The Hidden Costs of Recycling: Lead Exposure and Student Learning

Gerald Ipapa[†]

July 2026

Abstract

I examine the impact of lead exposure from newly opened lead-acid battery recycling plants on student learning in Kenya. Using individual test scores from a national end-of-primary exam, I estimate difference-in-differences and event-study models exploiting this natural experiment. Lead exposure reduces test scores by 0.18 standard deviations, with losses twice as large at schools closest to the plants, exceeding the gains from common school-based learning programs. Boys are more affected than girls, and suggestive evidence points to larger losses for private-school students; I find no rural-urban difference. Lead exposure poses a significant threat to human capital in sub-Saharan Africa.

Keywords: lead exposure, battery recycling, environmental health, human capital, student achievement, learning crisis, sub-Saharan Africa

JEL Codes: Q53, I25, I15, I21, I24, O15

[†]Ipapa: Georgetown University, gje11@georgetown.edu. I am grateful to Adrienne Lucas, Francisco Costa, Sabrin Beg, Kimberly Oremus, Parag Mahajan, Olga Gorbachev, Isaac Mbiti, Andrew Brudevold-Newman, Owen Ozier, Bryan Stuart, Manaswini Rao, Pamela Jakiela, Andrew Zeitlin, the participants of the OSWEET Energy and Environment Economics Workshop, Philadelphia Federal Reserve seminar and the University of Delaware Economics Department seminar participants for helpful comments. The findings, interpretations, and conclusions expressed herein are entirely my own. All errors and omissions are my own.

1 Introduction

Human capital is a critical driver of economic growth, making the prevalent lack of learning in sub-Saharan Africa (SSA) a major concern. The World Bank (2018) identifies four key determinants of learning in schools: learner preparedness, teacher skills and motivation, school inputs, and school management and governance. These areas have traditionally been the focus of school-based programs aimed at improving learning in low- and middle-income countries (LMICs) over the last two decades, yet extensive research examining their impact has yielded mixed results (Banerjee and Duflo, 2007; Duflo, 2001; Lucas et al., 2014; Glewwe et al., 2009). Environmental factors, however, are excluded from this framework. This exclusion is concerning given the ongoing learning crisis and rising pollution levels (particularly lead) in SSA (Gottesfeld et al., 2018).[1] Despite its significance, the impact of lead exposure in LMIC schools remains understudied.

In this paper, I study the impact of recycling-driven lead exposure on student learning outcomes in Kenya. Leveraging a natural experiment generated by the opening of 26 used lead-acid battery (ULAB) recycling plants in 2007, which triggered a sudden increase in lead emissions, I exploit the variation in students' exposure to these emissions to identify the impacts. The level of exposure depends on two sources of variation: (i) the schools' proximity to the plant, which determines whether a student is in the exposure zone, and (ii) the fraction of a student's primary schooling years that overlap with the plant's operation, which captures the duration of exposure.[2] I measure exposure intensity as the product of a binary proximity indicator and this continuous timing measure. Together with a cross-sectional harmonized secondary school qualifying exam, the Kenya Certificate of Primary Education (KCPE), I estimate difference-in-differences (DiD) regressions to identify the effects of lead exposure on student learning.

I find that exposure to lead emissions reduces students' test scores by 0.18 standard deviations, and the effect roughly doubles for students attending the schools closest to the plants, consistent with a dose-response relationship. This is equivalent to a reduction of 1.17 Learning-Adjusted Years of Schooling[3] (1.17 LAYS) per child, an effect size that exceeds the gains from common school-based learning programs. For example, Angrist et al. (2020) show that a well-known deworming program in Kenya delivered 0.05 LAYS and merit scholarships for girls in Kenya delivered 0.34 LAYS.[4] I

[1] Over the past two decades, LMICs have experienced a significant surge in pollution levels. This rise is primarily attributed to the expansion of manufacturing, mining, and recycling sectors as governments in these countries grapple with and respond to the high levels of unemployment and a growing youthful population.

[2] For example, a student whose school is near a plant that opened in 2007 and who took the KCPE exam in 2010 was exposed for 3 out of 8 primary school years (fraction = 3/8), while a student at the same school who took the exam in 2007 was exposed for only 1 out of 8 years (fraction = 1/8).

[3] Learning-adjusted years of schooling for a given country (or LAYS) are the product of years of schooling and a measure of schooling quality (Angrist et al., 2020).

[4] In constructing the LAYS from my lead impact estimates, I follow Angrist et al. (2020) and take the lead impact estimate, divide by a high-performance benchmark learning rate of 0.8 standard deviations per year, and then multiply by the duration, in years, that the impact is expected to persist. Following Angrist et al. (2020), I set the persistence horizon equal to the duration of the evaluation, here taken as five years to match the persistence assumptions used for comparable interventions in the Angrist et al. (2020) database. These choices imply that 1 LAYS corresponds to 0.157 SD (0.8/5). For example, the main estimate of -0.183 SD at the 4 km radius converts to $(0.183/0.8) \times 5 = 1.17$ LAYS, or equivalently $0.183/0.157 = 1.17$ LAYS. I use this fixed benchmark rather than a country-specific rate

also find that the negative effects of lead emissions are more pronounced for boys than girls and for students attending private schools than public schools, suggesting heterogeneity by student gender and by school type but not by school location type.

I also examine and test three potential threats to identification. First, the timing and location of plant operations may be endogenous, which raises concerns about potential pollution sorting.[5] While my research design does not require recycling plants to be exogenously located, I address these concerns with controls for local time-varying shocks at the district level. I also report entropy-balancing matching estimates following Hainmueller (2012) as a robustness check.[6] The balancing is based on observable school characteristics including school location and school type, as well as school quality and student socioeconomic status.

Second, residential sorting may occur due to households moving away in response to a plant opening nearby or for employment reasons. I run a regression of the log of the number of exam takers on the exposure intensity measure to determine whether the number of students taking the KCPE exam decreases or increases after a plant is opened. I find no discernible relationship between lead exposure and the number of exam takers, suggesting limited net migration in response to the plant openings. While this test cannot definitively rule out offsetting migration flows that leave total numbers unchanged but alter student composition, institutional factors (poverty, limited infrastructure, and lack of public information about the plants) make such compositional sorting unlikely in this setting.

Finally, the recent difference-in-differences literature shows that estimates from TWFE estimators can be severely biased and may even be incorrectly signed (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020). The common trends assumption holding, the TWFE estimator linearly combines the treatment effects across treated units. However, when most units are treated in later periods, some treated observations may be assigned negative weights in the calculation of the estimated treatment effects. This can be a source of bias if the treatment effects are heterogeneous across units. Because my study setting does not exhibit staggered lead exposure, this is not a concern. Nevertheless, in my last robustness check, I run a simple diagnostic test proposed by Jakiela (2021) to assess the extent to which negative weighting on treated units is a concern and whether homogeneity of the treatment effect assumption is plausible. I find that the relationship between the lead exposure measure and the standardized test scores is consistent with the assumption required for TWFE estimation. Beyond these checks, a falsification test that assigns a placebo opening to the non-lead toxic sites in the same inventory produces null estimates, and wild cluster bootstrap inference confirms that the results are not an artifact of the modest number of plant clusters.

This paper contributes to the literature on environmental justice and human capital develop-

because it provides a stable, externally defined reference point that does not depend on any particular country's potentially variable performance data.

[5]Pollution sorting is when owners of the battery recycling plants place them in areas with lower land prices (to minimize their costs), which also tend to be locations with households of lower socioeconomic status.

[6]See Basri et al. (2021) for a recent application of entropy balancing to matching.

ment. Previous research on lead exposure in high-income countries (HICs) has demonstrated its harmful effects on various outcomes for children, including reduced test scores (Aizer et al., 2018; Hollingsworth et al., 2020). Even at low levels, lead impacts student proficiency in reading and math (Needleman et al., 2004; Lanphear et al., 2000). However, comparable evidence for LMICs is limited for several reasons. First, data constraints, specifically the lack of lead exposure monitoring among children, make causal research designs difficult. Second, identifying the impacts of lead exposure in LMICs is not straightforward. The complex interconnectedness of the drivers of learning (or lack thereof) in schools, and the fact that the lack of learning is due to numerous factors such that the actual impact of lead exposure alone could be negligible or small, make it difficult to discern.[7] Due to these constraints, in their recent work, Crawfurd et al. (2023) use a meta-analysis of lead studies in HICs to simulate learning outcomes in LMICs under a policy that completely eliminates lead to the levels observed in HICs. Given the increasing phenomenon of pollution havens, where pollution shifts to LMICs as HICs tighten regulations, it is critical to address these challenges to produce empirical evidence tailored to the unique settings in LMICs, ultimately informing regulations in the recycling sector.

Additionally, this paper contributes to the literature on lead exposure in LMICs by providing direct evidence of how recycling-driven lead exposure affects student learning in Kenya. Using administrative test score data and a research design that capitalizes on the opening of battery recycling plants in 2007, it effectively controls for potential biases. The findings also provide insights relevant to LMICs due to contextual similarities with Kenya. Moreover, the paper reveals the hidden costs of recycling that governments and policymakers must consider when creating policies for the recycling sector.

The rest of the paper proceeds as follows. Section 2 describes the setting of the study, the Kenyan education system, and battery recycling plants. Section 3 describes the conceptual framework, and Section 4 provides the empirical framework and discusses the empirical and identification strategy. Section 5 discusses the data used in the study and reports summary statistics. Section 6 presents the empirical findings, followed by a battery of robustness checks, mechanisms, and limitations. Section 7 concludes.

2 Background

2.1 Schooling in Kenya

During the period relevant to this study, Kenya's education system mandated that students spend eight years in primary school, followed by four years in secondary school and four years at university. The main analysis uses KCPE data from 2000 to 2010; the harmonized five-subject exam format was introduced in 2003 and remained consistent through 2010.[8] To progress to the next level,

[7] For example, the effects of a policy that trains teachers to improve students' quality of learning may be dampened if parents withdraw their support in response to the policy.

[8] I use data from 2000–2002 to establish pre-trends in the event study analysis. Because the exam format changed from seven subjects to five subjects in 2003, all specifications include exam-year fixed effects, which absorb level

students were required to take standardized national exams, known as the KCPE and the Kenya Certificate of Secondary Education, at the culmination of their primary and secondary education, respectively.

Up until 2002, students in grades 4–8 had a seven-subject workload during the school year, and therefore eighth graders were required to take seven-subject exams as part of the KCPE, with scores ranging from 0 to 700 points. However, this format changed in 2003 and stayed the same for the period of the study. During the study period, students in grades 4–8 had a five-subject workload during the school year, and therefore eighth graders were only required to take five-subject exams for the KCPE, with scores ranging between 0 and 500 points. This KCPE exam holds significant importance as it not only serves as a certification of primary education completion but also determines admission into secondary schools. It is thus considered a high-stakes assessment for students in Kenya.

2.2 Lead Exposure

Lead, like many other metals, is a natural element of the earth's crust. Highly immobile, its traces can be found in water, plants, and soil. However, it becomes highly toxic when mined and used by humans (Meyer et al., 2003; WHO, 2015). Lead has been mined for thousands of years, and according to estimates, approximately 300 million tons have been released into the environment due to mining and ore processing activities over the past 500 years (WHO, 2015).

Lead's properties make it valuable in a wide range of products, such as protective equipment (e.g., x-ray shields and fire protection gear), gasoline (as an anti-knocking agent, a use now close to being phased out globally), and ammunition. On the other hand, the detrimental effects of lead have been known for centuries. In recent decades, research has linked lead exposure to a broad spectrum of negative laboratory findings and adverse health effects.

Children are particularly vulnerable to lead exposure for two main reasons. First, they often engage in hand-to-mouth behavior and have a natural inclination to explore their surroundings. Their tendency to touch and taste the objects and substances they come into contact with increases the risk. Furthermore, habits such as chewing nails or sucking fingers can further increase the risk of ingesting lead-contaminated soil or dust particles. The act of mouthing or swallowing objects or substances containing lead or coated with lead can significantly contribute to the accumulation of lead in their bodies (WHO, 2015).

Second, children's organ systems are not fully developed, and their mechanisms for metabolizing lead are relatively ineffective. Infants and children have much higher rates of lead absorption in the gastrointestinal tract than adults: children absorb up to 40 percent of ingested lead, compared to approximately 10 percent in adults. Nutritional deficiencies, particularly in iron and calcium, which are more prevalent in children, further increase the absorption rate of lead. Furthermore, the transport of lead through children's immature blood-brain barrier increases the risk of adverse neurological and developmental effects (WHO, 2015; Landrigan and Bellinger, 2021)

differences across exam formats; test scores are standardized over the analysis sample.

Broadly, there are two sources associated with lead exposure among children in LMICs. The first is via non-industrial sources, which includes gasoline, paint, and piped water. Within this category, gasoline was a major source of chronic exposure to lead in LMICs until 2006, when a majority of SSA countries prohibited leaded gasoline. The second source is cottage industries or the informal sector, which includes ULAB recycling plants, mining, electrical repairs, and jewelry. Thus, the threat of children being exposed to lead remains incredibly high in LMICs (Meyer et al., 2003).

Despite the overwhelming evidence showing the significantly negative effects of lead in HICs, particularly the irreversible nature of the damage to a child's brain, pollution levels in LMICs have significantly surged over the past two decades. This rise is primarily attributed to the expansion of manufacturing, mining, and recycling sectors (Gottesfeld et al., 2018) as governments in these countries grapple with and respond to high unemployment levels and a growing youthful population, driven by global dynamics. On the one hand, stricter environmental regulations in HICs and the absence or laxity of regulations in LMICs may lead to the relocation of environmentally hazardous production through trade.[9] In the realm of policy, such displacement effects are often downplayed as innocuous (World Bank, 2020). However, in a recent study examining the recycling of used ULABs, Tanaka et al. (2022) provide empirical evidence showing a direct impact of tightened environmental regulations in the United States, which led to the relocation of polluting activities to Mexico.[10] Another factor contributing to this expansion is the increased demand for raw materials used in the computer chip and automobile industries. For example, lead is employed in the production of the ULABs used in the automotive sector. In 2003, recycled lead accounted for 45 percent of the global lead supply, with the majority coming from used ULABs (WHO, 2015).

In Kenya, a toxic site identification program (TSIP), commissioned by the UN Industrial Development Organization (UNIDO) and implemented by Pure Earth, identified 81 pollution sites across the country. Among them, 26 sites (ULAB recycling plants) had high levels of toxic lead emissions, while the remaining sites were contaminated with various forms of mercury (UNIDO, 2009). One example is the Owino Uhuru plant in Mombasa that opened in 2007, which used a special furnace to extract lead from used batteries and subsequently exported the purified lead (see Figure 1). Like the other sites, Owino Uhuru emitted pollution in the form of dust, fumes, and effluent filled with lead particles (Ministry of Health, 2015). The plant was shut down in 2014, and between 2015 and 2018, media reports indicate that 10 more toxic waste smelters closed following pressure from environmental activists.

[9] This situation could allow investors to reduce production costs and maximize individual benefits by relocating production to LMICs without adequately addressing the associated social costs, resulting in adverse health effects in the destination countries.

[10] Anecdotal evidence from the Kenya case finds that some of the ULAB plants were owned by Indian investors following tightening environmental regulations in India.

3 Conceptual Framework

To investigate how lead exposure impacts student learning and contributes to the learning crisis in SSA, I introduce a simple conceptual framework, building on Glewwe and Kremer (2006). I then discuss the nuances unique to the LMICs settings that make estimation complex. A useful assumption to make here is that subject to constraints, a household in LMICs maximizes its (life-cycle) utility function. The main arguments in the utility function are consumption (including leisure) at different points in time and each child's academic achievement. The constraints faced are the following: the child's learning production function; the household's information set; a life-cycle budget constraint; and either some credit constraints or an agricultural/retail production function (for which child labor is one possible input), or both. The child's learning production function is a structural relationship that can be depicted as

$$A_i = a(S, H, Q, C, EP, I), \tag{1}$$

where A denotes cognitive skills produced for child i (achievement), S denotes years of schooling, H is a vector of household characteristics, Q is a vector of school and teacher characteristics (quality), and C is a vector of child characteristics (including the child's innate ability). EP denotes the institutional arrangements that affect the efficiency of education provision, and I is a vector of school inputs under the control of parents, such as children's daily attendance and purchases of textbooks and other school supplies. For the purposes of this framework, little is lost (and some simplicity is gained) by treating A as a scalar even though children may acquire different skills in school, suggesting that equation (1) should have multiple outputs and A should be a vector.

Assume that all elements in the vectors C and H are exogenous. Examples of such variables are credit constraints, parental tastes for schooling, parental education, and children's genetic endowments of ability. Some child characteristics affecting education outcomes (such as health) could be endogenous; such variables can be treated as elements of I, all of which are endogenous. Another important set of variables is prices related to schooling, denoted by the vector P. These prices can include school fees, prices for school supplies purchased by parents, and even wages paid for child labor. P does not appear in equation (1) as it has no direct effect on learning; its effect works through decisions made for the endogenous variables S and I.

In the simplest scenario, assume that only one school is available to each household and parents cannot change the school's characteristics. Thus, all variables in Q and P are exogenous to the household. Parents choose S and I (subject to the above-mentioned constraints) to maximize household utility, which implies that years of schooling S and schooling inputs I can be expressed as general functions of the four vectors of exogenous variables:

$$S = f(Q, C, H, P), \tag{2}$$

$$I = g(Q, C, H, P). \tag{3}$$

Inserting (2) and (3) into (1) gives the reduced-form equation for (A):

$$A = h(Q, C, H, P, EP). \tag{4}$$

While this reduced-form equation is a causal relationship, it is not a production function as it reflects household preferences and includes prices among its inputs.

The more realistic assumption that households can choose from more than one school implies that Q and P are endogenous even if they are fixed for any given school. In this scenario, households maximize utility with respect to each schooling choice and then choose the school that leads to the highest utility. Conditional on choosing that school, they choose S and I, as in the case where there is only one school from which to choose.

Consider a change in one element of Q, denoted as Q_j. Equation (1) shows how changes in Q_j affect A when all other explanatory variables are held constant, and thus provides the partial derivative of A with respect to Q_j. In contrast, equation (4) provides the total derivative of A with respect to Q_j because it allows for changes in S and I in response to the change in Q_j. Households may respond to better teaching by increasing their provision of educational inputs such as textbooks. Alternatively, if they consider better teaching a substitute for those inputs, they may decrease them. In general, the partial and total derivatives in this case could be quite different.

Now recall that lead poisoning or exposure among children occurs via inhalation and/or ingestion.[11] Consider a sudden change in the lead levels to which a child is exposed, an element of C denoted as C_j. Equation (1) shows how changes in C_j affect A when all other explanatory variables are held constant, and thus provides the partial derivative of A with respect to C_j. In contrast, equation (4) provides the total derivative of A with respect to C_j because it allows for changes in S and I in response to the change in C_j.

However, in the short run, households cannot immediately respond to changes in the lead exposure levels their children face due to two important reasons unique to the LMIC context. The first is a lack of awareness: households, educators in schools, and healthcare providers have limited awareness of the effects of lead exposure, delaying necessary support. The second is limited lead screening and testing capacity, resulting in lead-related health issues going undetected or untreated. Because of these constraints, in the short run the partial and total derivatives are likely to be similar, as households do not adjust S and I in response to the shock. Over time, however, as exposure persists and symptoms manifest, household behavioral responses can drive the two derivatives apart.

In the long run, as exposure persists, the child becomes symptomatic, which can cause a household subject to a budget constraint to respond by reallocating resources away from supporting the child's education to seek healthcare for them.[12] The initial impact on the child's innate ability and health is compounded by extended untreated exposure, their absence from the classroom, and the

[11] For a detailed discussion of why children are vulnerable, see Section 2.

[12] Households also reduce their labor supply to care for their child, further limiting their budget.

household's resource reallocation, further affecting the child's production of cognitive skills.

When examining the impact of policies on academic skills, A, equation (4) is the policy-relevant quantity: it shows what will actually happen to A after a change in one or more elements of Q, C, or P, inclusive of household responses. The partial derivative from equation (1) remains of interest because it can better approximate welfare effects: at an optimum, the marginal welfare consequences of households' optimizing responses are second-order by the envelope theorem.

With the increase in children's lead exposure and households not being able to respond immediately, the two equations deliver similar effects in the short run. In the long run, households respond. Following a reduction in C due to increased lead levels, households reduce educational inputs I as they reallocate resources toward healthcare expenditure. The reduced-form impact (total derivative) reflects the drop in A due to both the direct effect of C and the reduction in I. By contrast, the structural impact measured in equation (1) captures only the direct effect of the lead shock, holding S and I constant. While these approaches differ in principle, both provide informative estimates of the impact of lead exposure on achievement. Since this paper ultimately estimates the reduced-form effect, β captures the total impact of lead exposure on test scores, inclusive of any household behavioral responses over the study period.

Using equation (4), a learning crisis emerges when the total negative effect of lead exposure on achievement, combining the direct cognitive damage and any reduction in educational inputs, exceeds the gains from school-based programs that improve school quality. To allow for better synthesis of this evidence, following Angrist et al. (2020), I express the set of outcomes A_i (or the estimates from the reduced-form estimator) in terms of LAYS. I can then find which share of the learning gap is explained by lead exposure.

Thus far, the framework has settled on reduced-form effects in equation (4) that reflect both the partial derivatives of increased lead exposure (reduction in C) and agents' optimizing responses as suitable for examining lead impacts. While this framework is a useful guide for empirical evaluation work in LMICs, it has two limitations. First, it assumes a unitary household model and thus abstracts from bargaining among household members regarding education decisions (Glewwe and Kremer, 2006).[13] Second, it abstracts from the program's general equilibrium effects in the long run.[14] Given these limitations, it is crucial to reevaluate how they might influence empirical assessments, particularly in environments with persistent high lead levels as observed in LMICs.

Although empirical evidence has clearly demonstrated the negative effects of lead in HICs, LMICs still have incredibly high lead levels. Is it possible that researchers consider the structural relationship in equation (1), thus underestimating the impacts? Most empirical studies on lead impacts in both HICs and LMICs use administrative data, based on data generated by ordinary (non-experimental) variation across schools and households. Even when using the reduced-form equations, vanilla OLS regressions using retrospective data deliver correlational relationships that

[13]Education decisions can be affected by household bargaining both between men and women and between parents and children within a household.

[14]Changes in education inputs may eventually result in the change in the supply of educated individuals, which may then change the returns to education and thus the demand for education.

may not be useful. Perhaps in countries with high levels of lead, the actual marginal effects of lead might be relatively small due to confounding factors. For example, polluting ULAB recycling plants are often placed in areas with lower land prices, which also tend to be locations with households of lower socioeconomic status. Therefore, endogenous program placement is the first issue.

The second issue is omitted variables. Using the production function (1), researchers may estimate the marginal effect of increasing lead exposure while holding all other factors constant. A common problem with this approach is that it is generally very difficult to control for all the relevant variables, resulting in potentially biased estimates. Third, the presence of measurement errors in the observed covariates is another difficult estimation issue, as random measurement error typically leads to underestimation of the true underlying impacts, while nonrandom measurement error could lead to biases in either direction. Compounding this issue is that household and school survey data collected in LMICs contain a substantial amount of error.

Correcting for these biases in an OLS regression with non-experimental data is difficult such that a correlational output is not helpful. Natural experiments (in this case generated by an environmental shock) together with techniques such as entropy balancing, which ensures that $E(u|Treatment) = E(u|Control)$, offer a tremendous opportunity to address these issues and to correctly identify lead impacts in LMIC settings.

4 Empirical Framework

4.1 Empirical Strategy

In my empirical strategy, I leverage the environmental shock induced by the 2007 opening of the ULAB recycling plants, using a difference-in-differences design complemented by an event study to compare learning outcomes for children attending schools near these plants. Specifically, I analyze differences between children at schools exposed to more severe versus less severe emissions and between those who were younger versus older when the toxic emissions began. Figure 7 shows the distribution of the plants across the country. The level of students' lead exposure depends on their schools' proximity to the plant and the timing of when it opened.

For proximity (first difference), I draw donut-like bins around the ULAB recycling plants. Students attending schools within a 4 km radius of a plant (the buffer) are the treated group, while those attending schools located 10–20 km from a plant (the donut area) are the control group. Figure 2 illustrates this design for the Owino Uhuru plant in Mombasa. The timing variation (second difference) captures how long a student is exposed to lead during the school period: the number of years between when the plant started operating and when the student took the exam, divided by the total number of years spent in primary school (typically eight). I measure the intensity of lead exposure, R_{ist}, as the product of a binary proximity indicator and this continuous timing measure. This variation in exposure intensity allows me to run a two-way fixed effects specification (5) at the student level:

$$y_{ist} = \beta \cdot R_{ist} + \gamma_s + \delta_t + \delta_{d(s),t} + \Gamma X'_{ist} + \varepsilon_{ist}, \tag{5}$$

where y_{ist} is an outcome for student i in school s in exam year t, and the exposure intensity measure R_{ist} is defined as:

$$R_{ist} = \textit{Recycling Plant}_s \times \frac{\max(t - E_s,\ 0)}{8}, \tag{6}$$

where *Recycling Plant*$_s$ is an indicator equal to one if school s is within 4 km (6 km, 8 km, or 10 km) of a recycling plant, E_s is the year the nearest plant opened (2007 for all plants in this study), t is the exam year, and the denominator of 8 reflects the eight years of primary schooling. For students in treated schools who took the exam before the plants opened ($t < E_s$), $R_{ist} = 0$. For students in treated schools who took the exam after the plants opened, R_{ist} captures the fraction of their primary schooling during which they were exposed to lead emissions. For all students in control schools, $R_{ist} = 0$ regardless of timing. The coefficient of interest is β, which captures the effect of a one-unit increase in exposure intensity, equivalent to being exposed for the entirety of one's primary schooling while attending a school within the treatment radius.

X'_{ist} is a set of student-level controls to address potential observable confounders, including gender. Gender is important here as differentiated gender roles are still common in LMICs. For example, one would expect household chores and social activities to differ for boys and girls, with girls spending more time engaging in indoor household chores (e.g., dishes) and boys spending more time on outdoor chores (e.g., working on the farm) and playing outdoors. Those involved in outdoor chores or sports are likely to be exposed to lead emissions from plants near the playgrounds. Controlling for gender accounts for potential differential trends across gender that influence individual test scores and spuriously correlate with a sharp increase in lead exposure levels for students attending schools close to the recycling plants.

δ_t denote exam year fixed effects and account for common annual shocks across all schools, and γ_s is a set of school fixed effects that accounts for time-invariant differences across schools. $\delta_{d(s),t}$ are district-by-year fixed effects that absorb local time-varying shocks at the district level $d(s)$ in which school s is located, such as commodity price movements, weather, or county-level political events. All 19 districts represented in the estimation sample contain both treated and control schools, so the treatment effect remains identified within every district after conditioning on district-by-year fixed effects. I do not include school-specific linear time trends in the preferred specification because they are mechanically collinear with the linear-in-t treatment, and report estimates with school trends as a robustness check. Robust standard errors are clustered at the recycling plant level.[15]

For specification (5) to provide unbiased estimates of the parameter β, I make the identifying assumption that the exposure intensity measure R_{ist} is independent of the regression error term,

[15] While standard errors are clustered at the school level (the unit of treatment assignment), I also cluster them at the recycling plant level to be conservative. Only the latter standard errors are reported.

conditional on the explanatory variables. In other words, specification (5) is well identified if, conditional on the controls, there are no remaining omitted variables that vary within a school or over time and are correlated with both student test scores and the intensity of lead exposure.

$$E[R_{ist} \cdot \varepsilon_{ist} \mid X'_{ist}, \gamma_s, \delta_t, \delta_{d(s),t}] = 0. \tag{7}$$

A violation of this assumption occurs if, for example, the opening of a plant responded to unobserved time-varying shocks (such as migration or sorting).[16]

$$R_{ist} = \alpha + \phi \cdot \Delta Z_{st} + \gamma_s + \delta_t + \eta_{st}, \tag{8}$$

where ΔZ_{st} is the change in the average school-level test score since the last time students took the exam. The result of estimating equation (8) is discussed in detail in Section 6.1. Prior changes in the test scores are uncorrelated with the launch of the recycling plants,[17] indicating that the estimates from my preferred specification correctly identify the causal effect of lead exposure.

To trace the dynamics of the effect and assess pre-trends, I complement the difference-in-differences specification with an event-study design, given by equation (9):

$$y_{ist} = \sum_{\tau \neq 2006} \beta_\tau [\textit{Recycling Plant}_s \times \mathbb{1}[t = \tau]] + \gamma_s + \delta_t + \delta_{d(s),t} + \Gamma X'_{ist} + \varepsilon_{ist}. \tag{9}$$

where *Recycling Plant*$_s$ is the binary proximity indicator defined above and $\mathbb{1}[t = \tau]$ is a year indicator. The reference category is the exam year just before the ULAB recycling plants were opened (the omitted year is $\tau = 2006$, i.e., event time -1). The rest of the controls remain as in equation (5). Robust standard errors are clustered at the plant level. Each coefficient β_τ estimates the effect of attending a school within the treatment radius in year τ relative to the omitted year. If β_τ is statistically insignificant for all years before the plant rollout ($\tau < 2007$), this lends support to the validity of the parallel trends assumption.

Figures 3–6 plot β and 95 percent confidence intervals for each year of lead exposure. Figure 3 shows the effect on lead exposure when a 4 km radius is chosen for the buffer zone. All the figures illustrate that before the plant rollout, there are no statistically significant differences in the outcomes across the treatment and control schools, lending support to the parallel trends assumption.

[16]The assumption may also be violated if changes in a school's ability to limit a recycling plant's operations are correlated with the school's effectiveness. This is inconsistent with my findings as it would create pretrends, which I do not observe.

[17]Using the full sample, the regression of exposure intensity on all these changes simultaneously has an F-stat of 0.49 with a p-value of 0.6112, indicating that changes in the school performance variable (test score average) do not significantly predict exposure intensity.

5 Data

This study combines data on lead exposure with primary school data and student outcomes. The three sources of data in the analysis are: lead exposure (TSIP) data from Pure Earth (Pure Earth, 2021); KCPE exam data from the Kenya National Examinations Council (Kenya National Examinations Council, 2000–2010); and Kenyan primary school data from the World Bank (Kenya Ministry of Education, 2007).

5.1 Lead Exposure Data

The TSIP dataset provides publicly available data on all existing and past recycling and manufacturing sites with pollutants. The variables include the ULAB recycling plants' geocoordinates, the counties in which the plants are located, and soil pollution levels. I use the plant geocoordinates and the school geocoordinates, from the 2007 Kenya primary school dataset, to construct the distance between the plants and the schools around them. Using these distances, I classify schools as treated (within the buffer radius) or control (in the 10–20 km donut area), as described in Section 4.1. The exposure intensity measure R_{ist}, defined in equation (6), combines this binary proximity classification with the fraction of each student's primary schooling years that overlap with the plant's operation.

5.2 Main Outcome Data

The KCPE is a standardized secondary school qualification exam taken by all students completing primary education to determine entry into secondary school. The KCPE dataset contains students' test scores from 2000 to 2010, year of birth, gender, year of exam, location or exam center, and exam center county. The sample consists of students attending schools within a 20 km radius around any of the 26 recycling plants, who took their national exams in 2000–2010. These individuals were between 13 and 19 years old when they took the exam.

5.3 Other Data

The Kenya primary school dataset contains school-level covariates, including the number of classrooms, pupil-teacher ratios, student enrollment rates, the number of teachers, type of school (private versus public, rural versus urban), and school composition (mixed boys and girls, girls only, or boys only). I use this dataset to check for balance on observables between treated schools and those within the control zone, as described in Section 6.1.

I combine all three datasets into one dataset that I use for my analysis. Using this dataset, I draw donut-like bins around the ULAB recycling plants: schools within the buffer radius are treated,[18] and schools located 10–20 km from a plant (the donut area) form the control group. The 4 km buffer radius is motivated by the physical dispersion of lead from small-scale recycling operations:

[18]Schools within the buffer zone are considered as treated.

Tanaka et al. (2022) measure airborne lead from battery recycling plants traveling approximately a 2-mile radius, and the ULAB plants in Kenya are small relative to those in HICs.[19] I report results at 6, 8, and 10 km buffer radii as sensitivity checks, and at a 2 km radius as an intensity check.

Table 1 compares treated and control schools in the estimation sample on observable characteristics. The two groups differ on several dimensions: treated schools are less likely to be rural, enroll more students, employ more teachers, and are more likely to be private than control schools, a pattern consistent with plants locating in peri-urban areas. These differences indicate that the placement of the plants may not have been random, suggesting potential pollution sorting.[20] Level differences of this kind do not by themselves threaten the design: school fixed effects absorb all time-invariant differences between schools, and identification requires only that treated and control schools would have trended similarly in the absence of the plants, an assumption supported by the flat pre-trends documented in Section 6. Nevertheless, the preferred specification addresses local time-varying confounders with district-by-year fixed effects, and as a robustness check I additionally apply the entropy-balancing method of Hainmueller (2012) to construct reweighted control schools matching the treated schools on observable characteristics including region/county, location type, school type, and other indicators of school quality and student socioeconomic status (pupil-teacher ratios, total enrollments, and number of classrooms). Appendix Table A1 demonstrates the successful entropy balancing for the 6 km buffer specification with the 10–20 km donut as the control group. In Section 6 I provide support for the absence of the pre-plant rollout difference in the test score trends between students in the treatment and control schools.

The ULAB plant start dates are critical for the setup of this study, as they determine the initial exposure to lead by students and the length of time they were exposed before taking their secondary school qualifying exam.[21] While these data are unavailable in the TSIP dataset, media reports indicate that most of the plants were rolled out in 2007, and therefore I assume that all the ULAB plants opened in that year. If some plants in fact opened earlier, the affected pre-period years would be partially treated and would surface as pre-trends in the event study; the flat pre-2007 coefficients in Figures 3–6 provide indirect support for the assumed timing. Furthermore, I also assume that schools located closer to the newly launched plants experienced more severe exposure relative to those further away.

In my analysis, the sample does not include schools in Nairobi county for two reasons. First, Nairobi has multiple ULAB recycling plants that are so close to each other that schools treated by one plant may be controls for another plant; thus, there is no clean comparison. Second, in

[19] The literature on lead has established and consistently used a 10 km radius around a lead-emitting source as an exposure zone; therefore, I consider subjects within the 10 km radius as being potentially exposed to lead, and the control donut begins at 10 km.

[20] ULAB recycling plants are usually placed in areas with lower land prices, which also tend to be in areas with households of lower socioeconomic status.

[21] For example, a student who was in grade 5 in a school that started experiencing exposure in 2007 (because a plant was built next to it) and took their KCPE exam in 2010 was exposed for three years out of the eight years they spent in primary school. However, a student who was in grade 8 and took their exam in 2007 was only exposed for one year out of the eight years in primary school. A student in this school who took their exams in 2006 was not exposed.

large cities in LMICs where leaded gasoline was in use before 2006, automobile leaded gasoline combustion accounted for 80 to 90 percent of airborne lead emissions (Meyer et al., 2003). Not only does Nairobi fit this description, but the city also experienced (and still does) severe traffic. In 2006, the country prohibited leaded gas from being sold to the general public in Kenya. Impacts of the ban are expected to be felt across the board but are much larger in cities with initially severe traffic levels, such as Nairobi. The distorted impact of the ban, which acts in the opposite direction, adds complexity to the analysis.

6 Results

The main results in Figure 3 suggest that exposure to lead emissions from the ULAB plants in 2007 reduced students' test scores.[22] I next check how the event-study pattern varies with the choice of the buffer radius by re-estimating equation (9) at 6 km and 8 km around the plants (2 km and 4 km wider than the preferred 4 km radius, respectively). Figures 4 and 5 show post-2007 point estimates that are negative but attenuated relative to the 4 km case, with confidence intervals that include zero in most years.

Repeating the same exercise at the 10 km radius commonly used in the literature, the post-2007 point estimates attenuate further (Figure 6). Taken together, Figures 3–6 show that the negative effect on test scores is concentrated within the 4 km buffer and weakens as the buffer expands, consistent with airborne and soil-borne lead emissions decaying steeply with distance from the source rather than with a threshold below which lead becomes harmless. Schools beyond the 4 km buffer are likely included in the exposure zone only nominally and dilute the measured average effect.

Table 2 presents difference-in-differences results based on estimating equation (5). Column (4) presents the average treatment effect (ATE) estimate of lead exposure on learning when the buffer is restricted to a 0–4 km radius around the plants. Columns (1), (2), and (3) show the results when the buffer radius is 10 km, 8 km, and 6 km away from the plants, respectively. Exposure to lead emissions reduces student test scores by 18.3 percent of a standard deviation at the 4 km radius (column 4, plant-clustered $p \sim 0.05$). The effect is concentrated within this radius: at 6, 8, and 10 km the magnitudes fall to 7.1, 8.5, and 6.8 percent of a standard deviation respectively, with none of these estimates statistically distinguishable from zero.

The dose-response pattern sharpens when the buffer is tightened further. Restricting the treated group to schools within 2 km of a plant, keeping the same 10–20 km control donut, roughly doubles the estimated effect to -0.383 SD (SE 0.130; Table 3). Together with the attenuation at 6–10 km, the estimates trace a monotonic gradient in distance from the plants: the closer a school is to a

[22] These results are from an event study based on equation (9), which restricts the buffer zone (containing the treated schools) to a 4 km radius around the plants. Recall that the omitted category is $\tau = 2006$ (event time -1), the year before the plants opened. A joint F-test of pre-treatment event-study coefficients does not reject parallel trends (F=1.68, p=0.18 at 4 km; F=1.89, p=0.13 at 6 km).

plant, the larger the loss in test scores. This is the pattern one would expect if proximity to plant emissions, rather than school unobservables, drives the estimated effect.

To allow for better synthesis of this evidence, following Angrist et al. (2020), I express the estimates of lead exposure in terms of LAYS. I find a reduction of 1.17 LAYS per child at the 4 km radius, though the 95 percent confidence interval ranges from approximately 0 to 2.4 LAYS, reflecting the marginal statistical significance of the underlying point estimate. At the 2 km radius the loss is 2.44 LAYS per child, with a 95 percent confidence interval that excludes zero. The point estimates exceed the gains from common school-based learning programs, including the 0.34 LAYS delivered by merit scholarships for girls in Kenya and the 0.05 LAYS delivered by deworming (Angrist et al., 2020).

6.1 Identification Threats and Robustness Checks

Three concerns typical to LMIC settings may raise identification problems: residential (Tiebout) sorting, pollution sorting, and negative weight assignment in the presence of treatment effect heterogeneity.

Residential (Tiebout) Sorting. Residential sorting, known as Tiebout sorting, occurs when individuals or households choose their residence based on their preference, income, location characteristics, and amenities. Identification is threatened when these are unobserved and correlated with both learning and pollution levels. For instance, high-income families often gravitate toward less polluted neighborhoods, potentially relocating from areas with higher pollution levels. In this example, household income and other unobserved factors may be correlated with both pollution exposure and student learning. While cases of within-country migration may exist, they are uncommon in Kenya. For example, a Kenya Life Panel Survey of a deworming study in western Kenya tracked over 80 percent of its study participants back to their native homes almost 20 years later.

To assess the extent of residential sorting associated with lead exposure in this study, I examine the relationship between exposure intensity and the number of students taking the KCPE exam following the plant rollout. To do this, I run a regression of the log of the number of exam takers on the exposure intensity measure R_{ist}. Table 4 shows no discernible relationship between lead exposure and the number of exam takers, suggesting limited net migration in response to the plant openings. While this test cannot rule out offsetting migration flows that leave total numbers unchanged, the limited application of Tiebout sorting in the Kenyan context (attributable to factors such as poverty, lack of infrastructure, and limited transportation options, which restrict individuals' ability to relocate) makes compositional sorting unlikely.

A related selection concern is that test scores are observed only for students who sit the KCPE exam. If lead exposure pushes some students to repeat grades or drop out before grade 8, as the absenteeism mechanism discussed in Section 6 suggests, the estimated effect is identified from the students who remain. The null effect on the number of exam takers limits the scope of this concern, since large exposure-induced changes in the test-taking population would appear in the counts.

To formalize how much selection the data allow, I construct trimming bounds in the spirit of Lee

(2009). A school-year-level regression of the log number of exam takers on exposure intensity under the preferred specification yields a coefficient of -0.135 (SE 0.178) per unit of R_{ist}, statistically indistinguishable from zero; at observed exposure levels, this point estimate implies differential attrition of at most 5 percent of takers by 2010. Trimming the control-group score distribution by the year-specific implied attrition share bounds the 4 km estimate between -0.328 (SE 0.093) and $+0.041$ (SE 0.092).[23] The two endpoints are not equally plausible. The upper bound requires that every student kept from the exam by lead exposure would otherwise have scored at the top of the distribution, the opposite of the absenteeism mechanism, under which health-impaired and academically struggling students are the ones who fail to reach the exam. Under the plausible direction of selection, the relevant bound is the lower one, and the main estimate understates the true loss.

Additionally, the capacity of LMICs to provide a diverse range of public goods and services is often constrained, offering limited options for households seeking to choose their neighborhoods based on desired public goods and services. Last, it is unlikely that residents would be aware of plans for ULAB plants to be established in their vicinity, as public participation was infrequent and not mandated during the study period.

Pollution Sorting. The criteria used by plant owners for siting or allocating plants are unclear, raising potential pollution sorting concerns. Recycling plants may be placed in areas with lower land prices, which also tend to be neighborhoods with households of lower socioeconomic status. The key identifying assumption is that, conditional on school fixed effects, year fixed effects, and district-by-year fixed effects, plant siting is uncorrelated with time-varying unobservables that affect student test scores. The event study pre-trends (Figures 3–6) and the formal test based on equation (8) provide direct evidence supporting this assumption. I additionally report entropy-balancing matching estimates following Hainmueller (2012) as a robustness check.[24] This method constructs reweighted control schools matching treated schools within the buffer area, based on observable characteristics such as region/county, location type, school type, and other indicators of school quality and student socioeconomic status (e.g., pupil-teacher ratio, total enrollments, and the number of classrooms; see a detailed description of this process in Section 5). Appendix Table A1 presents balance checks after entropy matching, and the Specification Sensitivity subsection below reports estimates that include entropy weights.

Negative Weight Assignment and Treatment Effect Heterogeneity. Recent advances in the difference-in-differences literature highlight that estimates from TWFE estimators can be severely biased and may even be incorrectly signed (Goodman-Bacon, 2021; de Chaisemartin and

[23] Using instead the most negative attrition consistent with the 95 percent confidence interval of the takers regression (a worst case of roughly 18 percent attrition by 2010) widens the bounds to $[-0.713, +0.393]$. The trimming follows Lee (2009): because treated schools (weakly) lose takers, the control group contains excess marginal test takers; trimming the bottom (top) of the control distribution bounds the estimate under the assumption that exposure-induced non-takers would have scored at the bottom (top) of the distribution.

[24] Entropy balancing is a preprocessing procedure that consists of a reweighting scheme that assigns a scalar weight to each sample unit such that the reweighted group(s) satisfy a set of balance constraints imposed on the sample moments of the covariate distributions. The balance constraints ensure that the reweighted groups match exactly on the specified moments.

D'Haultfœuille, 2020). With the common trends assumption holding, the TWFE estimator linearly combines the treatment effects across treated students. However, when most students are treated in later periods, some treated observations may be assigned negative weights in the calculation of the estimated treatment effect. This would be a source of bias if the treatment effects are heterogeneous across units.

However, my study setting does not exhibit staggered lead exposure, eliminating this concern. Typical to LMICs, due to data constraints, information about the plants' specific opening dates is unavailable. However, because media reports indicate a start date of 2007 for most plants, I assume that all ULAB recycling plants rolled out in 2007. This implies that all students attending schools within the buffer area were simultaneously exposed to lead, with no staggered treatment rollout. Furthermore, my research design clearly distinguishes the control group (students attending schools located 10–20 km from a plant) so that students in the control schools are always untreated. Students in the treatment buffer zone receive treatment simultaneously and continue to be treated throughout the study period (no exiting treatment once treated as long as the recycling plant remains open).[25]

Nonetheless, to further investigate this concern, I conduct a diagnostic test proposed by Jakiela (2021) to assess the extent to which negative weight placement on treated units may be a concern and whether the homogeneity of treatment effect assumption is plausible. Appendix Figure A1 presents the results from this negative weights assessment.

Negative weighting of observations in the treatment group is appropriate if the treatment effects are homogeneous but can introduce bias when they change over time within the treated units. Under the assumption of homogeneous treatment effects (and common trends), the residualized outcome variable is a linear function of a residualized treatment variable, and the slope does not differ between the treatment and control groups. Therefore, I first examine the functional form visually using a local polynomial regression of residualized test scores on the residualized lead exposure variable. Appendix Figure A2 presents a scatter plot of these residuals. While the flexible local polynomial suggests some non-linearity in the overall relationship, the key test for the TWFE assumption is whether the *linear* slope differs between treated and control groups, which would indicate problematic heterogeneity in treatment effects. I find no evidence of such a difference (p = 0.40; see Appendix Table A2). The observed non-linearity in the pooled sample is consistent with differences in the distribution of the residualized treatment variable across groups, rather than heterogeneous treatment effects. Consequently, the relationship between the lead exposure measure and standardized test scores is consistent with the assumption required for TWFE estimation.

Tanaka et al. (2022) measure how far from a battery recycling plant lead may travel while in the air (an approximate 2-mile radius). While I use this distance in my main analysis, I also run the event-study estimation at buffer radii of 6 km, 8 km, and 10 km. Figures 4–6 show event-study patterns consistent with the attenuation reported in Section 6, with effects most pronounced at the

[25] Also, as discussed in Section 2, infants and children relative to adults take longer periods to shed lead absorbed into their system and organs.

4 km buffer and diminishing at larger radii.

As discussed in Section 4, Figures 3–6 plot β and 95 percent confidence intervals for each year of lead exposure across the 4, 6, 8, and 10 km buffer radii. Figure 3 shows the effect at the 4 km buffer. All the figures show that pre-plant rollout, there are no statistically significant differences in outcomes across the treated and control schools, lending support to the parallel trends assumption.

I conduct an additional test in support of the pretrends assumption, testing for correlation between prior trends in school performance (changes in schools' mean test scores) and exposure intensity (the independent variable of interest). The results of this exercise, based on regression equation (8), show no evidence that the exposure intensity is driven by trends in school performance. A regression of exposure intensity on all these changes simultaneously has an F-stat of 0.49 with a p-value of 0.6112, indicating that changes in the school performance variable (average school test score) do not significantly predict exposure intensity. This result, together with the previous pretrends tests, shows that the estimates from my preferred specification correctly identify the causal effect of lead exposure.

6.2 Inference with Few Clusters

Standard errors throughout the paper are clustered at the recycling plant level, which yields 26 clusters. With few clusters, asymptotic cluster-robust standard errors can over-reject the null hypothesis (Cameron et al., 2008). To assess whether the reported inference is an artifact of the small number of clusters, I compute wild cluster bootstrap p-values imposing the null, with 9,999 replications and Rademacher weights, using the boottest implementation of Roodman et al. (2019); Webb weights give nearly identical results.

Appendix Table A3 reports the results. The bootstrap leaves my paper's conclusions unchanged. The 2 km estimate remains significant at the 5 percent level (analytic $p = 0.007$, bootstrap $p = 0.017$), and the main 4 km estimate moves only marginally (analytic $p = 0.056$, bootstrap $p = 0.062$), with a bootstrap 95 percent confidence interval of $[-0.394, 0.011]$. The estimates at 6–10 km remain statistically indistinguishable from zero. Among the heterogeneity interactions at the 4 km radius, the gender interaction remains significant at the 5 percent level under the bootstrap ($p = 0.012$), the private-school interaction strengthens slightly ($p = 0.006$), and the rural interaction remains null. The same patterns hold at the 6 km radius.

6.3 Sensitivity to Parallel-Trends Violations

The pre-trend tests reported above cannot rule out violations of parallel trends that are too small to detect, and with 26 clusters their power is limited. I therefore report the sensitivity analysis of Rambachan and Roth (2023), which constructs confidence sets for the treatment effect that remain valid when post-treatment violations of parallel trends are allowed to be up to $\bar{M}$ times the largest pre-treatment violation (the relative magnitudes restriction). The target parameter is the average of the four post-2007 event-study coefficients at the 4 km radius, estimated under the preferred specification.

Appendix Table A4 reports the robust confidence sets. Under exact post-period parallel trends ($\bar{M} = 0$), the confidence set for the average post-2007 effect is $[-0.097, -0.002]$, matching the conventional event-study inference. The breakdown value lies below $\bar{M} = 0.25$: once post-period violations even one quarter the size of the largest pre-period violation are allowed, the confidence set includes zero. This is the expected outcome for an estimate at the margin of conventional significance, and I report it for transparency rather than dispute it. Two features of the exercise inform its interpretation. First, the pre-period coefficients are small in absolute terms (none exceeds 0.022 SD), so the violations entertained at $\bar{M} = 1$ are modest; even there, the confidence set of $[-0.218, 0.133]$ remains centered below zero and is consistent with effects considerably larger than the point estimate. Second, the exercise evaluates the event-study evidence in isolation. The case for a causal interpretation in this paper does not rest on the event study alone: a confounding differential trend near plants would have to reproduce the dose-response gradient across radii, including the doubling of the effect at 2 km, while leaving no imprint on schools near non-lead toxic sites and no imprint on the number of exam takers. The sensitivity analysis accordingly bounds what the event-study evidence can establish by itself, under (approximately) exact parallel trends, while the convergent design-based checks carry the remaining identifying weight.

6.4 Specification Sensitivity

The preferred specification (Equation 5) reflects two consequential modeling choices: the inclusion of district-by-year fixed effects to absorb local time-varying shocks, and the exclusion of school-specific linear time trends, which are mechanically collinear with the linear-in-t treatment intensity. To assess how each choice shapes the main estimate, I report a battery of variations at the 4 km buffer radius.

Removing district-by-year fixed effects from the preferred specification yields a substantially larger point estimate of -0.430 (SE 0.192); see Appendix Table A5 for the main variations. District-by-year fixed effects absorb regional shocks (commodity prices, weather, county-level political events) that would otherwise be attributed to lead exposure, so the larger magnitude without them reflects regional confounding rather than a more credible identification of lead's effect. Adding school-specific linear time trends to the preferred specification attenuates the point estimates toward zero (-0.273, SE 0.204 without district-by-year fixed effects, untabulated; 0.009, SE 0.170 with both school trends and district-by-year fixed effects, column 3 of Appendix Table A5). The attenuation reflects the mechanical collinearity between R_{ist} (linear in t for treated schools after 2007) and $\gamma_s \cdot t$ (also linear in t): the trend term absorbs a meaningful share of the treatment-effect variation by construction.

I also report the entropy-balanced specification used in earlier drafts of this paper. With school-specific time trends and entropy weights, the point estimate is -0.166 (SE 0.170), qualitatively similar in direction to the preferred specification but smaller in magnitude and less precise.

Appendix Figure A3 plots the corresponding entropy-balanced event-study coefficients at the 4 km and 6 km buffers. Two methodological caveats apply to the entropy-balanced version. First, balancing on time-invariant covariates is partially redundant under school fixed effects: any time-invariant difference between treated and control schools is already absorbed by γ_s, so the matching primarily reweights the sample (changing the estimand) rather than addressing omitted-variable bias. Second, entropy balancing without pre-treatment outcomes in the balancing set does not address selection on pre-treatment levels.

6.5 Tighter Control Donut

The control group in the preferred specification consists of schools located 10–20 km from the nearest plant, balancing two competing concerns: schools too close to plants risk being treated themselves (control contamination), while schools too far away may differ systematically from treated schools on unobserved dimensions. To probe sensitivity to this geographic extent, I rerun the specification using a tighter 10–15 km control donut (Appendix Table A5, column 4). Tightening the donut reduces the unweighted point estimate by roughly 20–35 percent (the magnitude depends on whether district-by-year fixed effects are included) but does not flip the sign of the coefficient. The qualitative story is robust to the geographic restriction: bringing the control group closer to the treatment buffer does not change the direction of the estimated effect, suggesting that the point estimate is not an artifact of comparing treated schools to systematically different control schools far from the plants.

6.6 Plant Overlap

A potential concern in this setting is that schools may be within the treatment radius of multiple plants, creating SUTVA-violation contamination if matched control schools are themselves treated by a different plant. A diagnostic analysis confirms that the nearest-plant matching rule that defines the control donut (schools 10–20 km from the *nearest* plant) prevents control contamination by construction: zero percent of nominally-control schools fall within the 4 km treatment radius of any other plant in the dataset.

Among treated schools, however, roughly 5–6 percent are within the treatment radius of two or more plants. These multi-plant-treated schools experience higher actual exposure than the binary indicator captures. As a robustness check, I rerun the preferred specification dropping these multi-plant schools. The main estimate is essentially unchanged: at the 4 km radius, the point estimate moves from -0.183 (SE 0.091) to -0.176 (SE 0.098). The robustness of the estimate to dropping multi-plant exposure indicates that the result is not driven by these schools.

I also construct a continuous lead-weighted version of R_{ist} that sums treatment intensity across all plants within radius, weighted by each plant's measured soil lead level. The coefficient on this lead-weighted measure is small (less than 2 percent of a standard deviation) and statistically indistinguishable from zero. Three readings are consistent with this result: (i) measurement error in soil lead concentrations attenuates the coefficient toward zero, (ii) the relationship between lead exposure and learning is non-linear in concentration so that any meaningful proximity produces similar harm, or (iii) the operative exposure pathway is not purely soil lead but also airborne lead, water, or household goods. All three readings remain consistent with the binary-proximity specification serving as the appropriate main specification.

6.7 Falsification: Non-Lead Toxic Sites

The TSIP inventory that identifies the 26 ULAB recycling plants also records 51 toxic sites whose key pollutants are not lead, primarily arsenic, mercury, and chromium contamination at artisanal mining operations and dump sites. These sites provide a falsification test of the interpretation of the main estimates. If the main result reflected generic proximity to toxic or industrial facilities, or spurious post-2007 trends in test scores near such locations, then assigning a placebo 2007 "opening" to the non-lead sites should reproduce the pattern observed around the ULAB plants.

I replicate the preferred design around the non-lead sites: schools within 4 km of a non-lead site form the placebo treated group, schools 10–20 km from the nearest non-lead site form the control group, and the placebo exposure intensity interacts placebo proximity with the same post-2007 timing measure. The specification mirrors equation (5), with standard errors clustered at the non-lead site. To keep the placebo sample free of actual lead exposure, schools within 10 km of any ULAB plant are excluded.

Appendix Table A6 reports the results. The placebo estimates are statistically indistinguishable from zero at both the 4 km radius (-0.177, SE 0.311, wild cluster bootstrap $p = 0.65$) and the 6 km radius (-0.111, SE 0.241, bootstrap $p = 0.67$). Two caveats apply. First, fewer schools lie near non-lead sites, so the placebo estimates are less precise than the main estimates. Second, arsenic and mercury are themselves neurotoxic, which makes this a conservative falsification: any true harm from the non-lead sites would push the placebo estimate away from zero, yet no significant effect appears. The contrast between the null placebo estimates and the sharp, dose-responsive pattern around the ULAB plants supports attributing the main estimates to plant emissions rather than to generic toxic-site proximity.

6.8 Treatment Heterogeneity by Gender

Because ATEs can mask substantial differences in impacts across treated groups, I explore treatment heterogeneity by gender. I estimate a pooled version of equation (5) that adds an interaction between R_{ist} and a Female dummy, taking Male as the baseline category (Table 5, Panel A). At the 4 km radius, the baseline effect of R_{ist} on male students is -0.274 SD (SE 0.098, p $<$ 0.01), and the differential effect for female students ($R_{ist}\times$ Female) is $+0.183$ SD (SE 0.058, p $<$ 0.01), implying that female students are less adversely affected by lead exposure than male students. The 6 km radius shows the same qualitative pattern, with the interaction still significant at the 1 percent level ($+0.154$ SD, SE 0.053).

The mechanism for this difference is consistent with differentiated gender roles in the household.[26] Boys' greater outdoor exposure plausibly translates into greater contact with airborne and soil-borne lead emissions from the recycling plants. As discussed in Section 2, lead exposure affects brain development regardless of gender, impairing cognition, attention, and short-term memory even at low exposure levels (see Aizer et al., 2018; Needleman et al., 2004, and Lanphear et al., 2000), so the observed difference operates through exposure rather than through differential biological response.

6.9 Treatment Heterogeneity by School Location Type

I next examine heterogeneity by school location type. I run a pooled version of equation (5) interacting R_{ist} with a Rural dummy, taking Urban as the baseline (Table 5, Panel B). The differential effect for rural relative to urban schools ($R_{ist}\times$ Rural) is -0.110 SD (SE 0.230) at 4 km and -0.038 SD (SE 0.210) at 6 km; neither is statistically distinguishable from zero, and I cannot reject equality of effects between rural and urban schools.

The absence of a rural-urban gap is consistent with lead emissions from the ULAB plants being a sufficiently large pollution shock that the marginal effect on test scores is detectable even where urban congestion exposes students to competing sources of pollution. The rural composition of the analysis sample also matters for interpretation: rural schools account for roughly 93 percent of student-year observations within 4 km of plants, so the ATE in column (4) of Table 2 is effectively driven by rural-school students.

6.10 Treatment Heterogeneity by School Type

A third dimension of heterogeneity is school type. I run a pooled version of equation (5) interacting R_{ist} with a Private dummy, taking Public as the baseline (Table 5, Panel C). The baseline effect for public-school students is small and statistically indistinguishable from zero (-0.028 SD, SE 0.076 at 4 km), and the differential effect for private-school students ($R_{ist}\times$ Private) is -0.667 SD (SE 0.262, $p = 0.017$) at 4 km and -0.671 SD (SE 0.239, $p < 0.01$) at 6 km. The negative sign indicates that private-school students are substantially more adversely affected by lead exposure than public-school students.

Three caveats temper the strength of this finding. First, the private-school subsample is small: roughly 86,000 student-year observations versus 738,000 for public schools at the 4 km radius. Second, the differential effect is imprecisely estimated (SE 0.262 at 4 km), and the interaction's significance is driven in part by the tightness of the public-school baseline coefficient near zero rather than by precise estimation of the private effect. Third, private schools in Kenya may be spatially concentrated in urban or peri-urban areas closer to ULAB plant locations on average, which would amplify their measured exposure relative to public schools. While these caveats argue against

[26] Boys are known to spend more of their time outdoors, either in playgrounds or the streets (in urban areas) or in farms (in rural areas). Girls, on the other hand, may spend more time indoors working on house chores, including cooking in rural areas.

treating the private-school effect as a primary finding, the magnitude and sign of the interaction warrant further investigation when richer school-level data become available.

6.11 Treatment Heterogeneity by Age Group

In my analysis, I am unable to carry out treatment heterogeneity analysis by age group due to the unavailability of relevant data. As previously discussed in Section 2, children are particularly vulnerable to lead exposure for several reasons, which highlights their increased risk relative to adults. Yet, the literature suggests that lead exposure significantly impacts learning outcomes for students of all ages, including adults. However, there are few detailed studies focusing on the differential impacts of lead among various child age groups. Consequently, this paper does not make specific claims about differential impacts among children due to these data constraints. Similarly, it is difficult to assess treatment heterogeneity analysis across students' initial performance distributions; for example, whether students initially at the bottom of the performance distribution are more adversely affected than those at the top remains an open question. This lack of detailed data points to significant gaps in the research and underscores the need for further studies when more comprehensive data become available.

6.12 Mechanisms

Lead exposure can impair learning in several ways, but two main channels stand out: health and time use. The effects observed in this study may be mediated through either of these channels individually or through a combination of both. A thorough mediation analysis to support any one of these requires several pieces of data, including school attendance records, health or student medical records, and local labor market information. However, these data are unavailable, and thus I cannot carry out the mediation analysis. Nevertheless, in this section I discuss the likely pathways given the knowledge of the local conditions.

6.12.1 Health Channel

There are three possible pathways through which learning could be impacted via the health channel. First, lead exposure compromises students' neurological features, which is a common pathway discussed in the literature. Following exposure, a child's brain development is affected, impairing their cognitive abilities and diminishing their capacity to understand classroom teaching and keep up with learning. These effects become worse if exposure remains undetected and untreated, as is the case in LMICs.

The second pathway is absenteeism. Even at moderate blood lead levels, chronic exposure can produce fatigue, headaches, irritability, and difficulty concentrating, symptoms that may plausibly increase school absences across a broad portion of the treated population. At higher levels of exposure, more severe symptoms emerge: gastrointestinal abnormalities (anorexia, nausea, vomiting, stomach cramps) are likely to occur at levels greater than or equal to 50μg/dl, while acute, high concentrations from 70μg/dl to 100μg/dl can produce severe kidney abnormalities. Coma and eventually death can occur at very high blood lead concentrations of 100μg/dl and above.

At blood levels greater than or equal to 40μg/dl, lead affects the activity of enzymes needed for the proper functioning of red blood cells, which can result in anemia. The life span of red blood cells also decreases, further contributing to the development of anemia. Hemolytic anemia occurs when high levels of lead exposure occur over a short period of time. Additionally, levels at or below 30μg/dl affect vitamin D metabolism and calcium levels, leading to abnormal growth and bone development (WHO, 2015).

If these symptoms remain undiagnosed, as is often the case in LMICs, then the child will stay home instead of going to school. The longer they are absent from school, the more they will remain behind in their studies, leading to reduced or no learning. Currie et al. (2009) support this narrative, finding that higher pollution levels over six-week attendance periods are associated with more student absences.

The third pathway is resource reallocation. Following these major medical complications, the household may reallocate resources away from supporting the child's education to seek healthcare for them. A household may also reduce its labor supply (affecting income and thus resources available to the child) to care for the child.

6.12.2 Time Use/Reallocation Channel

The launch of the ULAB plants may provide some employment opportunities for local residents. If earnings from employment in the plants are spent locally, this could modestly improve local economic conditions. With such conditions in place, children may shift their time away from school or schoolwork to work in the plants or sell goods, as the opportunity cost of schooling rises. If this channel is operative, it would imply that the estimated treatment effect captures the *net* impact of a negative health shock and a potentially positive income shock, meaning the health-related effect of lead exposure on learning may be even larger in magnitude than the reduced-form estimate suggests.

7 Conclusion

In this paper, I examine the effects of lead exposure on student learning, leveraging a temporal and spatial variation generated by an environmental shock induced by the 2007 opening of ULAB recycling plants. I find that relative to students in control schools, those in lead-exposed schools experienced a reduction in test scores of 18.3 percent of a standard deviation at the 4 km buffer radius, with effects attenuating at larger radii. The estimated impacts are more pronounced for boys than for girls and for students attending private rather than public schools, with no detectable rural-urban gap. These estimated impacts exceed the effects of traditional interventions aimed at improving learning in LMICs.

My findings advance our understanding of the learning crisis puzzle: as long as the effect of lead on students outweighs efforts to improve learning, the crisis will persist. Better data on individual-student blood lead levels could further enrich research in this area, potentially clarifying how much of the learning crisis is driven by lead exposure. Lead exposure undermines students' ability to learn by irreversibly damaging neurological features crucial for cognitive development. There is also evidence that these effects not only persist into adulthood but also influence peers, indicating a more complex and extensive problem than previously recognized.

A limitation of this study is its reliance on administrative data assembled for purposes other than research; in particular, the plant opening dates rest on media reports rather than registry records, and no direct measures of children's blood lead levels exist for this period. Given the significance of this issue, collecting data on students' blood lead levels should be a priority. This would allow for structured, long-term study on the impact of lead in SSA, as well as more research on treatment heterogeneity across different student groups, especially across performance distributions.

Furthermore, these findings reveal the hidden costs of recycling and show that lead poisoning in LMICs is a serious threat to the human capital of future generations, a critical driver of economic growth. The acceleration of lead-based pollution exacerbates these concerns. Given the permanent damage caused, such pollution can worsen the economic growth prospects of LMICs in SSA. Thus, in the long run, the cost of tolerating uncontrolled lead-based pollution from recycling and manufacturing plants could lead to stagnated economic growth.

Data Availability

The lead-exposure data come from the Toxic Sites Identification Program (TSIP) inventory (Pure Earth, 2021), which is publicly available from Pure Earth at `https://www.contaminatedsites.org/`. The 2007 Kenya primary school locations and characteristics (Kenya Ministry of Education, 2007) are publicly available from the World Bank Data Catalog at `https://datacatalog.worldbank.org/search/dataset/0038039/kenya-schools`. The Kenya Certificate of Primary Education (KCPE) examination records (Kenya National Examinations Council, 2000–2010) are restricted-use administrative microdata maintained by the Kenya National Examinations Council; I obtained them from a third party under an agreement that does not permit redistribution, though researchers may request them directly from the Kenya National Examinations Council. All code required to reproduce the results from these sources is provided in the replication package.

References

Aizer, A., J. Currie, P. Simon, and P. Vivier (2018): "Do Low Levels of Blood Lead Reduce Children's Future Test Scores?" *American Economic Journal: Applied Economics*, 10, 307–341.

Angrist, N., D. Evans, D. Filmer, R. Glennerster, F. H. Rogers, and S. Sabarwal (2020): "How to Improve Education Outcomes Most Efficiently? A Comparison of 150 Interventions using the New Learning-Adjusted Years of Schooling Metric," Policy Research Working Paper 9450, World Bank.

Banerjee, A. V. and E. Duflo (2007): "The Economic Lives of the Poor," *Journal of Economic Perspectives*, 21, 141–168.

Basri, M. C., M. Felix, R. Hanna, and B. A. Olken (2021): "Tax Administration versus Tax Rates: Evidence from Corporate Taxation in Indonesia," *American Economic Review*, 111, 3827–71.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008): "Bootstrap-Based Improvements for Inference with Clustered Errors," *The Review of Economics and Statistics*, 90, 414–427.

Crawfurd, L., R. Todd, S. Hares, J. Sandefur, and R. S. Bonnifield (2023): "How Much Would Reducing Lead Exposure Improve Children's Learning in the Developing World?" CGD Working Paper 650, Center for Global Development.

Currie, J., E. A. Hanushek, E. M. Kahn, M. Neidell, and S. G. Rivkin (2009): "DOES POLLUTION INCREASE SCHOOL ABSENCES?" *The Review of Economics and Statistics*, 91, 682–694.

de Chaisemartin, C. and X. D'Haultfœuille (2020): "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 110, 2964–96.

Duflo, E. (2001): "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment," *The American Economic Review*, 91, 795–813.

Glewwe, P. and M. Kremer (2006): "Schools, Teachers, and Education Outcomes in Developing Countries," in *Handbook of the Economics of Education*, ed. by E. Hanushek and F. Welch, Elsevier, vol. 2 of *Handbook of the Economics of Education*, chap. 16, 945–1017.

Glewwe, P., M. Kremer, and S. Moulin (2009): "Many Children Left Behind? Textbooks and Test Scores in Kenya," *American Economic Journal: Applied Economics*, 1, 112–135.

Goodman-Bacon, A. (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225, 254–277, themed Issue: Treatment Effect 1.

Gottesfeld, P., F. H. Were, L. Adogame, S. Gharbi, D. San, M. M. Nota, and G. Kuepouo (2018): "Soil contamination from lead battery manufacturing and recycling in seven African countries," *Environmental Research*, 161, 609–614.

Hainmueller, J. (2012): "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies," *Political Analysis*, 20, 25–46.

Hollingsworth, A., M. Huang, I. J. Rudik, and N. J. Sanders (2020): "Lead Exposure Reduces Academic Performance: Intensity, Duration, and Nutrition Matter," Working Paper 28250, National Bureau of Economic Research.

Jakiela, P. (2021): “Simple Diagnostics for Two-Way Fixed Effects,” .

Kenya Ministry of Education (2007): “Kenya – Schools (Primary School Locations and Characteristics),” Washington, DC: World Bank Data Catalog, Dataset ID 0038039, https://datacatalog.worldbank.org/search/dataset/0038039/kenya-schools (accessed 2021).

Kenya National Examinations Council (2000–2010): “Kenya Certificate of Primary Education (KCPE) Examination Records,” Nairobi: Kenya National Examinations Council, restricted-use administrative microdata; obtained from a third party under a non-redistribution agreement.

Landrigan, P. and D. Bellinger (2021): “It’s Time to End Lead Poisoning in the United States,” *JAMA Pediatrics*, 175.

Lanphear, B. P., K. Dietrich, P. Auinger, and C. Cox (2000): “Cognitive Deficits Associated with Blood Lead Concentrations le than 10 mug/dL in US Children and Adolescents,” *Public Health Reports (1974-)*, 115, 521–529.

Lee, D. S. (2009): “Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects,” *The Review of Economic Studies*, 76, 1071–1102.

Lucas, A. M., P. J. McEwan, M. Ngware, and M. Oketch (2014): “IMPROVING EARLY-GRADE LITERACY IN EAST AFRICA: EXPERIMENTAL EVIDENCE FROM KENYA AND UGANDA,” *Journal of Policy Analysis and Management*, 33, 950–976.

Meyer, P. A., M. A. McGeehin, and H. Falk (2003): “A global approach to childhood lead poisoning prevention,” *International Journal of Hygiene and Environmental Health*, 206, 363–369.

Ministry of Health (2015): “Report on the Health Assessment of Communities Living around Battery Recycling Plants,” Tech. rep., Republic of Kenya, Ministry of Health.

Needleman, H. L., P. J. Landrigan, and M. J. Brown (2004): “Discussion of Susan M. Bernard, Should the Centers for Disease Control and Prevention’s Childhood Lead Poisoning Intervention Level Be Lowered?.” *American Journal of Public Health*, 94, 8 – 9.

Pure Earth (2021): “Toxic Sites Identification Program (TSIP) Database: Kenya,” New York: Pure Earth (formerly Blacksmith Institute); program commissioned by the United Nations Industrial Development Organization, https://www.contaminatedsites.org/ (accessed 2021).

Rambachan, A. and J. Roth (2023): “A More Credible Approach to Parallel Trends,” *The Review of Economic Studies*, 90, 2555–2591.

Roodman, D., M. Ø. Nielsen, J. G. MacKinnon, and M. D. Webb (2019): “Fast and Wild: Bootstrap Inference in Stata Using boottest,” *The Stata Journal*, 19, 4–60.

Tanaka, S., K. Teshima, and E. Verhoogen (2022): “North-South Displacement Effects of Environmental Regulation: The Case of Battery Recycling,” *American Economic Review: Insights*, 4, 271–88.

UNIDO (2009): “Kenya: Toxic Site Identification Program,” Tech. rep., United Nations Industrial Development Organization, implemented by Pure Earth (formerly Blacksmith Institute).

WHO (2015): *Lead exposure in African children: contemporary sources and concerns*, World Health Organization. Regional Office for Africa.

World Bank (2018): *World Development Report 2018: Learning to Realize Education's Promise*, World Bank: the International Bank for Reconstruction and Development.

World Bank, T. (2020): *World Development Report 2020: Trading for Development in the Age of Global Value Chains*, World Bank: the International Bank for Reconstruction and Development.

Table 1: Balance on Observables: Treated versus Control Schools (Estimation Sample)

	Treated (≤4 km) (1)	Control (10–20 km) (2)	Standardized difference (3)
Female student	0.489	0.484	0.010
Rural school	0.778	0.975	−0.474
Private school	0.180	0.073	0.279
Mixed-gender school	0.980	0.984	−0.027
Boys-only school	0.014	0.011	0.026
Pupil-teacher ratio	38.97	42.02	−0.126
Pupil-classroom ratio	41.76	42.08	−0.018
Male government teachers	4.50	5.58	−0.303
Female government teachers	11.92	5.88	0.646
Total enrollment	773.5	534.4	0.452
Number of classrooms	17.39	12.66	0.611
Number of teachers	18.85	12.47	0.661

Notes: Columns (1) and (2) report means over student-year observations in the estimation sample at the 4 km treatment radius: treated schools lie within 4 km of a ULAB plant, control schools 10–20 km from the nearest plant. Column (3) reports the standardized difference (the difference in means divided by the treated-group standard deviation). School covariates are from the 2007 Kenya primary schools dataset. Treated and control schools differ on several dimensions; Section 5 discusses why these level differences do not threaten the design, and Appendix Table A1 reports covariate balance after entropy reweighting.

Table 2: Average Treatment Effects of Lead Exposure on Learning Outcomes

	Dep. Var.: Standardized Test Scores (Z)			
	10 KM RADIUS	8 KM RADIUS	6 KM RADIUS	4 KM RADIUS
	(1)	(2)	(3)	(4)
Exposure Intensity (R)	-0.0678	-0.0854	-0.0706	-0.1834^{*}
	(0.093)	(0.098)	(0.088)	(0.091)
Observations	1179182	1026831	900709	823199
Plant clusters	26	26	26	26

Notes: Each column reports estimates from a difference-in-differences regression as specified in equation (5). The dependent variable is the standardized test score. The main regressor is the exposure intensity measure R_{ist}, defined in equation (6) as the product of a binary proximity indicator and the fraction of primary schooling years during which the student was exposed to lead emissions. All specifications include controls for gender, school fixed effects, exam-year fixed effects, and district-by-exam-year fixed effects. Standard errors are in parentheses and clustered at the plant level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Intensity Check: Effects at the 2 km Treatment Radius

	Dep. Var.: Standardized Test Scores (Z)
	2 km radius
	(1)
Exposure Intensity (R)	-0.3825^{***}
	(0.130)
Observations	734,036
Plant clusters	26

Notes: Same specification as Table 2, with the treatment buffer tightened to schools within 2 km of a plant and the control group unchanged (schools 10–20 km from the nearest plant). The dependent variable is the standardized test score. Standard errors are in parentheses and clustered at the plant level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Residential Sorting Test: Lead Exposure and the Number of KCPE Exam Takers

	Dep. Var.: Log of Number of Test Takers			
	10 km radius	8 km radius	6 km radius	4 km radius
	(1)	(2)	(3)	(4)
Exposure Intensity (R)	−0.107	−0.041	−0.055	−0.135
	(0.125)	(0.131)	(0.136)	(0.178)
	[0.398]	[0.756]	[0.688]	[0.457]
Observations (school-years)	30,322	26,676	23,608	21,528
Schools	3,519	3,097	2,750	2,516
Plant clusters	26	26	26	26

Notes: Each column reports a school-year-level regression of the log number of KCPE exam takers on the exposure intensity measure R_{ist}, with school fixed effects, exam-year fixed effects, and district-by-exam-year fixed effects as in equation (5). The sample in each column consists of treated schools within the indicated radius and control schools 10–20 km from the nearest plant. Standard errors, in parentheses, are clustered at the recycling plant level; p-values in brackets.

Table 5: Treatment Effect Heterogeneity by Gender, Location, and School Type

	Dep. Var.: Standardized Test Scores (Z)	
	(1)	(2)
	4 km radius	6 km radius
Panel A. Gender (baseline group: Male)		
R_{ist}	-0.274^{***}	-0.146
	(0.098)	(0.091)
$R_{ist}\times$ Female	0.183^{***}	0.154^{***}
	(0.058)	(0.053)
Panel B. School Location (baseline group: Urban)		
R_{ist}	-0.102	-0.041
	(0.147)	(0.146)
$R_{ist}\times$ Rural	-0.110	-0.038
	(0.230)	(0.210)
Panel C. School Type (baseline group: Public)		
R_{ist}	-0.028	0.055
	(0.076)	(0.080)
$R_{ist}\times$ Private	-0.667^{**}	-0.671^{***}
	(0.262)	(0.239)
Observations	823,199	900,709
Plant clusters	26	26

Notes: Each panel reports the indicated coefficients from a single pooled difference-in-differences regression (Equation 5) at the given radius, run with R_{ist}, its interaction with a moderator dummy (Female / Rural / Private), and the moderator's main effect. The Rural and Private main effects are absorbed by the school fixed effects because they are time-invariant school-level characteristics; the Female main effect is identified within school. R_{ist} measures the effect of exposure intensity for the baseline group (Male / Urban / Public); the interaction $R_{ist}\times$ Moderator gives the differential effect for the listed subgroup. A positive interaction coefficient therefore indicates that the subgroup is *less* adversely affected by lead exposure than the baseline group, and a negative coefficient indicates the opposite. All regressions include student gender, school fixed effects, exam-year fixed effects, and district-by-exam-year fixed effects. Plant-clustered standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 1: Owino Uhuru battery recycling plant in Mombasa

The lead-acid battery smelter, visible in the background of this photo, lead to a mass poisoning in Owino Uhuru, a village in Momb

Photo credit: Phylis Omido

Figure 2: Illustration of the Donut Design: Schools Around the Owino Uhuru Plant

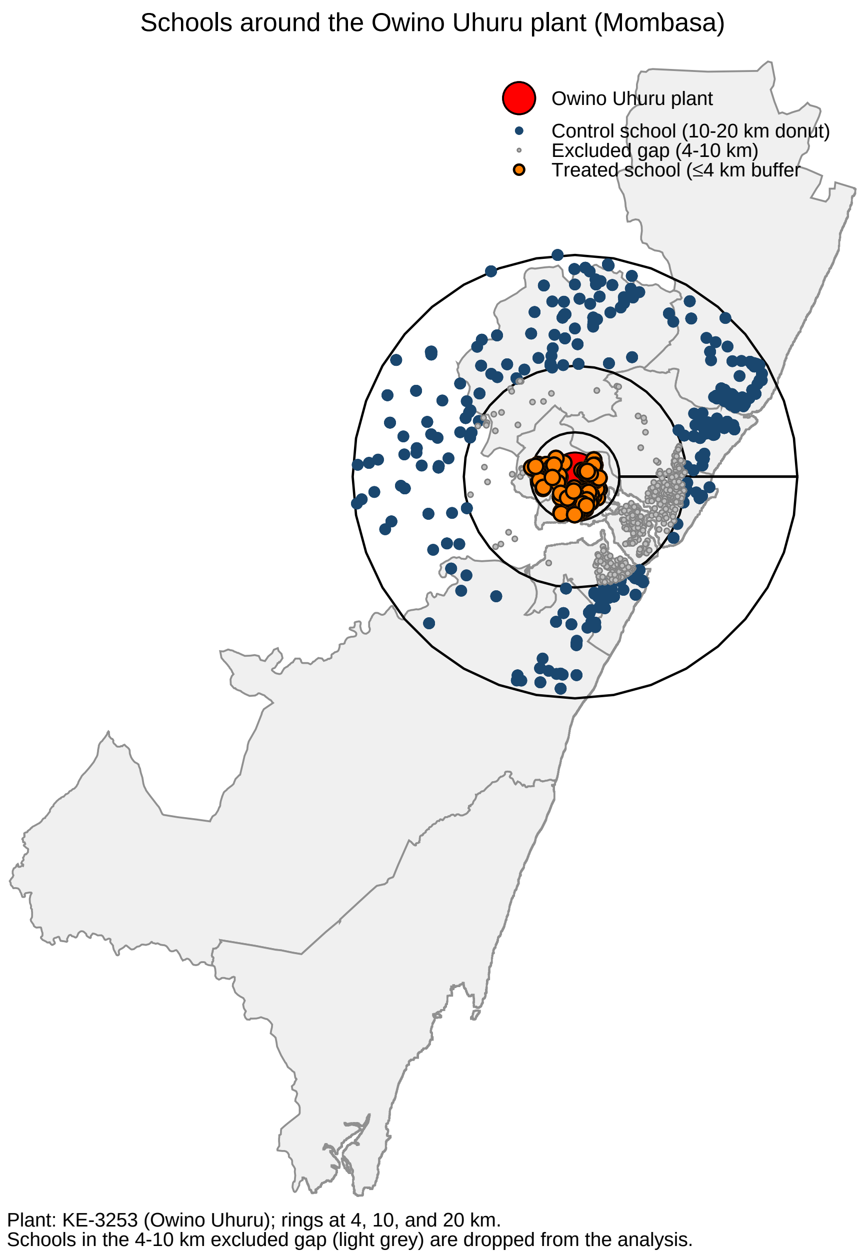

Notes: Concentric rings drawn at 4 km, 10 km, and 20 km from the Owino Uhuru ULAB plant (KE-3253) in Mombasa. Treated schools (orange) lie within the 4 km buffer; control schools (navy) lie in the 10–20 km donut; schools in the 4–10 km excluded gap (light grey) are dropped from the analysis to limit control contamination. Base map: Mombasa, Kwale, and Kilifi sub-counties from GADM 4.1.

Figure 3: Event Study Results (with 4 km radius): Effects of Lead Exposure on Learning

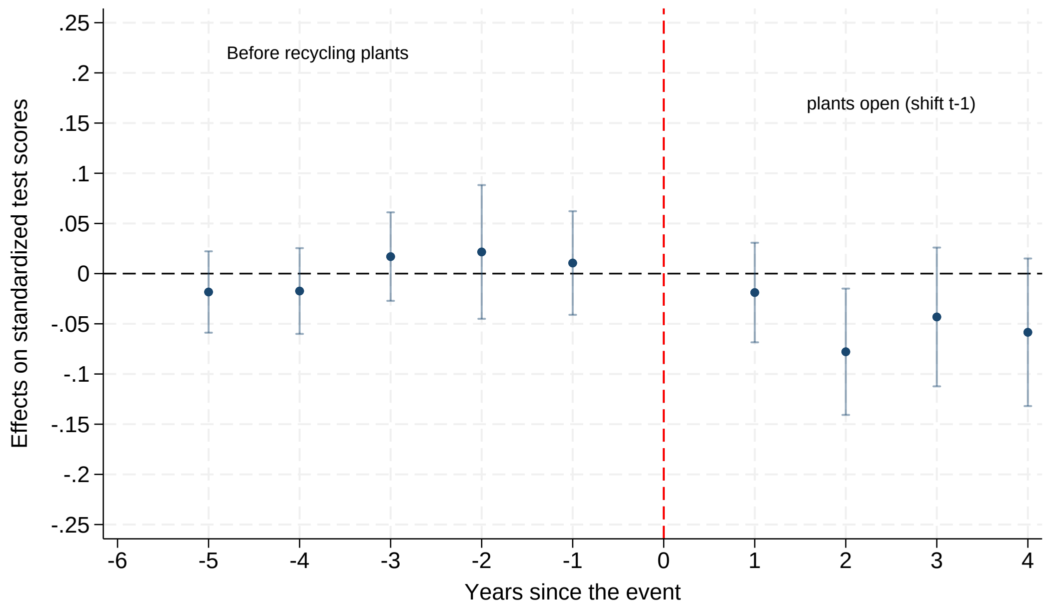

Figure 4: Sensitivity Analysis (with 6 km radius): Effects of Lead Exposure on Learning

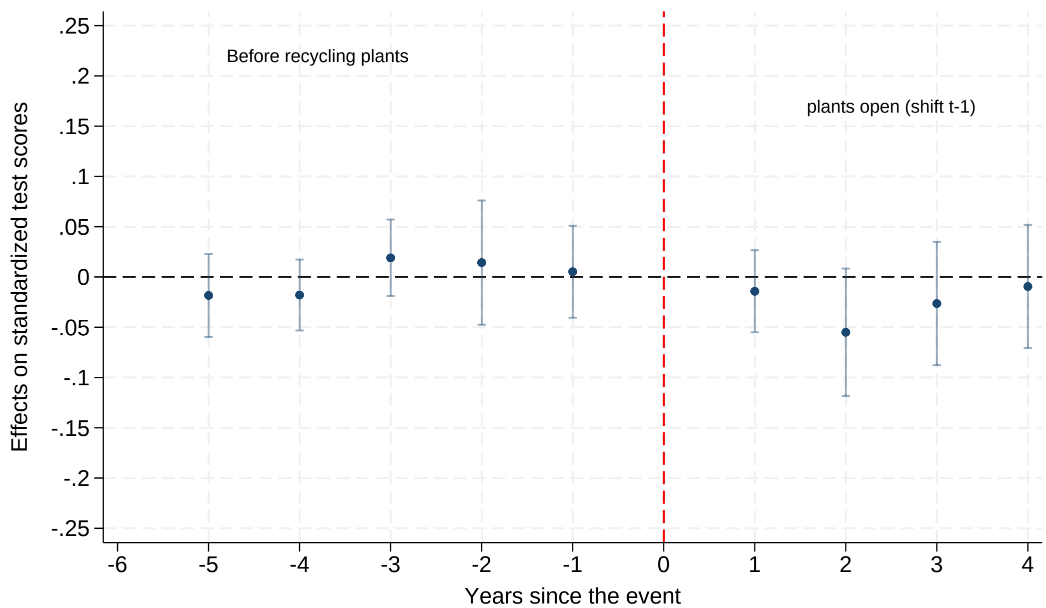

Figure 5: Sensitivity Analysis (with 8 km radius): Effects of Lead Exposure on Learning

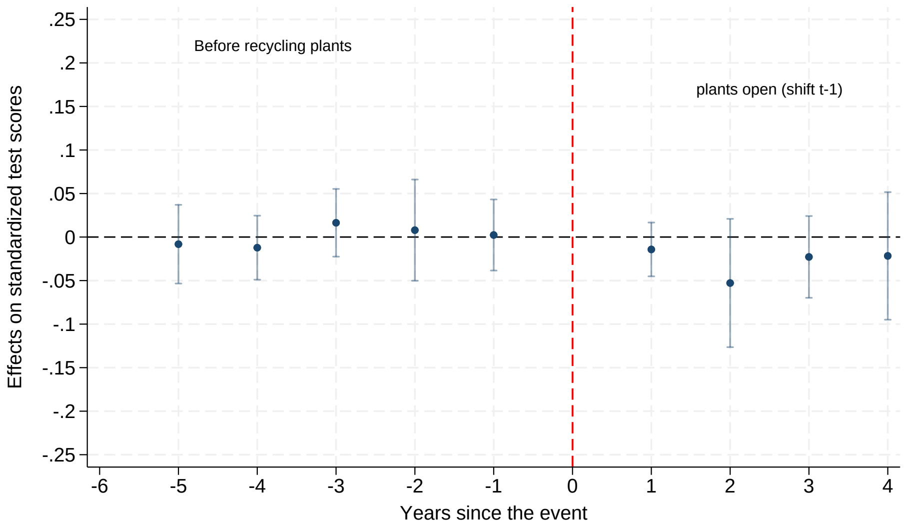

Figure 6: Sensitivity Analysis (with 10 km radius): Effects of Lead Exposure on Learning

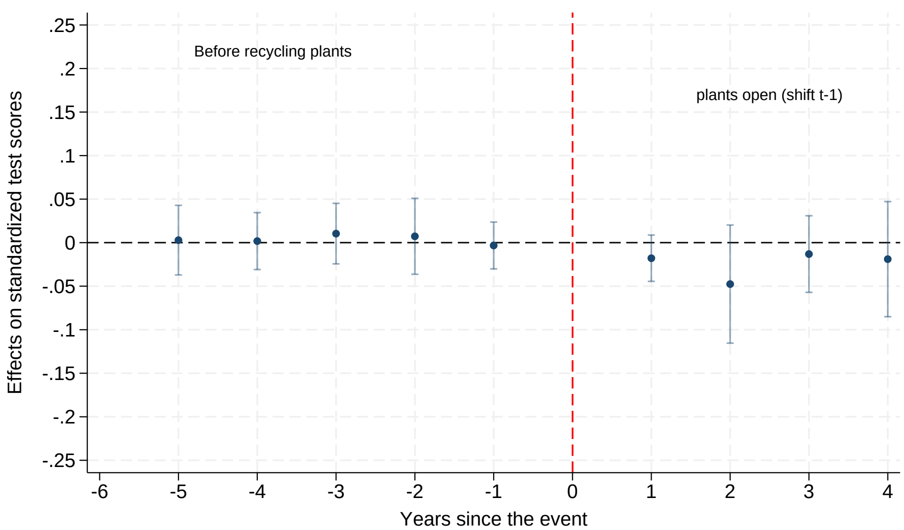

Figure 7: Distribution of the 26 ULAB Recycling Plants Across Kenya

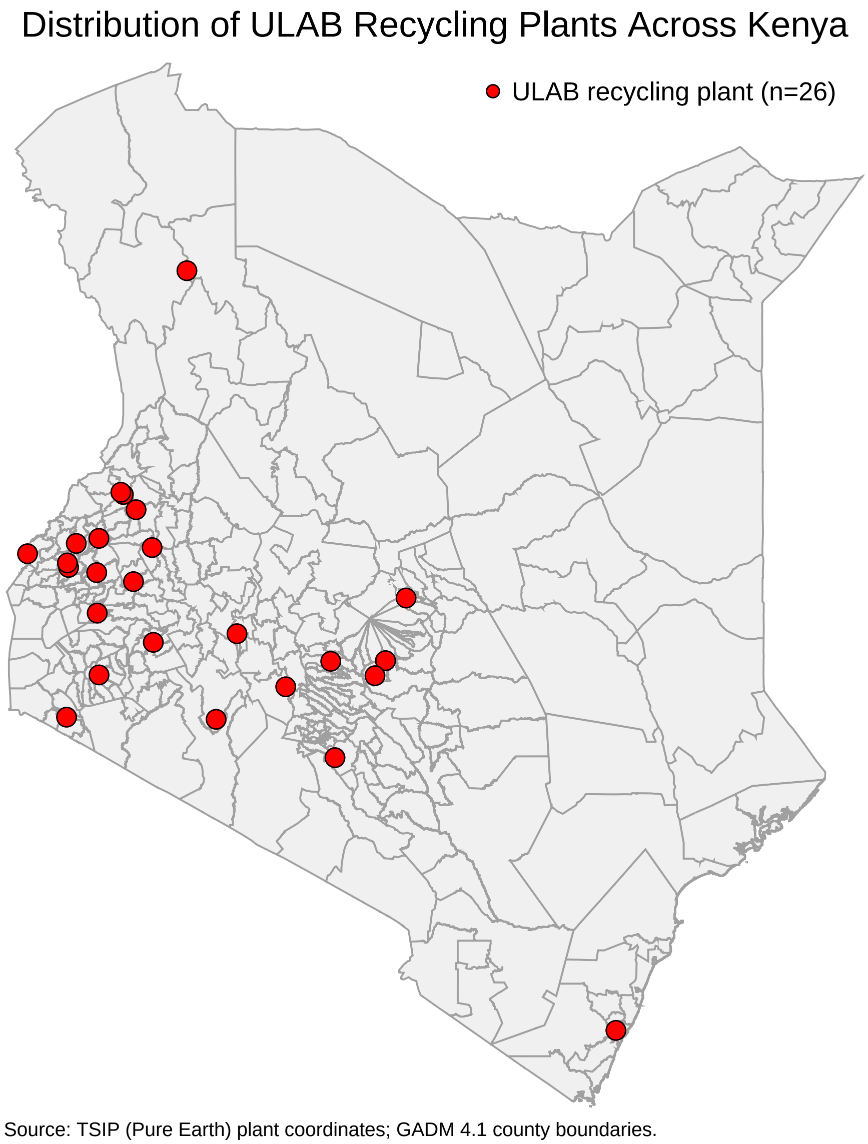

Notes: Red dots mark the 26 used lead-acid battery (ULAB) recycling plants identified by the toxic site identification program (TSIP) administered by the UN Industrial Development Organization and implemented by Pure Earth. Plants are most heavily concentrated in western Kenya near Lake Victoria; one operates on the coast near Mombasa (the Owino Uhuru plant). County boundaries are GADM 4.1.