who moved once during childhood. I analyze how moving one year earlier influences children's outcomes by estimating how the expected outcomes from those in the origin neighborhood converge to those in the destination neighborhood. Specifically, the model captures the rate at which a child's outcomes converge toward the outcomes of permanent residents in the destination area, with each additional year spent in the destination.[4]

[4]For a comprehensive and formal introduction to this identification strategy, see Chetty and Hendren (2018a).

3.1 Step 1: Estimating Neighborhood-Level Predicted Outcomes

In the first step, I generate predicted outcomes for children growing up in different neighborhoods using data from permanent residents. To do this, I estimate the relationship between parental income ranks and children's CITO scores within each municipality. This relationship is captured by the following linear regression model:

$$CITO_i = \alpha_c + \pi_c p_i + \epsilon_i, \tag{1}$$

where $CITO_i$ denotes the child's CITO score, p_i is the parental income rank, α_c represents the neighborhood (municipality) fixed effect, and π_c measures the effect of parental income rank p_i on child outcomes within each neighborhood. The error term ϵ_i accounts for unobserved factors that may influence a child's outcomes.

In my baseline specification, I focus on the effects of parental income rank, omitting variation by birth cohorts, unlike Chetty and Hendren (2018a). This is because, unlike other educational systems that group students by birth year, school entry in the Netherlands is determined by developmental readiness, making cohort distinctions less relevant. In my robustness analysis, however, I account for neighborhood predicted outcomes within birth calendar years to ensure the robustness of my findings. Consistent with Chetty and Hendren (2018a), I use rank-based measures to avoid issues related to attenuation and life-cycle bias.

After estimating the model, I calculate two key predicted outcomes for each child in the sample of movers: $\bar{y}_{op}$ denotes the predicted outcome if the child grew up entirely in their origin neighborhood. And $\bar{y}_{dp}$ denotes he predicted outcome if the child grew up entirely in their destination neighborhood. These predictions will later be used to compute the "expected gains" from moving between neighborhoods by tracking the change from outcomes expected in the origin neighborhood to those expected in the destination neighborhood.

Figure 1 maps children's test scores by municipalities for children whose parents are at the 25th income percentile. Children's outcomes vary significantly across municipalities. For example, among children with parents at the 25th percentile, CITO scores are approximately 15 percentage points higher in municipalities at the top (95th percentile) of the mean test score distribution than those at the bottom (5th percentile).

3.2 Step 2: Estimating the Impact of Neighborhood Moves on Child Outcomes

In the second step, I use the sample of movers to assess how the age at which children move between neighborhoods affects their outcomes. Specifically, I test whether children who move at younger ages demonstrate more of the predicted difference between their origin and destination neighborhoods. To quantify this, I estimate a child's eventual CITO score as a function of two key factors: their predicted outcome in the origin neighborhood ($\bar{y}_{op}$) and the predicted change in outcomes resulting from the move to the destination neighborhood ($\Delta odp = \bar{y}_{dp} - \bar{y}_{op}$), interacted with the child's age at the time of the move (m).

Consistent with Chetty and Hendren (2018a), the model is estimated using a semi-parametric approach, represented by the following specification:

$$CITO_i = \sum_{m=0}^{15} [\alpha_m + \phi_m p_i + \zeta_m \bar{y}_{op} + b_m \Delta odp] + \epsilon_i, \tag{2}$$

where $CITO_i$ is the eventual CITO score of child i, α_m is an intercept that varies with the age at which the child moves, ϕ_m is the age-specific effect of parental income rank p_i, and ζ_m is the age-specific coefficient on the predicted outcome in the origin neighborhood ($\bar{y}_{op}$).

The key parameters of interest are the b_m coefficients, which measure how much of the predicted change in outcomes from moving to a new neighborhood (Δodp) is realized by children moving at different ages. In other words, these coefficients capture the degree to which children's outcomes shift toward the outcomes predicted for their destination neighborhood, depending on their age at the time of the move. The differences between these coefficients—such as $b_m - b_{m+1}$—are interpreted as the impact of additional exposure to a neighborhood with higher predicted outcomes.

The rate of convergence to the outcomes of permanent residents in the destination neighborhood is used to measure the effects of neighborhood exposure. In this context, the model identifies how much of the difference between origin and destination outcomes is absorbed by children, depending on when the move occurs.

3.3 Identification Assumptions

The coefficient b_m captures both selection and exposure effects. Selection effects arise because families who move to better neighborhoods may differ systematically from those who stay, even after controlling for observed socio-economic factors such as income or education level. These differences may be driven by unobserved characteristics, such as preferences for education or aspirations for their children, which could bias the estimated effects of the neighborhood. In contrast, exposure effects reflect the causal impact of the time spent in a new neighborhood on a child's educational outcomes.

To separate the exposure effect from the selection effect, I make two key identification assumptions following Chetty and Hendren (2018a): **Constant Selection Effects** assumes that selection effects do not systematically vary with the child's age at the time of the move. In other words, the characteristics driving a family's decision to move are assumed to be constant across different ages. This allows me to attribute differences in b_m for children who move at different ages primarily to differences in exposure time rather than unobserved differences in family characteristics. **Linearity of Exposure Effects** suggests that the effect of neighborhood exposure grows linearly with the time a child spends in the destination neighborhood. Under this assumption, each additional year spent in the new neighborhood contributes equally to the child's outcomes. The improvement in outcomes from moving one year earlier can then be interpreted as a constant exposure effect.

Given the first assumption, the difference between b_m values for children who move at different ages can be interpreted as the causal effect of neighborhood exposure. The difference $b_{m+1} - b_m$ reflects the annual exposure effect, which measures how much a child's outcomes improve with each additional year spent in the new neighborhood. This framework allows me to estimate how outcomes converge toward those of permanent residents in the destination neighborhood over time.

Assuming exposure effects are linear, as posited in the second assumption, I can further parametrize equation (2) as follows:

$$\begin{aligned}
CITO_i = &\sum_{m=0}^{15} I(m_i = m)[\alpha_m + \phi_m \hat{T}_o + \zeta_m p_i] \\
&+ K(m_i \leq 12)(\gamma' + \gamma(12 - m_i))\Delta_{odp} \\
&+ C(m_i > 12)(\rho' + \rho(m_i - 12))\Delta_{odp} + \epsilon_i. \qquad (3)
\end{aligned}$$

In this specification, the coefficient γ represents the annual exposure effect for children who move at age 12 or earlier. It captures the average effect of moving one year earlier to a neighborhood where permanent residents score one percentile higher on educational outcomes. Similarly, the coefficient ρ measures the corresponding slope for children who move after age 12, capturing how exposure to a new neighborhood affects older children's outcomes.

By estimating these coefficients, I can calculate the annual exposure effect for different age groups and infer the overall impact of neighborhood exposure on long-term outcomes. This approach allows me to introduce additional controls such as family fixed effects.

Figure 2 displays estimates of b_m from Equation 2, revealing two main patterns: selection effects after age 12 and exposure effects before age 12. The positive values of b_m for ages $m > 12$ clearly indicate selection effects, as moves after age 12 cannot causally affect CITO test scores, which are obtained at age 12. This finding suggests that children in families who move to better neighborhoods often have favorable unobservable attributes. Furthermore, the degree of selection remains constant across ages above 12, as evidenced by a statistically insignificant slope of 0.001 when regressing b_m on m. This stability aligns with the assumption that selection does not significantly vary based on the child's age at the time of moving.

Figure 2 also shows a steady decline in b_m estimates with the age at move (m) for $m < 12$. According to the first assumption, this downward trend provides evidence of exposure effects, meaning that moving to a better neighborhood earlier in childhood results in greater benefits. The linear relationship between b_m and the age at move (m) for ages below 12 suggests that the exposure effect remains relatively constant across different ages. A regression of b_m on m for ages below 12 estimates an average annual exposure effect of 0.025, indicating that children's outcomes improve and align with those of permanent residents at a rate of 2.5% per year of exposure up to age 12.

3.4 Validation of Identification Assumptions

One key threat to identification is the possibility that families who move at different times do so for different unobservable reasons. For instance, families moving when their children are young may prioritize long-term educational opportunities, while those moving when their children are older might do so for reasons related to employment or financial constraints. This would violate the first Assumption. The result could bias estimates of the neighborhood effects, as variation in b_m across ages may capture differences in family characteristics and not just differences in exposure time.

To address the concern of family selection effects, I incorporate family fixed effects into the model following the approach of Chetty and Hendren (2018a). This approach allows me to compare siblings who moved at different ages but share the same (observed and unobserved) family background, including the learning environment in the family and the educational preferences and aspirations of the family. Put differently, by examining within-family variation, I can isolate the effects of neighborhood exposure from family-specific selection effects. The key insight is whether the difference in school outcomes between siblings is proportional to their difference in exposure time, once family-level confounders are held constant.

However, time-varying factors, such as a parent's new job or changes in household financial circumstances coinciding with the move, could still introduce bias. To account for this, I also implement outcome-based placebo tests designed to test whether neighborhood exposure specifically affects the outcome of interest rather than unrelated outcomes. The detailed Dutch data allows me to examine subject-specific test scores, providing a novel test of whether the neighborhood's advantage in a particular subject, such as math, translates into stronger gains in that same subject for children exposed to the neighborhood. If the causal model holds, a child's math test scores should correlate more strongly with the neighborhood's math performance than unrelated subjects like the Dutch language. This approach ensures that the observed effects are specific to the neighborhood's influence rather than being driven by unobserved time-varying factors or general improvements unrelated to subject-specific advantages.

I implement these robustness tests in Table 2. Column 1 presents estimates of the average annual exposure effect of 3%, which are robust across various specifications and outcome definitions. Column 2 shows that controlling for maternal education and immigration background yields similar results. Columns 3 and 4 show that even when analyzing the data by subject, the convergence effects persist. Column 5 provides estimates where permanent residents' outcomes are measured based on their birth year. Column 6 shows results using variation in age at moves within families when parents relocate. Across these various robustness tests, the results remain consistent, reinforcing the robustness of the exposure effect estimates.

3.5 Relevance of School Performance for Long-Term Educational Outcomes

The findings demonstrate substantial effects of childhood exposure to higher-quality neighborhoods on school performance, as measured by CITO test scores. The remaining question is whether these early improvements in school performance translate into better long-term educational outcomes, such as years of schooling or overall educational attainment. This issue is examined in detail in *Appendix A*, where the relevance of school performance at age 12 for predicting educational attainment at age 24 is analyzed. The methodology builds on Rothstein (2019), who explores how schools mediate the intergenerational transmission of income, focusing on educational outcomes as a mediator for long-term socioeconomic mobility. Compared to Rothstein (2019), who decomposes income variation at the commuting zone level, a key advantage of the Dutch data is the ability to link end-of-primary-school performance to educational attainment at age 24 directly.

The mediation analysis in *Appendix A* decomposes the total effect of parental income on children's long-term educational attainment into direct and indirect effects. The indirect effects focus on primary school performance as a mediator, allowing us to assess whether improvements in CITO test scores—driven by exposure to higher-quality neighborhoods—significantly contribute to educational attainment, measured in years of schooling. This analysis clarifies how school performance is a crucial mechanism linking neighborhood quality to children's long-term educational trajectories.

The results suggest that differences in school performance account for approximately

40% of the variation in educational attainment across neighborhoods. This highlights that improvements in CITO test scores, resulting from better neighborhood environments, have significant implications for long-term outcomes. The role of schools in my context appears larger than in Rothstein (2019) (He estimated skills mediate 11% of the spatial income variation). This may be due to two reasons: First, my analysis focuses on primary school performance and its impact on educational attainment, whereas Rothstein (2019) examines income as the outcome. Second, the school tracking system in the Netherlands makes primary school performance more decisive for future educational success, as early school performance determines secondary school placement, shaping long-term educational trajectories. Therefore, improvements in primary school performance have more immediate and far-reaching consequences within the Dutch educational system.

4 Mechanisms

To determine the mechanisms underlying my findings, I identify neighborhood effects more granularly by evaluating the impact of moves across *buurten*, the smallest administrative units in the Netherlands. I then explore the spatial decay of these effects, investigating how geographic proximity influences neighborhood outcomes. Finally, I examine the role of educational institutions in explaining neighborhood effects.

4.1 Buurt-Level Exposure Effects

I first replicate the exposure effects analysis from Section 3, focusing on children who moved across different *buurten* while staying within the same municipality. *Buurten* are highly localized areas, with approximately 1,400 residents on average and a range of around 500–2,000 residents. Hence, *buurten* are much smaller than municipalities; in urban areas, they can be as small as a few city blocks.

Each *buurt* typically contains one or two primary schools, depending on the density and size of the population, although some smaller *buurten* may not have any schools within their boundaries. *Buurten* may also differ substantially in their availability of local amenities such as churches, small parks, playgrounds, or community centers.

Given the diversity and the small scale of *buurten*, they provide a granular view of interactions that operate locally compared with the larger, more heterogeneous municipalities. Because I focus on moves across *buurten* within the same municipality, the municipal policies—such as educational reforms, public services, or economic initiatives—should affect both the origin and destination *buurten* in a similar way. Therefore, any outcome differences after a move are likely driven by local within-*buurten* factors, such as peer interactions, social networks, local norms, or environmental characteristics.

By focusing on such small-scale geographic units, I can better isolate the influence of local dynamics and community interactions that may not be captured when analyzing larger areas such as municipalities. The fact that these moves often occur over distances of only a few kilometers emphasizes the importance of neighborhood-level social and environmental factors, as opposed to broader municipal policies.

The estimation results, shown in Figure 3, reveal that moves across *buurten* also exhibit clear exposure effects, with patterns similar to those observed in inter-municipality moves. Even with highly localized moves, often within a very short distance, children experience significant gains in standardized test scores. This suggests that neighborhood effects operate highly granularly, driven by localized social and environmental dynamics within each *buurt*, rather than broader municipal-level interventions.

4.2 Spatial Decay of Neighborhood Effects

The available data allow me to investigate to what extent neighborhood effects are within a *buurt* rather than across *buurten* located close to each other. My analysis is similar to that of **?** at the census-tract level in the U.S. To capture the spatial decay of neighborhood effects, I estimate a regression in which the predictions for the origin and destination neighborhoods ($\hat{y}_{op}$ and $\hat{y}_{dp}$) are replaced with predictions for the child's specific *buurt*. In addition to these immediate neighborhood effects, I include interactions between the child's age at the move and the average observed outcomes in the 10 closest *buurten* to the origin and destination neighborhoods. The selection of these ten closest buurten is based on the distance between their geographic centers.[5]

The decision to include the 10 closest buurten is motivated by the fact that, in the

[5]The geographic distance between *buurten* is the straight-line distance between their geographical centers. The 10 closest *buurten* are then ranked based on proximity, from nearest to farthest.

Netherlands, most children attend schools and participate in activities within a radius of 1 to 3 kilometers from their homes. Therefore, nearby buurten often share key community resources, such as schools and recreational facilities, and children frequently interact across these small boundaries. By incorporating this surrounding *buurten*, I capture potential spillover effects from neighboring environments, such as shared peer groups, social networks, or common access to public goods.

The results, presented in Figure 3, reveal a clear pattern of spatial decay. Moving to a higher-performing buurt—where permanent residents achieve test scores one standard deviation higher—has the strongest effect in the child's own buurt. However, the impact of nearby buurten rapidly diminishes with distance. By approximately 1 to 1.5 kilometers away, the effect becomes negligible. These findings suggest that policy interventions aimed at improving child development in specific areas should be targeted at the micro-level, within the *buurten* of the children. Broader municipal or regional policies may not have any effects.

Figure 4 presents the results, showing that relocating to a neighborhood with higher test scores earlier in childhood significantly improves a child's educational performance. In contrast, moving to a neighborhood where only the surrounding neighborhoods have higher test scores, without a corresponding improvement in the child's immediate neighborhood, does not significantly affect outcomes. This finding suggests that the beneficial effects of high-performing neighborhoods are hyperlocal and directly tied to the child's immediate environment.

4.3 Separating School and Neighborhood Effects

When children move to a new neighborhood in the Netherlands, they experience changes in school quality and non-school neighborhood factors. To understand the impacts of these changes, I decompose the total effect of moving into two components: a school-related effect, which captures differences in school quality across locations, and a non-school-related effect, reflecting other neighborhood factors such as income composition, public services, and social networks.

The approach I adopt here closely follows the methodology developed by Laliberté (2021), who introduced a framework to decompose the effect of neighborhood moves into separate school and non-school components. This approach is particularly well-suited to the Netherlands due to its unique education system and neighborhood structure. In the Netherlands, students can access schools beyond their immediate residential neighborhood owing to a relatively liberal school choice policy. As a result, students from the same neighborhood may attend different schools, and students attending the same school may come from different neighborhoods. This variation provides the necessary identification to separate school and non-school effects. In many other countries, the primary school children attend is often tightly linked to their neighborhood (and the quality of a school is tightly linked to the neighborhood's affluence). In the Netherlands, the decoupling of residential location and school attendance allows me to separate the contributions of school quality and neighborhood characteristics to children's educational outcomes.

I estimate the effects of school and neighborhood quality separately for permanent residents—those who remain in the same neighborhood throughout their schooling. Using a similar specification to Laliberté (2021), the model for permanent residents is given by:

$$CITO_{it} = \Omega_{s(i)} + \Lambda_{l(i)} + X_i\beta + \epsilon_i, \tag{4}$$

where $CITO_{it}$ is the primary school performance at time t of child i, who lives in neighborhood $l(i)$ and attends school $s(i)$. $\Omega_{s(i)}$ captures the quality of the schools attended by children who live in l_i, while $\Lambda_{l(i)}$ captures the separate effect of the neighborhood where children live. X_i includes observable characteristics of children, and ϵ_i accounts for unobserved factors. The model is identified because students from the same neighborhood often attend different schools, and students in the same school often come from different neighborhoods, allowing me to separate the effects of school and neighborhood on student outcomes.

For students who move between neighborhoods, I estimate the total exposure effect of moving on their educational outcomes. The total difference in outcomes between permanent residents of the origin neighborhood o and destination neighborhood d is represented as $\Delta_{od} = \bar{y}_o - \bar{y}_d$.

The primary estimating equation for movers is the following:

$$y_{imod} = \gamma(m_i\Delta_{od}) + \beta X_{imod} + \alpha_{od} + \alpha_m + \epsilon_{imod}, \tag{5}$$

In this model, y_{imod} represents the educational outcome of student i, who lived in neighborhood o (origin) at baseline and moved to neighborhood d (destination) at age m. The coefficient of interest, γ, captures the annual rate at which the outcomes of movers converge to those of permanent residents in their destination neighborhood. Unlike equation (3), which includes income-specific effects, equation (5) focuses purely on the spatial exposure effects by incorporating origin-by-destination fixed effects (α_{od}) to control for differences across neighborhoods. Additionally, unobserved differences between children who move at different ages—such as disruption costs from the move—are controlled through age-at-move fixed effects (α_m). The model fundamentally compares children who began in the same origin neighborhood and moved to the same destination neighborhood, but at different ages, to isolate the effect of neighborhood exposure independent of income-specific factors.

Following Laliberté's (2021) decomposition approach, the total effect of moving is separated into contributions from school quality and non-school factors. The difference in outcomes between origin and destination neighborhoods Δ_{od} can be written as:

$$\Delta y_{od} = \Delta\Omega_{od} + \Delta\bar{y}_{od}^{ns}, \tag{6}$$

where $\Delta\Omega_{od}$ represents the difference in the expected quality of the schools children attend in neighborhoods o and d. This is calculated as a weighted average of school quality estimated in equation 4, with weights corresponding to the share of permanent residents' children attending different schools. This specification implicitly assumes that movers will make similar school choices to those of the permanent residents, allowing us to estimate the Intention-to-Treat (ITT) effects. Meanwhile, $\Delta\bar{y}_{od}^{ns}$ captures the difference in expected non-school (ns) factors, including both neighborhood fixed effects and family composition differences.

Because of the unique structure of the Dutch education system, I can separately estimate the following two models.

$$y_{imod} = \gamma_s(m_i\Delta\Omega_{od}) + \beta X_{imod} + \alpha_{od} + \alpha_m + \epsilon_{imod}, \tag{7}$$

and

$$y_{imod} = \gamma_n s(m_i \Delta \bar{y}_{od}^{ns}) + \beta X_{imod} + \alpha_{od} + \alpha_m + \epsilon_{imod}. \tag{8}$$

The variation in school attendance and neighborhood residency patterns in the Netherlands enables estimation of these two effects separately. This flexibility is key to disentangling school quality from broader neighborhood characteristics.

The total exposure effects can now be written as the sum of two effects:

$$\begin{aligned}\gamma &= \frac{\mathrm{cov}^r(y_{imod}, \Delta y_{od})}{\mathrm{var}^r(\Delta y_{od})} \\ &= \underbrace{\frac{\mathrm{cov}^r(y_{imod}, m_i \Delta y_{od}^{ns})}{\mathrm{var}^r(m_i \Delta y_{od}^{ns})}}_{\gamma_{\mathrm{ns}}} \frac{\mathrm{var}^r(m_i \Delta y_{od}^{ns})}{\mathrm{var}^r(m_i \Delta y_{od})} + \underbrace{\frac{\mathrm{cov}^r(y_{imod}, m_i \Delta \Omega_{od})}{\mathrm{var}^r(m_i \Delta \Omega_{od})}}_{\gamma_{\mathrm{s}}} \frac{\mathrm{var}^r(m_i \Delta \Omega_{od})}{\mathrm{var}^r(m_i \Delta y_{od})}\end{aligned} \tag{9}$$

where cov^r and var^r refer to the covariance and variance obtained using the residuals of regressions on the controls employed when estimating school and non-school neigborhood effects.

The relative contributions of school and non-school factors to the total effect of moving can now be obtained following as:

$$F^{\mathrm{school}} = \frac{\gamma_s \mathrm{var}^r(m_i \Delta \Omega_{od})}{\gamma \mathrm{var}^r(m_i \Delta y_{od})}, \tag{10}$$

and

$$F^{\mathrm{non\text{-}school}} = \frac{\gamma_{ns} \mathrm{var}^r(m_i \Delta \bar{y}_{od}^{ns})}{\gamma \mathrm{var}^r(m_i \Delta y_{od})}. \tag{11}$$

Table 3 presents the results. The upper panel highlights statistically significant total exposure effects, ranging from 0.023 to 0.033, which closely mirror the estimation results derived from the specifications outlined in Section 3. The breakdown of these effects reveals notable disparities in the contributions of school and non-school factors. In particular, approximately 39% of the total exposure effect (0.023) is attributed to school-related factors, indicating a significant influence of educational institutions on overall test scores. For math performance, γ_s is estimated at 0.288, and around 65% of the observed changes after relocation can be ascribed to differences in the quality of schools attended by children in different neighborhoods. Non-school factors, represented by $\gamma_n s$, contribute to a lesser extent account for approximately 35% of the total exposure effect. A different pattern is observed in Dutch proficiency. Here, school factors explain roughly 26% of the total change in Dutch proficiency across neighborhoods, while non-school factors contribute to 74% of the total exposure effect.

These findings reveal the role of schools in shaping educational outcomes following moves, particularly in cognitive development. The effect sizes associated with school-related factors emphasize the importance of investing in and improving educational resources and opportunities within neighborhoods, particularly for subjects such as math.

5 Conclusion

This study provides a comprehensive analysis of how neighborhoods affect educational opportunities, showing the crucial role of primary school performance in mediating children's educational attainment. The findings emphasize the significant variation in educational outcomes across neighborhoods and illustrate that moving to a more advantageous neighborhood during childhood can substantially improve test scores and educational attainment.

By examining administrative data from the Netherlands, I obtain several key insights. First, Every additional year a child spends in a neighborhood with higher average test scores improves their test score rank by approximately 2.5% per year of childhood exposure, up to age 12. The longer a child is exposed to a better-performing neighborhood, the closer their test score converges to those of children who have always lived in that neighborhood.

Second, the mediation analysis demonstrates that primary school performance accounts for roughly 40% of the variation in educational attainment across neighborhoods. The decomposition analysis further supports this, revealing a correlation coefficient of 0.785 between primary school performance and educational attainment. This finding emphasizes improving primary school quality to bridge educational disparities.

The analysis at the smallest administrative unit, the *buurt*, yields similar exposure effects, indicating that neighborhood impacts are highly localized. Even moving within a municipality can, therefore, significantly affect educational outcomes.

Decomposing the total exposure effect on educational attainment into school and non-school factors reveals that school quality contributes to approximately 39% of the improvement in test scores at the end of primary school. In math, 65% of the exposure effect is attributable to school quality, indicating the critical role of schools in reducing educational inequalities. Non-school neighborhood amenities account for the remaining 35%.

These findings imply that improving primary school quality in disadvantaged neighborhoods can help reduce educational disparities by shaping early cognitive development and promoting better educational trajectories. When it comes to non-school neighborhood effects, it is important to recognize that these are highly localized. As a result, it is important that policies narrowly target disadvantaged neighborhoods. More broadly, a better understanding of the impact of neighborhood exposure and the underlying mechanisms should allow policymakers to target resources better to reduce educational disparities and promote socioeconomic mobility across diverse communities.

References

Aizer, A. and Currie, J. (2019). Environmental factors in disadvantaged neighborhoods and behavioral and cognitive outcomes. *American Economic Review*.

Card, D., Domnisoru, C., and Taylor, L. (2018). The role of school quality in regional differences in upward mobility. *Journal of Economic Perspectives*, 32(3):167–194.

Chetty, R., Friedman, J. N., Hendren, N., Jones, M. R., and Porter, S. R. (2020). The opportunity atlas: Mapping the childhood roots of social mobility. *Nature*, 577:462–468.

Chetty, R. and Hendren, N. (2018a). The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. *The Quarterly Journal of Economics*, 133(3):1107–1162.

Chetty, R. and Hendren, N. (2018b). The impacts of neighborhoods on intergenerational mobility ii: County-level estimates. *The Quarterly Journal of Economics*, 133(3):1163–1228.

Chyn, E. and Katz, L. F. (2021). A summary of recent literature on neighborhood effects. *Journal of Economic Literature*.

Deutscher, N. (2020). Place, peers, and the teenage years: Long-run neighborhood effects in australia. *American Economic Journal: Applied Economics*, 12(2):220–49.

Epple, D. and Romano, R. (2003). Neighborhood schools, choice, and the distribution of educational benefits. In Hoxby, C., editor, *The Economics of School Choice*, pages 227–286. University of Chicago Press. Accessed: 2024-10-05.

European-Commission (2019). Education and training monitor 2019 - netherlands. Accessed: 2024-10-05.

Gibbons, S. and Silva, O. (2018). Urban density and pupil attainment. *Journal of Urban Economics*, 103:96–115.

Guryan, J., Heller, S. B., Ludwig, J., and Pollack, H. A. (2021). The effect of parental resources and neighborhood violence on cognitive skills. *American Economic Review*.

Jackson, C. K. (2020). Public school funding and student outcomes. *Brookings Papers on Economic Activity*, 2020(1):65–138.

Laliberté, J.-W. (2021). Long-term contextual effects in education: Schools and neighborhoods. *American Economic Journal: Economic Policy*, 13(2):336–377.

OECD (2023). *Education at a Glance 2023: OECD Indicators*. OECD Publishing. Accessed: 2024-10-05.

Patrinos, H. A. (2011). School choice in the netherlands. *CESifo DICE Report*, 9(2):55–59.

Rossin-Slater, M. (2018). The importance of early childhood education. *Annual Review of Economics*, 10:269–295.

Rothstein, J. (2019). Inequality of educational opportunity? schools as mediators of the intergenerational transmission of income. *Journal of Labor Economics*, 37(S1):85–123.

Sykes, B. and Musterd, S. (2011). Examining neighbourhood and school effects simultaneously: What does the dutch evidence show? *Journal of Urban Studies*, 48(4):1–18.

Webbink, D., ter Weel, B., and Odding, C. (2023). On the estimation of neighbourhood exposure effects. Unpublished manuscript.

Wodtke, G. T. (2018). Parental resources and children's cognitive development. *Sociology of Education*, 91(1):32–53.

6 Tables and Figures

Table 1: Summary Statistics

	Mean	Std. Dev.	Median	Number of Obs.
A. Total				
Native Parents	0.803	0.398	1	1,732,497
Maternal Schooling	16.03	3.173	16	1,107,395
Parental Income	49,228	45,570	42,941	1,731,210
CITO Std. Scores	535.4	9.802	537	1,732,497
Child Schooling	17.27	2.104	17	631,896
B. Permanent Residents				
Native Parents	0.808	0.394	1	1,334,623
Maternal Schooling	15.90	3.166	16	820,807
Parental Income	48,996	40,823	43,253	1,334,623
CITO Std. Scores	535.3	9.792	537	1,334,623
Child Schooling	17.24	2.078	17	485,443
C. Movers				
Native Parents	0.783	0.412	1	397,874
Maternal Schooling	16.41	3.160	17	286,014
Parental Income	50,007	58,754	41,682	397,874
CITO Std. Scores	535.7	9.829	537	397,874
Child Schooling	17.37	2.187	17	146,453

Notes: This table presents summary statistics for the analysis sample used in this study. The sample is restricted to children born between 1994 and 2007, whose moving histories are fully observable from age 1 onward, and who have valid CITO test scores from 2006 to 2019. The sample is divided into three groups: (1) The "Total Sample" includes all children with complete residential and educational records throughout the observation window, encompassing both movers and permanent residents. (2) "Permanent Residents" refers to children who continuously lived in the same neighborhood from age 1 to age 15, providing the baseline for neighborhood-level predicted outcomes as these children experienced constant exposure to a single neighborhood throughout their childhood. (3) "Movers" include children who relocated between neighborhoods at least once before age 15. This group is further analyzed based on their age at the time of the move, differentiating between those who moved at or before age 12 and those who moved after age 12. This differentiation allows for investigating the timing of neighborhood exposure on educational outcomes. Parental income is measured as the average disposable income of the household when the child is between 9 and 12 years old, adjusted for inflation and standardized across years. CITO test scores are standardized nationally and converted into percentile ranks to facilitate comparison across cohorts. All other variables, including parental education and child schooling years, are similarly standardized to account for variations in measurement across different data sources and years.

Table 2: Estimates of Test Score Exposure Effects

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	Overall CITO	Overall CITO	Math	Dutch	Overall CITO	Overall CITO
Exposure Effect γ	0.0304***	0.0274***	0.0223***	0.0348***	0.0242***	0.0250*
	(0.00704)	(0.00874)	(0.0103)	(0.00889)	(0.00450)	(0.0137)
ρ	-0.0208	-0.0433	-0.0169	-0.0744	-0.0646*	-0.110
	(0.0465)	(0.0545)	(0.0655)	(0.0540)	(0.0326)	(0.0720)
Observations	250,159	231,112	231,112	231,112	106,102	106,102
Controls	no	yes	yes	yes	no	no
Within Year of Birth	no	no	no	no	yes	yes
Family Fixed Effects	no	no	no	no	no	yes

Note: This table reports estimates of annual childhood exposure effects on children's test score ranks at age 12. The estimates can be interpreted as the impact of spending one year of childhood in a municipality where the test scores of permanent residents are one percentile point higher. Each column presents estimates from a regression of a child's test score ranks at age 12 on the difference between permanent residents' predicted ranks in the destination versus the origin neighborhoods, which interacted with the child's age at the time of the move (m). I allow for separate linear interactions for $m \leq 12$ and $m > 12$ and report both coefficients. Standard errors are shown in parentheses.

Table 3: School Shares

Dependent Variables	(1) CITO	(2) Math	(3) Dutch
Total Exposure Effects			
γ	0.0255***	0.023***	0.033***
	(0.00655)	(0.00635)	(0.00757)
School and Non-School Components			
γ_s	0.0215**	0.0288***	0.0191*
	(0.00965)	(0.00885)	(0.01101)
γ_{ns}	0.0263***	0.013	0.0338***
	(0.00826)	(0.00792)	(0.00852)
School Shares(s^{school})	39%	65%	26%
Observations	43,640	43,640	43,640

Robust standard errors are in in parentheses.
*** $P < 0.01$, ** $P < 0.05$, * $P < 0.1$

Figure 1: Mean Test Scores for Children of Permanent Residents at the 25th Income Percentile

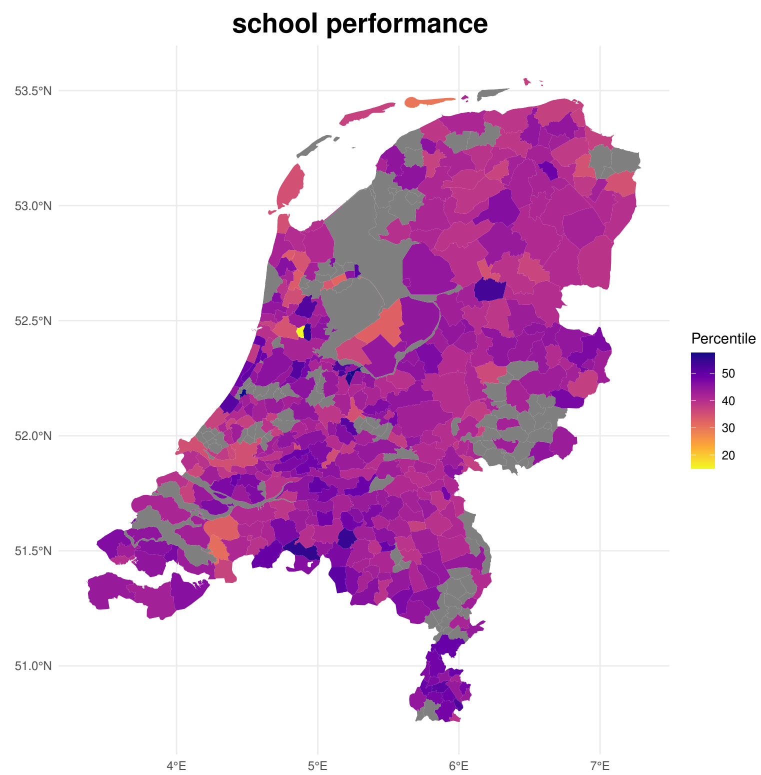

Note: This plot illustrates children's mean percentile ranks in test scores, conditional on having parents at the 25th income percentile. Darker-shaded colors correspond to higher outcomes for children, while gray indicates areas with fewer than 40 children, where data are insufficient to estimate outcomes. The sample includes all children in the analysis sample whose parents are permanent residents (i.e., those who do not move across municipalities before the child turns 16). To create these estimates, I first regress children's test score ranks on a constant and their parents' family income ranks separately for each municipality. I then calculate the predicted income rank for children with parents at percentile p in municipality c as the intercept $+p$ times the slope of this regression.

Figure 2: Childhood Exposure Effects on Test Scores at Age 12

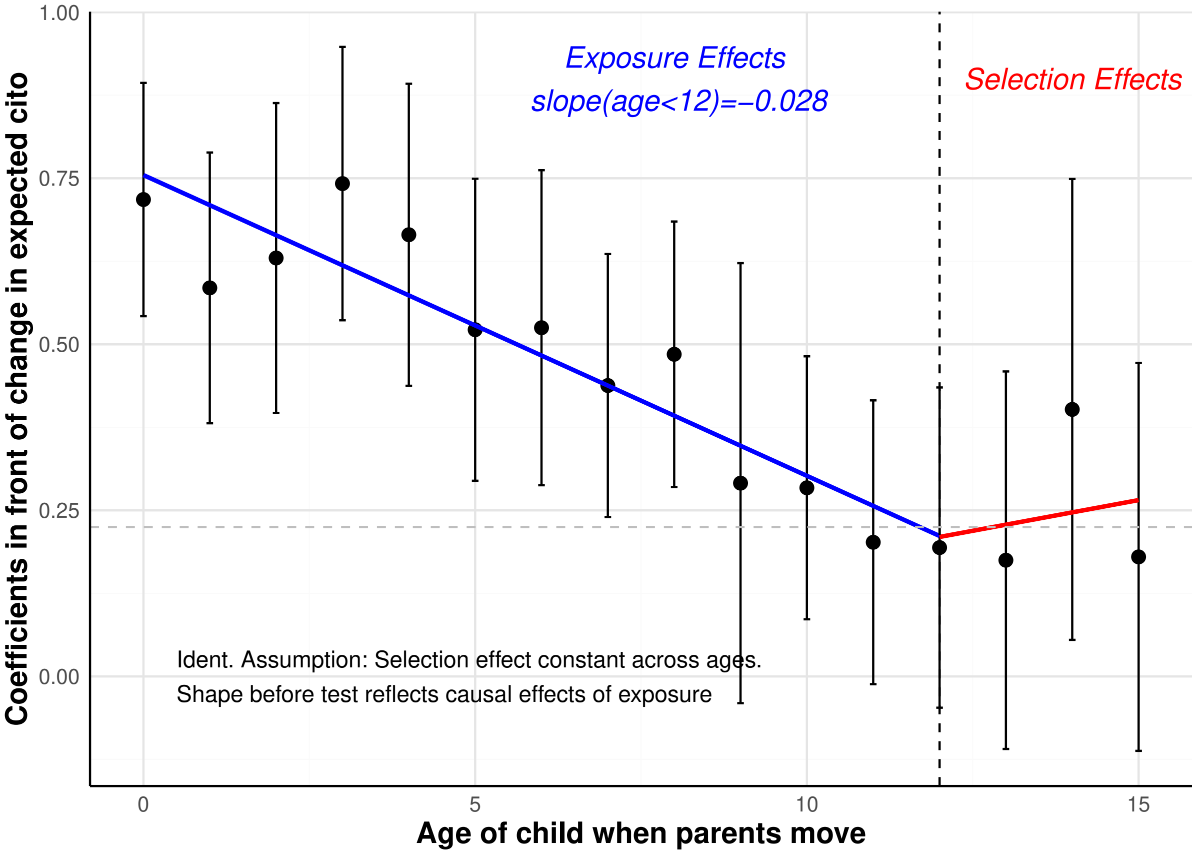

Note: This figure illustrates the relationship between the coefficients b_m and the child's age at the time of moving (m), using the semi-parametric model outlined in Equation 1. The analysis evaluates children's test scores at age 12, and the sample consists of children in the primary analysis who moved exactly once between 1994 and 2007. The b_m coefficients represent the effect of moving to a neighborhood where permanent residents achieve one percentile higher outcomes at a given age (m). These coefficients are derived by regressing a child's test score rank (T_i) on $\Delta odp = T_{op} - T_{dp}$, which is the difference in predicted ranks between permanent residents of origin and destination neighborhoods, interacted with each age of the child when they moved (m). Dashed vertical lines distinguish two groups: children who moved at ages $m \leq 12$ and those who moved at ages $m > 12$. The best-fit lines are calculated using unweighted OLS regressions of the b_m coefficients on m for each age group. The slopes of these regression lines—reported with standard errors in parentheses—approximate annual childhood exposure effects for children who moved at ages $m \leq 12$. This analysis assumes that selection effects do not vary based on the child's age at the time of moving (m).

Figure 3: Exposure Effects Estimation Using Within-Municipality Moves

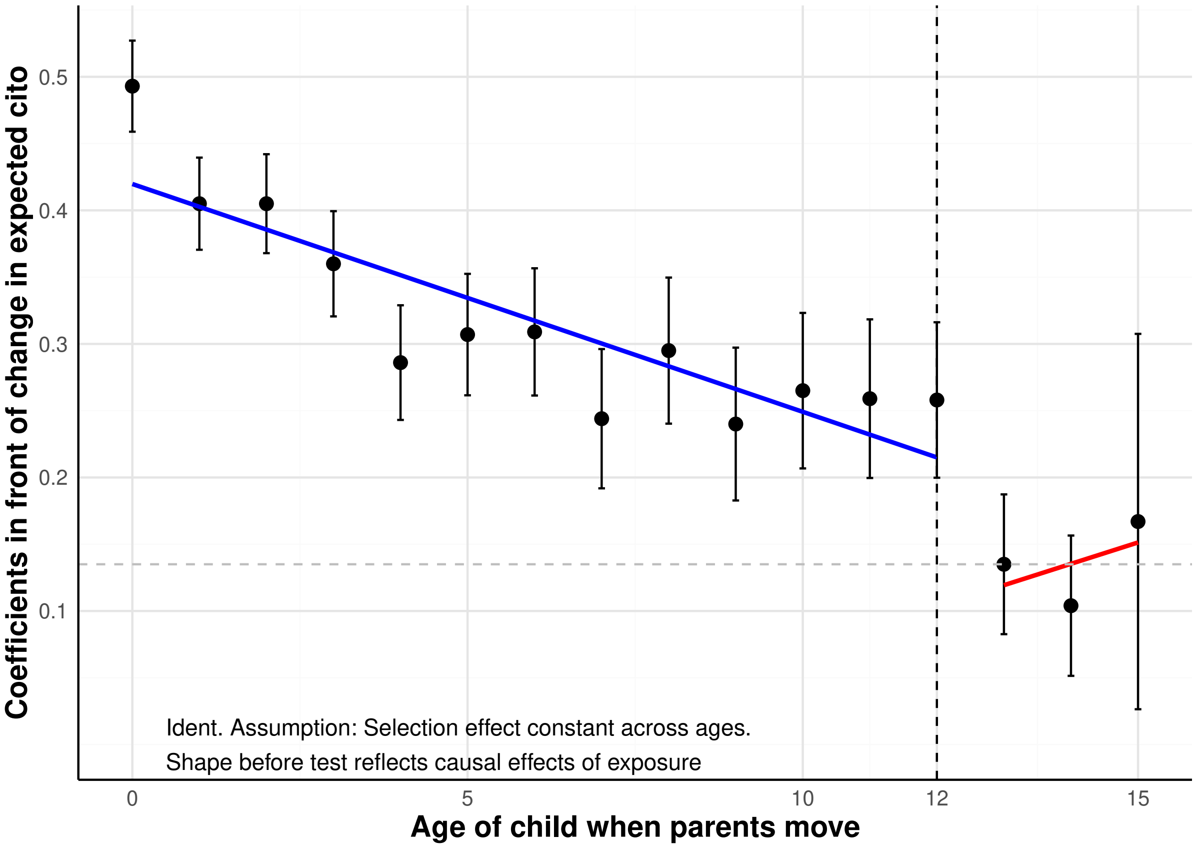

Note: This figure illustrates the relationship between the coefficients b_m and the age at which a child moves (m), using the semi-parametric model in Equation 1. The analysis evaluates children's test scores at age 12, and the sample includes children in the primary analysis who move exactly once within a municipality between 1994 and 2007. The b_m coefficients represent the effect of relocating to a neighborhood where permanent residents achieve one percentile higher outcomes at a given age (m). These coefficients are derived by regressing a child's test score rank T_i on $\Delta_{odp} = T_{op} - T_{dp}$, the difference in predicted ranks between permanent residents of origin and destination neighborhoods, interacted with each age at which the child moved (m). Dashed vertical lines distinguish two groups: children who moved at ages $m \leq 12$ and those who moved at ages $m > 12$. The best-fit lines are calculated using unweighted OLS regressions of the b_m coefficients on m for each age group. The slopes approximate annual childhood exposure effects for children who moved at ages $m \leq 12$. This analysis assumes that selection effects do not vary based on the child's age at the time of moving (m).

Figure 4: Spatial Decay

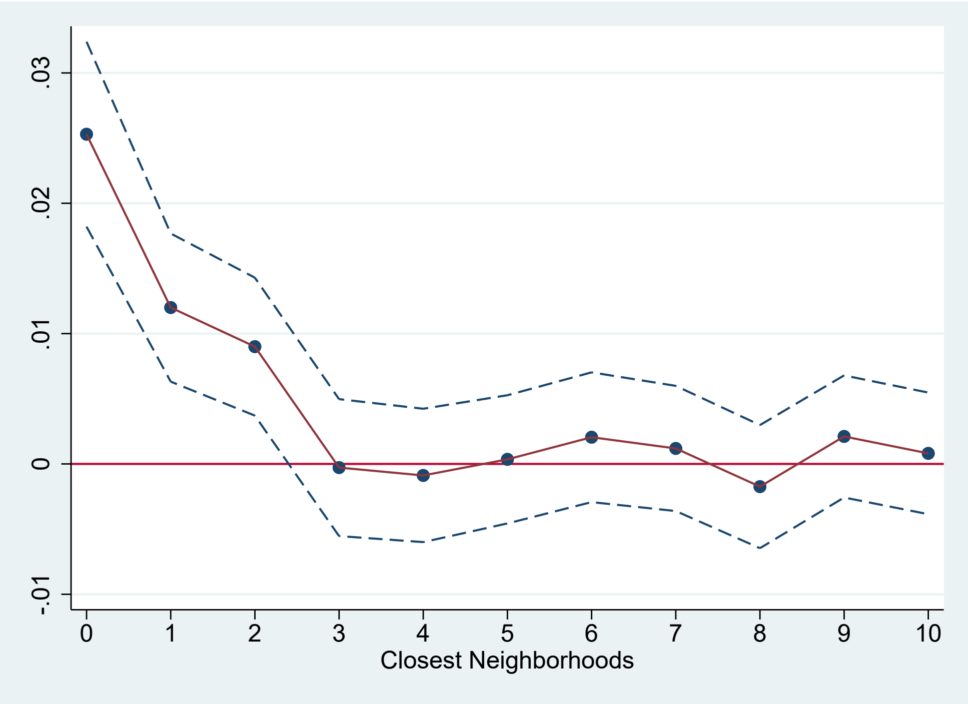

Note: The plot visualizes 11 coefficients of interaction terms between the child's age at the time of moving and neighborhood outcomes. This analysis is based on a mover design incorporating the mean observed outcomes of permanent residents from the 10 closest *buurten* to the origin and destination *buurt*.

Appendix A: Mediation Analysis of School Performance and Educational Attainment

This appendix provides a deeper exploration of the role of school performance, measured by CITO test scores, as a mediator in the relationship between neighborhood quality and long-term educational attainment. Specifically, I aim to assess whether improvements in school performance at age 12 can serve as a path by which neighborhood characteristics, such as total years of schooling, affect long-term educational attainment.

A.1 Effects on Educational Attainment

A key policy question is whether neighborhoods that enhance school performance at age 12 also positively influence long-term outcomes. In other words, does improving CITO scores translate into higher educational attainment? To investigate this possibility, I replace the outcome variable in the main analysis with the educational attainment of movers, measured in years of schooling, and estimate the following equation:

$$S_i = \sum_{m=0}^{15} (\alpha_m + \phi_m + \zeta_m p_i + b_m \Delta odp) + \epsilon_i \tag{12}$$

where S_i represents the educational attainment of movers, measured in years of schooling; b_m measures the effect of moving to a neighborhood where test scores are one percentile point higher on total years of schooling; Δodp captures the difference in neighborhood quality (as measured by test scores) between the origin and destination neighborhoods; p_i represents the child's parental income rank; and other terms such as α_m and ϕ_m capture age-specific effects and additional controls.

The results, illustrated in Figure A.1, show a steady decline in b_m estimates with the age at move (m) for children younger than 12, which provides evidence of neighborhood exposure effects on long-term outcomes. The findings indicate that moving to a neighborhood with better school performance earlier in childhood leads to more significant gains in educational attainment. The magnitude of this effect suggests that a one standard deviation increase in CITO test scores leads to an approximate gain of 0.025 years of schooling. If a child was born in such a neighborhood, the cumulative improvement would be around 0.6 years, or roughly 35% of a standard deviation in schooling.

Figure A.1: Reduced-Form Effects on Educational Attainment

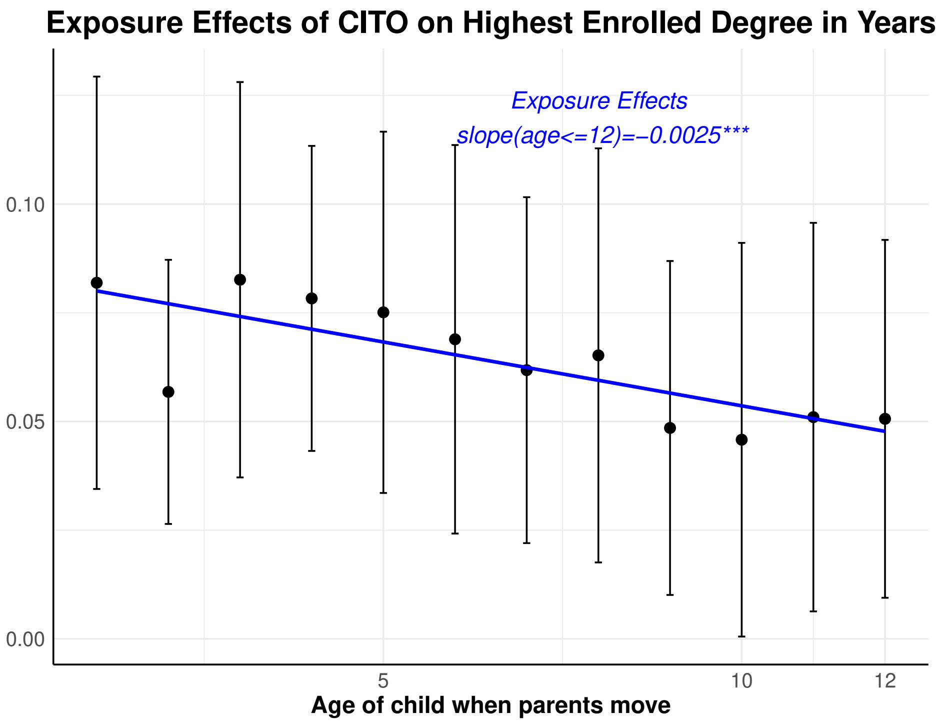

Note: This figure shows the relationship between the coefficients b_m and the child's age at the time of moving (m), using the semi-parametric model outlined in Equation 12. The analysis evaluates children's educational attainment up to age 24, and the sample consists of children who moved exactly once between 1994 and 1998. The b_m coefficients represent the effect of relocating to a neighborhood where permanent residents achieve CITO test scores that are one percentile higher at a given age (m). These coefficients are derived by regressing a child's years of schooling (S_i) on $\Delta odp = T_{op} - T_{dp}$, which is the difference in predicted ranks between permanent residents of origin and destination neighborhoods, interacted with the child's age at the time of the move (m). The best-fit lines are calculated using unweighted OLS regressions of the b_m coefficients on m. The slopes approximate annual childhood exposure effects for children who moved at ages $m \leq 12$.

A.2 Mediation Analysis

The mediation analysis focuses on breaking down the total effect of parental income on long-term educational outcomes into direct and indirect effects. In this context, *indirect effects* refer to the role of primary school performance (CITO test scores) as a mediator that links parental income to long-term educational attainment. In contrast, *direct effects*capture the influence of parental income on educational outcomes that bypass school performance, reflecting other pathways such as family resources, social capital, or access to information.

This section aims to quantify the extent to which primary school performance can explain the relationship between parental income and children's educational attainment. This will provide insights into whether school performance serves as an important mechanism or pathway by which socioeconomic status is transmitted across generations.

The following system of equations models the relationships between parental income, school performance, and educational attainment:

1. Reduced-form transmission of parental income to educational attainment

$$S_i = \phi_c + \theta_c p_i + \xi_i \tag{13}$$

   In this equation, S_i represents the child's long-term educational attainment; p_i is the child's parental income rank; θ_c measures the total effect of parental income on children's educational attainment, which includes both direct and indirect effects; and ξ_i captures unobserved factors.

2. Transmission of parental income through school performance

$$S_i = \kappa_c + \lambda_c T_{ic} + \mu_c p_i + u_i \tag{14}$$

   In this equation, T_{ic} represents the child's primary school performance (CITO test scores); λ_c captures the extent to which primary school performance contributes to the child's long-term educational attainment; μ_c measures the direct effect of parental income on educational attainment, independent of school performance; and u_i captures unobserved factors specific to the child.

The key parameter of interest here is λ_c, which quantifies the impact of school performance on long-term outcomes, holding parental income constant. This allows us to estimate how much parental income's total effect on educational attainment operates through school performance.

By substituting Equation 14 into Equation 13, the total effect of parental income rank on long-term educational attainment becomes:

$$S_i = (\kappa_c + \lambda_c \alpha_c) + (\lambda_c \pi_c + \mu_c) p_i + (\lambda_c \epsilon_i + u_i) \tag{15}$$

Thus, the reduced-form transmission of parental income to children's educational attainment can be expressed as a combination of direct and indirect effects:

$$\theta_c = \lambda_c \pi_c + \mu_c \tag{16}$$

A.3 Results and Interpretation

The mediation analysis results suggest that, on average, 40 percent of the variation in long-term educational attainment across neighborhoods can be attributed to differences in primary school performance. This finding indicates that school performance is a meaningful mediator in the transmission of socioeconomic status from parents to children.

As shown in Figure A.2, there is a strong positive correlation between test score transmission (π) and educational attainment transmission (θ), with a correlation coefficient of 0.785. This result suggests that improvements in CITO test scores are strongly linked to gains in long-term educational attainment. The decomposition analysis further reveals that primary school performance accounts for approximately 40% of the variation in educational attainment across neighborhoods, emphasizing the critical role of early school performance in shaping long-term outcomes.

A.4 Implications of the Findings

These findings underscore the importance of primary school performance as a key mechanism through which neighborhoods influence children's future outcomes. By improving primary education in disadvantaged neighborhoods, policymakers can potentially reduce

Figure A.2: Correlation Between Parental Income Rank and Educational Transmission

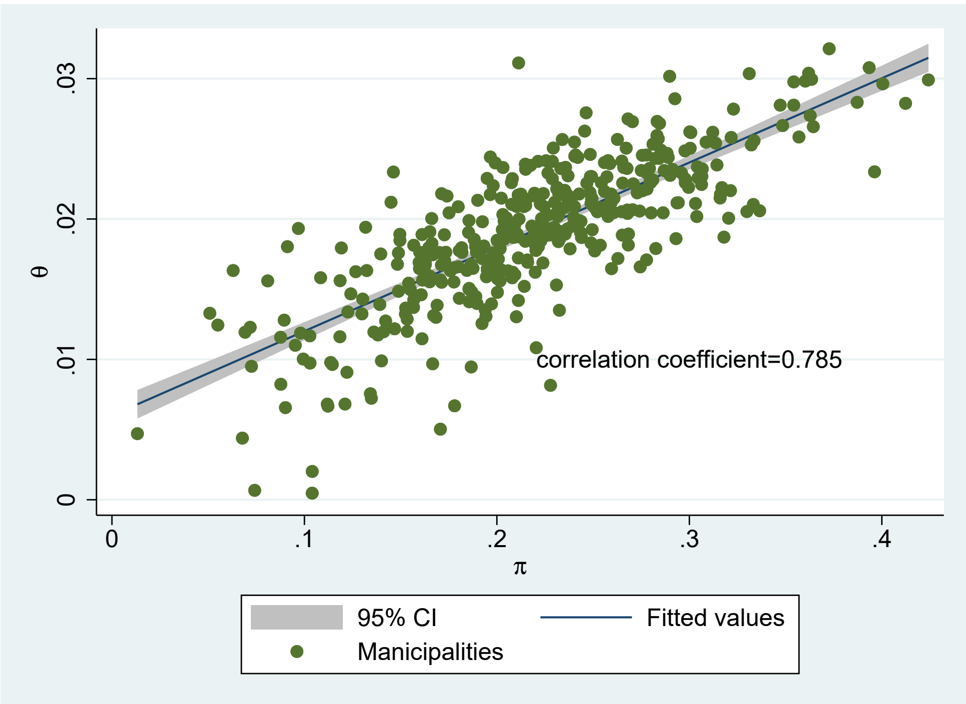

Note: This figure illustrates the relationship between coefficients θ (parental income rank to schooling years) and π (parental income rank to test scores). Each point represents a municipality. The fitted line and confidence interval show an unweighted regression of θ on π, and the text indicates the correlation coefficient between the two coefficients across municipalities.

educational disparities and promote upward mobility. Furthermore, the mediation analysis highlights that while school performance is a critical pathway, other neighborhood factors also play a role in shaping long-term outcomes, suggesting the need for a multifaceted approach to improving childhood environments.