

Essays on Air Toxics: Distribution, Drivers, and Consequences

By

Irene Baldwin Jacqz

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Agricultural and Applied Economics)

at the

UNIVERSITY OF WISCONSIN-MADISON

2020

Date of final oral examination: 06/22/2020

The dissertation is approved by the following members of the Final Oral Committee:

Daniel Phaneuf, Professor, Agricultural and Applied Economics
Corbett Grainger, Associate Professor, Agricultural and Applied Economics
Sarah Johnston, Assistant Professor, Agricultural and Applied Economics
William Provencher, Professor, Agricultural and Applied Economics
Jeffrey Smith, Professor, Economics

Acknowledgments

This dissertation represents, among other things, an accumulation of support and input from others, without which the work would not have been possible.

I am grateful to Dan Phaneuf, for his unconditional encouragement when I proposed scrapping my third-year paper and shifting my research agenda towards environmental inequality, and who has been an excellent sounding board and advocate for my work since then and throughout the degree.

My committee in Agricultural and Applied Economics, Corbett Grainger, Sarah Johnston, and Bill Provencher, engaged carefully and seriously with my research, and provided much-appreciated mentorship in teaching. I felt valued as a student and researcher in a collegial and supportive department, and am especially thankful for the work of Laura Schechter, Barbara Forrest, and Mary Treleven.

I would like to acknowledge the welcoming interdisciplinary community I found at the Institute for Research on Poverty, and its support of this research. I am grateful to Jeff Smith in particular, who was thoughtful in his feedback and generous with his time.

My cohort in environmental economics and on the job market made the hard things easier, the confusing things clearer, and the absurd things funnier. I am grateful for Nicholas Hudson, Jared Hutchins, Matt Klein, Adam Theising, and Justin Winikoff.

I would be nowhere without my family, Susan, Christian, Sarah, and Henry, who are constantly teaching me what it looks like to be hardworking, curious, caring, and wise.

Above all, I am grateful for Luke Seaberg, whose partnership has been a constant source of love, support, patience, perspective, and joy.

Abstract

Environmental amenities—such as access to lead-free water, proximity to green space, or the absence of unhealthy levels of smog—vary widely across the United States and can affect mental and physical health. This dissertation asks how firms affect disparities in exposure to toxic air pollution, and whether exposure early in childhood diminishes later cognition.

In my first essay, I study the effects of heightened toxic air pollution across cohorts at public elementary schools in the United States, linking the Risk Screening Environmental Indicators model of the Toxics Release Inventory to standardized testing data from EDFacts. I find harmful effects on test scores from exposure to airborne chemicals in early childhood, and particularly chromium.

I build on these themes in the second essay, asking how toxicity shocks in infancy affect cognitive development throughout early childhood. I use the individual-level Birth Cohort of the Early Childhood Longitudinal Survey, and exploit variation in the timing of children's births within zip codes to identify causal effects. I test whether household behaviors and characteristics mitigate or exacerbate that initial exposure. While I find no behavioral response to pollution exposure, the magnitude of the effect of neonatal toxicity exposure on later cognition varies with household income. This finding suggests removing spatial inequalities in pollution exposure does not entirely alleviate the contribution of environmental harm to socioeconomic status-associated inequalities in outcomes.

My third essay turns to the role of firms in the disproportionate exposure of disadvantaged neighborhoods to environmental hazards. I link Toxics Release Inventory facilities to their parent companies and to corporate merger and acquisition activity. This allows me to estimate the effect of acquisitions on both facility-level high-risk air pollution and its firm-level distribution, using variation in the timing of acquisition among acquired facilities. I find evidence emissions fall following an acquisition and are redistributed to facilities in more-disadvantaged neighborhoods.

Taken together, these essays highlight the continued importance of understanding both the causes and consequences of inequality in exposure to toxic air pollution.

Contents

Acknowledgments	i
Abstract	ii
1 Toxic Test Scores: The Impact of Chemical Releases on Standardized Test Performance Within U.S. Schools	1
1.1 Introduction	1
1.2 Background	3
1.3 Data	6
1.3.1 EDFacts	6
1.3.2 Toxics Release Inventory	9
1.4 Empirical approach	13
1.4.1 Identification	15
1.5 Results	19
1.5.1 Aggregate toxicity	19
1.5.2 Metal toxicity	20
1.6 Conclusion	21
1.7 Figures and tables	24
1.A Additional robustness tables	37
2 The Cognitive Cost of Toxic Chemicals in Early Childhood	46
2.1 Introduction	46
2.2 Background and theory	51

2.2.1	Toxic pollution and human capital formation	53
2.3	Data	57
2.3.1	Birth Cohort of the Early Child Longitudinal Survey	57
2.3.2	Toxics Release Inventory and Risk-Screening Environmental Indicators model	60
2.3.3	Particulate matter	63
2.4	Empirical analysis	64
2.4.1	Exogeneity of TRI exposures	68
2.4.2	Household responsiveness and heterogeneity	70
2.5	Results	72
2.5.1	Comparison to PM	73
2.5.2	Behavior and heterogeneity	74
2.6	Conclusion	76
2.7	Figures and tables	78
2.A	Production of human capital	96
2.B	Exogeneity and toxicity of TRI exposures	98
2.C	Risk-Screening Environmental Indicators model	105
2.D	Exposure attribution and timing	107
3	Pollution and Acquisition: The Environmental Justice Effects of Mergers	119
3.1	Introduction	119
3.2	Data	125
3.2.1	Toxics Release Inventory	125
3.2.2	Linked merger, business, and Census tract data	127
3.2.3	Summary statistics	129
3.3	Empirical analysis	130
3.3.1	Two way fixed effect estimates	131
3.3.2	Event study estimates	131
3.3.3	Interaction-weighted estimates	132

3.3.4	Facility closures	134
3.3.5	Distributional outcomes	135
3.4	Results	138
3.4.1	Target facility effects	139
3.4.2	Distributional effects	142
3.4.3	Acquiring facility effects	144
3.5	Conclusion	145
3.6	Figures and tables	147
3.A	Additional figures and tables	168

Chapter 1

Toxic Test Scores: The Impact of Chemical Releases on Standardized Test Performance Within U.S. Schools

Abstract

This paper estimates the effects of exposure to toxic chemicals, and individual airborne metals in particular, in early childhood on school-level standardized test performance a decade later. The analysis links the cohort-level proficiency of primary school students born in the United States in the early 2000s to Toxics Release Inventory exposure, exploiting variation in the timing and magnitude of toxicity risk within public school catchment areas to estimate the impact of early exposure on educational outcomes. Estimates of airborne toxicity may better correspond to the human health risks from the varied composition of particulate pollution. One standard deviation higher aggregate airborne toxicity in the catchment area during the year in which most students were born causes cohorts to perform 0.02 standard deviations worse on statewide tests. Finally, airborne chromium is identified as a driver of this effect.

1.1 Introduction

The hypothesis that individuals build skills in each period of life, compounding human capital from investments in earlier periods, is informed by theory in Heckman (2007) and Cunha and Heckman (2016). The most extreme version of the theory that early childhood matters disproportionately for a child's trajectory through life is the fetal

origins hypothesis, which argues conditions from conception to birth can have lasting consequences for health and cognition (Almond, 2006; Black et al., 2007; Royer, 2009). It follows that policies minimizing threats to cognition when children are most vulnerable could have an outsized influence in ensuring all children have an equal chance at success.

Many environmental threats to fetal and infant health are well established. Increases in air pollution have negative consequences for contemporaneous health and cognitive performance (Chay and Greenstone, 2003; Currie and Neidell, 2005; Currie and Schmieder, 2009; Agarwal et al., 2010; Currie et al., 2011, 2015; Sanders and Stoecker, 2015; Ebenstein et al., 2016; Knittel et al., 2016; Persico et al., 2016; Persico and Venator, 2019). Recently, researchers have begun investigating the more difficult question of how maternal and early-childhood exposure to airborne pollution could be affecting longer-term outcomes. If exposure to air pollution during early stages of development matters for health and human capital formation over a lifetime, protecting vulnerable populations from extreme pollution becomes an issue of social mobility in addition to public health.

This paper is the first national-scale, empirical examination of the effects on later-life cognitive performance from birth-year exposure to the many airborne chemicals in the Toxics Release Inventory (TRI). I ask whether variation in local toxicants in the early 2000s diminishes student performance on standardized tests taken nearly a decade later, from 2010 to 2015, by comparing the relative proficiency of cohorts of children passing through U.S. schools over six years. I pair a toxicity-weighted measure of the risk from all TRI emissions with the U.S. Department of Education’s ED Facts, a national panel of school-level standardized test proficiency. I aggregate gridded pollution data using the School Attendance Boundary Survey, a geospatial dataset of public school catchment areas.

Identification of effects on student standardized test proficiency relies on local, temporal variation in toxicity-weighted concentrations of airborne chemicals during the years in which successive cohorts of students were born. This extends previous work exploiting plant openings and closings over time (for example, Currie et al., 2015), while exploiting toxicity-weighted concentrations of airborne pollutants better allows for identifying

variation to come from less salient changes in environmental amenities.

These national-level data also provide sufficient coverage and variation to estimate the effects of exposure to specific chemicals. I refine my findings by replacing the measure of aggregate toxicity with estimated concentrations of individual airborne metals, and find robust negative effects on later standardized test performance from chromium emissions in the first year of life. This analysis does not rely on the EPA’s inhalation toxicity weights, but a significant effect from airborne chromium exposure corroborates the toxicity rankings.

That worse educational performance may driven by exposure to airborne toxic chemicals nearly a decade earlier during vulnerable periods of development can inform the regulation of sources of toxic chemicals, and reinforce our understanding of the ways high-stakes standardized testing outcomes reflect underlying environmental inequalities. The study of specific chemical releases calls attention to the importance of accounting for the composition of particulate matter, and the geographic breadth of the data, which cover a large majority of public elementary schools in the United States, make the findings immediately relevant to national education policy.

1.2 Background

Many chemical compounds found at many different doses via several exposure pathways comprise environmental threats to health and cognition. Particulate matter (PM) has been studied extensively as a driver of adverse outcomes from conception through old age, but without widespread measurement of the *constituents* of PM, there has been little work in economics that identifies the population-wide effects of the toxicity of air emissions on long-term neurological development.

Both health generally and cognitive health in particular are implicated in the hypothesis that prenatal and neonatal environmental exposures matter. Poor health in childhood is strongly correlated with socioeconomic status (Case et al., 2002), and low-income and minority households are also more likely to live in polluted places (Ard, 2015).

Outcomes such as premature birth, low birth weight, or a diagnosis of asthma are indicators of poor health, and could affect a child’s performance in school through a channel such as reduced attendance, as in Neidell (2004), which finds a causal link between CO and asthma hospitalizations, or Currie et al. (2009), which finds CO reduces attendance for a sample of school districts in Texas. Diminished cognitive (or behavioral) skills in kindergarten means a child starts school already behind peers—and with Heckman-style dynamic complementarity of skill-building over the course of formal education, closing those gaps becomes more costly with time.

The importance of any finding that early exposures matter for future outcomes is heightened by a large body of literature correlating environmental exposures with other forms of socioeconomic disadvantage (for example, Tessum et al. (2019) find excess pollution burden, relative to consumption, among nonwhite racial-ethnic groups). This means children growing up breathing the most polluted air are likely also experiencing other insults of poverty and minority status. Household investments in health, including diagnosis, treatment, and avoidance behavior, are all costly, which means low-SES families are not in a position to mitigate the adverse impacts of environmental harm, or avoid such harm in the first place. For example, Ferrie et al. (2012) find heterogeneity in the magnitude of damage to IQ from waterborne lead exposure across income levels: the lowest-SES individuals suffered the greatest harm, while higher-SES households may have been better able to attenuate the cognitive harm from lead poisoning.

The notion that differences in environmental quality contribute to variation in educational performance is grounded in findings that exposure to a wide array of chemical toxicants harms human cognition. Recent work has focused on lead, arguably the most prominent environmental neurotoxin. Ferrie et al. (2012) highlights the interaction of labor market outcomes with lead using historic data on drinking water, and Sampson and Winter (2016) uses longitudinal blood tests to connect racial inequality in lead exposure to socioeconomic disadvantage in Chicago. Aizer et al. (2018) and Evens et al. (2015) connect lead and test scores. The first looks at policy changes that reduced blood lead levels in Rhode Island and resulting decreases in the racial disparities in test scores in

the state, and the second finds an association between blood lead levels and lower standardized test scores, controlling for some other health measures and demographics. In an examination of heavy metal exposure, Rau et al. (2015) find reduced standardized test scores from proximity to a site containing deposits of hazardous mining waste. Persico et al. (2016) looks more generally at airborne toxics and childhood cognitive outcomes in Florida using superfund sites, while Ebenstein et al. (2016) link low performance on a standardized exam to days with contemporaneous airborne particulate matter. Finally, the recent finding of Isen et al. (2017) suggests the cognitive cost from airborne pollution is disproportionately determined in the first year of life, but persists well into adulthood, with lower labor market participation and earnings decades later.

To examine the long-run impacts of birth-year exposure to pollution, Sanders (2012) follows the identification strategy of Chay et al. (2003), and Isen et al. (2017) follows Chay and Greenstone (2003). Instrumenting for pollution with the 1980s recession, the former finds negative impacts of TSP on high school test scores, and the later finds reduced labor force participation and earnings at age 30 for cohorts born before and after the 1970 CAA. Bharadwaj et al. (2017) links fetal exposure to CO and PM to lower math and language skills in fourth grade for a panel of siblings in Santiago born between 1992 and 2001. In a paper using the ECLS-K, Marcotte (2017) finds unequal early childhood exposure to PM and pollen may contribute to school readiness gaps in kindergarten. Rosales-Rueda and Triyana (2018) find the most persistent health effects from exposure to wildfires in Indonesia among children exposed in utero. Collectively, these previous efforts find convincingly negative effects from early childhood exposure to particulate matter. But most PM data do not allow researchers to examine the heterogeneity of airborne pollution; hazardous constituents, such as heavy metals, may have disproportionate effects on cognition.

Currie (2009) and Currie et al. (2015), use the Toxics Release Inventory to show negative impacts on infant health and housing values from toxicant releases. The earlier study regresses gestation time, birth weight, and infant deaths at the county level on the pounds of chemicals released in the previous year (also at the county level) for a subset of

the toxicants reported in the TRI. The later paper exploits plant openings and closings by linking restricted business-level data from the US Census, which the authors argue provides cleaner identification, but abstracts away from estimating the marginal effect of changes in the levels of different pollutants. More evidence of short-run evidence for adverse cognitive outcomes is in Persico and Venator (2019), who find the years in which TRI facilities opened were linked to lower standardized test performance among nearby students. In contrast, this study concentrates not on contemporaneous changes in TRI exposure, but the conditions of airborne toxicity when children were very young, and identifies the effects of particular chemicals, accounting for the widespread heterogeneity in the health risks from airborne emissions.

This study includes individual toxic metals—airborne elements like arsenic, chromium, lead, and mercury—because these compounds are widely identified as threats to human health and cognition by the EPA and the World Health Organization (Tchounwou et al., 2012) and because their releases are empirically both widespread and generally uncorrelated. Airborne metals tend to be among the most fine particulates (Ravindra et al., 2008), and smaller particle sizes are associated with most of the health impacts from PM inhalation as they penetrate furthest into the body (Pope and Dockery, 2006). Lead poisoning has received particular attention for causing cognitive harm in children, but less is known about population-wide exposure to most other airborne metals, even though several are ranked even more highly than lead by the EPA for inhalation toxicity risk. The EDFacts proficiency data are nearly national, and therefore provide sufficient coverage and variation in exposures to begin to understand the potential for neurotoxic effects. This attention to the composition of airborne pollution, rather than its mass concentration, foregrounds the risks from PM comprised of highly toxic chemicals (Lippmann, 2010; Kelly and Fussell, 2012).

1.3 Data

I pair standardized test score data from the U.S. Department of Education’s EDFacts with the U.S. Environmental Protection Agency’s Risk-Screening Environmental Indicators (RSEI) to form a six-year panel in which public elementary schools are the unit of observation. The standardized test data come from assessments given in the 2009–2010 through 2014–2015 school years, and the environmental toxicant data are from 2001–2006; both sources cover the entire United States.

1.3.1 EDFacts

The EDFacts data provide proficiency percentages on annual, statewide, federally-mandated standardized tests. Each state reports school-level proficiency outcomes on its standardized tests for each combination of year (2009–2015), grade (3–8), and subject (math and reading). This measures the aggregate performance of students within each school, grade level, subject area, and year; individual students’ scores are not provided, though there are counts of students taking each test, and breakdowns for some subgroups, including by gender and major racial and ethnic groups.¹ The U.S. Department of Education compiles the states’ reports, performs quality review on the submitted data, and releases EDFacts to the public (U.S. Department of Education, 2015). A restricted use license from the National Center for Education Statistics allows me to include proficiency percentages that are suppressed or rounded in the public version of the data.

Though administration of standardized tests for certain grades and subjects is federally mandated, states’ standardized tests vary in content, difficulty, and proficiency cutoffs across grades and from year to year (described further in Fahle et al., 2017). There are two reasonable approaches to make meaningful comparisons in spite of those differences. First, an attempt can be made to standardize scores within each subject, grade, and state; this is the approach in, for example, the Stanford Education Data Archive’s

¹Counts and proficiency scores for these subgroups are consistently, but not perfectly, reported. For example, the EDFacts category for the percent of students who are “Two or more races” is missing in 7 percent of schools. Counts for the subgroup “non-white” are constructed as the difference between the total number of students and the number of white-only students, who are the most consistently measured.

district-level measures of proficiency, which link to National Assessment of Educational Progress (NAEP) scores. This sort of standardization involves extensive interpolation, and the SEDA scores are only available at the level of a school district—which may mask important school-level heterogeneity in neighborhood characteristics.

A second approach is to recognize that school-level student proficiency is an important policy outcome in its own right. Students, teachers, schools, and school districts are all held accountable by high-stakes standardized testing policies, which aim to incentivize performance and can determine, for example, student placement or school funding. Parents and researchers can easily observe school-level standardized test scores, making it a frequent measure of school quality for both home buyers and in the large literature estimating parents' willingness to pay for school quality (Black, 1999; Bayer et al., 2007). The No Child Left Behind Act (NCLB), which was in effect for all the years in my test score data, was a prominent example of federal incentives to improve student performance on standardized tests. Under NCLB, schools missing targets for student proficiency could face state and federal sanctions.² Though the use of high-stakes standardized testing has been criticized (Afflerbach, 2005; Haladyna, 2006; Wiliam, 2010), it remains influential: the Every Student Succeeds Act (ESSA), which replaced the NCLB in 2015, de-emphasizes but does not eliminate standardized testing for students in grades 3 through 8.

This paper presents several outcome variables, showing results are generally robust to either interpretation of the EDFacts proficiency data. I consider (1) the raw percent proficient reported by EDFacts at the school, grade, and subject level; (2) school-level proficiency in standard deviations relative to the test (i.e., within a state, subject, and grade level); and (3) the ranked proficiency for each school, grade, and subject among their cohort.

I standardize the reported percentage of students scoring proficient in school i , year t , subject s , and grade level g using the mean μ and standard deviation σ among all students in state S who took the exact same test (where the calculation of μ and σ use a

²See more information about NCLB here: <https://www.edweek.org/ew/section/multimedia/no-child-left-behind-overview-definition-summary.html>.

balanced panel of schools who report statewide tests in each year of the EDFacts data):

$$StandardizedScore_{itsg} = (PctProficient_{itsg} - \mu_{Stsg}) / \sigma_{Stsg}.$$

I interpret this standardized score for a group of students in a specific school, subject, and year as their performance relative to peers in the same cohort.

The third outcome is a school's ranked proficiency on a standardized test. This ordinal measure does not use the mean or spread of proficiency in a given year, but instead simply ranks schools by the percentage of students who passed, relative to all other schools with students who took the same test. Improvement in this metric requires a school to move up in the statewide rankings, which can in itself be a high-stakes outcome for teachers and schools. To compare the rankings across states with very different numbers of schools, I standardize ranks to a percentile, following

$$PercentileRank_{itsg} = 100 * (Rank_{itsg} - 1) / (N_{Stsg} - 1),$$

so the lowest-performing school is at 0 and the highest performing school is at 100. For this outcome, I reduce the sample of schools to a balanced panel (so the denominator, which counts schools in a state-subject-grade cell, is fixed over time).

1.3.2 Toxics Release Inventory

The U.S. Environmental Protection Agency's Toxics Release Inventory (TRI) is a national, annual assessment reporting the location, type, and quantity of individual chemical releases by facilities above a certain size, as mandated by the EPA, from 1988 to the present. Releases to air, water, and soil from a wide range of industries are covered, including mining, power generation, manufacturing, and hazardous waste facilities; I focus on air releases in this study.³

³The TRI is a unique source for a detailed panel of industrial releases at the national level, but it does omit some sources of airborne toxicants that have been linked to cognitive outcomes, notably non-point sources such as highways, railroads, seaports, and airports. A relatively short panel of six years ensures variability in pollution is not being driven by long-term changes in TRI reporting requirements.

The Risk-Screening Environmental Indicators (RSEI) model transforms reported TRI releases by location, quantity, and pathway into a product called “Geographic Micro-data,” which includes estimates of ambient airborne concentrations from TRI releases over 810-square-meter grid cells covering the United States. To compare risk to people over space due to hundreds of different exposures, the toxicity-weighted concentrations scale each estimated concentration of a chemical by its relative estimated inhalation toxicity risk, producing a unitless measure of aggregate potential harm from airborne TRI releases at each grid cell.

These RSEI toxicity-weighted concentrations account for transport and residence time of chemicals, where local dispersion depends on factors like weather, stack heights, and atmospheric chemistry. Exposure to toxicants is modeled rather than measured (a contrast with air quality monitor data), and the modeling process makes simplifying assumptions—such as constant, uniform release of a chemical by a facility over the course of a year—but doesn’t rely on a symmetric radius around a point source of toxic releases (Ash and Fetter, 2004), and doesn’t have to match locations to the nearest air quality monitor, which is sometimes quite far away. For this study, I use the EPA’s inhalation-specific toxicity weights, and since TRI reporting requirements have changed over time, I use subsets of chemicals for which reporting requirements were consistent for all cohorts of a given age. The RSEI model is described in detail, including the development of inhalation toxicity weights and the atmospheric dispersion of chemical releases, in its documentation (U.S. EPA, 2018), and in Ash and Boyce (2018).

Inhalation toxicity weights likely encompass some chemicals that are not neurotoxic to infants, and some that may be strongly so. In this sense, the aggregate toxicity measure—which accounts for long term risk for cancers as well as reproductive, developmental, and respiratory effects—is noisy relative to my hypothesized pathway. Conversely, the sum may be greater than its parts, if the cumulative insult of many toxic releases matters more than relatively high levels of individual toxicants.

To examine the potential contribution of individual toxicants to diminished learning outcomes, I use estimated concentrations of airborne metals reported to the TRI and

available in the RSEI “Disaggregated Microdata”. In general, causal interpretation is tricky with these disaggregate data: the individual chemical-level RSEI data are noisier than the aggregate toxicity scores, many listed compounds are chemically similar or highly collinear with others, and many chemicals are only reported in a small subset of states. However, the metals reported in the TRI generally do not have these drawbacks. They are prevalent, appearing in most states in the data (Table 1.1) and releases of each are generally uncorrelated with others, so detecting individual causal effects is plausible (Figure 1.1). This approach does not depend on the EPA’s assigned toxicity scores, but I do standardize the estimated airborne concentrations of each chemical using its distribution in the sample for comparability and interpretation (significant results are robust to alternatives).

Data accuracy is a potential issue for any study using reported TRI releases. Accurate reporting is legally required by the EPA, and the EPA conducts quality review on submitted data. Marchi and Hamilton (2006) investigate the question of TRI accuracy and find more evidence of random noise (likely due to imprecise estimation at facilities) in the data than strategic under-reporting, suggesting data inaccuracy is introducing classic measurement error rather than bias. The greatest potential threat to identification in this paper will be time-varying misreporting correlated with determinants of standardized test performance, but if these trends are in response to state-level policies, a state-by-year fixed effect should alleviate those concerns. Ultimately, while widespread and frequent monitoring—rather than modeling—of specific hazardous air pollutants would be a great benefit to empirical research, for now the TRI are one of the best available sources for national coverage of point-source industrial pollutants.

In order to attribute toxicant exposure to schools, I use school attendance boundaries from a publicly available shapefile, the School Attendance Boundary Survey from the National Center for Education Statistics, which provides substantial (though not complete) coverage of school attendance boundaries in the United States. I average the higher-resolution pollution data, weighting for grid-level population, up to the attendance boundaries for primary schools. For a school i with J grid cells within its school

attendance boundary, where a given cell j has toxicity-weighted concentration $ToxConc$, RSEI-reported population Pop , and $PctCoverage$ percent of the cell overlapping the school catchment area, the aggregate annual measure for toxicant exposure risk is

$$ToxExp_{it} = \frac{\sum_{j=1}^J ToxConc_{jt} * Pop_{jt} * PctCoverage_j}{\sum_{j=1}^J Pop_{jt} * PctCoverage_{ij}}.$$

This captures both the toxicity and the proximity of TRI releases to the population living within the attendance boundary for a school. For lack of national data on how school attendance areas have changed over time, I hold constant the 2015 school attendance boundaries.

Use of the School Attendance Boundary Survey (SABS) for this analysis allows me to calculate a population-weighted average toxicant exposure over each school’s geographic catchment area. However, coverage of schools in the SABS is incomplete (see Figure 1.2): public schools only appear in the SABS if their school district responded to a survey conducted by the National Center for Education Statistics in 2013–14, which may imply a sample selection process that compromises the external validity of the findings. About 72 percent of schools for which I have third grade test scores and geographic coordinates in the EDFacts data are represented in the School Attendance Boundary Survey. Comparing observable sociodemographic characteristics across the schools included and excluded from the SABS reveals excluded schools are, on average smaller, with comparable racial composition and county-level income (Table 1.2). Schools included in the SABS have fewer extreme proficiency outcomes than those excluded (Figure 1.3).

Changes in toxicity-weighted concentrations over time arise, for example, when local firms increase or decrease production, change pollution control practices, or experience accidental releases. The existence of sufficient time-varying exposure is supported anecdotally in Figure 1.4, which shows annual toxicant scores for schools in the Los Angeles area in the years when third graders in the EDFacts data were born. Los Angeles is representative of many urban areas in the country in that it is comprised of hotspots recording persistently high toxicity scores, large areas with consistently low toxicant exposure, and also many places which show variability over time.

States in the U.S. generally require that students are five years old in August, September, or October to start kindergarten (some states allow school districts to set their own cutoffs, but provide similar guidance); this creates a year-long window, most often starting on September 1st, during which students in the same cohort in school were born. Following the pattern $BirthYear = TestYear - Grade - 6$ allows me to attribute lagged toxicant exposure for each cohort; for example, the majority of third graders taking tests in the spring of 2013 were born in 2004. This assumes a school-level observation is comprised of typical students who do not get held back or skip a grade. The length of the panel is constrained by available test score data: a public school with complete observations between 2010 and 2015 has six cohorts of third graders with assessments in two subjects between 2010 and 2015 (Table 1.3). There are inevitably students in my data who take standardized tests in a different location than where they were born; the historic pollution exposure for part of a school-cohort is mis-attributed because of these students, but this is not a threat to identification unless unobserved characteristics of those students are correlated with local shifts in toxic emissions; this and other threats to identification are discussed in Section 1.4.1.

A descriptive preview of the data suggests high levels of environmental toxicants typically accompany educational and economic disadvantage, while affluent neighborhoods enjoy both strong public schools and high environmental quality (Figure 1.5). This means environmental harm is frequently an additional dimension of poverty for students experiencing other challenges. If toxic environments have a lasting effect on students' educational attainment, this cognitive externality should be taken seriously by schools and school districts evaluating standardized test performance in heavily polluted places. Summary statistics for the data are presented in Table 1.2.

1.4 Empirical approach

An appropriate regression design to isolate the effect of TRI exposure on cognitive ability depends on the hypothesized pathway from local toxic releases to students' performance

on tests. Toxicant harm could be cumulative (damage in each year of life matters, possibly in an additive way), contemporaneous (toxic releases to the environment in the year of testing matter), or age-dependent (toxic releases in specific, sensitive years of life matter). Building on previous findings in the literature I test the final pathway, examining a narrow window during which exposure to toxicants likely affects cognition in the long run. My reduced form analysis is agnostic about the exact pathway from chemical releases in the year of birth to standardized test performance; direct cognitive harm or reduced school attendance due to physical health issues are both plausible.

My main regressions isolate the causal effect of toxicant exposure during a cohort's year of birth on its future performance on standardized tests, relative to peers in the same school, grade, and subject area:

$$Y_{itg} = \theta \text{ToxExp}_{itsg} + \beta_g X_{itsg} + \alpha_{isg} + \alpha_{Sts} + \epsilon_{its} \quad (1.1)$$

where a unit of observation is a cohort of students in public school i , year t , subject area s , and grade g .

The outcome variables Y , as described in Section 1.3, measure the performance of students relative to peers taking the same statewide assessment, in either un-transformed proficiency percentages, standard deviations of the proportion of proficient students, or the ranked performance of a cohort of students. The EDFacts report proficiency separately for math and reading assessments. If I think each standardized test is a noisy measure of students' overall cognitive abilities, it makes sense to control for the individual test with fixed effects (and sometimes test-specific standardization) and then pool these scores into one regression. If, however, the effects of early exposure to environmental toxicants are meaningfully different for the development of math and reading skills, it makes sense to estimate separate coefficients θ for each subject. The literature is not definitive on this point, so I present both.⁴

The independent variable ToxExp is the inhalation toxicity-weighted concentration

⁴In a study of particulate matter exposure among adults in China, Zhang et al. (2018) find larger effects on verbal assessments than math, and write this is because "air pollution has a stronger effect on white matter (required more by verbal tests) than on gray matter (required more by math tests)".

of airborne pollutants, averaged over the school’s catchment area. The RSEI toxicity measures are unitless, only indicating relative risk to human health, and their distribution is strongly right-skewed (Panel A of Figure 1.6). I transform this variable twice in order to make results readily interpretable in terms of typical changes in airborne toxicity risk across the schools in the panel by first taking a log of the toxicity scores (Panel B), and finally converting this more-symmetric distribution to Z -scores (Panel C). Modeled concentrations (in $\mu\text{g}/\text{m}^3$) of the individual airborne metals are meaningful, but still hard to interpret, so I again standardize relative to the observed distribution of each chemical. Robustness checks in Appendix 1.A show the results are robust to alternate specifications of the independent variables. Coefficient θ is the effect of increases in Z -scores of school catchment-level toxicity when a cohort is in infancy on that cohort’s standardized test proficiency in elementary school.

The controls X are observable characteristics of a cohort as reported in EDFacts, quadratics in both county temperature and precipitation in the pollution year, and county-level economic indicators in the test year. I use 5-year ACS county estimates for median household income, unemployment rate, and population in order to keep smaller counties in the data. The coefficient θ is allowed to vary by subject. the fixed effects α_{isg} and α_{Stsg} are for individual schools by grade level and subject, and for states by subject, year, and grade, respectively.

Because my analysis links individual schools with local pollution nearly a decade ago, it will be useful to focus on third grade students: first, elementary schools have the smallest geographic footprint, and therefore more precise pollution attribution, and second, third graders are the youngest students in the EDFacts data, and therefore least likely to have moved since birth. (Additionally, including standardized tests for additional grades doesn’t include many more individuals, but instead mostly the same cohorts at older ages – see Table 1.3.)

I expect both pollution and standardized test scores to be correlated both within schools over time and across schools, so I cluster standard errors at the school district level. I weight regressions using the number of students in each test-taking cohort, since

student proficiency is informed by a larger number of students, but results are robust to omitting these weights.

1.4.1 Identification

I expect air pollution where a household lives will be related to the household's education, occupation, income, health, and preferences over correlated neighborhood attributes (such as density, crime, and property values). In this context, claiming a causal effect of toxic exposure on test score performance relies on controlling for correlated, household- and neighborhood-level effects that also determine student ability. Put differently, annual variability in the lagged concentrations of airborne environmental toxicants within a school attendance boundary must enter this setting as exogenous conditional on included controls.

Uncontrolled-for variables that correlate with TRI releases could introduce bias that either attenuates or exaggerates the effect of exposure to toxicity. Much work has been done identifying correlations between low socioeconomic status (SES) and high exposure to pollution in the United States; this type of correlation might bias upwards the magnitude of an effect of toxicant exposure on student ability. Conversely, higher levels of pollution may correlate with local economic activity that generates higher income for families and perhaps funding for schools; this type of effect could bias downward the magnitude of an estimate for cognitive health impacts.

The panel fixed effects design exploits the intertemporal variability in toxic releases within the catchment area for individual schools, over a short panel: identification rests on changes in the levels of toxicant exposure not affecting the households' (and in particular, new parents') mobility decisions in a way that causes toxic releases to be correlated with unobserved differences in student ability across otherwise-similar cohorts nearly a decade later. A state-by-year fixed effect, present in all regressions, absorbs average changes in reporting behavior over time. Including county-level weather in the year of airborne emissions should help control for the effect of temperature and precipitation on the dispersion of air pollution, and including county-level population, income, and unem-

ployment should help control for time-varying differences in regional economic activity. A relatively short panel also mitigates the possibility that longer-term changes (such as deindustrialization and gentrification) are driving results.

Household turnover is likely strongest when changes in airborne pollution levels are salient, substantial, and perceived as permanent, as when a plant opens or closes in a neighborhood; Currie et al. (2015) find a housing market response to the opening and closing of TRI plants, though only among houses inside half a mile of the site. The authors find significant concentrations of pollutants at a wider radius than the sorting response—up to a mile from toxic plants—so there may be many changes in exposures outside of the salience radius that affect children but do not drive a sorting response (the RSEI dispersion model estimates atmospheric concentrations within a 30-mile radius of a plant). This echoes the finding in Spencer Banzhaf and Walsh (2008), who present evidence of demographic change over a decade within a half mile of TRI openings and closures in California. Unlike those papers, this analysis exploits changes in facility output rather than openings, and most plants remain in operation.

Still, any evidence of neighborhood composition changing in response to variation in toxicant exposures presents a possible threat to identification, and the assumption that most students don't move is violated in the worst case if moves are correlated with unobserved characteristics of students that drive test scores (for example, if an increase in toxic releases causes affluent families to leave, or conversely an increase in environmental quality drives gentrification and more wealthy families).

The fact that I use a long lag, rather than contemporaneous pollution exposure, diminishes the confounding effect of household sorting in response to changes in environmental quality. This is because the timing of exposure is particularly important for identification: I pair cohorts with the ambient pollution just in their year of birth, and compare them to cohorts born at most few years earlier and later, who take tests in the same public school. Households therefore have nearly a decade to “vote with their feet” after birth-year emissions are realized and before standardized tests are taken—relocation does not introduce bias unless households with children born exactly in the year of higher releases move at

a higher rate than households with kids just a few years older and younger. A more realistic process of household sorting is that the families who move in and out of a school catchment area—even potentially in response to changing environmental quality—have children at many ages. Put differently, the identifying assumption is that the effect of TRI releases on household sorting is constant across all cohorts within in a school after nearly a decade, but the effect of TRI releases on cognition is uniquely harmful for the cohort exposed in their birth year.

It is possible to test for violations of the “no correlated household mobility” assumption. In particular, if the demographic characteristics across waves of elementary school students in the EDFacts data are correlated with chemical releases in their years of birth, there would be compelling evidence cohorts of students within one school are not good controls for one another. Regressing observable attributes of students (as reported to EDFacts) on lagged local toxicant releases shows how school-level panel fixed effects control for correlations between household demographics (race and income) and exposure to environmental harm (columns (3) and (6) of Table 1.4); this is consistent with the assumption that other unobserved child attributes are unrelated to variation in the birth year toxicant exposure experienced across cohorts.

Within a state for a given year, the birth-year toxicant exposures of cohorts are correlated with available demographic characteristics (proportion of minority groups, class size, economic disadvantage, and limited English proficiency); this is expected and consistent with widespread patterns of household sorting over space. Within school districts, environmental justice correlations persist—schools reporting more nonwhite students tended to have higher exposure to airborne toxicity within their district— but these correlations are not present with the inclusion of school fixed effects. The finding that class size is not correlated with lagged toxic releases also addresses a possible concern that toxicant exposure drives students to start school late, or repeat grades—if this was the case, smaller third grade class sizes might correlate with high-pollution birth years.

Beyond threats to identification, there are some unavoidable sources of measurement error that come from using elementary schools as the unit of observation. I only observe

traditional public schools, so students in private, magnet, or charter schools, which do not have geographic catchment areas, are not included in the data. I also have to approximate the age of students in each grade: for example, a typical individual in third grade in the spring of 2010 was born in 2001, but some of her classmates were actually born at the end of 2000. I attribute the year 2001 as the birth year exposure for this entire class of third graders, recognizing some individuals are already infants by the start of 2001. However, if the effect I’m looking for—cognitive damage from environmental toxicants impacts school performance, then students in polluted districts might be more likely to be held back and be older than the typical student in a grade level. In this case, there is incorrect attribution of exposure for the most affected students. Additionally, should the effect of toxicant exposure truly be substantial and negative, then a finding of decreased performance on standardized tests—which are, in many places, high-stakes for the teacher as well as the students—is net of any compensating effort by parents and teachers to bring under-performing students up to grade level.

1.5 Results

1.5.1 Aggregate toxicity

I examined whether greater local toxicity during infancy contributed to lower standardized test performance in the EDFacts proficiency scores for more-exposed cohorts. Main results for the average effect are presented in Table 1.5 for three outcomes: the reported, school-level student proficiency rate; the test-normalized percentage of students scoring proficient; and a statewide test-specific percentile relative to other schools. The independent variable is in Z -scores for interpretability. Across these three outcomes, a change in airborne toxicity risk that increases airborne toxicity by one standard deviation within a school catchment area when a cohort is in its infancy is expected to reduce that cohort’s collective standardized test proficiency in third grade by 0.3 percentage points (e.g., one fewer student passes in a cohort of 300 students), 0.02 standard deviations, or 0.4 percentiles, all relative to their peers in other schools taking the same test (by state, year,

grade level, and subject area).

Estimates are about twice as large among third graders than for the panel that includes grades 4 and 5. This may be because households move, and even if moves are uncorrelated with pollution nearly a decade ago, the presence of children taking standardized tests in a location different from where they were born reduces the proportion of the cohort with correctly-attributed pollution exposure. As kids get older, their likelihood of a household move since birth increases.⁵ Third graders also typically attend smaller elementary schools than middle schools, with correspondingly smaller catchment areas; this means the attribution of air pollution to their third grade school is more geographically precise. Finally, it's also possible that some of the observed decrease is because the effect of early exposure to air pollution on early standardized test performance doesn't persist to later grades as students learn new skills or get additional help from parents and teachers.

The average effect of an increase in airborne toxicity-weighted concentrations is to lower performance across both subject areas; there is mixed evidence suggesting this negative effect is significantly larger for math tests (the interaction terms in Table 1.6 indicate estimated coefficients for reading tests are always smaller in magnitude, but not always significantly different from those estimated for math tests). EDFacts proficiency percentages are also reported for subgroups across students' race and gender, providing an opportunity to examine the heterogeneity of the effects of high toxicant exposure across groups of students. Using these subgroup proficiency percentages as outcome variables, and modifying the set of fixed effects so student proficiency in each group is compared only to other students of their race or gender, I find no statistically significant differences in the magnitude of the effect (Table 1.7).

⁵To give a sense of the possible scale of this measurement error, in 2012 the Census estimated the rates of 5–9 year olds who moved in the past five years was 45 percent, but the great majority of these were in the same county (there is no more precise estimate for shorter-distance moves), and moves outside the county still tended to be under 50 miles. See <https://www.census.gov/prod/2012pubs/p20-567.pdf>.

1.5.2 Metal toxicity

I next extended the estimating equation from Section 1.4 to the case of multiple, distinct airborne metals in two ways: estimating separate regressions for concentrations of each airborne metal, and then jointly in one regression; because airborne releases of metals are not strongly correlated across the SABS catchment areas (recall Figure 1.1), coefficients are expected to be similar across both approaches. As with the aggregate RSEI toxicity scores, I standardized modeled concentrations of each metal to Z -scores (relative to the distribution of each) for interpretability. Coefficients from separate regressions of proficiency outcomes on each airborne metal are plotted in Figure 1.7, and coefficients from the combined regression are presented in Table 1.8.

Among the metals reported in the TRI, early-childhood exposure airborne chromium strongly drives diminished standardized test proficiency in third grade. Coefficients on the estimated birth-year concentrations are significant both individually and in the regression with all airborne metals, across different transformations of estimated concentrations (Table 1.A3), and after correcting for multiple hypothesis testing. A standard deviation increase in chromium exposure has an estimated effect of reducing a cohort's ranking by about 0.08 percentage points, 0.09 percentiles, or 0.006 standard deviations relative to peers. This finding corroborates the EPA's toxicity weights used above: among the 18 metals included in this study, chromium is ranked first for inhalation toxicity.⁶

The lack of significant findings for other metals does not mean they are not toxic, but instead that there is not sufficient variation in the data to detect effects in this study, which is quite saturated with controls. For cognitive risk imposed by less common or correlated toxic compounds, including nonmetallic chemicals, which were not analyzed individually, the aggregate toxicity measure remains a useful representation of risks across exposures from industrial releases.

⁶All releases containing chromium are reported in a single category in the TRI but only one form of the chemical, hexavalent chromium (chromium-6), is classified as carcinogenic by the EPA. Chromium-6 is used across a variety of industries, including ore refining, chemical processing, textile dyeing, wood preserving, stainless steel production, and electroplating; it is most frequently linked to lung cancer from occupational exposure. Using facility-level records from the TRI, about 830,000 pounds of chromium were released per year into the air in the United States in the 2000s, and almost every state had some level of activity.

1.6 Conclusion

This study brings recent, national data to the question of long-term harm to educational proficiency from cohorts' early childhood exposure to airborne toxicants. I estimate the effect on cohort-level standardized test performance from modeled exposures nine years prior, and my findings suggest early environmental harm may have lasting impacts on high-stakes educational outcomes much later in life.

Using an estimate of inhalation toxicity-weighted concentrations of TRI chemicals provides a clear link from facility releases to human health risk, and is motivated by findings that particulate matter can have substantial spatial and temporal heterogeneity. While the EPA primarily measures and regulates particulate pollution by particle size, the toxicity of its composition also matters for human health.

Across elementary schools with a third grade, a standard deviation increase in the EPA toxicity score in the year in which most students were born is predicted to lower that cohort's overall passing rate by about three tenths of a percentage point, and that class is predicted to rank about 0.4 percentile points lower statewide. Since similarly-aged cohorts of students in a school catchment area experience much of the same toxicity exposure, but at different developmental stages, this study highlights the importance of infancy as a period of particular vulnerability to environmental toxicants.

The extensive coverage of both the RSEI and the EDFacts data also provide a context to study the harm from individual airborne metals without making assumptions about their relative toxicity. Airborne metals have been associated with numerous adverse health outcomes, and are prevalent in the TRI data. Among the 18 metals studied here, chromium exposure in infancy was linked to diminished cohort proficiency in third grade: a one standard deviation increase in exposure in infancy is predicted to decrease a cohort's passing rate by 0.08 percentage points.

In light of potential measurement error, the findings in this paper suggest a lower bound on the impact local environmental toxicity can have on children's cognitive development. First, there is inevitable attenuation bias from using elementary schools, rather than individuals, as the unit of observation: any students who move into a local elemen-

tary school area after birth “miss out” on the toxic exposure history which is attributed to their cohort. Second, there is likely some mismatch around the age of students and their toxicant exposure, as both are approximated to a year but will not always overlap.

Finally, it is important to re-contextualize the magnitude of these findings within the aims of the identification strategy. In the U.S., there is much more variation in exposure to toxicants across places than within them over time. This means a within-school fixed effect strategy has good causal purchase but weaker variability off of which to identify effects. The effects here—worse proficiency rates among students born in years when local pollution is higher—should be interpreted as a small suggestion of the potential for persistent cognitive harm in communities where the average level of toxicant exposure is always substantially higher than in other places. My finding that within-school increases in toxicant exposure diminish performance on standardized tests relative to other students attending the exact same school tells us nothing about how much worse students who live in places that are always polluted do relative to students in places that are always clean, all else equal. Future work should further examine the magnitude of this fundamental inequality.

1.7 Figures and tables

Table 1.1: Exposure to airborne metals

Toxicity rank	Symbol	Chemical	Places with nonzero modeled TRI concentrations		
			Schools	Districts	States
1	Cr	Chromium	35,210	9,340	50
2	Co	Cobalt	26,400	6,510	47
3	As	Arsenic	20,750	4,800	48
4	Be	Beryllium	7,370	1,540	35
5	Cd	Cadmium	13,320	3,060	42
6	Ni	Nickel	35,090	9,150	51
7	Pb	Lead	38,110	10,520	51
8	Sb	Antimony	23,830	5,730	48
9	Mo	Molybdenum	13,630	2,910	40
10	Tl	Thallium	3,780	1,080	29
11	Mn	Manganese	35,650	9,580	51
12	Ba	Barium	31,520	8,000	50
13	Cu	Copper	35,300	9,500	50
14	Al	Aluminum	21,220	4,950	42
15	Ag	Silver	12,910	2,650	42
16	Se	Selenium	11,370	2,620	46
17	V	Vanadium	26,470	6,520	48
18	Zn	Zinc	35,370	9,440	51

Metals are among the most consistently reported TRI chemicals. Counts are of schools (with a third grade), districts and states reporting positive concentrations of each toxicant in at least one year. Relative inhalation toxicity is from the EPA.




Figure 1.1: Pairwise correlations among school-level concentrations of airborne metals

Unlike many of the reported TRI compounds, estimated concentrations of particulate metals are generally uncorrelated with one another.




Figure 1.2: School Attendance Boundary Survey coverage of schools with a third grade.




Figure 1.3: Comparison of distributions of scores at schools included in and omitted from the School Attendance Boundary Survey

Table 1.2: Summary statistics for grades 3–5 in the panel by SABS coverage

	Full EDFacts	Not in SABS	In SABS
	(1)	(2)	(3)
Students per cohort	72 (48)	59 (47)	77 (47)
Reported percent proficient	65 (24)	62 (27)	66 (23)
Standardized proficiency	-0.1 (1.2)	-0.3 (1.5)	-0.1 (1.1)
Proficiency percentile	50 (0.3)	50 (0.3)	50 (0.3)
Pct. Asian	0.04 (0.10)	0.04 (0.09)	0.05 (0.10)
Pct. Black	0.16 (0.26)	0.20 (0.30)	0.15 (0.24)
Pct. Hispanic	0.22 (0.27)	0.19 (0.26)	0.23 (0.28)
Pct. Native American	0.02 (0.10)	0.03 (0.15)	0.02 (0.07)
Pct. White	0.54 (0.35)	0.53 (0.36)	0.54 (0.34)
Pct. ECD	0.57 (0.29)	0.58 (0.30)	0.57 (0.28)
County med. HH inc.	53,790 (14360)	53,560 (14020)	53,900 (14480)
County unemp.	8.85 (2.80)	9.06 (3.00)	8.77 (2.72)
County population	727800 (1336700)	677800 (1243200)	745700 (1368100)
Birth year toxicity			16070 (445030)
Z-score log(Birth year toxicity)			0.00 (1.00)
Grades	3–5	3–5	3–5
Distinct schools	61640	19480	42160
Mean years in panel	5.73	5.38	5.85

*Rounding dictated by EDFacts data disclosure rules

The EDFacts “ECD” category corresponds to economically disadvantaged students. Proficiency outcomes and toxicity exposure are defined in Section 1.3. Schools not participating in the SABS appear to be slightly smaller and score a bit lower on standardized tests.




Figure 1.4: School attendance boundary-level toxicity-weighted concentrations for the Los Angeles Area

Darker red areas represent higher toxicity-weighted airborne concentrations among all schools with a third grade, 2001–2006.




Figure 1.5: School-level toxic exposure and standardized test scores

Each point is a school; toxicity scores are unitless and truncated at the 99th percentile for presentation. As median household income rises, fewer and fewer schools are exposed to the highest levels of airborne toxicity.




Figure 1.6: Distribution of relative airborne toxicity risk

Airborne inhalation toxicity risk scores are transformed following the three stages of distributions shown. Results are robust to the choice of log- or IHS-transformed toxicity data.

Table 1.3: Test score data in EDFacts

		Cohort										
		1	2	3	4	5	6	7	8	9	10	11
Grade by year	2010	8	7	6	5	4	3					
	2011		8	7	6	5	4	3				
	2012			8	7	6	5	4	3			
	2013				8	7	6	5	4	3		
	2014					8	7	6	5	4	3	
	2015						8	7	6	5	4	3

The EDFacts data provide standardized test proficiency for all grades taking standardized tests in each school from 2010–2015. Within each school, the data contain fewer cohorts than proficiency observations, since most cohorts appear multiple times at different ages.

Table 1.4: Correlation between lagged toxicant exposure and cohort demographics

	<i>Dependent variable:</i>					
	<i>Z-score log(Toxicity)</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
Pct. Male	0.026 (0.027)	−0.005 (0.011)	0.015 (0.010)	0.020 (0.026)	−0.002 (0.011)	0.016 (0.010)
Pct. ECD	−0.605*** (0.053)	0.036 (0.025)	0.005 (0.021)	−0.506*** (0.047)	0.021 (0.024)	0.004 (0.021)
Pct. White	−1.020*** (0.053)	−0.139*** (0.033)	0.019 (0.019)	−1.031*** (0.248)	−0.100 (0.062)	0.021 (0.051)
Pct. Asian				0.234 (0.238)	0.003 (0.062)	−0.063 (0.057)
Pct. Black				0.060 (0.252)	0.031 (0.067)	0.021 (0.055)
Pct. Hispanic				−0.192 (0.275)	0.075 (0.060)	−0.003 (0.052)
Pct. Native American				−1.286*** (0.326)	0.011 (0.098)	0.091 (0.074)
log(Class size)	0.133*** (0.016)	0.010 (0.008)	−0.002 (0.007)	0.143*** (0.016)	0.010 (0.007)	−0.002 (0.007)
Grades	3–5	3–5	3–5	3–5	3–5	3–5
State-year-subject-grade FE	Yes	Yes	Yes	Yes	Yes	Yes
District-subject-grade FE	No	Yes	No	No	Yes	No
School-subject-grade FE	No	No	Yes	No	No	Yes
Cluster	District	District	District	District	District	District
Observations	1,277,480	1,277,480	1,277,480	1,259,480	1,259,480	1,259,480

Note:

*p<0.1; **p<0.05; ***p<0.01

School-level panel fixed effects minimize correlation between lagged pollution exposure and observable demographics. Columns (4)–(6) include additional covariates, which together often become collinear with reported white students where diversity is low. Within states and even within school districts, minority populations are correlated with higher toxicant exposure, but panel (school FE) variation in pollution lags are uncorrelated with cohort demographics. Results are the same for the 3rd grade-only subset of the data.

Table 1.5: Main findings

	<i>Dependent variable:</i>					
	Pct. proficient		SD pct. prof.		Percentile	
	(1)	(2)	(3)	(4)	(5)	(6)
Z-score log(Toxicity)	−0.307** (0.133)	−0.187* (0.100)	−0.018** (0.008)	−0.011* (0.005)	−0.442* (0.229)	−0.205 (0.164)
Grades	3	3-5	3	3-5	3	3-5
Observations	428,690	1,266,690	428,690	1,266,690	366,060	1,070,470

Note:

*p<0.1; **p<0.05; ***p<0.01

Across all three standardized test proficiency outcomes, higher birth-year toxicity drives lower cohort performance, particularly among younger students; possible explanations for this are discussed in the main text.

Here—and in all regression tables going forward unless otherwise noted—included fixed effects are for state-subject-grade-year, and school-subject-grade groups. Standard errors are clustered at the school district level, and regression weights are the size of each cohort of students.

Table 1.6: Heterogeneity by subject area

	<i>Dependent variable:</i>					
	Pct. proficient		SD pct. prof.		Percentile	
	(1)	(2)	(3)	(4)	(5)	(6)
Z-score log(Toxicity)	−0.407** (0.169)	−0.229* (0.126)	−0.025*** (0.009)	−0.014** (0.007)	−0.650** (0.279)	−0.267 (0.198)
... × reading	0.200 (0.135)	0.083 (0.091)	0.015* (0.008)	0.006 (0.006)	0.407* (0.243)	0.121 (0.172)
Grades	3	3-5	3	3-5	3	3-5
Observations	428,690	1,266,690	428,690	1,266,690	366,060	1,070,470

Note:

*p<0.1; **p<0.05; ***p<0.01

I find suggestive evidence math scores (the baseline category in each regression) are more affected by toxicity exposure in regressions that allow the effect to vary by subject area.

Table 1.7: Heterogeneity by subgroup

	<i>Dependent variable:</i>			
	Subgroup pct. proficient			
	(1)	(2)	(3)	(4)
<i>Z</i> -score ihs(Toxicity)	−0.283** (0.141)	−0.194* (0.107)	−0.328* (0.180)	−0.188 (0.142)
... × Male	−0.042 (0.088)	0.011 (0.062)		
... × White			0.030 (0.181)	0.023 (0.148)
Grades	3	3–5	3	3–5
Observations	861,040	2,543,970	795,040	2,342,040
<i>Note:</i>		*p<0.1; **p<0.05; ***p<0.01		

The baseline category in columns (1) and (2) is proficiency outcomes for female students; in columns (3) and (4) the baseline category is non-white students. In these regressions, the outcome variable is each subgroup's proficiency percentage, weights used are the number of students in each school-year-subject-grade subgroup, and separate fixed effects for each subgroup are included (so proficiency for male students, for example, is relative to other cohorts of male students in the same stage, subject, grade, and year, and male students in the same school, subject, and grade). There are not meaningful differences in the effect size by students' race or gender.




Figure 1.7: Outcomes regressed on individual airborne metals

Coefficients are estimated from separate regressions of each standardized test outcome on standardized concentrations of each airborne metal. Metals shown are those from Table 1.1 with more than 15,000 schools providing estimating variation; those omitted are not consistently statistically significant. Error bars are $1.96 \times \text{SE} \times \text{SD}$.

Table 1.8: Outcomes regressed on all airborne metals

Toxicity rank		<i>Dependent variable:</i>		
		Pct. proficient	Percentile	SD pct. prof.
		(1)	(2)	(3)
1	Chromium	−0.083*** (0.015)	−0.092*** (0.024)	−0.006*** (0.001)
2	Cobalt	−0.062 (0.055)	−0.049 (0.075)	−0.002 (0.003)
3	Arsenic	−0.045 (0.028)	−0.155 (0.078)	−0.005 (0.003)
4	Beryllium	−0.054** (0.017)	−0.420 (0.322)	−0.003*** (0.001)
5	Cadmium	−0.051 (0.042)	−0.063 (0.032)	−0.003 (0.002)
6	Nickel	0.021 (0.015)	0.003 (0.021)	0.001 (0.001)
7	Lead	0.010 (0.023)	0.013 (0.035)	0.001 (0.002)
8	Antimony	0.049*** (0.010)	0.060** (0.021)	−0.001 (0.002)
9	Molybdenum	0.042 (0.053)	0.072 (0.111)	0.004 (0.005)
10	Thallium	−0.0004 (0.015)	0.034 (0.019)	0.001 (0.001)
11	Manganese	0.032 (0.021)	0.061 (0.047)	0.005** (0.002)
12	Barium	−0.066 (0.139)	−0.037 (0.127)	−0.007 (0.006)
13	Copper	−0.025 (0.019)	−0.093 (0.081)	−0.003** (0.001)
14	Aluminum	0.011 (0.051)	0.185 (0.140)	0.003 (0.004)
15	Silver	−0.042*** (0.012)	0.423 (0.530)	−0.003*** (0.0004)
16	Selenium	−0.119 (0.076)	−0.193*** (0.033)	−0.006 (0.004)
17	Vanadium	0.042 (0.052)	0.024 (0.077)	0.001 (0.003)
18	Zinc	0.017 (0.025)	0.012 (0.033)	−0.0004 (0.001)
Grades		3	3	3
Observations		423,560	362,010	423,560

Note change in significance cutoffs:

*p<0.01; **p<0.006; ***p<0.001

Using modeled airborne concentrations of each metal reported in the TRI, significant and negative effects generally corroborate the EPA's toxicity rankings and the main findings presented above. The full set of controls and weights used throughout were included. Only Chromium exposure is robust to the choice of outcome variable at the $p < 0.006$ significance level (stricter cutoffs are used in the presence of multiple hypotheses of interest; 0.006 corresponds to the Bonferroni correction $\alpha/N = 0.1/18$). Significant effects that are not robust to the choice of dependent variable are among those least frequently reported in the TRI (Table 1.1).

Appendix

1.A Additional robustness tables

Additional robustness checks are presented in this appendix. Table 1.A1 shows the inclusion of contemporaneous (to the testing year) ambient toxicity does not change the main results. Table 1.A2 shows the results are not dependent on the use of z -scores as the dependent variables, and Table 1.A3 presents similar robustness for concentrations of airborne chromium. Table 1.A4 replicates the main findings (Table 1.5) using the z -scores of the IHS-transformed RSEI toxicity instead of the log transformation; results are attenuated slightly but not meaningfully different. Table 1.A5 replicates the main specifications without weighting by cohort size; results are again similar.

Table 1.A1: Main findings with contemporaneous toxicity

	<i>Dependent variable:</i>					
	Pct. proficient		SD pct. prof.		Percentile	
	(1)	(2)	(3)	(4)	(5)	(6)
Z -score log(Toxicity)	−0.251** (0.122)	−0.251** (0.122)	−0.020** (0.008)	−0.020** (0.008)	−0.501** (0.238)	−0.501** (0.238)
Z -score log(Contemp. toxicity)		−0.004 (0.164)		0.006 (0.009)		−0.034 (0.267)
Grades	3	3	3	3	3	3
Observations	354,540	354,540	354,540	354,540	304,540	304,540

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 1.A2: Main findings with alternate functional forms

<i>Dependent variable:</i>								
	Pct. proficient			SD pct. prof.			Percentile	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
								(9)
log(Toxicity)	-0.108** (0.047)			-0.006** (0.003)			-0.156* (0.081)	
Toxicity/10,000		-0.001*** (0.0002)	-0.002** (0.001)		-0.0001** (0.00003)	-0.0001 (0.0001)	-0.002*** (0.0004)	-0.003** (0.002)
(Toxicity/10,000) ²			0.00000 (0.00000)			0.000 (0.000)		0.00000 (0.00000)
Grades	3	3	3	3	3	3	3	3
Observations	428,690	429,690	429,690	428,690	429,690	429,690	366,060	366,790
<i>Note:</i>								
*p<0.1; **p<0.05; ***p<0.01								

Table 1.A3: Alternate transformations of metal exposure

<i>Dependent variable:</i>								
	Pct. proficient			Percentile			SD pct. prof.	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8) (9)
Chromium [conc. in $\mu\text{g}/\text{m}^3$]	-7.236*** (1.304)			-8.004*** (2.096)			-0.554*** (0.080)	
Chromium [IHS(conc.)]		-10.052*** (2.297)			-11.698** (3.660)			-0.711*** (0.161)
Chromium [z-score IHS(conc.)]			-0.089*** (0.020)			-0.103** (0.032)		-0.006*** (0.001)
Grades	3	3	3	3	3	3	3	3
Observations	423,560	423,560	423,560	362,010	362,010	362,010	423,560	423,560
<i>Note change in significance cutoffs:</i>								
*p<0.01; **p<0.006; ***p<0.001								

Coefficients on chromium exposure in a cohort's birth year. All specifications include all airborne metals, as in Table 1.8, and the full set of controls and weights. The IHS transformation is preferred to the log since there are many more zeros in the data than with the aggregate toxicity measure.

Table 1.A4: Main findings with Z -scores for IHS-transformed toxicity

	<i>Dependent variable:</i>					
	Pct. proficient		SD pct. prof.		Percentile	
	(1)	(2)	(3)	(4)	(5)	(6)
Z -score ihs(Toxicity)	−0.293** (0.136)	−0.173* (0.103)	−0.016** (0.008)	−0.010* (0.006)	−0.397* (0.229)	−0.187 (0.165)
Grades	3	3-5	3	3-5	3	3-5
Observations	429,690	1,270,370	429,690	1,270,370	366,790	1,073,110
<i>Note:</i>				*p<0.1; **p<0.05; ***p<0.01		

Table 1.A5: Main findings without weighting

	<i>Dependent variable:</i>					
	Pct. proficient		SD pct. prof.		Percentile	
	(1)	(2)	(3)	(4)	(5)	(6)
Z -score log(Toxicity)	−0.292** (0.124)	−0.145 (0.097)	−0.020*** (0.008)	−0.010* (0.006)	−0.449** (0.211)	−0.139 (0.165)
Grades	3	3-5	3	3-5	3	3-5
Observations	428,690	1,266,690	428,690	1,266,690	366,060	1,070,470
<i>Note:</i>				*p<0.1; **p<0.05; ***p<0.01		

Bibliography

- Afflerbach, Peter**, “National Reading Conference Policy Brief : High stakes testing and reading assessment,” *Journal of Literacy Research*, 2005, 37 (2), 151–162.
- Agarwal, Nikhil, Chanont Banternghansa, and Linda T.M. Bui**, “Toxic exposure in America: Estimating fetal and infant health outcomes from 14 years of TRI reporting,” *Journal of Health Economics*, July 2010, 29 (4), 557–574.
- Aizer, Anna, Janet Currie, Peter Simon, and Patrick Vivier**, “Do Low Levels of Blood Lead Reduce Children’s Future Test Scores?,” *American Economic Journal: Applied Economics*, January 2018, 10 (1), 307–341.
- Almond, Douglas**, “Is the 1918 influenza pandemic over? Long-term effects of in utero influenza exposure in the post-1940 U.S. population,” *Journal of Political Economy*, August 2006, 114 (4), 672–712.
- Ard, Kerry**, “Trends in exposure to industrial air toxins for different racial and socioeconomic groups: A spatial and temporal examination of environmental inequality in the U.S. from 1995 to 2004,” *Social Science Research*, 2015, 53, 375–390.
- Ash, Michael and James K Boyce**, “Racial disparities in pollution exposure and employment at US industrial facilities,” *Proceedings of the National Academy of Sciences of the United States of America*, October 2018, 115 (42), 10636–10641.
- **and T. Robert Fetter**, “Who Lives on the Wrong Side of the Environmental Tracks? Evidence from the EPA’s Risk-Screening Environmental Indicators Model,” *Social Science Quarterly*, June 2004, 85 (2), 441–462.
- Bayer, Patrick, Fernando Ferreira, and Robert McMillan**, “A Unified Framework for Measuring Preferences for Schools and Neighborhoods,” *Journal of Political Economy*, 2007, 115 (4), 588–638.
- Bharadwaj, Prashant, Matthew Gibson, Joshua Graff Zivin, and Christopher Neilson**, “Gray Matters: Fetal Pollution Exposure and Human Capital Formation,” *Journal of the Association of Environmental and Resource Economists*, June 2017, 4 (2), 505–542.
- Black, Sandra E**, “Do Better Schools Matter? Parental Valuation of Elementary Education,” *The Quarterly Journal of Economics*, 1999, 114 (2), 577–599.
- Black, Sandra E., Paul J. Devereux, and Kjell G. Salvanes**, “From the cradle to the labor market? The effect of birth weight on adult outcomes,” *Quarterly Journal of Economics*, February 2007, 122 (1), 409–439.

- Case, Anne, Darren Lubotsky, and Christina Paxson**, “Economic status and health in childhood: The origins of the gradient,” *American Economic Review*, November 2002, *92* (5), 1308–1334.
- Chay, Kenneth, Carlos Dobkin, and Michael Greenstone**, “The Clean Air Act of 1970 and Adult Mortality,” *Journal of Risk and Uncertainty*, October 2003, *27* (3), 279–300.
- Chay, Kenneth Y. and Michael Greenstone**, “The impact of air pollution on infant mortality: Evidence from geographic variation in pollution shocks induced by a recession,” *Quarterly Journal of Economics*, 2003, *118* (3), 1121–1167.
- Cunha, Flavio and James J Heckman**, “Formulating , Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation,” *Journal of Human Resources*, 2016, *43* (4), 738–782.
- Currie, Janet**, “Healthy, Wealthy, and Wise: Socioeconomic Status, Poor Health in Childhood, and Human Capital Development,” *Journal of Economic Literature*, 2009, *47* (1), 87–122.
- **and Johannes F Schmieder**, “Fetal exposures to toxic releases and infant health,” *American Economic Review*, April 2009, *99* (2), 177–183.
- **and Matthew Neidell**, “Air pollution and infant health: What can we learn from California’s recent experience?,” *Quarterly Journal of Economics*, August 2005, *120* (3), 1003–1030.
- **, Eric A. Hanushek, E. Megan Kahn, Matthew Neidell, and Steven G. Rivkin**, “Does pollution increase school absences?,” *Review of Economics and Statistics*, November 2009, *91* (4), 682–694.
- **, Lucas Davis, Michael Greenstone, and Walker Reed**, “Environmental health risks and housing values: Evidence from 1,600 toxic plant openings and closings,” *American Economic Review*, February 2015, *105* (2), 678–709.
- **, Michael Greenstone, and Enrico Moretti**, “Superfund cleanups and infant health,” *American Economic Review*, May 2011, *101* (3), 435–441.
- Ebenstein, Avraham, Victor Lavy, and Sefi Roth**, “The long-run economic consequences of high- stakes examinations: Evidence from transitory variation in pollution,” *American Economic Journal: Applied Economics*, October 2016, *8* (4), 36–65.
- Evens, Anne, Daniel Hryhorczuk, Bruce P Lanphear, Kristin M Rankin, Dan A Lewis, Linda Forst, and Deborah Rosenberg**, “The impact of low-level lead toxicity on school performance among children in the Chicago Public Schools: a population-based retrospective cohort study,” *Environmental Health*, December 2015, *14* (1), 21.
- Fahle, E. M., B. R. Shear, D. Kalogrides, S. F. Reardon, R. DiSalvo, and A. D. Ho**, “Stanford Education Data Archive: Technical Documentation,” <https://purl.stanford.edu/db586ns4974>, 2017.

- Ferrie, Joseph P., Karen Rolf, and Werner Troesken**, “Cognitive disparities, lead plumbing, and water chemistry: Prior exposure to water-borne lead and intelligence test scores among World War Two U.S. Army enlistees,” *Economics and Human Biology*, June 2012, *10* (1), 98–111.
- Haladyna, Thomas M.**, “Testing Concerns Perils of Standardized Achievement Testing,” *Educational Horizons*, 2006, *85* (1), 30–43.
- Heckman, James J.**, “The economics, technology, and neuroscience of human capability formation,” *Proceedings of the National Academy of Sciences of the United States of America*, August 2007, *104* (33), 13250–13255.
- Isen, Adam, Maya Rossin-Slater, and W. Reed Walker**, “Every breath you take—every dollar you’ll make: The long-term consequences of the clean air act of 1970,” *Journal of Political Economy*, June 2017, *125* (3), 848–902.
- Kelly, Frank J. and Julia C. Fussell**, “Size, source and chemical composition as determinants of toxicity attributable to ambient particulate matter,” *Atmospheric Environment*, December 2012, *60*, 504–526.
- Knittel, Christopher R., Douglas L. Miller, and Nicholas J. Sanders**, “Caution, drivers! children present: Traffic, pollution, and infant health,” *Review of Economics and Statistics*, May 2016, *98* (2), 350–366.
- Lippmann, Morton**, “Targeting the components most responsible for airborne particulate matter health risks,” *Journal of Exposure Science & Environmental Epidemiology*, 2010, *20* (2), 117–118.
- Marchi, Scott De and James T. Hamilton**, “Assessing the accuracy of self-reported data: An evaluation of the toxics release inventory,” *Journal of Risk and Uncertainty*, 2006, *32* (1), 57–76.
- Marcotte, Dave E.**, “Something in the air? Air quality and children’s educational outcomes,” *Economics of Education Review*, 2017, *56*, 141–151.
- Neidell, Matthew J.**, “Air pollution, health, and socio-economic status: The effect of outdoor air quality on childhood asthma,” *Journal of Health Economics*, 2004, *23* (6), 1209–1236.
- Persico, Claudia, David Figlio, and Jeffrey Roth**, “Inequality Before Birth: The Developmental Consequences of Environmental Toxicants,” *National Bureau of Economic Research Working Paper Series*, May 2016, No. 22263.
- Persico, Claudia L and Joanna Venator**, “The Effects of Local Industrial Pollution on Students and Schools,” *Journal of Human Resources*, 2019.
- Pope, C. Arden and Douglas W. Dockery**, “Health Effects of Fine Particulate Air Pollution: Lines that Connect,” *Journal of the Air & Waste Management Association*, 2006, *56* (6), 709–742.

- Rau, Tomás, Sergio Urzúa, and Loreto Reyes**, “Early Exposure to Hazardous Waste and Academic Achievement: Evidence from a Case of Environmental Negligence,” *Journal of the Association of Environmental and Resource Economists*, 2015, 2 (4), 527–563.
- Ravindra, Khaiwal, Marianne Stranger, and Rene Van Grieken**, “Chemical characterization and multivariate analysis of atmospheric PM_{2.5} particles,” *Journal of Atmospheric Chemistry*, July 2008, 59 (3), 199.
- Rosales-Rueda, Maria and Margaret Triyana**, “The Persistent Effects of Early-Life Exposure to Air Pollution: Evidence from the Indonesian Forest Fires,” *Journal of Human Resources*, April 2018, pp. 0117–8497R1.
- Royer, Heather**, “Separated at girth: US twin estimates of the effects of birth weight,” *American Economic Journal: Applied Economics*, January 2009, 1 (1), 49–85.
- Sampson, Robert J and Alix S Winter**, “The racial ecology of lead poisoning,” *Du Bois Review: Social Science Research on Race*, 2016, 132, 1–23.
- Sanders, Nicholas J.**, “What doesn’t kill you makes you weaker: Prenatal pollution exposure and educational outcomes,” *Journal of Human Resources*, 2012, 47 (3), 826–850.
- **and Charles Stoecker**, “Where have all the young men gone? Using sex ratios to measure fetal death rates,” *Journal of Health Economics*, May 2015, 41, 30–45.
- Spencer Banzhaf, H. and Randall P. Walsh**, “Do people vote with their feet? An empirical test of tiebout’s mechanism,” *American Economic Review*, May 2008, 98 (3), 843–863.
- Tchounwou, Paul B, Clement G Yedjou, Anita K Patlolla, and Dwayne J Sutton**, “Heavy Metal Toxicity and the Environment,” in Andreas Luch, ed., *Molecular, Clinical and Environmental Toxicology: Volume 3: Environmental Toxicology*, Basel: Springer Basel, 2012, pp. 133–164.
- Tessum, Christopher W, Joshua S Apte, Andrew L Goodkind, Nicholas Z Muller, Kimberley A Mullins, David A Paoella, Stephen Polasky, Nathaniel P Springer, Sumil K Thakrar, Julian D Marshall, and Jason D Hill**, “Inequity in consumption of goods and services adds to racial-ethnic disparities in air pollution exposure,” *Proceedings of the National Academy of Sciences of the United States of America*, March 2019, 116 (13), 6001–6006.
- U.S. Department of Education**, “State Assessments in Reading/Language Arts and Mathematics - School Year [2010–2015] EDFacts Data Documentation,” <http://www.ed.gov/edfacts>, 2015.
- U.S. EPA**, “Risk-Screening Environmental Indicators Methodology Document,” Technical Report January, Environmental Protection Agency, Washington, DC 2018.
- Wiliam, D**, “Standardized Testing and School Accountability,” *Educational Psychologist*, 2010, 45 (2), 107–122.

Zhang, Xin, Xi Chen, and Xiaobo Zhang, “The impact of exposure to air pollution on cognitive performance,” *Proceedings of the National Academy of Sciences of the United States of America*, 2018, *115* (37), 9193–9197.

Chapter 2

The Cognitive Cost of Toxic Chemicals in Early Childhood

Abstract

I estimate the effect of exposure to toxic chemicals in infancy on both child cognition in the critical years before kindergarten and parent responsiveness to this shock to child cognition. Linking data from the Early Childhood Longitudinal Study and emissions from the Toxics Release Inventory, I exploit within-zip code variation in children’s months of birth. I find higher exposure substantially diminishes cognition—moving from the 60th percentile to the median toxicity hazard improves early reading and math skills by 0.08 standard deviations—but I find no evidence parents alter behavior in response. Children from more affluent households are less affected by airborne chemical exposure, suggesting heterogeneity in early skill investments, which are associated with household resources, may exacerbate environmental inequality at the neighborhood level.

2.1 Introduction

A growing empirical literature has established the vulnerability of later-life outcomes to exposure to air pollution in childhood (Almond and Currie, 2011; Currie et al., 2014).¹ Much of this work draws from and affirms a theoretical framework in which skills are complementary across multiple dimensions and compound over the life course, so early

¹As a “shock” to early human capital, air pollution fits into a much larger context of risk factors for children, such as malnutrition, maternal stress, disease, recession, and natural disasters; these shocks are reviewed in detail by Almond et al. (2018).

skill development has cascading benefits for later success in learning (Heckman, 2008). Distilling early childhood into a series of endowments and investments, this theory predicts parents choose their allocations of time and money in response to their children’s human capital in early childhood. Empirically, observed harm to cognition from neighborhood pollution is net of household decisions, a product of both biological and behavioral mechanisms (Almond et al., 2018). When environmental exposures are disproportionately concentrated among low-income and minority neighborhoods (Ash and Boyce, 2018; Banzhaf et al., 2019), this channel from environmental shocks to cognitive outcomes may contribute to observed school readiness gaps and perpetuate socioeconomic inequality.

This paper finds a substantial and harmful effect of exposure to airborne toxic releases during infancy, a critical period of neurodevelopment, on child cognition several years later. Individual observations from a detailed survey of early childhood also allow me to study the relationship between household choices and environmental shocks to cognition in the formative years before kindergarten: I estimate both the responsiveness and protectiveness of parenting behaviors, and find suggestive evidence of socioeconomic status (SES)-associated resilience to local environmental harm. This highlights the importance of households for mitigating children’s experiences of adverse neighborhoods.

I pair the Birth Cohort of the Early Childhood Longitudinal Study (ECLS-B) with the Risk-Screening for Environmental Indicators (RSEI) data on toxic air pollution. The ECLS-B surveys a nationally representative cohort of children born in the United States in 2001, and includes multiple, standardized, direct assessments of child cognition and parent attention, as well as detailed household responses to surveys. These repeated records of inputs to, and measurements of, human capital help identify the mechanisms through which environmental harm has persistent effects on development in the period from birth to kindergarten entry.

The Toxics Release Inventory (TRI) data are records of releases from industrial sources in the United States, providing a year-, chemical-, and facility-specific measure of both permitted and accidental air emissions for hundreds of common, harmful compounds. The RSEI model estimates and aggregates the dispersion of these chemicals into a measure

of relative, neighborhood-level variation in ambient inhalation toxicity hazard over time. This better accounting for the relative toxicity of the constituents of airborne pollution improves on the use of particulate matter (PM; often measured as PM2.5 or PM10), which has a heterogeneous composition across both time and space. Inhalation presents a clear pathway from firm releases to infants' health. Relative to the toxicity-weighted index, using annual satellite PM2.5 measurements in an otherwise-identical specification results in estimates that are attenuated towards zero by about half, which is consistent with PM2.5 being an imprecise proxy for cognitive risks.

The process of households sorting into neighborhoods differentiated by amenities, from labor markets to school quality to clean air, confounds a cross-sectional evaluation of the effect of local environments on skills in early childhood; instead, I exploit the timing of children's births within zip codes. Identification rests on a child's month of birth being unrelated to annual variation in toxic releases, and zip code fixed effects controlling for constant, unobserved, neighborhood-level determinants of children's cognition. I create a month of birth-weighted average of RSEI concentrations in 2001 and 2002, assigning younger children a higher proportion of releases from the later year, and demonstrate that this variation in environmental exposure is uncorrelated with child and household characteristics, household moves, and sample attrition.

Decreased exposure to toxic chemicals in infancy predicts higher cognitive skills throughout early childhood, using as an outcome variable children's scores on standardized ECLS-B assessments. Moving from the 60th percentile of toxicity to the median corresponds to an increase in cognitive skills of about 0.08 standard deviations, and this harm from exposure in infancy persists until kindergarten entry. The relatively small change in toxicant exposure hazard associated with the 60th to 50th percentile change (0.5 units) is in line with neighborhood-level panel variation in exposure (0.3 units), and much smaller than a standard deviation (about 3 units). Its predicted effect is comparable in magnitude to the 0.08 to 0.12 standard deviation increase in early reading and math skills Lee et al. (2014) estimate from participation in Head Start for the same ECLS-B cohort, indicating reductions in exposure to environmental toxicants could have widespread benefits for

children in the United States.

The broader potential benefits to public health implied by these estimates depend on the extent to which early cognitive scores are predictive of future labor market outcomes. In standard deviations of cognitive performance relative to peers, I find no evidence the effects of exposure to toxicants in infancy diminish across the three waves when the same cognitive assessments were given (roughly, preschool age through kindergarten entry). I do not observe how the ECLS-B cohort is doing today, now old enough to be graduating high school, but other research concludes the foundations for inequalities in educational attainment are established well before children enter a kindergarten classroom (Currie and Stabile, 2003; Chetty et al., 2011). The early math and reading skills with which children enter kindergarten tend to be the strongest predictors of later achievement in school across a wide range of characteristics (Duncan et al., 2007; Claessens et al., 2009). Gaps in cognitive ability among five year olds often follow familiar patterns of disadvantage along race and class lines, and are strong predictors of later performance in school, which means targeting investments in children’s health, learning, and environments at very young ages can be a particularly important channel for reducing intergenerational inequality (Heckman, 2007).

A persistent and causal link from toxicity-weighted air pollution exposure in infancy to children’s prekindergarten cognition is an important finding in its own right, and expands on literature that has largely focused on prenatal exposures to particulate matter. It is also an average effect, net of parental investments and home environments that may exacerbate or remediate a biological pathway of harm. If the impact of environmental pollution in early childhood depends on a household’s available resources, then avoidance behavior, healthcare, childcare, or extra time spent reading are all ways in which higher-SES parents can attenuate adverse environments (Case et al., 2002; Ferrie et al., 2012); conversely, adverse shocks early on may be invariant to household SES if the biological effect dominates the behavioral (Figlio et al., 2014). The quasi-exogenous variation in environmental harm in infancy allows me to estimate whether parental investments respond to children’s skills in the formative years before kindergarten.

The role of responsive investments matters for the interpretation of reduced-form estimates, and in turn for policies regulating airborne toxicants near residential neighborhoods with the goal of improving child outcomes or reducing environmental inequality. If household investments are endogenous with respect to environmental shocks to cognition, then it is possible to estimate the extent to which decreases in cognitive skills are explained by reinforcing parental investment, or conversely, to what extent harm from toxicant exposure is understated due to compensating investments. I test for responsiveness using parents' choices of reading-related activities and videotaped assessments of parent engagement, asking whether households whose children were exposed to relatively more toxic emissions subsequently report higher or lower levels of investment. I rule out compensatory or reinforcing investments as a large part of the net effect of toxicity exposure on cognition, and this null finding may reflect the context: in low-information settings, where parents do not have clear signals about shocks to children's human capital, responsive behavior may not be plausible.

Despite this, I find toxicity exposure in infancy is more harmful among children from lower-income households than among children from higher-income households. I rule out within-zip code, SES-associated differences in exposure by testing for, and failing to find, a similar gradient in income for motor skills. This pattern is consistent with observed underlying heterogeneity in investments by income, so one potential explanation of this finding is that early parental investments modulate early harm from exposure to environmental toxicants. Put differently, if household SES is protective against harm from toxicity, then even if pollution exposure were distributed evenly across the population (it is not: Boyce et al. (2016) find environmental inequality is starker than income inequality in the United States), it would still cause more harm among poorer households.

These findings contribute to a broader understanding of the ways parents allocate time and effort toward children in adverse circumstances, highlighting two distinct channels through which exposure to toxic air pollution may contribute to school readiness gaps: unequal exposures across neighborhoods, and unequal resources across households. Across the United States, the most-polluted neighborhoods are disproportionately home

to children facing other dimensions of disadvantage. Across households, household resources may then mitigate external threats in a way that reinforces inequality, if families with more resources are able to provide costly investments protective to cognition. This second channel illustrates the need for more research on how household characteristics and choices modulate harmful neighborhoods in the critical developmental years before kindergarten, and reinforces the importance of well-resourced and widely accessible early childhood programs.

Studying the consequences of toxic air pollution also has clear policy implications. Understanding the potential harm from air toxicity motivates both better monitoring of the composition of particulate pollution across time and space, and careful consideration of facility siting that exposes neighborhoods to harmful emissions. As the current EPA has considered weakening rules that count co-benefits among the gains from regulating emissions and minimizing the role of communities to appeal pollution permits in their neighborhoods,² this study presents evidence that industrial pollutants are a persistent threat to children’s early cognitive development.

2.2 Background and theory

Empirical studies of the effects of air pollution on health and cognition implicitly or explicitly speak to several dimensions of treatment: the relative toxicity of the pollutant; the magnitude, or dose, of exposure; and the timing of the exposure (in the context of early childhood, this is the prenatal or neonatal period). Estimated effects range from contemporaneous or short-term outcomes (mortality or birthweight) to long-term behavioral, educational, and labor market outcomes. Much recent research has examined criteria pollutants that are measured by widespread air quality monitors (e.g., carbon monoxide, ozone, and particulate matter). The effects of these relatively well-measured emissions on early childhood outcomes have been studied across a wide range of doses, on outcomes throughout the life course. Conclusions drawn from both county- and individual-level

²See <https://www.epa.gov/environmental-economics/increasing-consistency-and-transparency-considering-costs-and-benefits> and <https://www.nytimes.com/2019/07/12/climate/epa-community-pollution-appeal.html>.

data generally indicate adverse long-term effects on measures like test scores and earnings (Sanders, 2012; Black et al., 2013; Bharadwaj et al., 2017; Isen et al., 2017; Voorheis, 2017; Rosales-Rueda and Triyana, 2018).

Particulate matter (e.g., PM_{2.5} or PM₁₀), while a commonly used measure of air pollution, reflects the density of all the constituent parts of particulate pollution in air, and this makeup can vary over both time and space in ways that could matter substantially for human health. Valavanidis et al. (2008) find strong associations between the chemical composition and toxicity of fine PM, and Kelly and Fussell (2012) argue that “identifying and quantifying the influences of specific components or source-related mixtures [of PM] on measures of health-related impacts” is an important step toward reducing the disease burden of air pollution. This study uses a measure of airborne toxicity in order to better attribute the threat to human health from industrial pollution.

Several existing studies link exposure to toxic air pollution with negative outcomes for fetal and infant health using the TRI. It is the most comprehensive record of toxic chemicals released from point sources into the environment in the United States, with mandated reporting from firms above a certain size in major industrial sectors such as manufacturing, mining, chemical production and treatment, and power generation. Currie and Schmieder (2009) find shorter gestation length and birth weight from prenatal exposure to county-level toxic releases, using the total weight of releases normalized by county area. Agarwal et al. (2010) use a similar explanatory variable and find negative effects on county-level infant, but not fetal, mortality rates. These studies benefit from the large sample sizes provided by natality data, but inevitably average out the substantial spatial heterogeneity in exposure to toxicants within cities and counties. Currie et al. (2015) exploit firm openings and closings to estimate the effects of toxic releases in the immediate neighborhood of TRI-reporting facilities and find an increase in the incidence of low birthweight within a mile of firms.

Persico and Venator (2019) study the effect of TRI facility openings and closings in the proximity of Florida public schools, identifying effects of cumulative pollution exposure on student standardized test outcomes and school rankings. They find this discrete change

in exposure reduces test scores by 0.024 standard deviations for schools within a mile of TRI facilities, relative to schools one to two miles from a TRI site; their smaller effect size from exposure among students in grade 3 or above is consistent with reduced sensitivity of cognitive development at older ages.

Two other recent papers provide useful points of comparison for long-term cognitive outcomes from household proximity to sites contaminated with hazardous waste. Rau et al. (2015) find short-term cognitive harm of 0.09 and 0.07 standard deviations for math and language skills, respectively, for children living a kilometer nearer to a site containing deposits of mining waste (particularly lead) in northern Chile, and larger long-term effects from increases in blood lead levels associated with exposure to the site. Studying mothers living within two miles of Florida Superfund sites (areas flagged by the EPA for deposits of toxic waste), Persico et al. (2016) find differences in test scores of 0.06 standard deviations for children conceived before and after the sites were cleaned up, as well as increases in grade repetition and school suspensions for exposed siblings.

This study introduces parent behavior as a potential modulator of environmental shocks, providing generalizable insights about the role of household heterogeneity beyond the study context of toxic pollution exposure. These findings also reinforce our understanding of the ways harmful environments in the earliest stage of life may have persistent effects years later, and are based on a nationally representative sample of children with a toxicity measure that accounts for the heterogeneity in potential health impacts of chemicals released by the TRI.

2.2.1 Toxic pollution and human capital formation

The hypothesis that infancy is an important window of sensitivity—that is, a period during which toxic chemicals are disproportionately harmful for neurodevelopment—to air pollution is grounded in findings across epidemiology, pharmacology, and toxicology.³

³After birth, the body's organs and metabolic processes to eliminate toxicants mature rapidly; this maturation is nearly complete by age one (Renwick, 1998; Scheuplein et al., 2002). Ünüvar and Büyükgebiz (2012) write, “Children and especially newborns are more sensitive to environmental toxins compared to adults. ... The ability of the newborn to metabolize, detoxify and eliminate many toxins is different from that of the adults. Although exposures occur during fetal or neonatal period, their effects may sometimes be observed in later years.” Describing the potential for unique vulnerability during

Heft-Neal et al. (2018) also find strong effects from PM2.5 exposure on infant mortality (that is, during the 12 months after birth) in sub-Saharan Africa. Following these other literatures, I examine airborne toxic exposures during this period of postnatal neurological development.

That air toxics matter for contemporaneous infant mortality and morbidity does not mean exposure causes long-term skill gaps: parents' behavioral responses in subsequent years may mitigate the initial exposure. When parental investments counteract the role of cognitive harm from environmental insults, net effects undercount the cost from exposure to air toxics. This echoes the intuition that failing to measure costly avoidance behavior as a response to pollution results in undermeasurement of pollution's damages. But unlike avoidance behavior, parental behavior can also go the other direction and exacerbate net effects. Theory linking early shocks to cognition and parental investments is formalized by Almond et al. (2018) and modified here for intuition. Other recent papers such as those by Heckman (2007), Del Boca et al. (2014), Bharadwaj et al. (2018), and Attanasio et al. (2019) also explore these themes.

Consider a multiperiod model of early childhood, where circumstances of birth endow a child with unidimensional human capital H_0 . Early childhood development, the accumulation of skills at each age t , is approximated as the production of human capital, where parent investments X and skills in previous periods are inputs to human capital in each period:

$$H_t = f(H_{t-1}, X_t).$$

Parents get utility U from their children's human capital, but the investments X is costly and comes out of the household's total budget Y , which is otherwise spent on consumption C . If parents know the human capital production function f , and observe H_{t-1} , they can

which organs are much more sensitive than later in life, Bruckner (2000) write that infants' anatomy and physiology are quite different from adults, including their nervous systems. As those systems develop, "relatively small disruptions in these processes may set individuals on trajectories that have subtle effects in early years and profound effects later in life" (Block et al., 2012). At very young ages, children's metabolic rates are higher, meaning their relative oxygen consumption is relatively higher than adults (Bearer, 1995; Bateson and Schwartz, 2008); it follows that when the air is polluted, infants breathe in a proportionally larger share of the pollution.

choose an optimal investment X^* as part of the optimization

$$\max_{X_t} U(H_t, C_t) \text{ s.t. } Y \geq X_t p_x + C_t.$$

Environmental harm enters as a shock $S < 0$ to children’s cognition, decreasing X at time t :

$$H_t = f(H_{t-1}, X_t + S_t).$$

Since f is increasing in both inputs H and X , S unambiguously decreases contemporaneous human capital. We also know the reduced form effect from a growing body of literature: a measurable long-term decrease in later-life human capital can be linked to many environmental shocks in childhood ($dH_{t+N}/dS_t < 0$, where N is many periods in the future). In theory, this decrease is relative to the same child’s counterfactual human capital in the absence of the shock; empirically, researchers estimate the “baseline” using siblings or peers who are not exposed or are less exposed.

The reduced form effect dH_{t+N}/dS_t embeds two possibilities. With compensatory investment, parents allocate additional time and effort into children’s skills in order to make up for the shock in the previous period ($dX_{t+1}^*/dS_t < 0$); with reinforcing investment, parents reduce their inputs in light of the negative shock ($dX_{t+1}^*/dS_t > 0$). Understanding whether the investment response magnifies or minimizes environmental shocks is directly useful for understanding (and mitigating) the source of long-term, reduced-form harm.

The functional form f has been intentionally left unspecified, but it—and the functional form of parents’ preferences U —is necessary for a theoretical prediction about the direction of parental investment.⁴ In particular, the extremes of perfect complementarity and perfect substitutability of endowments of skills and parental investments produce opposite predictions. Three functional forms are presented for intuition in Appendix 2.A.

Any production function for human capital generating responsive parental investment implies something about the ways a shock persists across periods when human capital is an input to its own production. In the case of perfect substitutes, the model implies

⁴Cunha and Heckman (2016) and Cunha et al. (2010), among others, provide some insight toward estimating this production function.

compensating investments: extra time and effort in the current period can “make up” for diminished human capital in the previous period. When endowments and investments are perfect complements in the production function for skills, however, we expect reinforcing investments. Intuitively, if the returns to investment are limited by diminished skills in the previous period, then parents correspondingly divest, reallocating instead toward consumption. In the absence of any other inputs or shocks to human capital (such as preschool), these two extremes predict very different trajectories for skills over the course of early childhood. Since there exist functional forms consistent with either outcome, the direction of responsiveness remains an empirical question.

The theoretical assumption that parents observe (and respond) to child human capital provides intuition about parent responsiveness, but it may not hold up to empirical scrutiny. Many studies of parental investment examine responses to birthweight or nutrition shocks by exploiting sibling and twin fixed effects models to isolate differences in parents’ allocation of time or resources among siblings (Yi et al., 2015; Bharadwaj et al., 2018). It may be unrealistic, however, to frame those findings as descriptive of the ways parents with a single child recognize and respond to shocks in other contexts. Having a twin for comparison is not useful just for econometricians: if relative, rather than absolute, ability informs investment allocations, then this behavioral signal is not available to parents of singleton children. Birthweight is also an easily and universally observed dimension of human capital; it seems less plausible parents precisely know whether their child’s cognitive skills are age appropriate (Dizon-Ross, 2019). In fact, for the ECLS-B cohort, parents do not even appear to be strongly responsive to birthweight: Royer (2009) finds little evidence of compensatory or reinforcing investments in early medical care across twins, and Lynch and Brooks (2013) find no evidence of responsive longer-term parenting investments, after controlling for cross-sectional associations of poor child health and lower investment behaviors. This lack of parent responsiveness is echoed in the empirical findings here, which do not identify responsive parent behavior in the presence of environmental harm.

Finally, household environments likely matter. In the abstract, this model imposes

a budget constraint on investment, so resources, not parents’ abilities or preferences, are the relevant constraint over allocations of time and effort. This informs the choice to study empirical heterogeneity households of different SES. Even if parents do not actually observe cognitive shocks, income heterogeneity is plausible: household wealth, education, health insurance, and habituated behaviors all affect resources available to children. Because I expect parents’ investment decisions—and access to resources more generally—will vary with income, it will be useful to think about “protective environments,” or characteristics of households that mitigate the effects of adverse environmental shocks, in addition to directly responsive investments.

2.3 Data

This paper pairs a survey of early childhood development with local estimates of toxicity hazard from airborne chemical emissions. Individual microdata on child outcomes and parental choices provide detailed outcome and control variables with which to examine the role of shocks in environmental exposure, which I link to children using residential zip code and month of birth.

2.3.1 Birth Cohort of the Early Child Longitudinal Survey

The Early Childhood Longitudinal Study – Birth Cohort (ECLS-B), designed by the National Center for Education Statistics, follows a cohort of children born in the United States in 2001 from birth through kindergarten entry. The 10,700 children were sampled from birth certificates, and the first wave of the survey was conducted when most children were nine or ten months old.⁵ The second wave collected information from children and families around age two, and waves three through five were conducted between ages 4 and 7. Attrition reduced the sample size to about 8,800 students, four-fifths of whom entered kindergarten in 2006 and the remainder the next year, in 2007. Each wave of data collection provides an evolving picture of child development over these early

⁵To protect child privacy, this and all sample sizes going forward are rounded to the nearest 50, as per NCES privacy rules.

years, alongside rich data on individual measures and family controls, such as parent age, education, attitudes, and SES. Unfortunately, the nature of this sample precludes sample designs used elsewhere; in particular, the study does not collect information about the focal cohort's siblings, so sibling fixed effects and other multiyear comparisons are not an option.

I restrict the sample for estimation to neuro-typical children (identified by the “special needs” variable in the ECLS-B) who are living with their birth mother; this is 96 percent of the entire sample. The main outcome variable of interest is the performance of children on direct cognitive assessments given in each wave of the study. Assessments are standardized in the ECLS-B to have a unit variance for each wave, so changes in this outcome variable will be in standard deviations of performance relative to peers, or age-appropriate cognitive development. To get at medium-term effects of contemporaneous exposure from airborne toxicity in the first year of life, I do not use the wave 1 cognitive score, collected when children were nine months old, as an outcome variable.

Assessments in wave 2 tested early language and problem solving using the Mental Scale of the Research Edition of the Bayley Short Form, and in waves 3 through 5 tested early mathematics, and early language and literacy, using a range of assessments, including some from the ECLS-K and other proprietary (but commonly used) assessments. In order to have a continuous outcome variable from waves 2 through 5, I average the math and reading scores to create a single cognitive score for each child in these later waves, and also consider each subject area separately. Empirical analysis is designed to estimate cognitive differences across children at the same age with respect to their history of exposure to toxic pollution, so cognitive scores are compared within waves rather than across them. This alleviates some concerns that the assessments given in each wave of the ECLS-B measure different sets of skills at each age.

A secondary outcome is parental time investments in child skills, which are measured through two channels: parent-reported investment activities in reading, and the videotaped “Two Bags Task.” Literacy investments are revealed in survey responses to questions about how many times per week the parent reads to the child and how many

children’s books parents have in the house.⁶ Because both of these are a recall question, I expect responses to be a noisy measurement of the underlying variability in parent effort. In order to minimize hypothesis testing and improve precision, I combine the measures into an index by adding the z -scores of each, relative to responses in the wave in which the responses were collected.

The Two Bags Task measured the quality of parent-child interactions during a “semi-structured play activity” at two years and preschool age (waves 2 and 3 of the ECLS-B). Videotapes of the sessions, which asked parents to play with toys and read a book to their children, were graded using a standardized rating scale across several dimensions meant to characterize the quality of the interactions. The parent rating scales assessed sensitivity to children’s behaviors, stimulation of cognitive development, and expressions of positive regard, as well as some negative responses, which I do not include. Steps were taken to ensure the task was administered in a standardized way across all ECLS-B home visits, and ratings were uniformly applied across the coders who watched the video tapes (the collection and measurement of these data are discussed further by Andreassen C and Fletcher P (2007) and Najarian et al. (2010)). As with the parent-reported reading activities, I sum z -scores of the individual measures into an index reflecting the positive aspects of parental engagement during the activity.

The two measures of parent investment are correlated, but not strongly (the correlation coefficient is about 0.32 for waves 2 and 3). While the index of reading activities focuses specifically on how parents foster early literacy, the Two Bags index is a more general assessment of parent attentiveness. The relationship between reading sessions and the Two Bags index is visualized in Figure 2.1.

I focus on income as the main dimension of household heterogeneity with respect to child sensitivity to environmental toxics, recognizing income can also be considered a proxy for household SES more broadly. Reports of household income were collected in

⁶ “Number of kids’ books” was an open-ended response, and has a positive skew. Results are generally invariant to top-coding the largest responses. Responses for reading sessions were collected in four bins: “0 times per week,” “1–2 times per week,” “3–6 times per week,” or “Every day”; I treat the middle of each interval as a number. Unfortunately, two other potential measures of literacy investment were not collected consistently: responses for “telling stories” were collected only in waves 1 and 2, while responses for “minutes spent reading” were collected only in waves 3–5.

14 bins in the survey; I average the endpoints of each bin to create a continuous variable.

Across the children in the ECLS-B, the associations of child skills, parent investment, and parent resources are apparent: wealthier parents tend to invest more time and effort into children and have higher-quality interactions, and children from wealthy families tend to score higher on cognitive assessments. The effect of extra parental effort, therefore, cannot be easily separated from the lucky circumstance of being born into a household with more material resources at its disposal. Figure 2.2 shows the stratification of both early childhood cognition and parent investments by terciles of income in the ECLS-B (low-, middle-, and high-income households).

Summary statistics are presented in Table 2.1, which also previews the balance necessary for identification across zip code toxicity exposure (discussed in the next section). Identifying variation relies on within-zip code differences in ambient air pollution: for internal validity, it is important children exposed to more and less pollution are comparable on observable and unobservable characteristics; for external validity, it is important children in zip codes not contributing to identifying variation are comparable to “treated” children on observable and unobservable characteristics.

The ECLS-B cohort initially lives in 4,000 distinct zip codes in over 2,300 cities across the United States. Children in the survey were geographically clustered at the first wave to facilitate data collection, but the ECLS-B cohort is nationally representative with sampling weights, which are used throughout empirical analysis. The survey provides household zip codes, which I use to link neighborhood-level environmental conditions. I assume the household zip code reported at the first wave of the survey, conducted when children are just nine months old, is where a child has lived since birth.

Anticipating the variation that will identify the effect of toxicant exposure, 1,950 zip codes have children born in more than one month of the year, and these zip codes contain 76 percent of the ECLS-B children (Table 2.2). Household mobility is substantial (29 percent of children moved between waves 1 and 2, and 45 percent report a different zip code at least once in the survey), which makes it difficult to control for children’s zip codes in later waves of the study, since most ECLS-B families who move end up in

a zip code with no ECLS-B peers. I test that moves after wave 1 are uncorrelated with exposure to air toxicity at wave 1 of the survey, and control only for zip code of birth.

2.3.2 Toxics Release Inventory and Risk-Screening Environmental Indicators model

The Toxics Release Inventory (TRI) is a federally mandated reporting program for chemical releases in the United States. Since 1988, manufacturing firms have been required to report their releases of hundreds of chemicals directly to the EPA, which compiles the data and makes it publicly available. In 1994, over 200 chemicals were added to the TRI, and in 1998, a major expansion brought additional industry sectors under mandated TRI reporting, including metal and coal mining, electric utilities, chemical distributors, and chemical treatment and disposal facilities. Federal facilities have always been required to report, regardless of sector. The reporting thresholds for persistent bioaccumulative toxic chemicals—toxicants which remain in the human body and the environment over long periods of time—were lowered for the 2000 reporting year, which means the TRI years used here contain more accurate information about smaller releases of highly toxic chemicals. These expansions improve the reports as a representation of neighborhood-level environmental conditions—in 2001, nearly 650 distinct chemicals were required to be reported, and total releases were in the billions of pounds—but the TRI is not exhaustive. The data do not account for mobile sources of pollution (such as airports and highways), nor firms falling outside of covered sectors or below reporting thresholds (which are based on number of employees and weight of chemicals used).

Reported toxic releases typically come from regular, permitted operation of industrial facilities. These fluctuate year to year due to changes in the type or quantity of activity at a facility (for example, changes in demand for a product, changes in production inputs or processes, or improvements in pollution control technology); an additional source of annual variation at a facility is accidental or one-time releases, any instances of which must be included in TRI reporting.

This paper primarily uses TRI data from the 2001 and 2002 reporting years to at-

tribute environmental toxicity in infancy, and over these years the TRI maintained consistent rules across sectors and chemical compounds; this means changes in toxicity (and estimated effects) are not being driven by changes in reporting requirements. A single reporting change (the threshold for lead and lead compounds was lowered for 2001) affects comparisons from 2000 to 2001, so in Appendix 2.D where zip code-level concentrations from 2000 are included, I recreate an inhalation toxicity-weighted concentration that omits releases of lead to ensure consistency across TRI years.

To reflect the potential harm to people in the neighborhood of TRI facilities, it is important to consider the size of the area around a firm affected by TRI releases, and the relative toxicity of individual chemicals. Currie et al. (2015) identify a uniform radius around TRI facilities affected by releases and exploit plant openings and closings for variation in local exposure. This paper instead takes advantage of a dispersion model providing traction on the quantities of chemicals released, the fate and transport of those chemicals, and their relative toxicities. Two features of the RSEI model make the health risks implied by reported TRI releases more tractable: dispersion and toxicity. These aspects are discussed briefly here, and at length in Appendix 2.C.

An EPA dispersion model called AERMOD⁷ estimates the fate and transport of each release in the TRI, producing ambient concentrations of the chemical in the area around the point location of a facility. These concentration estimates account for site-specific quantities, release pathways (stack or fugitive), and weather; and compound-specific dispersion characteristics. (Over half of ECLS-B zip codes do not contain any TRI facilities but are affected by facilities in neighboring zip codes via the RSEI dispersion model.)

Toxicity weights are the other component of the RSEI data, and reflect the fact that the threat to human health from TRI releases varies widely among the chemicals reported, from relatively benign (for example, sodium nitrite, used as a food preservative) to known carcinogens and neurotoxicants (like asbestos and chromium). I scale concentrations with the RSEI's inhalation-specific toxicity weights, which reflect the EPA's collected research about relative risk for adverse health consequences from inhalation of each chemical

⁷AERMOD is documented at www.epa.gov/scram/air-quality-dispersion-modeling-preferred-and-recommended-models.

compound.⁸

Using estimated concentrations and toxicity weights, the RSEI model aggregates all TRI releases up to a standardized grid, so each cell reflects the total, relative levels of toxicity-weighted airborne hazard from the contributions of each chemical from each individual facility affecting the location in a given year. I link the RSEI grid, which covers the entire United States at a resolution of 810 square meters, to ECLS-B zip codes, assigning airborne toxicity-weighted concentrations by joining the centroid of each zip code to its nearest grid cell. Zip codes do not represent geographic areas (as do counties, census tracts, or school attendance boundaries) so it is not possible to average RSEI exposure over an entire zip code. Consequently, I expect rural zip codes, which cover residential addresses over a large area (typically the most rural places) to have the highest potential for measurement error. Because environmental releases tend to be low in rural places, I do not expect this source of measurement error to meaningfully affect findings. Because ambient concentrations are modeled, there is considerable smoothing in the RSEI grid, which means neighboring grid cells are spatially correlated.

The TRI present uniquely detailed toxicity data, but the quantities of chemicals are ultimately self-reported, and in many cases estimated rather than measured by firms. However, there are legal requirements report accurately to the EPA, and penalties for firms failing to do so. The accuracy of TRI reporting is discussed by Marchi and Hamilton (2006), who find more evidence of random noise (possibly due to inaccurate guessing) than strategic underreporting. The exogeneity of these releases with respect to unobservable attributes of the ECLS-B cohort is discussed in Section 2.4.1.

Since the RSEI toxicity-weighted concentrations are unitless, they are best interpreted as a measure of relative risk for chemical exposure across a nationally representative sample. Changes in airborne toxicity exposure can be expressed in terms of quantiles of the sample-weighted distribution (e.g., moving from the 75th percentile to the median), which are independent of monotonic transformations of the underlying data. For reference, with

⁸Coefficients from regressions using the all-around toxicity weights are slightly attenuated toward zero, corroborating the assumption that inhalation is the most immediate pathway by which TRI releases affect infants; issues of accumulation in water or soil are left for future study.

the inverse hyperbolic sine-transformed independent variable I use for analysis, moving from the 60th percentile of toxicity risk to the median is about 0.5 units, and from the 90th percentile to the median is about 2.2 units. As would be expected, panel variation is smaller than cross-sectional variation: the median absolute within-zip code change from 2001 to 2002 for the RSEI measure is about 0.3 units. Summary percentiles are presented in Table 2.3, and the distribution is visualized in Figure 2.3.

2.3.3 Particulate matter

In order to evaluate whether the use of toxicity-weighted concentrations from the RSEI model leads to substantively different conclusions from the use of PM measurements, I replicate the main analysis with PM data from two sources.

The first is an intermediate estimate from the RSEI data. I sum the modeled concentrations of airborne particulates from TRI emissions without using the inhalation toxicity weights, to produce a gridded value for the total concentration of PM from all industrial sources reporting to the TRI. Of course, non-TRI sources are still missing from these data, but particulates of all sizes are included.

The second is the NASA Global Annual PM_{2.5} Grids from van Donkelaar et al. (2018), which combine multiple satellite measurements to estimate average near-surface PM_{2.5} with dust and sea salt removed at a resolution of 0.01 degrees and in units of micrograms per cubic meter. This measurement of the mass concentration of ambient particulates includes non-point and non-industrial sources of air pollution, such roads, airports, and other sites that do not report to the TRI, but omits larger particles. PM_{2.5} is preferred to PM₁₀ since finer airborne particulates have been more strongly linked to adverse health outcomes. Since the satellite data is also an annual measure, and very similar in resolution to the RSEI grid (0.01 degrees is 850 meters at 40 degrees north), the comparison to the main results is empirically straightforward. As above, I link the centroid of each ECLS zip code to the gridded PM data in each year.

For each zip code, satellite PM_{2.5} is almost always greater than TRI PM (on average by about 10 micrograms per cubic meter), which is to be expected as the satellite mea-

sure includes many sources of airborne PM not captured by the TRI. Empirically, the correlation coefficient of the toxicity hazard index and the satellite PM_{2.5} is only 0.35 for exposures in infancy among the ECLS–B children (Figure 2.4). That is, there are many locations where toxicity is relatively high but PM is relatively low, and vice versa, and so PM may be a noisy measure of relative health risk in this context.

2.4 Empirical analysis

Empirical analysis will determine whether child cognition in the years leading up to kindergarten is sensitive to airborne toxicant exposure in infancy, and test whether parental behavior and household environments mitigate or exacerbate environmental shocks to early human capital. I discuss the exogeneity of the independent variable after introducing the main specification.

At the first wave of the survey, many zip codes are shared by multiple ECLS-B households. I exploit this geographic clustering and focus on differences in outcomes among children born in different months in the same zip code. Since the ECLS-B cohort was born entirely in 2001, and TRI releases are reported only at annual intervals, I create zip-level weighted averages of the RSEI toxicity-weighted concentrations for the 2001 and 2002 reporting years. This generalizes the intuition that children born in January experience ambient pollution throughout 2001 during their first year of life, while children born in December are exposed to air toxics released in 2002 during their infancy. For any child i born in month B living in zip code z , toxicity risk in the first 12 months of life is a weighted average of the value Tox_z at the matched RSEI grid cell for 2001 and 2002:

$$Tox_{z(i)} = Tox_{z,2001} * (12 - B_i)/12 + Tox_{z,2002} * (B_i)/12. \quad (2.1)$$

Since only one RSEI grid cell is used per zip code, variation in estimated exposure risk among children born in the same zip is exclusively derived from variation in month of birth. This attribution based on timing of birth for an ECLS cohort is comparable to the strategy used by both Marcotte (2017) and Heft-Neal et al. (2018).

Treatment at the zip code level, then, is the sign and magnitude of the change in RSEI airborne toxicity risks from 2001 to 2002 combined with a child's month of birth. Using rolling averages avoids potentially large, discrete breaks in pollution exposure attributed to children born, or assessed, only a few months apart, but it also inherits the assumption built into the RSEI model of uniform air release rates over the course of a year. Finally, because the RSEI toxicity-weighted concentrations have an enormous right skew, I transform each zip-level RSEI concentration Tox_z with the inverse hyperbolic sine to prevent outliers in the far right tail from overinfluencing regression estimates.⁹ The spatial distribution of the average, unitless ambient toxicity risk across 2001 and 2002 is shown in Panel A of Figure 2.5, while the zip code-level changes from 2001 to 2002 in the RSEI measure are shown in Panel B.

Using this variation in local airborne toxicity during infancy among children born in the same zip code in 2001, a fixed effect regression of cognitive assessments on ambient airborne toxicity and observable household characteristics gives the causal effect of interest,

$$Cog_{iw} = \beta_1 Tox_{z(i)} + \gamma X_{iw} + \theta_w A_{iw} + \mu_{iw} + \eta_{r(i), B_i} + \zeta_{z(i)} + \epsilon_{izw}, \quad (2.2)$$

where Cog is child i 's cognitive score measured in wave w at age A ; the fixed effect μ captures the month and year in which the child was assessed. $Tox_{z(i)}$ approximates ambient airborne toxicity during infancy (following equation 2.1), and ζ is a zip code of birth fixed effect. X is a matrix of observed, household-level controls: child sex, race, number of siblings, and twin status; mother's age at birth, marriage status, and dummies for drinking or smoking during pregnancy; and household income.

It is possible children's birth months are correlated with other socioeconomic determinants of child cognition. To account for this possibility I include 48 region by month-of-birth fixed effects ($\eta_{r(i), B_i}$); this assumes any drivers of selection into birth month are constant across larger geographies than neighborhood-level pollution exposure. Birth-

⁹Recall that the RSEI toxicity-weighted concentrations are unitless, and meaningful only as ordinal measures of relative risk. This means a linear-log functional form should not be interpreted as a dose-response relationship (which would imply diminishing harm from unit changes at higher levels of toxicity).

weight has been closely examined as both an outcome of exposure to toxic pollution in utero, and also as a determinant of later-life outcomes in its own right; because birthweight may be influenced by local pollution, I do not include birthweight and gestational length as controls but instead explore the relationship between prenatal and neonatal exposures for this sample in Appendix 2.D—to preview that result, I find that for air toxicity risk, neonatal exposures matter relatively more for later cognition.

Coefficient β_1 is the parameter of interest, and represents the effect on cognition at age 2 and older of exposure to toxic air releases in the first year of life. For β_1 to be identified, with a zip code fixed effect, the air toxicity during the first year of life, must matter. This is because by March of 2002, for example, the entire cohort had been born and experienced the ambient pollution of the zip code, but some children were 3 months old and some were 15 months old.

I cluster standard errors at the child’s city of birth, even in later waves of the survey. The clearest need for this comes from the RSEI model, which estimates exposure smoothly around facilities, so nearby grid cells are spatially correlated. This means neighboring zip codes, particularly those very close to one another, experience correlated pollution exposure.

It is also useful to relax the zip code fixed effect specification, instead controlling for the local level of ambient toxicity for 2001 and 2002. I include this pollution “baseline” with $Tox_{z,2001}$ and $Tox_{z,2002}$, as above:

$$Cog_{iw} = \beta_2 Tox_{z(i)} + \beta_3 Tox_{z,2001} + \beta_4 Tox_{z,2002} + \gamma X_{iw} + \theta_w A_{iw} + \mu_{iw} + \eta_{r(i),B_i} + \alpha_{s(i)} + \epsilon_{izw}. \quad (2.3)$$

This cross-sectional specification treats the observed environmental quality of a zip code as a proxy for determinants of early childhood cognition that lie outside the household, which cannot be controlled for with observed attributes of children and their families, and are correlated with local air toxics. After controlling for average pollution in residential zip codes, the timing of births should be quasi-exogenous. I add a state-of-birth fixed

effect to this regression, $\alpha_{s(i)}$, and keep the same set of household and region by month-of-birth controls.

The average toxicity specification is less plausible for causal inference, but provides greater estimating variation, since zip codes with only one child (or all children born in the same month of 2001; column (2) of Table 2.1) contribute to the estimate of the effect in this specification but not the specification with zip code fixed effects. Point estimates turn out to be quite similar across the zip code fixed effect and average toxicity specifications. This reinforces the external validity of specification (2.2), where we might otherwise be concerned the limited estimating variation distorts the national representation of the sample. The average toxicity specification will also have the advantage of more identifying variation in household characteristics (relative to individual zip codes) when later regressions estimate heterogeneity in effects by SES.

2.4.1 Exogeneity of TRI exposures

Causal identification of the effects of toxic exposure in infancy requires that month of birth for children born in the same zip code in 2001 is unrelated to variation in RSEI concentrations between 2001 and 2002; if this is the case, then local toxicant exposure can be considered a quasi-exogenous shock. Much has been made of household sorting and neighborhood demographic turnover in response to changes in environmental amenities, which in turn correlates unobservable determinants of child cognition with local environmental conditions. “Household sorting” in this context could take the form of relocation decisions (e.g., some households opt to leave a neighborhood when environmental conditions deteriorate), but also the timing of parenthood (e.g., some families delay having children in response to air pollution). This type of sorting requires both salience of environmental shocks and time for households to change behavior. For example, Spencer Banzhaf and Walsh (2008) find compositional changes in the immediate neighborhood of TRI facilities when they open or close, but they identify these relatively small changes over a decade.

Several empirical checks confirm both characteristics of children, and their attrition

from the sample, are not significantly related to their relative (zip-level) “treatment” into toxicant exposure. I show balancing tests of assigned exposure to toxicity-weighted concentrations are uncorrelated with household characteristics within zip codes (Table 2.4), and controlling instead for average toxicity combined with state or city fixed effects does nearly as well on balance as zip code fixed effects. In Table 2.5, I show toxicant exposures are unrelated to the likelihood a child’s household moves during ECLS-B data collection, the number of household zip codes reported by a child’s household, or the number of cognitive assessments collected for each child. Additional discussion of the exogeneity of toxicity exposure, and robustness checks of the main empirical results, are presented in Appendix 2.B, which includes specifications that drop observations that had larger (and therefore, potentially salient) changes in the number of TRI facilities both within or affecting the household zip code from 2001 to 2002; show removing household characteristics does not meaningfully change estimates (consistent with as-random assignment to pollution); estimate placebo effects for randomized child zip codes, months of birth, or years of exposure.

Finally, some descriptive characteristics of this study and the TRI make it less vulnerable to household sorting than longer-term and primarily cross-sectional analyses. The ECLS-B cohort was born within a single calendar year, minimizing the potential for substantial demographic turnover that would make children compared to one-another within a zip code bad counterfactuals. TRI data is available to the public more than a year after each reporting year, making it impossible for households to respond to the information specifically contained in the TRI, and further, the RSEI measure reflects variation in the levels of emissions from continuously operating facilities, so most changes are not being driven by firm openings and closings. And since “many toxic pollutants are colorless, odorless, and not well monitored, making them less salient than other negative externalities” (Currie et al., 2015), and most of the zip-level, year-over-year changes in the levels of toxicants released from 2001 to 2002 are small, households may not accurately perceive levels of toxic pollutants affecting their neighborhood.

The specification presented so far restricts the effect of toxicant releases on cognition

to be perfectly persistent over time: the harm to cognition β_1 is invariant over the waves of the survey in which cognition was assessed. This is a useful first step to establish the existence of a relationship between toxic exposure and cognition, but unnecessarily restrictive. Does the relative performance of children who were exposed to higher toxicity fall relative to peers as the cohort ages, or does the magnitude of harm revert toward zero?

Interacting toxicant exposure in the first year of life with the wave of the survey in which children are assessed, or the age at which they are assessed, helps show the (reduced form) persistence of the relative harm to cognition for children more exposed in early childhood. Modifying equation (2.2) only slightly, I estimate specifications where the effect of toxicity is flexible with respect to the wave of the survey (w is waves 2–5), and alternately where I include an interaction with assessment age.

It is not conceptually obvious whether the harmful effect of exposure in infancy is decreasing or increasing in magnitude with age. If compensating investments better describe how parents adapt to children’s environmental harm in infancy, the coefficient of an interaction between early exposure and age should be positive, while the opposite sign would be consistent with growing inequality in cognitive skills, relative to peers, as children age.

2.4.2 Household responsiveness and heterogeneity

The reduced form effect estimated in equation (2.2) is net of parent investments in child cognition, which are an important component of early development. If parental behavior is compensatory or reinforcing with respect to shocks to children’s human capital, there should be evidence parents shift resources toward child development. To examine this, I omit children’s cognitive scores, and estimate the causal effect of neonatal toxic exposure directly on investment choices, circumventing the simultaneity arising from parents’ choices of investments and children’s endowments of skills. I assume the effect of quasi-exogenous toxic shocks in infancy on parents’ behavior only operates through the channel

of children's cognition:

$$Inv_{iw} = \beta_1 Tox_{z(i)} + \gamma X_{iw} + \theta_w A_{iw} + \mu_{iw} + \eta_{r(i), B_i} + \zeta_{z(i)} + \epsilon_{izw} \quad (2.4a)$$

$$Inv_{iw} = \beta_2 Tox_{z(i)} + \beta_3 Tox_{z(i)} * \log(Inc_{i1}) + \gamma X_{iw} + \theta_w A_{iw} + \mu_{iw} + \eta_{r(i), B_i} + \zeta_{z(i)} + \epsilon_{izw} \quad (2.4b)$$

where $\log(Inc_{i1})$ is household income in the first wave of the survey, and Inv_{iw} is reported reading activity or parent engagement assessed with the Two Bags task. As always, I include controls for household characteristics. The coefficient β_1 reflects the average responsiveness of parents' investments (reading activities or the Two Bags assessment) to neonatal toxicity shocks, and β_2 and β_3 allow that responsiveness to vary with household income. These equations allow me to test whether, for example, parents whose children were more exposed to toxicity than their peers in the zip code also read to their children more than other parents. More precisely, when β_1 is positive, investments are compensatory—higher toxicity exposure diminishes cognitive skills and parents respond with increased effort; the opposite sign implies reinforcing investments. Estimates consistent with responsive investments will depend on both a strong relationship between early toxicity exposure and children's later skills, and also that parents actually adjust allocations of time and effort in response to their perceptions of children's skills.

I limit the empirical analysis in this section to waves 2 and 3, since later waves have substantially greater attrition, the Two Bags task was only administered in these waves, and by wave 4 the older children had already entered kindergarten, which likely has large effects on parents' investment decisions. If household incomes are relatively homogeneous within zip codes (zip codes are segregated by affluence or poverty), then I expect the specification allowing cross-sectional comparisons will provide greater identifying variation.

Investment behaviors have the potential to mitigate adverse circumstances in early childhood, even if they are not responsive. And of course, parents' choices of effort are constrained by the total household budget. If investment choices do not respond to parent's perceptions of children's skills, then the potential endogeneity bias from including

investment activities on the right hand side of equations (2.4a) and (2.4b) is avoided.

Empirically, I use two descriptive regressions to ask whether there is a link between the magnitude of harm from exposure to neonatal environmental toxics and household resources (equation 2.5a) or investments in reading activities or the two bags assessment (equation 2.5b) (this has precedent in Currie and Hyson (1999)). These regressions, by now familiar, involve an interaction of toxicity exposure with household characteristics and behaviors:

$$Cog_{iw} = \beta_1 Tox_{z(i)} + \beta_2 Tox_{z(i)} * \log(Inc_{i1}) + \gamma X_{iw} + \theta_w A_{iw} + \mu_{iw} + \eta_{r(i), B_i} + \zeta_{z(i)} + \epsilon_{izw} \quad (2.5a)$$

$$Cog_{iw} = \beta_3 Tox_{z(i)} + \beta_4 Tox_{z(i)} * Inv_{iw} + \gamma X_{iw} + \theta_w A_{iw} + \mu_{iw} + \eta_{r(i), B_i} + \zeta_{z(i)} + \epsilon_{izw}. \quad (2.5b)$$

If β_2 is positive, household resources are “protective” against cognitive shocks — that is, the same magnitude of exposure to toxicity in early childhood predicts smaller decreases in performance relative to peers for children from more affluent households. When household investment decisions are the only way to mitigate neonatal toxic shocks, then rejecting the hypothesis $\beta_2 = 0$ is consistent with heterogeneous investments by household income, but for this to be a reasonable interpretation, “investment” needs to include much more than the econometrician can observe (such as avoidance behavior). The coefficient β_4 is not necessarily causal, since the baseline choice of investment is likely correlated with other determinants and outcomes of SES.

2.5 Results

Table 2.6 shows the reduced-form, causal effect of exposure to airborne toxicity in infancy on several outcomes: cognitive skills, motor skills, and health. I find a persistent, negative effect of exposure to air pollution in the first year of life on cognition in subsequent waves of the survey. Cognitive scores are 0.08 standard deviations lower for children exposed to birth-year toxicity at the 60th percentile versus the median, relative to peers, for

assessments given in waves 2 through 5 (using the reference percentiles presented in Table 2.3). In the later waves, where assessments differentiate math and reading skills, toxicity exposure seems more harmful for reading skills. It is possible the harmful effect on cognition is echoed in harm to other dimensions of human capital. Child motor skills may be affected by toxicity exposure, but child health (as reported by parents) does not seem to be, suggesting TRI chemicals are a particular threat to learning.

In general, results will be described for the specification using zip code fixed effects (Panel A). Estimates using the secondary, cross-sectional specification (regression equation (2.3), which uses zip-level average toxicity, rather than zip code fixed effects, to control for unobservables that correlate child cognition with ambient pollution risk) are presented in Panel B throughout, for comparison. The point estimates are generally consistent with the primary specification, and provide more precise estimates, especially when geographic variation is useful for identifying heterogeneity, as will be the case with varying effect sizes by household income.

The effect of early chemical exposure on cognition is persistent: children are not simply scoring worse than peers on cognitive assessments shortly after exposure, and then rebounding to the mean. Allowing the coefficient on neonatal toxicity exposure to vary with the age at which children are assessed or, more flexibly, with the wave of the survey in which they are assessed, Table 2.7 shows the harmful effect of early exposure is greatest at the age-2 wave of the survey, but remains negative through the age at which children enter kindergarten (note, though, the change in magnitude between waves 2 and 3 should not be over-interpreted, since the ECLS-B cognitive assessments changed between those waves). Between the ages of approximately 4 and 6, there is no evidence the overall magnitude of the initial shock is magnified or minimized with age. This medium-term persistence distinguishes these findings from studies of air pollution showing students perform worse on “bad air” days.

2.5.1 Comparison to PM

Using the concentration of airborne particulates, rather than a toxicity-weighted index, as the independent variable has the effect of diminishing the estimated magnitude of the effect of exposure in infancy on later cognition. This finding is consistent with the motivation that since the composition of PM varies over time and space, it is an imprecise measure of health risk. Two different proxies for PM informed this finding: the modeled RSEI concentrations, without toxicity weights, of particulates from TRI air emissions; and annual satellite measurements of near-surface PM_{2.5}.

Table 2.8 replicates column 1 of Table 2.6 using zip-level PM exposure. Estimated coefficients are smaller in magnitude (for a comparable change in exposure) and not statistically significant. Table 2.8 also shows the inclusion of either measure of PM in the the preferred specification has essentially no effect on the estimated magnitude of harm to cognition from toxicity.

2.5.2 Behavior and heterogeneity

The reduced-form analysis so far has established environmental exposure in infancy diminishes cognitive performance several years later, and these shocks are large enough that parent behavior plausibly adjusts in response. Table 2.9 shows estimates from regression equations (2.4a) and (2.4b), which directly test this. For the ECLS-B cohort, there is not much evidence for responsive investments in either direction (compensatory or reinforcing, a positive or negative coefficient on toxicity exposure): the average responsiveness to infant toxicity exposure is indistinguishable from zero.

Panel B of each table seems to suggest investment responsiveness may be heterogeneous with household income: across both outcomes and all three indicators for SES, the responsiveness of wealthier households to toxicity shocks is more compensatory. That is, holding constant average differences in levels of investment by income, upper-income households appear to increase reading activity significantly more in response to cognitive shocks in infancy. However, Figure 2.6 plots the responsiveness of households and shows no statistically significant effects at any income level; and in any case, any degree of

responsiveness implied by Table 2.9 is economically small, since the investment outcomes are standardized to have unit variance.

The degree of cognitive harm from early toxicity exposure is heterogeneous by the SES of households: Table 2.10 shows the mitigating effect of household income. Using household characteristics reported in wave 1 of the survey, an interaction between toxicity exposure and SES shows the impact of airborne toxic exposure on cognitive scores is smaller among children from wealthier households (this effect is consistent across three different indicators for household SES: logged household income, an indicator for whether household income is below 185 percent of the poverty line, and dummy variables for high-medium- and low-income households). These results are only significant for the average toxicity specification, which is reasonable if within-zip code variation in household income is low. The coefficient on toxicity exposure is about 30 percent lower for children in the top third of the income distribution relative to children from the bottom third of the income distribution.

This relationship is not necessarily causal, though it's consistent with a theory of human capital development where resources provide advantages protective against cognitive shocks: higher-income households may have better access to nutrition, childcare, or healthcare, and we know income is correlated with investments in cognition from the associations plotted in Figure 2.2. All of these hypothetical channels involve behaviors, intentional or not, that could help children's development catch up to peers' after environmental shocks.

This finding of income heterogeneity does not seem to be driven by wealthier children tending to be less-exposed within their zip codes: if this heterogeneity were being driven by income-correlated measurement error in toxicity exposure, which might arise from within-zip code sorting of wealthier households further from TRI facilities, or differential avoidance behavior by income (such as less time spent outdoors), then we would expect to see income (or the distance for which it may proxy) also moderates the effect of toxicity on motor skills. Table 2.11 and Figure 2.7 show this is not the case.

Table 2.12 considers whether children are less affected by toxicity in the medium-

term if their parents tend to invest more (which, following the results above, is plausibly not endogenous with respect to the toxicity shock). Across both specifications, there is evidence reading activities do just this, though not greatly: a standard deviation more reading investment reduces the estimated effect of toxicant exposure by about one tenth. To be clear, these effects should not be considered causal, since household reading activities are quite correlated with other developmental advantages determined by SES. There is no evidence the attentiveness measured with the parent scores on the Two Bags assessment is similarly protective, and neither investment is protective for children’s motor skills.

Because recall measures of reading investment and assessed quality of interactions are noisy signals for the underlying effort parents allocate toward early literacy and cognition for their children, I can’t rule out the possibility that environmental shocks drive responsive investments, but these estimates are not sufficiently precise. In particular, heterogeneous parent responsiveness is a mechanism consistent with the finding of lower-magnitude harm among children from higher-income households. The absence of strong evidence for responsiveness may simply mean cognitive shocks are not salient to parents, since in this setting, parents do not receive a clear signal about children’s human capital. More broadly, this study presents a case where behavior changes are probably not a large component of the net, reduced-form effects of toxicity on cognition. Instead, differences in resources across families, and the affordability of investments in early childhood, may be driving observed heterogeneity.

2.6 Conclusion

This study estimates a persistent effect from exposure in infancy to airborne toxic emissions on cognitive scores at kindergarten entry. It is the first to estimate the effect of TRI releases for cognitive outcomes using a nationally representative cohort of U.S. children, and highlights the particular sensitivity of cognitive development in infancy. Within-zip code variation in month of birth across children provides causal estimates of the effect

of toxicity in infancy, which is assigned using the RSEI dispersion model of ambient airborne toxicity-weighted concentrations. The average effect of reducing toxicity risk from the 60th percentile to the median has a comparable effect on early reading and math skills (0.08 standard deviations) to the effect of Head Start participation estimated by Lee et al. (2014), and to the effect of living a kilometer further from a land-based site contaminated with hazardous waste estimated by Rau et al. (2015). Since environmental toxicants are disproportionately imposed on less-advantaged neighborhoods, this harmful effect may exacerbate school readiness gaps through the channel of neighborhood sorting.

Evidence for a statistically significant, negative effect from exposure to airborne pollution in the first year of life contrasts the finding in Marcotte (2017), which reports no measurable effect from early exposure to particulate matter (PM_{2.5}) on school readiness in the ECLS–K. This absence of a statistically significant finding for PM is replicated here: using annual satellite-derived measures of PM_{2.5} shows that the RSEI inhalation toxicity-weighted index is more strongly linked to diminished early childhood cognition than the mass concentration of fine particulates. This highlights the empirical importance of accounting for the toxicity of particulate pollution, and should motivate greater monitoring of and attention to the composition of air pollution.

Detailed household-level data allowed me to consider the ways in which parents respond to environmental shocks, and whether children from more- or less-advantaged backgrounds are differentially affected by pollution. Despite the large average effect on cognition, in this context, where neither shocks to cognition nor endowments of skills are particularly salient, parents' behaviors are not responsive to environmental exposures during children's infancies. Yet, parent behavior may inadvertently modulate the effect of exposure to toxicant releases. A standard deviation increase in reading investments—which are strongly associated with economic resources available to the household—is associated with a 10 percent reduction in the harmful effects of exposure. Similarly, the harmful effect of toxicity on cognitive scores for the poorest third of children is about 40 percent larger than for children from the highest-third of the income distribution. That more-affluent children are less-harmed suggests an additional channel by which environ-

mental exposures may perpetuate inequality in educational outcomes. Importantly, this channel of SES-associated adaptation would exist even in the absence of larger-scale spatial inequality in exposure to toxic air pollution. The mechanisms behind these patterns, should be further investigated in future research.

2.7 Figures and tables




Figure 2.1: Two Bags scores by reported reading sessions

Parent scores on the standardized Two Bags index by reading sessions reported in waves 2 and 3. Parents who read more often to children typically have higher Two Bags scores, but this relationship is noisy.




Figure 2.2: Child cognitive scores, reading, and parents' Two Bags scores by wave and household income

Average levels of cognitive scores and investment (reading, Two Bags) are plotted by terciles of income in the ECLS-B, in standard deviations. Cognition was measured using different sets of assessments between wave 2 and waves 3–5. Time spent reading with the child is reported as sessions per week. The Two Bags Task, described in the main text, was only conducted in waves 2 and 3. Horizontal density of points reflects the number of ECLS-B children assessed at each age. Sample weights are used throughout.

Table 2.1: Summary statistics by wave 1 toxicity exposure

		<i>Zip-demeaned toxicity:</i>			
		Negative (1)	Zero (2)	Positive (3)	Difference (4)
<i>Exposure</i>	$Tox_{z(i)} - \overline{Tox_z}$	-0.11	0.00	0.10	0.215
		(0.21)	(0.00)	(0.17)	-0.040
	$Tox_{z(i)}$	8.31	7.90	8.49	0.178
		(2.47)	(2.76)	(2.52)	0.049
<i>Characteristics</i>	Pct. male	0.50	0.50	0.52	0.020
	Pct. Asian	0.03	0.02	0.03	0.001
	Pct. Black	0.16	0.11	0.14	-0.019
	Pct. Hispanic	0.27	0.19	0.28	0.012
	Pct. other nonwhite	0.05	0.04	0.04	-0.002
	Pct. mother unmarried	0.33	0.29	0.34	0.010
	Pct. twin or higher	0.03	0.03	0.03	0.001
	Pct. smoked during pregnancy	0.10	0.12	0.11	0.011
	Pct. alcohol during pregnancy	0.03	0.04	0.03	-0.002
	log(Income)	10.42	10.55	10.41	-0.005
		(0.93)	(0.90)	(0.90)	-0.032
	Mother's age	27.29	27.88	27.08	-0.215
		(6.17)	(6.18)	(6.18)	0.004
	Mother's education	12.84	13.20	12.75	-0.097
		(3.00)	(2.66)	(2.89)	-0.116
	Number of siblings	0.99	0.96	0.99	-0.004
		(1.13)	(1.10)	(1.09)	-0.039
<i>Attrition & moving</i>	Number of cognitive asmts.	3.38	3.32	3.34	-0.047
		(1.25)	(1.20)	(1.15)	-0.099
	Pct. ever moved	0.44	0.44	0.45	0.009
	Number of reported zips	1.58	1.57	1.57	-0.012
		(0.75)	(0.74)	(0.71)	-0.031
Number of children		3,850	2,000	4,250	400

Summary statistics are presented across three groups of children in the ECLS-B, split by the variation in toxicity exposure within zip codes (which is based on timing of birth): higher than zip-average (1), no zip-level change (2), and lower than zip-average (3). Differences between columns (1) and (3) are shown in column (4). After demeaning, there is similarity of children across higher and lower toxicity exposure. The similarity of columns (1) and (3) to column (2) informs the external validity of estimates identified using the part of the ECLS-B sample experiencing changes in neighborhood toxicity exposure; the children summarized in column (2) tend to be slightly whiter and more affluent. Weighted summary statistics are presented throughout.

Table 2.2: Counts of wave 1 ECLS-B zip codes by number of distinct birth months represented

Calendar months in zip	Number of zips	Number of children	$\overline{Tox}_z$
(1)	(2)	(3)	(4)
1	2000	2300	7.87
2	900	2050	8.36
3	450	1650	8.71
4	200	1100	8.47
5	150	850	8.65
6	50	450	7.91
7–8	100	950	8.37
9–10	50	500	8.53
11–12	0*	200	7.01
Full sample	3,900	10,100	8.30

*Nonzero but rounds to zero.

About four-fifths of the ECLS-B sample was born in a zip code shared by peers born in a different month. Birth month at the zip code level provides identifying variation in toxicity exposure in the main specification. Toxicity risk does not appear to be related to the density with which a zip code was sampled.

Table 2.3: Variation in Wave 1 toxicity exposure

Percentile	RSEI level	$Tox_{z(i)}$
(1)	(2)	(3)
0	0	0
10	56.04	4.72
25	627.71	7.14
50	3116.09	8.74
60	5175.56	9.24
75	12582.43	10.13
90	29297.99	10.98
100	53709310.31	18.49
Mean	20166.26	8.30
St. dev.	459033.7	2.56

Distribution of assigned zip-level toxicity for the ECLS-B cohort. Column (2) shows the right skew of the RSEI measure (before transformation using the inverse hyperbolic sine), and column (3) presents Wave 1 exposure for the transformed measure used in the analysis ($Tox_{z(i)}$). Sample weights are used for the reference percentiles, means, and standard deviations.




Figure 2.3: Distributions of toxicity exposure

Panel A shows the distribution of changes from 2001 to 2002 in RSEI toxicity for all the ECLS-B zip codes. Most zip codes experience small changes (that is, small panel variation) in toxicity risk. **Panel B** shows cross-sectional variation in toxicity exposure risk for the ECLS-B sample.




Figure 2.4: Correlation of toxicity-weighted concentrations and satellite PM2.5

The correlation coefficient for near-surface PM2.5 and toxicity-weighted concentrations (0.35) is positive but not large. Annual, gridded near-surface PM2.5 from satellite measurements are assigned using the same weighting formula across exposure in 2001 and 2002 as in the main analysis (equation 2.1); each point is an ECLS-B child.




Figure 2.5: RSEI toxicity-weighted risk, 2001–2002

Panel A shows unitless RSEI aggregate toxicity for the contiguous United States, transformed with the inverse hyperbolic sine and averaged between 2001 and 2002 (the two years identifying variation in neonatal exposure for the ECLS-B cohort); darker colors correspond to higher toxicity risk. **Panel B** maps the change in zip-level toxicity from 2001 to 2002: darker orange corresponds to larger increases, while darker blue corresponds to larger decreases. These year-over-year shifts are clustered locally, but their direction and magnitude are not distributed systematically over the United States. For the visualization, I mapped the zip code-level RSEI measure using a crosswalk to the U.S. Census Bureau's ZCTA geographies. Alaska and Hawaii are included in the ECLS-B sample and in the study. None of the shading is related to locations of the ECLS-B children.

Table 2.4: Toxicant exposure and household and child characteristics

	<i>Dependent variable: $Tox_{z(i)}$</i>				
	(1)	(2)	(3)	(4)	(5)
$T_{z,2001}$				0.491*** (0.037)	0.506*** (0.054)
$T_{z,2002}$				0.505*** (0.038)	0.502*** (0.055)
1(Male)	0.014 (0.059)	-0.016 (0.037)	-0.002 (0.011)	0.006 (0.008)	-0.002 (0.009)
1(Asian)	0.753*** (0.100)	0.164*** (0.056)	-0.025 (0.020)	-0.021 (0.014)	-0.026* (0.014)
1(Black)	0.903*** (0.125)	0.204*** (0.069)	-0.008 (0.020)	-0.025* (0.014)	-0.006 (0.013)
1(Hispanic)	0.879*** (0.103)	0.206*** (0.064)	-0.008 (0.022)	0.022 (0.021)	-0.009 (0.015)
1(Other nonwhite)	0.259* (0.144)	0.072 (0.067)	0.015 (0.034)	0.049* (0.025)	0.021 (0.024)
log(Income)	0.119** (0.054)	0.002 (0.026)	0.008 (0.010)	-0.007 (0.010)	0.004 (0.007)
1(Mother unmarried)	-0.021 (0.078)	0.038 (0.044)	-0.001 (0.015)	-0.003 (0.010)	-0.002 (0.011)
Mother's age	0.020*** (0.006)	0.010*** (0.003)	-0.002 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Mother's education	-0.013 (0.014)	-0.012* (0.006)	-0.002 (0.003)	0.002 (0.003)	-0.001 (0.002)
1(Smoked during pregnancy)	-0.521*** (0.110)	-0.166** (0.069)	0.006 (0.024)	0.009 (0.015)	0.001 (0.016)
1(Alcohol during pregnancy)	0.145 (0.168)	-0.029 (0.078)	0.032 (0.026)	-0.003 (0.019)	0.009 (0.022)
Number of siblings	-0.077** (0.033)	-0.027 (0.017)	0.007 (0.006)	0.003 (0.004)	0.004 (0.004)
1(Twin or higher)	0.134 (0.086)	0.055 (0.057)	0.013 (0.017)	-0.004 (0.010)	0.003 (0.012)
Number of cog. asmts.	-0.080*** (0.028)	-0.039*** (0.014)	-0.010 (0.007)	-0.004 (0.004)	-0.007 (0.005)
Waves	1	1	1	1	1
Fixed Effect	State	City	Zip	State	City
Observations	9,450	9,450	9,450	9,450	9,450

With state-of-birth fixed effects, characteristics of the ECLS-B children and households predict RSEI toxicity exposure (column 1); city fixed effects do not fully remove these correlations (column 2). Corresponding to equation 2.2, zip code fixed effects, which derive identifying variation exclusively from timing of birth, show no statistically significant correlations between household attributes and children's toxicity exposure (column 3). Controlling instead for zip-level ambient toxicity does nearly as well as zip code fixed effects (columns 4 and 5); the pairing of zip-level exposure with state fixed effects corresponds to equation 2.3.

Table 2.5: Sample attrition related to toxicant exposure

	<i>Dependent variable:</i>					
	Ever moved (1)	(2)	Reported zips (3)	(4)	No. cog. asmts. (5)	(6)
$Tox_{z(i)}$	0.031 (0.035)	−0.002 (0.018)	−0.015 (0.054)	−0.027 (0.026)	−0.176 (0.122)	−0.051 (0.051)
Zip FE	Yes	No	Yes	No	Yes	No
Zip-level tox.	No	Yes	No	Yes	No	Yes
State FE	No	Yes	No	Yes	No	Yes
Waves	1	1	1	1	1	1
Observations	9,600	9,600	9,600	9,600	9,600	9,600

Three outcomes related to attrition—whether the household ever moved during the ECLS-B data collection, the number of household zip codes given for each child during the survey, and the number of cognitive assessments collected for each child—are not significantly related to toxicity exposure using either the zip code FE (columns 1,3,5) or average-toxicity (columns 2,4,6) specifications.

Table 2.6: Net effect of neonatal toxicity exposure

	<i>Dependent variable:</i>				
	Cognition	Reading	Math	Motor	Health
A: Zip FE	(1)	(2)	(3)	(4)	(5)
$Tox_{z(i)}$	−0.162*** (0.052)	−0.162*** (0.060)	−0.103* (0.055)	−0.112* (0.063)	−0.038 (0.058)
Zip FE	Yes	Yes	Yes	Yes	Yes
B: Zip-level toxicity	(1)	(2)	(3)	(4)	(5)
$Tox_{z(i)}$	−0.097** (0.044)	−0.107** (0.046)	−0.065 (0.043)	−0.002 (0.030)	−0.033 (0.032)
Zip-level tox.	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes
Waves	2–5	3–5	3–5	2–5	2–5
Observations	23,600	15,400	15,400	23,250	25,250

These are the reduced form effects of exposure to toxicity risk in the first year of life on human capital across several measures, from assessments given in waves 2 through 5 (which occurred at least a year after neonatal toxicity exposure). A strong negative effect on cognition is apparent, and appears to be driven by diminished reading scores in the later waves. Motor skills may also be harmed by toxicity exposure, but child health does not appear to be affected. Estimates are quite similar across **Panel A**, which presents estimates from the zip code fixed effect specification and **Panel B**, showing the cross-sectional specification controlling for zip-level toxicity (explained in the main text).

NB: In this and all regression tables to follow, unless otherwise noted, included controls are date of assessment fixed effects; month of birth-by-region fixed effects; and household characteristics explained in the main text. Sampling weights are used and standard errors are clustered at the city of birth.

Table 2.7: Persistence of early toxicity exposure

	<i>Dependent variable:</i>					
	Cognition		Reading		Math	
A: Zip FE	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z(i)}$	-0.190*** (0.053)	-0.180*** (0.053)	-0.176*** (0.065)	-0.166*** (0.060)	-0.071 (0.060)	-0.101* (0.055)
$Tox_{z(i)} \times \text{age (years)}$	0.007*** (0.002)		0.003 (0.005)		-0.007 (0.005)	
$Tox_{z(i)} \times \text{Wave 3}$		0.028*** (0.008)				
$Tox_{z(i)} \times \text{Wave 4}$		0.028*** (0.009)		0.009 (0.007)		-0.004 (0.007)
$Tox_{z(i)} \times \text{Wave 5}$		0.023* (0.013)		0.006 (0.013)		-0.007 (0.013)
Zip FE	Yes	Yes	Yes	Yes	Yes	Yes
B: Average toxicity	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z(i)}$	-0.128*** (0.045)	-0.117*** (0.044)	-0.126** (0.052)	-0.110** (0.046)	-0.040 (0.048)	-0.061 (0.043)
$Tox_{z(i)} \times \text{age (years)}$	0.008*** (0.002)		0.004 (0.004)		-0.005 (0.004)	
$Tox_{z(i)} \times \text{Wave 3}$		0.030*** (0.007)				
$Tox_{z(i)} \times \text{Wave 4}$		0.027*** (0.008)		0.004 (0.006)		-0.009 (0.006)
$Tox_{z(i)} \times \text{Wave 5}$		0.034*** (0.013)		0.013 (0.012)		-0.003 (0.012)
Zip-level tox.	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Waves	2-5	2-5	3-5	3-5	3-5	3-5
Observations	23,600	23,600	15,400	15,400	15,400	15,400

The negative effect on cognition shown in table 2.6 is largest around age 2 (wave 2), and diminishes with child age, suggesting harm to human capital from neonatal exposure recovers somewhat as children get older. However, this attenuation is definitely not to zero. Allowing the effect of early exposure to vary more flexibly by survey wave, the effect holds constant across waves 3 through 5 (the left-out wave is 2 in column (2) and 3 in columns (4) and (6)). As in table 2.6, estimates are quite similar across **Panel A** and **Panel B**.

Table 2.8: Net effect of exposure to PM

	<i>Dependent variable:</i>					
	Cognition					
	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z(i)}$	−0.162*** (0.052)			−0.161*** (0.052)	−0.159*** (0.052)	−0.159*** (0.052)
$TRI\ PM_{z(i)}$		−0.104 (0.079)		−0.101 (0.077)		−0.103 (0.078)
$Satellite\ PM_{z(i)}$			−0.041 (0.051)		−0.032 (0.050)	−0.033 (0.050)
Zip FE	Yes	Yes	Yes	Yes	Yes	Yes
Waves	2–5	2–5	2–5	2–5	2–5	2–5
Observations	23,600	23,600	23,750	23,600	23,600	23,600

Note:

*p<0.1; **p<0.05; ***p<0.01

Replicating the specification from column (1) of Table 2.6 with two estimates of zip-level PM yields estimates that are smaller in magnitude and not statistically significant. Including PM concentrations in the toxicity specification has almost no effect on the estimate of the effect of toxicity. *TRI PM* is the RSEI modeled particulate concentration from TRI sources, and *satellite PM* is satellite-measured near-surface PM2.5; both are in units of micrograms per cubic meter. To interpret these magnitudes using the same framework as the main results, a change in satellite PM2.5 from the 60th to 50th percentile in these data is about 1 microgram per cubic meter—so the effect is about half as large, relative to the toxicity-weighted independent variable. This finding is consistent with PM as a noisy measure for the human health risks of airborne particulate pollution.

Table 2.9: Parent responsiveness by SES

DV: Reading index	A: Zip FE			B: Average toxicity		
	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z(i)}$	-0.027 (0.056)	-0.132 (0.099)	-0.011 (0.055)	-0.034 (0.058)	-0.030 (0.028)	-0.192*** (0.059)
$Tox_{z(i)} \times \log(\text{Income})$		0.010 (0.008)				0.016*** (0.005)
$Tox_{z(i)} \times \text{Below 185\% Poverty}$			-0.013 (0.015)			-0.026*** (0.009)
$Tox_{z(i)} \times \text{Income Q2}$				-0.0004 (0.015)		0.003 (0.009)
$Tox_{z(i)} \times \text{Income Q3}$				0.026 (0.019)		0.038*** (0.011)
Observations	17,250	17,250	17,250	17,250	17,250	17,250
DV: Two bags index	(1)			(2)		
	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z(i)}$	-0.010 (0.084)	0.021 (0.135)	-0.002 (0.082)	-0.011 (0.087)	-0.0002 (0.039)	-0.059 (0.069)
$Tox_{z(i)} \times \log(\text{Income})$		-0.003 (0.010)				0.006 (0.005)
$Tox_{z(i)} \times \text{Below 185\% Poverty}$			-0.001 (0.014)			-0.011 (0.008)
$Tox_{z(i)} \times \text{Income Q2}$				0.008 (0.017)		0.005 (0.011)
$Tox_{z(i)} \times \text{Income Q3}$				0.006 (0.018)		0.020** (0.010)
Fixed effect	Zip	Zip	Zip	Zip	State	State
Zip-level tox.	No	No	No	No	Yes	Yes
Waves	2:3	2:3	2:3	2:3	2:3	2:3
Observations	14,100	14,100	14,100	14,100	14,100	14,100

Columns (1) and (6) show the average effect of neonatal toxicity shocks on subsequent choices of investment activities (explained in the main text) is both statistically insignificant and economically small. All regressions include the un-interacted effect of household income on the outcome. When responsiveness of parent investments to toxicity shocks is allowed vary with income, testing the hypothesis that wealthier households are more sensitive to toxicity shocks and/or more able to adjust behavior in response, the effects are generally statistically insignificant (see Figure 2.6).




Figure 2.6: Coefficients and error bands on toxicity exposure as a function of household income

The estimated coefficient $\beta_1 + \beta_2 * \log(\text{Income})$ on neonatal toxicant exposure, where the outcome is parent investment, is plotted over levels household income. Each panel corresponds to the regressions presented in columns (2) and (6) of Table 2.9. Though the statistically significant interaction terms in the average toxicity specifications across both outcomes suggest meaningful differences by household income in the magnitude of responsive investment, the coefficient is significant outside the range of household incomes observed in the data. Possible explanations for the imprecision of these estimates are discussed in the main text. Error bands show a 10-percent significance level ($1.65 * \text{SE}$).

Table 2.10: Higher-SES children experience lower-magnitude effects from neonatal exposure to toxicants

	<i>Dependent variable:</i>							
	A: Zip FE				B: Average toxicity			
	Cognitive score							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Tox_{z(i)}$	-0.142*** (0.053)	-0.249*** (0.092)	-0.125** (0.052)	-0.145*** (0.052)	-0.094** (0.046)	-0.231*** (0.066)	-0.075 (0.050)	-0.101*** (0.047)
$Tox_{z(i)} \times \log(\text{Income})$		0.010 (0.007)				0.013*** (0.005)		
$Tox_{z(i)} \times \text{Below 185\% Poverty}$			-0.017 (0.013)				-0.020** (0.008)	
$Tox_{z(i)} \times \text{Income Q2}$				-0.004 (0.015)				0.010 (0.009)
$Tox_{z(i)} \times \text{Income Q3}$				0.023 (0.015)				0.031*** (0.010)
Fixed effect	Zip	Zip	Zip	Zip	State	State	State	State
Zip-level tox.	No	No	No	No	Yes	Yes	Yes	Yes
Waves	2:3	2:3	2:3	2:3	2:3	2:3	2:3	2:3
Observations	15,800	15,800	15,800	15,800	15,800	15,800	15,800	15,800

The harmful effect of increased toxic exposure on cognition measured in later ages appears to be lower among children who grow up in high-SES households: the harm to children's cognition from the same change in toxicity in infancy is about 30 percent smaller. This is finding is consistent across several measures of household resources, each using characteristics reported in wave 1: income, poverty status, and income terciles (the lowest and highest third of ECLS-B households by income are in "Q1" and "Q3," respectively). The average toxicity specification shown in **Panel B** produces lower standard errors but comparable point estimates to the zip code fixed effects specification in **Panel A**.

Table 2.11: Higher SES does not appear to mitigate the effect of toxicity on motor skills

	<i>Dependent variable:</i>							
	Motor score				B: Average toxicity			
	A: Zip FE							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Tox_{z(i)}$	-0.102 (0.072)	-0.029 (0.110)	-0.103 (0.071)	-0.090 (0.072)	-0.019 (0.030)	-0.015 (0.054)	-0.019 (0.031)	-0.011 (0.030)
$Tox_{z(i)} \times \log(\text{Income})$		-0.007 (0.007)				-0.0004 (0.004)		
$Tox_{z(i)} \times \text{Below 185\% Poverty}$			0.009 (0.013)				0.003 (0.008)	
$Tox_{z(i)} \times \text{Income Q2}$				-0.010 (0.014)				-0.014 (0.010)
$Tox_{z(i)} \times \text{Income Q3}$				-0.018 (0.015)				-0.002 (0.009)
Fixed effect	Zip	Zip	Zip	Zip	State	State	State	State
Zip-level tox.	No	No	No	No	Yes	Yes	Yes	Yes
Waves	2:3	2:3	2:3	2:3	2:3	2:3	2:3	2:3
Observations	15,550	15,550	15,550	15,550	15,550	15,550	15,550	15,550

In contrast to the findings in the previous table, there is no evidence household income mitigates the harm from toxicity exposure to motor skills. This suggests a measurement of household income is not acting as a proxy for household sorting within zip codes, or other forms of avoidance behavior reducing children's exposure to toxicity in the first place.




Figure 2.7: Coefficients and error bands on toxicity exposure as a function of household income

The estimated coefficient $\beta_1 + \beta_2 * \log(\text{Income})$ on neonatal toxicant exposure, where the outcome is the standardized cognitive score in waves 2 and 3, is plotted over levels household income. Each panel corresponds to the regressions presented in columns (2) and (6) of Tables 2.10 and 2.11. The significant gradient in the effect of toxicity on cognition is not replicated in the effect of toxicity on motor skills, which may suggest observed heterogeneity by income is not being driven by within-zip code differences in exposure by income. Error bands show a 10-percent significance level ($1.65 * \text{SE}$).

Table 2.12: Higher reading investment is associated with smaller effects from toxic exposure

	<i>Dependent variable:</i>					
	Cognitive score			Motor score		
	(1)	(2)	(3)	(4)	(5)	(6)
A: Zip FE						
$Tox_{z(i)}$	-0.142*** (0.053)	-0.141*** (0.051)	-0.147** (0.061)	-0.102 (0.072)	-0.100 (0.072)	-0.134** (0.065)
$Tox_{z(i)} \times \text{Reading index}$		0.012** (0.005)			0.007 (0.006)	
$Tox_{z(i)} \times \text{Two bags index}$			0.003 (0.006)			-0.003 (0.006)
Zip FE	Yes	Yes	Yes	Yes	Yes	Yes
B: Average toxicity	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z(i)}$	-0.094** (0.046)	-0.086* (0.044)	-0.084 (0.053)	-0.019 (0.030)	-0.017 (0.030)	-0.018 (0.033)
$Tox_{z(i)} \times \text{Reading index}$		0.013*** (0.004)			0.005 (0.004)	
$Tox_{z(i)} \times \text{Two bags index}$			0.003 (0.004)			0.001 (0.004)
Zip-level tox.	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Waves	2:3	2:3	2:3	2:3	2:3	2:3
Observations	15,800	15,550	15,800	15,500	13,500	13,300

Across the two measures of parent investment, only reading activities are correlated with attenuated effects of neonatal exposure several years later, and only for cognitive scores. The units of both indices of investment are standardized with mean zero for interpretability. Among households who read to their children one standard deviation above the mean, the estimated coefficient on neonatal toxicity risk is about 10 percent smaller. These effects are not causal, since household investment decisions may be endogenous with respect to early shocks to cognition, and correlated with unobserved or uncontrolled-for characteristics of households.

Appendix

2.A Production of human capital

This model illustrates the complementarity of endowments and investments, and modifies Almond et al. (2018). There, human capital is not an input to its own production in subsequent periods, so parental responsiveness relies on the intertemporal substitutability of investments across two periods. Here, the functional form for the production of human capital determines the predicted direction of parental investment in response to human capital shocks in preceding periods. Each of these functional forms oversimplifies the process by which children accumulate skills in early childhood, omitting inputs such as preschool and peers. Further, for both extreme cases (perfect substitutes and perfect complements), the predicted response is independent of household resources.

Following from the theory presented in the main text, if the shock S to X in period t reduces contemporaneous human capital H , then $\delta X_{t+1}/\delta S_t$ is equivalent to the more general $\delta X_{t+1}/\delta H_t$. For a Cobb-Douglas U , parents optimize $(1 - \alpha) \log C_t + \alpha \log H_t$, yielding choices that solve

$$\frac{1 - \alpha}{C_t^*} = \frac{\alpha}{f(H_{t-1}, X_t^*)} \frac{df}{dX_t^*} \frac{1}{p_x}.$$

1. Cobb-Douglas: $f(H_{t-1}, X_t) = H_{t-1}^\gamma X_t^{1-\gamma}$ implies optimization

$$\frac{1 - \alpha}{Y - X_t p_x} = \frac{\alpha(1 - \gamma)}{H_{t-1}^\gamma X_t^{1-\gamma}} \left(\frac{H_{t-1}}{X_t} \right)^\gamma \frac{1}{p_x},$$

with solution $X_t^* = \frac{Y}{p_x(1-\alpha\gamma)}$.

2. Perfect substitutes: $f(H_{t-1}, X_t) = \gamma H_{t-1} + (1 - \gamma)X_t$, implies

$$\delta X_t^* / \delta H_{t-1} = \frac{-\gamma(1 - \alpha)}{\gamma - 1} < 0,$$

so a shock to cognition reducing H_{t-1} is expected to induce compensating investments: X_t^* is greater than it would have been in the absence of the shock.

3. Perfect complements: $f(H_{t-1}, X_t) = \min(\gamma H_{t-1}, (1 - \gamma)X_t)$ implies

$$\delta X_t^* / \delta H_{t-1} = \frac{\gamma}{(1 - \gamma)} > 0$$

when f binds, and the Cobb-Douglas solution when the budget constraint binds, so harmful environmental shocks imply corresponding reductions in investment.

2.B Exogeneity and toxicity of TRI exposures

Understanding the estimated coefficient on neonatal toxicity exposure as causal requires that children’s exposure is quasi-exogenous, where the key identifying variation in exposure is within-zip codes, with respect to month of birth.

The disaggregated RSEI microdata model the dispersion of each TRI release to every neighboring grid cell for each year. This allows me to back out for each RSEI grid cell the number of firms and the number of distinct chemicals affecting the cell’s aggregate toxicity in each year, as well as the number of TRI facilities located within each ECLS-B zip code. Changes in these three measures might indicate the sorts of local economic shifts that drive household sorting. Restricting the sample to ECLS-B zip codes that never contained TRI facilities or experienced no change in the number of firms between 2001 and 2002 demonstrates firm openings and closings are not driving most of the year-to-year variation in toxic releases on which effects are identified (Table 2.B1). Similarly, imposing a cutoff on the year-over-year percent change in the number of distinct firms and distinct chemicals affecting each grid cell does not change results meaningfully. Since the RSEI model uses a 50-kilometer radius of potential influence around each facility, almost all ECLS-B zip codes experience some changes in sources of toxic exposure, though the estimated contributions from large distances are typically small; sensitivity to this cutoff is presented in Table 2.B2, and results are consistent with the main specification, except where sample sizes become very small.

Though I control for household demographics in all specifications, I also show their exclusion from the main specification does not meaningfully change results (Table 2.B3), which is consistent with random assignment to “treatment”. Put differently, there should be no correlations between, for example, parents’ incomes, races and ethnicities, or education levels, and the ambient airborne conditions when their children were born. I test this by regressing children’s ambient exposure risk in the first year of life on observable attributes of the household.

Several randomization exercises assign placebo month or location of birth. I first shuffled month of birth within all children born in the same zip code by sampling without

replacement from the actual months in which children were born, and reestimated equation (2.2) using toxicity exposure implied by children’s true locations and the randomized birth months. I separately shuffled household zip codes in the first wave of the survey (again, sampling without replacement, this time from all ECLS-B wave 1 zip codes), always keeping children from the same zip code together. I then reestimated equation (2.2) using the toxicity exposure associated with the new (placebo) location and children’s actual birth months.

Repeating each randomization test 3500 times produces two similar distributions of the estimated coefficients, both centered around zero, and shown in Figure 2.B1. For each distribution of placebo coefficients, about 0.7 percent of the time the estimated exposure to toxicity from randomized months or locations exceeded the effect estimated from the un-randomized data, which is well within what might be expected from random chance.

A final placebo check involves assigning neonatal toxicity exposure to the ECLS-B cohort as if the children had been born in a year other than 2001, without changing their home zip code or month of birth. The placebo years’ “birth year” exposures generally do not produce significant results (Table 2.B4), ruling out the possibility that past variation in TRI exposure is a proxy for contemporaneous variation in pollution.

Table 2.B1: Table 2.6, dropping any zip codes where the number of TRI facilities *located in* the zip code changed at all between 2001 and 2002

	<i>Dependent variable:</i>				
	Cognition (1)	Reading (2)	Math (3)	Motor (4)	Health (5)
A: Zip FE					
$Tox_{z(i)}$	-0.179*** (0.066)	-0.150** (0.070)	-0.105 (0.069)	-0.133* (0.069)	-0.036 (0.063)
Zip FE	Yes	Yes	Yes	Yes	Yes
B: Average toxicity	(1)	(2)	(3)	(4)	(5)
$Tox_{z(i)}$	-0.073 (0.052)	-0.067 (0.051)	-0.039 (0.052)	-0.017 (0.034)	-0.032 (0.035)
Zip-level tox.	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes
Drop open/close zips	Yes	Yes	Yes	Yes	Yes
Waves	2–5	3–5	3–5	2–5	2–5
Observations	19,300	12,600	12,600	19,000	20,600

Some of the ECLS-B zip wave 1 household zip codes contain TRI facilities. For these zip codes, changes in the number of facilities within in a zip code between 2001 and 2002 might be correlated with unobservable household sorting around salient changes in environmental amenities. The exclusion of the zip codes where the number of in-zip code TRI facilities changed between 2001 and 2002 does not meaningfully change estimates.

Table 2.B2: Column 1 of Table 2.6, where zip codes are dropped if the number of chemicals or facilities *affecting* the zip changed by more than a threshold percentage between 2001 and 2002

	<i>Dependent variable:</i>					
	Cognition					
	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z(i)}$	-0.162*** (0.052)	-0.154** (0.061)	-0.177*** (0.065)	-0.109* (0.059)	-0.082 (0.070)	-0.471* (0.267)
Zip FE	Yes	Yes	Yes	Yes	Yes	Yes
B: Average toxicity	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z(i)}$	-0.097** (0.044)	-0.123** (0.048)	-0.122** (0.050)	-0.076 (0.052)	-0.072 (0.058)	-0.597*** (0.217)
Zip-level tox.	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Drop over threshold	None	0.20	0.15	0.10	0.05	0.01
Waves	2–5	2–5	2–5	2–5	2–5	2–5
Observations	23,600	20,600	19,100	16,150	14,050	850

The RSEI model indicates the number of distinct firms and the number of distinct toxic chemicals contributing to the average toxicity at each zip code. Again, large changes in either count might be evidence of inconsistent TRI reporting, or particularly salient local economic shifts in industrial activity (though these totals do not account for the proximity of releases, so some changes may be driven by more remote facilities). Each column presents the estimated effect after dropping children from wave 1 household zip codes associated with successively lower thresholds for this change (from 20 percent down to 1 percent). This has the effect of reducing variation to changes in the quantity of releases from a more constant set of facilities, but since TRI facilities are opening and closing every year, it also dramatically reduces the sample.

Table 2.B3: Table 2.6, with variants of controls included

	<i>Dependent variable:</i>				
	Cognition (1)	Reading (2)	Math (3)	Motor (4)	Health (5)
$Tox_{z(i)}$	−0.008 (0.005)	−0.001 (0.006)	0.003 (0.006)	0.008* (0.004)	−0.011** (0.005)
Zip FE	No	No	No	No	No
Control variables	No	No	No	No	No
	(1)	(2)	(3)	(4)	(5)
$Tox_{z(i)}$	0.0001 (0.005)	0.004 (0.005)	0.008 (0.005)	0.008* (0.004)	−0.002 (0.004)
Zip FE	No	No	No	No	No
Control variables	Yes	Yes	Yes	Yes	Yes
	(1)	(2)	(3)	(4)	(5)
$Tox_{z(i)}$	−0.136** (0.062)	−0.133* (0.074)	−0.083 (0.063)	−0.106* (0.061)	−0.021 (0.053)
Zip FE	Yes	Yes	Yes	Yes	Yes
Control variables	No	No	No	No	No
	(1)	(2)	(3)	(4)	(5)
$Tox_{z(i)}$	−0.162*** (0.052)	−0.162*** (0.060)	−0.103* (0.055)	−0.112* (0.063)	−0.038 (0.058)
Zip FE	Yes	Yes	Yes	Yes	Yes
Control variables	Yes	Yes	Yes	Yes	Yes
Waves	2–5	3–5	3–5	2–5	2–5
Observations	23,600	15,400	15,400	23,250	25,250

Coefficients vary greatly with the inclusion of zip code fixed effects, which is consistent with household sorting across landscapes in a way that correlates unobservable attributes of households with toxicant exposure (though not always pairing high-SES and low-pollution), but the inclusion or exclusion of household demographic controls matters far less for the final estimates, which is consistent with quasi-exogenous treatment of toxicity exposure within zip codes: household covariates are included to improve efficiency, not reduce bias. Controls are child sex, race, household income, number of siblings, mother's age and marital status at birth, twin category, and drinking and smoking during pregnancy. Child age, assessment date, and birth month-by-region fixed effects are included in all specifications.




Figure 2.B1: Coefficients estimated from two randomization exercises

Two separate randomization checks show the importance of child zip code and month of birth for the estimated effects of toxicity exposure on early cognition. The distribution “Shuffle months” is generated by randomly shuffling (sampling without replacement) actual months of birth among all children born in the same zip code and using those placebo birth months to assign toxicity exposure. Similarly, “Shuffle zips” keeps all children from the same zip code together and randomly shuffles all the ECLS-B zip codes. I ran each data randomization 3500 times, and show here the distribution of estimated coefficients, which are quite similar, with means and medians of each effectively zero. About 0.2 percent of estimated coefficients from both distributions exceeded $\hat{\beta}$, the estimated coefficient from the nonrandomized data.

Table 2.B4: Placebo toxicity exposure by birth year

	<i>Independent variable: Placebo $Tox_{z(i)}$</i>						
	2001–02 (1)	2002–03 (2)	2003–04 (3)	2004–05 (4)	2005–06 (5)	2006–07 (6)	2007–08 (7)
A. Zip FE							
Coefficient	−0.150*** (0.052)	0.016 (0.070)	−0.043 (0.046)	0.127* (0.069)	−0.010 (0.070)	−0.017 (0.077)	−0.047 (0.063)
Zip FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
B. Average toxicity	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Coefficient	−0.091** (0.044)	−0.035 (0.042)	−0.023 (0.036)	0.041 (0.032)	−0.055 (0.041)	0.039 (0.042)	−0.040 (0.041)
Zip-level tox.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Waves	2–5	2–5	2–5	2–5	2–5	2–5	2–5
Observations	23,750	23,750	23,750	23,750	23,750	23,750	23,750

Column (1) uses RSEI toxicity exposure for the year in which the ECLS-B children were actually born; the other columns replicate the regression using toxicity exposure as if the cohort had been born in 2002 (averaging 2002 and 2003 RSEI years for placebo neonatal toxicity), 2003, and so on. Coefficients on this placebo exposure, in otherwise-identical regressions where the dependent variable is cognition, are generally indistinguishable from zero.

2.C Risk-Screening Environmental Indicators model

Modeled exposures to airborne industrial pollution used in this paper rely on the EPA’s RSEI model, which estimates the dispersion of chemical emissions from TRI facilities. This model intended to transform the point-source pounds of chemicals released at TRI facilities into a more useful measure of relative hazard to human health; it is discussed at length in the EPA documentation¹⁰, and summarized here.

The RSEI model starts by locating the reporting facility for each TRI release in space. Releases in pounds are reported separately for each distinct facility, chemical, release medium, and year. Releases to the air are differentiated by those that exit a facility at a point, through a stack or vent (“stack air emissions” in the RSEI model), and other emissions, such as leaks and evaporation (“fugitive”).

The dispersion of each TRI air release is mapped from its facility to an airborne concentration on an 810-meter grid covering the United States. The American Meteorological Society/EPA Regulatory Model (AERMOD) is a “steady-state Gaussian plume model used to estimate pollutant concentrations downwind of a stack or area source”; this means the neighborhoods where TRI releases end up depends on the local meteorology (such as wind speed and direction), chemical-specific air decay rates, and characteristics of the facility. Constant releases over the reporting year are assumed, as well as industry-specific stack heights, diameters, and velocities if facility-specific data are missing. The result is estimates of ground-level pollutant concentrations for grid cells up to 49 kilometers away from the facility (determined to be the distance from facilities at which concentrations are sufficiently close to zero under most conditions) (Figure 2.C1).

The toxicity index is constructed by summing the product of each chemical’s inhalation toxicity weight for each estimated concentration for each grid cell, producing an aggregate measure of chronic health risks over the 810-meter grid. Inhalation toxicity weights are chosen to reflect the EPA’s accumulated knowledge of the toxicity of chemicals, incorporating information from many offices and agencies (such as the Integrated

¹⁰See https://www.epa.gov/sites/production/files/2018-12/documents/rsei_methodology_v2.3.7.pdf for an overview, and for greater detail https://www.epa.gov/sites/production/files/2014-03/documents/rsei_analyses.pdf.

Risk Information System, Office of Air Quality Planning and Standards, Office of Pesticide Programs, Agency for Toxic Substances and Disease Registry, and California's EPA. Assigned values encompass risks for chronic, carcinogenic and non-carcinogenic outcomes resulting from exposure, and are down-weighted if there is uncertainty in the scientific evidence; acute health effects are not incorporated. The inhalation toxicity weights are unitless. Less than 1 percent of TRI releases are chemicals for which the EPA does not have a toxicity weight.

The full RSEI model also incorporates data from the U.S. census to estimate population exposures, but this analysis does not use those data.




Figure 2.C1: RSEI dispersion (Source: EPA)

The dispersion of each release reported to RSEI is a function of facility and chemical characteristics, and local meteorology, and is estimated for a 50-kilometer radius around TRI facilities. Figure from <https://www.epa.gov/rsei/modeling-air-releases-rsei>.

2.D Exposure attribution and timing

This paper focuses on exposures in infancy, in contrast to much of the extant literature, which focuses on fetal exposures. An exhaustive empirical approach to the question of exposure timing would require a panel of early childhood development spanning more than one birth year, but a discussion of in utero exposures, as well as simple tests of the sensitivity of results to changes in the independent variable, are discussed in this appendix.

TRI reporting requirements changed for lead (but nothing else) in 2001, and because these analyses use pollution data from TRI reporting year 2000, I recalculate the aggregate, zip-level toxicity-weighted concentrations using only the subset of chemicals for which reporting requirements were constant. This ensures that identifying variation is not being determined by changes in the lead reporting rule, but instead actual changes in firms' activity.

First, a difference-in-difference specification allows me to understand the relative harm from exposure during the fetal and infant periods without making any assumptions about the importance of each.

Consider a regression equation of the form:

$$\begin{aligned} Cog_{iw} = & \beta_0 T_{z,2000} \times B_i + \beta_1 T_{z,2001} \times B_i + \beta_2 T_{z,2003} \times B_i \\ & + \beta_3 T_{z,2000} + \beta_4 T_{z,2001} + \beta_5 T_{z,2003} + \gamma X_{iw} + \theta_w A_{iw} + \mu_{iw} + \eta_{r(i),B_i} + \epsilon_{izw}. \end{aligned} \quad (2.D1)$$

As in the main text, the outcome Cog is the standardized score on cognitive assessments for a child i born in zip code z in (integer) month of birth B . Fixed effects η_{r,B_i} flexibly absorb averages differences in cognitive scores across each region-by-month of birth; controls X and A are defined above. Departing from the main text, I include zip-level, annual, ambient toxicity T for 2000, 2001, and 2002—the years over which the entire ECLS-B cohort was both in utero and in infancy—rather than an average over just the neonatal years.

The coefficient of interest, β_1 , allows us to understand the relative importance of

exposure during the fetal and infant periods using variation in the 2001 cohort's months of birth. More precisely, the sign of β_1 suggests which window of toxic exposure matters relatively more for later cognition, and its magnitude reflects the net difference between the two effects. The coefficients β_3 , β_4 , and β_5 are not identifiable with zip code fixed effects, and so not the focus of this estimation, but for now consider each is somewhere between null (indistinguishable from zero) and harmful (negative).

Intuition follows from two extreme cases. First, assume *fetal* exposure is the only pollution that affects later cognition. Then a child born at the beginning of 2001 avoids that year's harmful exposure, having been in utero almost entirely during 2000. Children born in the same place progressively later in the year experience an increasingly negative penalty to cognition from toxic air pollution in 2001, as the share of their gestational time in 2001 increases. We should then expect $\beta_1 < 0$: the harm from toxicity in 2001 is increasing in month of birth. Similarly, $\beta_0 > 0$, the harmful effect of pollution occurring during 2000 diminishes toward zero as children are born later and later in 2001; and there is no prediction about β_2 . But if exposures in *infancy* are the only pollution that matters, then a child born at the beginning of 2001 experiences all of that year's harmful exposure during the critical window. Children born later in the year live through less of 2001 and more of 2002. It follows that we expect $\beta_1 > 0$: the magnitude of the cognitive penalty from exposure in 2001 falls toward zero with month of birth. Correspondingly, $\beta_2 < 0$ and there is no prediction about β_0 .

Of course, these two cases bound the more realistic scenario that *both* fetal and infant exposures matter, in which case β_1 captures the net effect of the timing of toxicant exposure during each crucial period of development. The sign of β_1 suggests which exposure window is more strongly related to adverse, medium-term cognitive outcomes, and its magnitude by how much. With respect to the other coefficients, even if the first case is "more" true, we might still expect $\beta_2 < 0$, or even if the second case dominates, it may still be that $\beta_0 > 0$.

Estimates for this regression are presented in Table 2.D1 for specifications that include either zip code or state fixed effects; the positive and significant coefficients on β_1 suggests

when airborne toxicity covering both the prenatal and neonatal periods are included in the same regression, neonatal exposures appear to matter relatively more than fetal exposures for cognition. This study, which combines a particular cohort of children and toxic air pollution, may be somehow unique, but these findings imply greater consideration should be given to a larger period of vulnerability than strictly fetal origins. More important than precise identification of “window of exposure”, researchers and policymakers could understand that early childhood comprises multiple vulnerable periods of development, and policy to reduce pollution for both mothers and young children has potentially large benefits for long-term learning.

For the TRI data, annual measures of releases further complicate identification, especially for children born in the middle months of 2001 likely have measurement error in the attribution of ambient toxicity risk in infancy than children born at either end of 2001: the distinction between pre- and neonatal ambient toxicity for a child born in the middle of the year is murkier than for a child born in January. For a child born in July, for example, some of the toxic releases reported for 2001 likely occurred while the child was in utero, and some occurred in infancy. Conversely, a January baby is exposed to all of the 2001 toxicant releases, and they all occur while the child is an infant. Mechanically, in the zip code fixed effects specification, if two children are born in adjacent months their attributed neonatal toxicity exposures are also quite similar, while children in a single zip code born at opposite ends of the year have substantially more identifying variation in ambient toxicity risk.

The solution to accurately attributing TRI exposure in Currie and Schmieder (2009) is to only examine births occurring in the first few months of the year. With only 10,700 ECLS-B children in the survey, restricting the sample by month of birth may diminish power more than it helps control measurement error. Reassuringly, results are fairly consistent whether the “middle-month” children are retained or dropped from the sample: Table 2.D2 presents coefficients estimated from equation (2.2) with only children from the end months of the year, indicating the main results are not being driven by very small differences in attributed exposure for children of nearly the same age.

Alternative constructions of the independent variables affect the magnitude of effects estimated from the main specifications. Focusing on within-zip code variation, the first year of life is when children most clearly experienced or avoided pollution exposure differently from their peers, so in the context of this study, the biological importance of the first year of life is not entirely separable from its econometric importance. The cross-sectional specification makes it clear other “windows of vulnerability” may also be associated with diminished cognitive outcomes, suggesting the potential for a broader consideration of early childhood vulnerability to air toxics than simply the fetal or neonatal periods. Regressions using different right-hand-side variables—ranging from the 6 months before birth and after birth to the 18 months after birth, as well as the 12-month window before the Wave 2 assessment (“contemporaneous” exposure using current household zip code)—are presented in table 2.D3. These different exposures weight combinations of the RSEI years 2000 through 2004.

Finally, though the main focus of this paper is the cognitive effect of exposure in infancy, and Table 2.D1 indicates that exposure matters relatively more for cognition in the case of airborne toxicity, it is possible to check how prenatal outcomes relate to toxicant exposure. However, the ECLS-B cohort is a selected sample with regards to prenatal health: since it draws from live births at 9 months, the most severe outcomes of fetal or infant death as a result of exposure to toxic air pollution are missing from the data. Table 2.D4 shows no measurable effect of neighborhood toxicant exposure on several measures of birthweight, several measures of gestational length, and five-minute APGAR scores. For these regressions, I use as the dependent variable the zip-level variation in toxicant exposure in the 12 months *before* birth, averaging the RSEI years 2000 and 2001 (denoted $Tox_{z(i)-12}$) and otherwise replicating equation (2.2).

Table 2.D1: Exposure-by-month of birth analysis

	<i>Dependent variable:</i>					
	Cognition		Reading		Math	
	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z,2000}$		0.023 (0.000)		0.015 (0.032)		−0.001 (0.030)
$Tox_{z,2001}$		−0.070** (0.000)		−0.083** (0.038)		−0.029 (0.032)
$Tox_{z,2002}$		0.039 (0.000)		0.065** (0.031)		0.029 (0.031)
$Tox_{z,2000} \times B_i$	−0.009** (0.004)	0.0001 (0.004)	−0.010* (0.005)	0.001 (0.004)	−0.006 (0.006)	0.002 (0.004)
$Tox_{z,2001} \times B_i$	0.018*** (0.005)	0.009** (0.004)	0.020*** (0.006)	0.010** (0.005)	0.012** (0.006)	0.005 (0.004)
$Tox_{z,2002} \times B_i$	−0.007** (0.004)	−0.008** (0.003)	−0.008** (0.004)	−0.010*** (0.003)	−0.003 (0.004)	−0.006 (0.004)
Zip FE	Yes	No	Yes	No	Yes	No
State FE	No	Yes	No	Yes	No	Yes
Cluster	City	City	City	City	City	City
Waves	2:5	2:5	3:5	3:5	3:5	3:5
Observations	23,750	23,750	15,500	15,500	15,500	15,500

A difference-in-differences specification (equation (2.D1), described in the text) suggests that higher levels of airborne toxicity are relatively more harmful for children currently in infancy than those still in utero. Coefficients on zip-level TRI toxicity risk interacted with children's months of birth indicate harm from pollution in 2001—the year over which the ECLS-B cohort was born—is diminishing in children's month of birth (the coefficient on $Tox_{z,2001} \times B_i$). This means higher pollution in 2001 causes children born at the beginning of the year, who experience that year as infants, to perform relatively worse on cognitive assessments than children born at the end of the year, who experience that year mostly in utero. Similarly, the coefficients on $Tox_{z,2002} \times B_i$ tend to suggest children born later in 2001, with a greater share of their infancy in 2002 than slightly older peers, perform relatively worse the more polluted their zip code in 2002.

Table 2.D2: Estimates from ECLS-B subsamples, by month of birth

	<i>Dependent variable:</i>					
	Cognition					
	(1)	(2)	(3)	(4)	(5)	(6)
$Tox_{z(i)}$	-0.162*** (0.052)	-0.177*** (0.057)	-0.201*** (0.063)	-0.217*** (0.072)	-0.171** (0.086)	-0.205 (0.154)
Birth months dropped	None	6-7	5-8	4-9	3-10	2-11
Zip FE	Yes	Yes	Yes	Yes	Yes	Yes
Waves	2-5	2-5	2-5	2-5	2-5	2-5
Observations	23,600	20,000	16,200	12,350	8,150	4,150

For the specification using a zip code fixed effect, identifying variability in neonatal toxicity exposure is stronger when children are born at opposite ends of the year. Very small differences across children born in the middle of year are not driving the main results: columns (2) through (6) reestimate the zip code fixed effect specification using increasingly restricted subsets of the data.

Table 2.D3: Regressions of cognitive scores on toxicant exposure over different critical periods

	<i>Independent variable: $Tox_{z(i)}$ over exposure window...</i>				
	± 6 months	9 months	12 months	15 months	18 months
A: Zip FE	(1)	(2)	(3)	(4)	(5)
Coefficient	0.002 (0.041)	-0.105*** (0.038)	-0.159*** (0.051)	-0.148** (0.053)	-0.124**
Zip FE	Yes	Yes	Yes	Yes	Yes
B: Average toxicity	(1)	(2)	(3)	(4)	(5)
Coefficient	-0.046 (0.034)	-0.085*** (0.033)	-0.087** (0.044)	-0.078* (0.046)	-0.074* (0.042)
Zip-level tox.	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes
Waves	2-5	2-5	2-5	2-5	2-5
Observations	23,600	23,600	23,600	23,600	23,600

Different windows of exposure affect the estimated magnitude of the effect of toxicity exposure on cognition. Column (3) corresponds to the window used in the main text. This sensitivity is likely driven both by the biological importance of the first year of life for cognitive development in early childhood, and the context of this study: differences in ambient toxicity during the first year of life are most-differentiated for children born in a single cohort.

Table 2.D4: Prenatal outcomes from toxicant exposure

	<i>Dependent variable:</i>						
	Bwt. (kg) (1)	LBW (0/1) (2)	VLBW (0/1) (3)	Gest. (wks) (4)	Prem. (wks) (5)	Prem. (0/1) (6)	APGAR (7)
$Tox_{z(i)-1}$	-0.019 (0.036)	-0.014 (0.018)	-0.003 (0.005)	0.217 (0.231)	-0.153 (0.116)	-0.006 (0.032)	0.053 (0.042)
Zip FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
B: Average toxicity	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Tox_{z(i)-1}$	-0.021 (0.023)	0.005 (0.007)	-0.0004 (0.002)	0.032 (0.115)	-0.044 (0.044)	0.007 (0.015)	0.040 (0.027)
Zip-level tox.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Waves	1	1	1	1	1	1	
Observations	9,950	9,950	9,950	9,950	9,950	9,950	7,724

Using panel and cross-sectional variation in toxicant exposure in the year before the ECLS-B children were born (combining RSEI years 2000 and 2001) shows higher exposure to TRI releases reduces birthweight.

Bibliography

- Agarwal, Nikhil, Chanont Banternghansa, and Linda T.M. Bui**, “Toxic exposure in America: Estimating fetal and infant health outcomes from 14 years of TRI reporting,” *Journal of Health Economics*, July 2010, 29 (4), 557–574.
- Almond, Douglas and Janet Currie**, “Human capital development before age five,” in “Handbook of Labor Economics,” Vol. 4 2011, pp. 1315–1486.
- , —, and **Valentina Duque**, “Childhood Circumstances and Adult Outcomes: Act II,” *Journal of Economic Literature*, 2018, 56 (4), 1360–1446.
- Andreassen C and Fletcher P**, “Early Childhood Longitudinal Study, Birth Cohort (ECLS-B): Psychometric Report for the 2-Year Data Collection.,” *National Center for Education Statistics*, 2007, (Methodology Report).
- Ash, Michael and James K Boyce**, “Racial disparities in pollution exposure and employment at US industrial facilities.,” *Proceedings of the National Academy of Sciences of the United States of America*, October 2018, 115 (42), 10636–10641.
- Attanasio, Orazio, Costas Meghir, and Emily Nix**, “Human Capital Development and Parental Investment in India,” 2019.
- Banzhaf, Spencer, Lala Ma, and Christopher Timmins**, “Environmental justice: The economics of race, place, and pollution,” *Journal of Economic Perspectives*, 2019, 33 (1), 185–208.
- Bateson, Thomas F. and Joel Schwartz**, “Children’s response to air pollutants,” *Journal of Toxicology and Environmental Health - Part A: Current Issues*, December 2008, 71 (3), 238–243.
- Bearer, C. F.**, “Environmental health hazards: how children are different from adults.,” *The Future of children / Center for the Future of Children, the David and Lucile Packard Foundation*, 1995, 5 (2), 11–26.
- Bharadwaj, Prashant, Juan Pedro Eberhard, and Christopher A. Neilson**, “Health at birth, parental investments, and academic outcomes,” *Journal of Labor Economics*, April 2018, 36 (2), 349–394.
- , **Matthew Gibson, Joshua Graff Zivin, and Christopher Neilson**, “Gray Matters: Fetal Pollution Exposure and Human Capital Formation,” *Journal of the Association of Environmental and Resource Economists*, June 2017, 4 (2), 505–542.

- Black, Sandra E, Aline Bütikofer, Paul J Devereux, and Kjell G Salvanes**, “This is Only a Test? Long-Run Impacts of Prenatal Exposure to Radioactive Fallout from Nuclear Weapon Testing,” Technical Report 18987 2013.
- Block, Michelle L., Alison Elder, Richard L. Auten, Staci D. Bilbo, Honglei Chen, Jiu Chiuan Chen, Deborah A. Cory-Slechta, Daniel Costa, David Diaz-Sanchez, David C. Dorman, Diane R. Gold, Kimberly Gray, Huei-wang Anna Jeng, Joel D. Kaufman, Michael T. Kleinman, Annette Kirschner, Cindy Lawler, David S. Miller, Srikanth S. Nadadur, Beate Ritz, Erin O. Semmens, Leonardo H. Tonelli, Bellina Veronesi, Robert O. Wright, and Rosalind J. Wright**, “The outdoor air pollution and brain health workshop,” *NeuroToxicology*, October 2012, *33* (5), 972–984.
- Boyce, James K., Klara Zwickl, and Michael Ash**, “Measuring environmental inequality,” *Ecological Economics*, 2016, *124*, 114–123.
- Bruckner, James V.**, “Differences in sensitivity of children and adults to chemical toxicity: The NAS panel report,” *Regulatory Toxicology and Pharmacology*, 2000, *31* (3), 280–285.
- Case, Anne, Darren Lubotsky, and Christina Paxson**, “Economic status and health in childhood: The origins of the gradient,” *American Economic Review*, November 2002, *92* (5), 1308–1334.
- Chetty, Raj, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan**, “How does your kindergarten classroom affect your earnings? Evidence from project star,” *Quarterly Journal of Economics*, November 2011, *126* (4), 1593–1660.
- Claessens, Amy, Greg Duncan, and Mimi Engel**, “Kindergarten skills and fifth-grade achievement: Evidence from the ECLS-K,” *Economics of Education Review*, 2009, *28* (4), 415–427.
- Cunha, Flavio and James J Heckman**, “Formulating , Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation,” *Journal of Human Resources*, 2016, *43* (4), 738–782.
- , —, and **Susanne M Schennach**, “Estimating the Technology of Cognitive and Noncognitive Skill Formation,” *Econometrica*, 2010, *78* (3), 883–931.
- Currie, Janet and Johannes F. Schmieder**, “Fetal exposures to toxic releases and infant health,” *American Economic Review*, April 2009, *99* (2), 177–183.
- and **Mark Stabile**, “Socioeconomic status and child health: Why is the relationship stronger for older children?,” *American Economic Review*, 2003, *93* (5), 1813–1823.
- and **Rosemary Hyson**, “Is the impact of health shocks cushioned by socioeconomic status? The case of low birthweight,” *American Economic Review*, May 1999, *89* (2), 245–250.
- , **Joshua Graff Zivin, Jamie Mullins, and Matthew Neidell**, “What Do We Know About Short- and Long-Term Effects of Early-Life Exposure to Pollution?,” *Annual Review of Resource Economics*, November 2014, *6* (1), 217–247.

- , **Lucas Davis, Michael Greenstone, and Walker Reed**, “Environmental health risks and housing values: Evidence from 1,600 toxic plant openings and closings,” *American Economic Review*, February 2015, 105 (2), 678–709.
- Del Boca, Daniela, Christopher Flinn, and Matthew Wiswall**, “Household choices and child development,” *Review of Economic Studies*, January 2014, 81 (1), 137–185.
- Dizon-Ross, Rebecca**, “Parents’ Beliefs about Their Children’s Academic Ability: Implications for Educational Investments,” *American Economic Review*, August 2019, 109 (8), 2728–2765.
- Duncan, Greg J., Chantelle J. Dowsett, Amy Claessens, Katherine Magnuson, Aletha C. Huston, Pamela Klebanov, Linda S. Pagani, Leon Feinstein, Mimi Engel, Jeanne Brooks-Gunn, Holly Sexton, Kathryn Duckworth, and Crista Japel**, “School Readiness and Later Achievement,” *Developmental Psychology*, 2007, 43 (6), 1428–1446.
- Ferrie, Joseph P., Karen Rolf, and Werner Troesken**, “Cognitive disparities, lead plumbing, and water chemistry: Prior exposure to water-borne lead and intelligence test scores among World War Two U.S. Army enlistees,” *Economics and Human Biology*, June 2012, 10 (1), 98–111.
- Figlio, David, Jonathan Guryan, Krzysztof Karbownik, and Jeffrey Roth**, “The effects of poor neonatal health on children’s cognitive development?,” *American Economic Review*, December 2014, 104 (12), 4205–4230.
- Heckman, James J.**, “The economics, technology, and neuroscience of human capability formation,” *Proceedings of the National Academy of Sciences of the United States of America*, August 2007, 104 (33), 13250–13255.
- , “Schools, skills, and synapses,” *Economic Inquiry*, 2008, 46 (3), 289–324.
- Heft-Neal, Sam, Jennifer Burney, Eran Bendavid, and Marshall Burke**, “Robust relationship between air quality and infant mortality in Africa,” *Nature*, July 2018, 559 (7713), 254–258.
- Isen, Adam, Maya Rossin-Slater, and W. Reed Walker**, “Every breath you take—every dollar you’ll make: The long-term consequences of the clean air act of 1970,” *Journal of Political Economy*, June 2017, 125 (3), 848–902.
- Kelly, Frank J. and Julia C. Fussell**, “Size, source and chemical composition as determinants of toxicity attributable to ambient particulate matter,” *Atmospheric Environment*, December 2012, 60, 504–526.
- Lee, Rae Hyuck, Fuhua Zhai, Jeanne Brooks-Gunn, Wen Jui Han, and Jane Waldfogel**, “Head start participation and school readiness: Evidence from the early childhood longitudinal study-birth cohort,” *Developmental Psychology*, January 2014, 50 (1), 202–215.
- Lynch, Jamie L. and Ryan Brooks**, “Low Birth Weight and Parental Investment: Do Parents Favor the Fittest Child?,” *Journal of Marriage and Family*, June 2013, 75 (3), 533–543.

- Marchi, Scott De and James T. Hamilton**, “Assessing the accuracy of self-reported data: An evaluation of the toxics release inventory,” *Journal of Risk and Uncertainty*, 2006, 32 (1), 57–76.
- Marcotte, Dave E.**, “Something in the air? Air quality and children’s educational outcomes,” *Economics of Education Review*, 2017, 56, 141–151.
- Najarian, Michelle, Kyle Snow, Jean Lennon, Susan Kinsey, and Gail Mulligan**, “Early Childhood Longitudinal Study, Birth Cohort (ECLS-B),” *Preschool-Kindergarten 2007 psychometric report*, 2010, (April), 2010–009.
- Persico, Claudia, David Figlio, and Jeffrey Roth**, “Inequality Before Birth: The Developmental Consequences of Environmental Toxicants,” *National Bureau of Economic Research Working Paper Series*, May 2016, No. 22263.
- Persico, Claudia L and Joanna Venator**, “The Effects of Local Industrial Pollution on Students and Schools,” *Journal of Human Resources*, 2019.
- Rau, Tomás, Sergio Urzúa, and Loreto Reyes**, “Early Exposure to Hazardous Waste and Academic Achievement: Evidence from a Case of Environmental Negligence,” *Journal of the Association of Environmental and Resource Economists*, 2015, 2 (4), 527–563.
- Renwick, A. G.**, “Toxicokinetics in infants and children in relation to the ADI and TDI,” *Food Additives and Contaminants*, January 1998, 15 (sup001), 17–35.
- Rosales-Rueda, Maria and Margaret Triyana**, “The Persistent Effects of Early-Life Exposure to Air Pollution: Evidence from the Indonesian Forest Fires,” *Journal of Human Resources*, April 2018, pp. 0117–8497R1.
- Royer, Heather**, “Separated at girth: US twin estimates of the effects of birth weight,” *American Economic Journal: Applied Economics*, January 2009, 1 (1), 49–85.
- Sanders, Nicholas J.**, “What doesn’t kill you makes you weaker: Prenatal pollution exposure and educational outcomes,” *Journal of Human Resources*, 2012, 47 (3), 826–850.
- Scheuplein, Robert, Gail Charnley, and Michael Dourson**, “Differential sensitivity of children and adults to chemical toxicity. I. Biological basis,” *Regulatory Toxicology and Pharmacology*, 2002, 35 (3), 429–447.
- Snow, Kyle, Lisa Thalji, Azucena Derecho, Sara Wheeless, Jean Lennon, Susan Kinsey, James Rogers, Melissa Raspa, and Jennifer Park**, “User’s Manual for the ECLS-B Longitudinal 9-Month–Preschool Restricted-Use Data File and Electronic Codebook,” Technical Report NCES Publication 2008024 December 2007.
- Spencer Banzhaf, H. and Randall P. Walsh**, “Do people vote with their feet? An empirical test of tiebout’s mechanism,” *American Economic Review*, May 2008, 98 (3), 843–863.
- Ünüvar, Tolga and Atilla Büyükgebiz**, “Fetal and neonatal endocrine disruptors,” *JCRPE Journal of Clinical Research in Pediatric Endocrinology*, June 2012, 4 (2), 51–60.

- U.S. EPA**, “Risk-Screening Environmental Indicators Methodology Document,” Technical Report January, Environmental Protection Agency, Washington, DC 2018.
- Valavanidis, Athanasios, Konstantinos Fiotakis, and Thomais Vlachogianni**, “Airborne particulate matter and human health: Toxicological assessment and importance of size and composition of particles for oxidative damage and carcinogenic mechanisms,” *Journal of Environmental Science and Health - Part C Environmental Carcinogenesis and Ecotoxicology Reviews*, December 2008, 26 (4), 339–362.
- van Donkelaar, A, R V Martin, M Brauer, N C Hsu, R A Kahn, R C Levy, A Lyapustin, A M Sayer, and D M Winker**, “Global Annual PM_{2.5} Grids from MODIS, MISR and SeaWiFS Aerosol Optical Depth (AOD) with GWR, 1998-2016,” 2018.
- Voorheis, John**, “Air Quality, Human Capital Formation and the Long-term Effects of Environmental Inequality at Birth,” *CARRA Working Paper Series*, 2017, pp. 1–22.
- Yi, Junjian, James J. Heckman, Junsen Zhang, and Gabriella Conti**, “Early Health Shocks, Intra-household Resource Allocation and Child Outcomes,” *Economic Journal*, November 2015, 125 (588), F347–F371.

Chapter 3

Pollution and Acquisition: The Environmental Justice Effects of Mergers

Abstract

I estimate the effect of mergers and acquisitions on both facility-level toxic air pollution and its firm-level distribution. Since acquisition is endogenous to the operation and emissions of polluting facilities, making target facilities substantially different from other polluting establishments, I use event study designs that exploit variation in the timing of acquisition among target facilities. When a plant continues to operate after acquisition (many facilities close entirely), I find evidence operations change, but not necessarily in ways that reduce the health risks from its air pollution. I also find evidence of increased inequality in emissions among target plants after acquisition, and shifts in pollution toward lower-income and less-white neighborhoods. These findings suggest consolidation in sectors with negative externalities may reduce the externality but increase inequality in its exposure.

3.1 Introduction

A large literature focuses on the human impacts of exposure to environmental harm, and particularly air pollution, on long- and short-run health, cognition, and behavior. These studies are empirically challenging because patterns of environmental inequality mean many dimensions of socioeconomic disadvantage are associated with higher pollution

burdens. Yet considerably less work has directly examined the determinants of these observed correlations between poor and minority neighborhoods and adverse environmental exposures.

This paper focuses on the supply of toxic emissions from facilities in the U.S., analyzing the ways in which firm behavior contributes to environmental inequality. Existing research on facility siting tends to focus on the relationship between individual plants and local characteristics. Wolverton (2009) argues “disproportionate siting”—the process by which firms target low-income and minority communities with heavily-polluting activities—likely matters less than local production costs, such as land, labor, and access to transportation. But we know relatively little about how corporations coordinate or allocate polluting activities across plants. This is a potentially important dimension of environmental inequality, since the majority of Toxics Release Inventory (TRI)-reporting facilities are owned by parent companies who control at least ten, and sometimes hundreds, of separate plants. Studying M&A-induced changes in operations at TRI facilities, most of which operate continuously over the study period, allows for a research design that is not limited to the set of facilities that open, close, or relocate.

This paper asks whether mergers and acquisitions (M&A) tend to exacerbate or mitigate environmental inequality, which refers to a distribution of pollution that puts disproportionate burden on low income and minority neighborhoods. I study corporate mergers and acquisitions as both a plant-level treatment, and a shift in the number and location of facilities already held by parent companies. I estimate the effect of a merger on emissions using event study designs, exploiting variation in the timing of acquisition among target facilities. This work does not explain why certain firms and not others are acquired, but instead takes as given the transactions that occurred in the past two decades and studies their effects.

Observations come from the EPA’s Toxics Release Inventory (TRI), a federal program mandating facilities report emissions of toxic chemicals across many industrial sectors. The number of firms and the pounds of emissions released in the US has fallen over the past two decades, but toxicity risk—the potential for adverse health effects from

exposure—has fallen by less, suggesting this source of air pollution has become increasingly harmful. Observing establishment-level emissions allows me to isolate the effect of an acquisition on emissions while controlling for fixed facility characteristics. I use corporate hierarchy data from Dun & Bradstreet, and the Thompson database of M&A activity. All of these sources cover both publicly- and privately-held U.S. companies, providing a detailed picture of the ownership of large point-source polluters in the US.

M&A activity is an important feature of the U.S. economy, with over \$2.5 trillion in announced deals in the first half of 2018. TRI-specific sectors are no different: in the global chemical manufacturing industry, for example, 2018 merger activity exceeded 72 billion dollars and 600 transactions, the bulk of which targeted U.S. companies.¹

Toxic air pollution is a local environmental externality, so firm choices over where, how much, and how cleanly to operate have meaningful impacts on the neighborhoods in which plants are located, and the effect of a merger on the quantity and distribution of this negative externality is not obvious. Acquiring companies tend to be larger and better-resourced, and may implement more efficient pollution control; these technological asymmetries and the transfer of environmental expertise between acquiring and target firms can motivate M&A activity (Berchicci et al., 2012; Kwon et al., 2018). This pathway resembles the strategy of target facilities adopting the technology and practices of the acquiring firm documented in Eliason et al. (2019). Consolidation-driven improvements in abatement technology is also consistent with observed trends in the manufacturing sector described in Shapiro and Walker (2018).

However, a merger may also allow firms to reduce or terminate risky or harmful emissions in the most advantaged communities, exacerbating environmental inequality; acquisitions may even be motivated by this opportunity. Banzhaf et al. (2019) proposes as a potential explanation for environmental justice correlations that firms locate hazardous activities across space and time to minimize costs. These costs may include location-specific expectations of enforcement, current and anticipated regulation, community oversight, and preferences for discrimination (Hamilton, 1995). This potential

¹New York times, July 2018; Deloitte's 2019 Global chemical industry mergers and acquisitions outlook.

side effect of mergers—the redistribution of the burden of pollution among facilities in a way that increases environmental inequality—adds another dimension to the tradeoff between increased market power and efficiency gains resulting from M&A activity.

Though much literature has focused on the effects of M&A activity on firm performance and market power (such as Blonigen and Pierce (2016)), there is no empirical evidence documenting the effect of M&A activity on environmental performance or the distribution of environmental externalities. Several relevant bodies of work consider the welfare effects of mergers in sectors with externalities, the role of trade in the environmental performance of firms, and the ways firms are sensitive to regulatory and community-based pressure in their choices of location and environmental performance.

Accounting for firm structure, ownership dynamics, and heterogeneity across plant locations is essential to understanding environmental performance. Grant and Jones (2003) finds subsidiary firms tend to pollute more heavily than their parent companies, and Akey and Appel (2017) finds greater liability protection for parent companies drives higher emissions among subsidiaries through disincentivizing investment in abatement technologies. Acquisitions may be motivated by technological asymmetries between the target and acquiring firms, or expansion into less-regulated (or less-scrutinized) locations. Cui et al. (2016) discusses the effects of firm structure and trade on environmental performance, finding larger and more productive facilities tend to have lower emission intensities for criteria pollutants. Similarly, Holladay (2016) finds toxic emissions are lower among exporting firms, and Li and Zhou (2017) associate lower toxic pollution at US plants with firms that import more from low-wage countries.

Fikru (2011) shows environmental policy can incentivize M&A activity among polluting firms, and Condliffe and Morgan (2009) finds stricter environmental regulations deter new plant openings, particularly in the most-polluting sectors. Among operating facilities, however, Becker (2011) does not find evidence higher environmental compliance costs diminish facility productivity. Fikru and Gautier (2020) provides a theoretical framework for incorporating environmental externalities in the welfare effects of mergers, proposing overall welfare effects depend on the size of the merger and the pollution intensity of the

firms involved. Simon and Prince (2016) finds increased competition reduces facility-level emissions through reduced output and increased abatement; however, that paper does not account for facility ownership or facility-specific heterogeneity, instead linking the HHI of each facility's industry. Park (2019) documents mixed effects of airline mergers on NOx emissions.

Community characteristics may inform both the perceived value of acquisitions and also subsequent management of target firms. De Silva et al. (2016) finds polluting facilities are more likely to locate in tracts with higher shares of nonwhite residents, and Millimet and Slottje (2002) finds federal regulation that increases environmental compliance costs acts to exacerbate disparities in air pollution exposure. Ash and Boyce (2011) document patterns of disproportionate exposure to toxic pollution within major firms in the US, and Kim et al. (2019) finds local facility ownership is associated with lower toxic chemical releases. Sanders (2012) and Powers (2013) estimate the effects of the information shock of the TRI itself on housing prices and plant births, and Wang et al. (2018) finds the 1989 TRI disclosures caused the relocation of polluting facilities towards more disadvantaged communities. Similarly, Kagan et al. (2003) attributes some heterogeneity in environmental performance across polluting firms to local community and activist pressure, and Campa (2018) finds a negative effect of press coverage on TRI emissions.

To understand the distributional impacts of M&A activity, I take as a conceptual framework the decisions made by a firm with multiple polluting establishments that perform similar functions and operate in multiple markets. Each firm decides each year how to allocate production across its set of facilities, minimizing costs at each facility, and constrained in the short run by total capacity. Each facility has its own access to local markets and supply chains, and faces local regulations and forms expectations of their enforcement. M&A activity occurs if it is perceived to increase value for the acquiring firm. As Berchicci et al. (2012) and Simon and Prince (2016) note, an acquisition may provide either the target or acquiring firm with new expertise, including for pollution control. The acquisition also increases the set of facilities available for production, and firms may

make changes at any of them after the acquisition. I hypothesize firms' increased scope to decide where to operate leads to an increased emissions share in more-disadvantaged neighborhoods (those that have lower median incomes and greater nonwhite populations). The firm's ability to reallocate production among facilities is constrained by the extent to which plants are substitutes for one another; I expect large firms to be the most likely to have multiple facilities with substitutable production, and similarly, that horizontal mergers are most likely to create redundancy and an opportunity to reallocate production.

This paper estimates the effects of a change in firm ownership on high-risk emissions at the facility level using the event study design of Sun and Abraham (2020), which allows for treatment effect heterogeneity and causal interpretation. I find the quantity of toxic chemicals released to the air at target facilities falls by about 24 percent in the decade following an acquisition, a much larger estimate than the descriptive decreases reported in Berchicci et al. (2012), and examine potential mechanisms for these effects. Reductions tend to be greatest when the target and acquiring firms operate in the same industrial sector, suggesting technology transfer is important. Facilities treated by an acquisition are also more likely to close down in the years immediately following the acquisition than never-acquired facilities of the same age and operating in the same sector.

I then estimate the heterogeneity of the effect of an acquisition by the demographic characteristics of the neighborhoods in which facilities are located, and characterize the changes in distribution and the socioeconomic composition of the population exposed to a firms' emissions using firm-level summary statistics, the exposure-weighted average (in the spirit of the firm-level distributional measure in Ash and Boyce (2011)), and the Gini coefficient. I show in the decade following an acquisition, pollution becomes less equally distributed among target firm plants, and lower-income and less-white neighborhoods end up relatively more exposed than before the acquisition. For acquiring firms there is not a clear and pre- and post-period, but I document a similar shift in the inequality of emissions across a fixed set of plants as the size of the firm increases.

This work informs our understanding of the supply side of neighborhood-level pollution—which has been widely established to threaten long-term health and cognition, particularly

for vulnerable populations—by explaining the distributional effects of mergers in heavily polluting industries. Consolidation in polluting sectors may have local environmental benefits but further entrench environmental inequality.

3.2 Data

An observation in the data consists of a TRI facility in a given year, and includes its annual toxic emissions, its parent company, industrial sector, and location. I link individual polluting facilities and their emissions from Environmental Protection Agency records to their parent firms using the Dun & Bradstreet business database and the Your Economy Time Series establishment-level data. I then use matched parent companies to identify facilities affected by corporate mergers and acquisitions using deals collected in the Thomson (formerly SDC Platinum) M&A database. These data cover the United States for the period 2001 through 2018.

3.2.1 Toxics Release Inventory

The facilities in this paper are all mandatory-reporters to the EPA’s Toxics Release Inventory (TRI), a federal environmental reporting program that has been in existence since 1988. Established as a “right to know” program, the TRI mandates all facilities operating in particular sectors above specific chemical usage and employment levels to report the quantity, type, and pathway for all chemicals released to the environment. Reporting sectors encompass major point sources of pollutants in the U.S. economy, from mining and manufacturing, to chemical production and disposal. Facilities must report if they have 10 or more employees and handle any listed chemical above its threshold—typically in the tens of thousands of pounds, though lower for a group known as persistent bioaccumulative (PBT) chemicals. Omitted from the TRI are non-point-source polluters, such as trucks and airports, and small establishments, such as dry cleaners.

Acute or long-term exposure to the chemicals covered under mandatory TRI reporting is known or suspected to threaten human or environmental health. The list currently

includes nearly 600 distinct chemicals and chemical groups, each of which must be reported separately, and has been expanded over the history of the program. Facility-level emissions must also be disaggregated by how they are released to the environment: air releases (through a smokestack or otherwise), releases to water, releases to land, and releases off-site for treatment or disposal. These different modes of release imply very different outcomes for local human and environmental health. Air releases, particularly those not emitted through a smokestack, escape into the vicinity of a facility at ground level, while water releases contaminate nearby streams, rivers, and lakes. Off-site transport, on the other hand, shifts the pollution burden to other neighborhoods. This paper focuses on air releases, which pose an immediate risk to neighborhoods surrounding the facility, and cannot be mitigated or remediated once they leave the facility.

Another way to understand the impact of toxic emissions is through the toxicity-weighted quantity of releases. Several billions of pounds of TRI chemicals are released to the air, water, and soil each year, but not all of these chemicals pose the same risk for health (for example, a pound of airborne mercury presents a much greater risk to human health than a pound of sodium nitrite). The EPA has established toxicity weights based on the available science in order to make more meaningful comparisons of health risk from emissions, which they define as exposure that increases the long-term (not acute or contemporaneous) likelihood of carcinogenic and non-carcinogenic adverse outcomes, such as respiratory, cardiovascular, and neurological harm.

I use as outcome variables both the pounds of toxic chemicals released and the total toxicity of releases. Changes in quantity in pounds of chemicals released to the air may be more salient, but changes in the toxicity-weighted releases from TRI facilities better reflect changes in health risks to the neighborhood surrounding a plant. Annual toxicity-weighted releases are the facility-level sum of each individual chemical release quantity times its toxicity weight; this number is unitless, and only meaningful in relative terms (higher toxicity-weighted releases represent greater health risks).

This analysis begins in 2001 because many chemicals were added for TRI reporting in 2000 and the thresholds for PBTs were lowered in that year; and shortly before, in

1998, many major industry sectors were included in the TRI for the first time. This cutoff ensures changes in release quantities reported are not due to changes in reporting requirements. I omit facilities owned by the federal government and those with zero reported emissions in all years between 2001 and 2018. Industries are assigned using the first 3- or 4-digits of the NAICS code, following the EPA's classifications.

Of course, in addition to varying emissions at the intensive margin, facilities start and cease to operate throughout the study period. Closures are an important possible outcome of M&A activity, because a polluting facility shutting down entirely matters for both local health and local employment. TRI facility reports do not identify the year in which a facility officially started or stopped operating, and closed facilities are not required to file reports. I consider the year after the last TRI report with positive emissions to be a meaningful zero in the data, and include the facility in the panel for 3 more years, with zero emissions. After that point, the facility drops out of the data (results are not sensitive to longer inclusion of zero-emission years in the panel). In specifications that keep these zero-emission years, closures inform the estimated average change in ambient air quality at acquired plants. To better understand just the intensive margin operations of TRI facilities, I also consider just the subset of facilities that are definitely operating (those that report nonzero emissions, to any release media, in a particular year).

3.2.2 Linked merger, business, and Census tract data

The Thomson M&A data list transactions in the U.S., providing the name of the target company (the company that was purchased), the name of the acquiring company (also referred to as the new parent company), the industrial sectors of each firm (which allows me to identify same-sector mergers), and the date of acquisition. Since the Thomson data include privately-held companies (which are not required to report mergers to regulators), they present one of the most complete pictures of M&A activity in the U.S. available to researchers.² I omitted transactions in the Thomson data where the listed purchaser is

²This paper leaves for future study the effect of an acquisition by a large multinational on offshoring pollution. Acquisitions by multinationals with production overseas appear in the Thomson data, but I only observe domestic facilities' emissions.

not another firm (instead, the same firm, a group such as bondholders, shareholders, or employee stock ownership), transactions where the target purchase was not the entire company (instead, a particular division, franchise, location, or unit), and transactions where either the target or acquiring firm is unknown.

The TRI is a very detailed facility-level reporting program, but it is often difficult and sometimes impossible to attribute corporate ownership of a plant, or to characterize how firm ownership has changed over time. Making accurate links from individual facilities to their parent firms allows correct assignment of mergers, which are recorded at the firm level. The relevant TRI reporting fields for facilities are an open-ended “Parent company”, an EPA-cleaned “Standardized parent company” that modifies the previous field, and the parent company’s Dun & Bradstreet number, which is a proprietary, EIN-like numeric identifier used primarily for corporate credit ratings. All of these fields contain numerous blanks, errors (typos or missing digits), and inconsistencies (small shifts from year to year in facility or firm names, including alternately listing subsidiaries and their parent companies) in the EPA data. I first cleaned facility and parent company names and removed small discrepancies and missing data from year to year within facilities.³ I then matched the standardized names of the target companies listed in the Thomson data to several cleaned and standardized fields in the TRI, prioritizing the “Standardized parent company” field, then the “Parent company” and finally the “Facility name” field.

I consider a TRI facility in the treatment group if it belongs to a firm that was acquired and has exactly one matched merger. This makes for distinct separation of pre- and post-acquisition periods, but excludes facilities that matched to multiple acquisitions (about 3 percent of all facilities). In addition to this, despite efforts to accurately categorize facilities, the plant-to-parent matching was intentionally quite conservative, which means

³Cleaning the TRI data fields is best explained with several representative examples. I removed all punctuation and common phrases like “UNKNOWN” and “SEE ABOVE”. I dropped uninformative strings like “LLC”, “PLANT” or “FACILITY”, as well the names of plants’ cities and counties, from the ends of strings, so “3M HARTFORD CITY PLANT” becomes just “3M”. Large companies that are commonly abbreviated in the data were identified, so “ADM” becomes “ARCHER DANIELS MIDLAND”. I standardize multiple common abbreviations for single words to one abbreviation, so consecutive years reported as “SOUTHERN MFG CO” and “SOUTHERN MANUFACTURING CO” register as the same ownership. These and all other rules were applied to the (already much more standardized) Thomson data as well as the TRI data.

there there are inevitably facilities that should be in the treatment group but are not (in particular, those with very inaccurate or frequently-missing parent company data). Both of these facts may compromise the representativeness of the sample, but minimize incorrect assignment of treatment.

About 60 percent of the facilities in the sample have matched records in the Your Economy Time Series (YTS) data,⁴ a database from the University of Wisconsin Business Dynamics Research Consortium, comparable to the National Establishment Time-Series database. These establishment-level records, which run from 1997–2019, provide figures for employment (about half of which is verified and half of which is estimated), and sales (which are always imputed from establishment NAICS codes and employment) for matched facilities. Because the TRI facilities are incompletely matched to the YTS data, and the YTS data are themselves often imputed, estimates relying these fields should be considered motivation for further investigation.

Finally, I link each TRI facility to its Census tract (ACS 5-year estimates since 2009) and county (ACS 1-year estimates since 2005) in order to represent the characteristics of the immediate neighborhoods and larger areas surrounding the facility. I focus on median household income, the percent of residents by racial and ethnic categories, and the population density. I primarily use Census tracts as an approximation of the neighborhood surrounding a TRI facility; they are typically several thousands of people.

3.2.3 Summary statistics

Figure 3.1 shows the largest share of on-site TRI facility releases over the past two decades have gone to land (such as waste piles, landfills, underground injection, and spills of liquids), while smaller shares of all releases are water discharges to streams, rivers, and lakes, and other bodies of water, and air releases. Figure 3.2 shows aggregate trends in the number of facilities reporting air emissions, total pounds released, and the toxicity of air releases over the study period. Across most industries, there have been declines in releases since the early 2000s, but per pound toxicity of emissions has risen.

⁴YTS staff performed the matching using a combination of TRI facility names, addresses, and geographic coordinates.

Many TRI facilities in the analysis open, or are acquired, or close down during the study period. Twenty-nine percent of TRI facilities in the sample had emissions from the start to the end of the study period, while the remainder opened and/or closed (Figure 3.3). Figure 3.4 shows all years have at least some merger activity, but the number of facilities acquired in each year, and the industrial composition of those acquisitions, varies between 2001 and 2018.

Summary statistics for the facilities used in the analysis are shown in Table 3.1. Of the 33,573 facilities in the sample that ever report air emissions, a great majority are never acquired: about 14 percent (4,560) have a single matched acquisition, and another 5 percent (1,680) have multiple matched acquisitions. Before their acquisition, facilities treated by a single merger tend to release more pounds of chemicals to the air than facilities never acquired. All three groups have comparable demographic characteristics.

3.3 Empirical analysis

The first goal of estimation is to understand whether (and by how much) an acquisition of a polluting facility affects its subsequent emissions. In this empirical setting, the panel comprises nearly two decades of emissions data for individual facilities, whose timing of treatment (acquisition) varies throughout the study period, and a large potential control group (more than 80 percent of facilities do not have a matched M&A event).

In light of the evolving literature laying out the ways the two-way difference-in-differences estimators may be biased when the timing of treatment varies and treatment effects vary over time, my preferred specification is an event study that omits never-acquired facilities. I use both the typical event study (ES) design allowing for dynamic treatment effects, and also show estimates using the Sun and Abraham (2020) estimator (which they call the interaction-weighted, or IW estimator), which allows for treatment effect heterogeneity by averaging dynamic treatment effects across cohorts treated at different points in time. I also present estimates for the static two-way fixed effect difference-in-differences (TWFE) for comparison. Identification is discussed in Section

3.3.3.

The second empirical goal is to understand whether the effects of mergers differ by plants' neighborhood characteristics, and whether acquisition leads to a reallocation of pollution among facilities within a treated firm. For the firm-level analysis, I use summary statistics that reflect relative changes in emissions exposures across the demographic groups in which facilities are located and the distributional inequality of emissions across target firm plants.

3.3.1 Two way fixed effect estimates

A two way fixed effects strategy to estimate the impact of a corporate acquisition on TRI facility emissions is as follows:

$$Y_{it} = \beta M_{it} + \alpha_i + \gamma_{jt} + \sigma_{st} + \epsilon_{it}$$

where the outcome Y represents the pounds, or toxicity-weighted quantity, of on- or off-site releases reported to the EPA for each facility i and year t , and can go to zero when the facility closes down. M is equal to one if the facility's parent company was the target of an acquisition, in all the years after the acquisition. Fixed facility-specific characteristics are captured with the fixed effect α_i and nonlinear, industry-specific time trends by the sector-by-year fixed effect γ_{jt} . Since environmental policies and macroeconomic conditions may also vary across states and years, σ_{st} is a state-by-year fixed effect. Standard errors are clustered at the industry level (EPA-assigned groups from the first three or four digits of the facility NAICS code, corresponding to sectors like "wood products", "chemical wholesalers", or "hazardous waste").

The coefficient of interest, β , reflects the average change in annual emissions at a TRI establishment after the establishment is the target of an acquisition. Identifying variation comes from both the average differences in emissions across treated and control facilities, and the timing of the merger relative to other acquired facilities. Following, for example, Goodman-Bacon (2018), the static estimate β is biased if treatment effects change over

time.

3.3.2 Event study estimates

The event study model,

$$Y_{it} = \sum_{l=-L}^{-1} \mu_l D_{it}^l + \sum_{l=1}^L \mu_l D_{it}^l + \alpha_i + \gamma_{jt} + \sigma_{st} + \epsilon_{it},$$

is analogous to the TWFE above, but treatment effects are estimated using a standardized lead or lag relative to the time of acquisition. The indicator $D_{it}^l = \mathbf{1}\{t - E_i = l\}$ corresponds to unit i being l periods away from year of treatment E in year t . The fixed effects are the same as in the specification above. If the μ_l where $l < 0$ are equal to zero, the event study is consistent with the assumption that acquired firms in the TRI do not ramp emissions up or down in the years leading up to acquisition. Similarly, if the μ_l where $l > 0$ are significantly different from zero, then mergers have short-term effects on emissions. For large lags and leads of the treatment, the μ_l are identified off of relatively few observations. The year of acquisition is $l = 0$, and its coefficient is not estimated in the regression. Since M&A activity happens throughout the year, and it may take some time for the effects of an acquisition to be felt at facilities after a change in ownership, it is plausible the first year after the merger is when its effects will show up in emissions.

3.3.3 Interaction-weighted estimates

The Sun and Abraham (2020) estimator builds on the dynamic event study model by estimating a different μ_l for each cohort e , which in this setting is a group of facilities that share a year of acquisition. The interaction-weighted estimator proceeds from the regression

$$Y_{it} = \sum_{e \notin C} \sum_{l \neq 0} \delta_{el} (\mathbf{1}\{E_i = e\} D_{it}^l) + \alpha_i + \gamma_{jt} + \sigma_{st} + \epsilon_{it},$$

where δ_{el} is the cohort-specific average treatment effect, and the other parameters are defined above. To avoid issues of multicollinearity, the omitted cohorts C are the first

cohort (2001), which is always treated in the sample, and the last cohort (2018), which is never treated in the sample.

The weighted average across cohorts of period-specific average treatment effects is $\sum_{e \notin C} w_{el} \hat{\delta}_{el}$, where weights are the sample shares of each cohort in each period, and is analogous to the event study coefficient μ_l , above. The weighted average across both periods and cohorts (for some fixed number of periods after treatment g) is $\frac{1}{|g|} \sum_{l \in g} \sum_e w_{el} \hat{\delta}_{el}$, and is analogous to the TWFE β , above. Standard errors for both are straightforward to derive.

In regression tables, I present both the TWFE estimates and the IW estimate for the average treatment effect in the post-period. Comparing both gives some sense of the extent to which the TWFE estimator departs from the IW estimator, which may be consistent with the “contamination” of estimates of particular leads and lags by effects in other periods described in Sun and Abraham (2020). In event study plots, I present both the “typical” and IW event study coefficients. My preferred specifications estimate effects for target facilities within 5- or 10- years of their acquisition, since longer horizons rely on estimates from fewer cohorts and so may not be representative.

Several assumptions are required for unbiased estimates of the effect of M&A activity on facility emissions.

First, the counterfactual must be reasonable. Since most of the TRI sample is not treated by a merger, there is a large potential control group against which to compare treated plants. However, this set of never-acquired TRI facilities is an implausible counterfactual if the plants have substantively different patterns of emissions, which is likely if the decision to acquire a plant or firm is related to its output and emissions control technologies. I address this concern by omitting never-acquired and multiply-acquired facilities from the sample.

A more complicated threat to identification is if earlier- and later-acquired facilities are not good counterfactuals for one another. The IW estimator requires that the treatment is homogeneous across cohorts—that is, the post-merger trajectory is the same across facilities acquired in 2005 as it is in 2015. However, this assumption may not be tenable

if, for example, effects vary across sectors and the composition of acquisition cohorts also varies across sectors (as Figure 3.4 shows). I include covariates for each industrial sector in each year, and assume treatment effects are homogeneous conditional on those. There may also be selection or timing effects within sectors. If mergers select for the most (or least) successful firms earlier in time, it is less plausible that effects are homogeneous across earlier and later acquisition years. Spillovers within industries may also lead an earlier acquisition to affect not-yet-acquired facilities, if those facilities have to adjust operations in order to stay viable. Changing regulatory and macroeconomic conditions (including the Great Recession) also threaten the assumption that post-treatment trajectories are constant over time.

A related concern is the presence of shocks that occur at the same time as the acquisition. This could be regulatory shocks that simultaneously affect emissions and the likelihood of acquisition, or local demand shocks. Consider the case of an estimated decrease in emissions resulting from acquisition. If mergers select very successful firms just when they would have ramped up production, the magnitude of the effect of the acquisition is underestimated. Conversely, if mergers select firms right before they are about to go out of business, so their continued operation results directly from the acquisition, the contribution of the merger to observed decreases in emissions is over-estimated. In order to avoid this source of bias, Arnold (2020) focuses on M&A activity among large firms, which also may be a relevant strategy here.

Finally, it is also necessary to assume treated units do not anticipate treatment. This is reasonable if the effects of the merger are primarily through transfer of management practices and technology that cannot affect target facilities until after the deal is complete, but could be problematic if target firms manipulate their plants' emissions as part of merger negotiations, or defer maintenance or upgrades until the change in ownership. Pre-merger event study coefficients showing facility emissions (or other statistics) are indistinguishable from zero in the years leading up to an acquisition can corroborate the assumption of no anticipation.

3.3.4 Facility closures

A facility closure is a meaningful outcome for local pollution and employment. When the regressions above include these zero-emission outcomes, the estimated effect of the merger captures both the extensive (closure) and intensive (operational) margin. In regressions that omit facilities after they close down, β just captures the average effect on facility operations.

In order to understand more directly the likelihood a firm acquired in year t closes, relative to all other never-acquired TRI facilities operating in the year of the acquisition, I estimate the model

$$Y_{it} = \beta A_{it} + \gamma_{ijt} + \sigma_{ct} + \epsilon.$$

The binary variable A is equal to one if the facility was the target of a merger. I use a cross section here, comparing each cohort of TRI facilities (those sharing a particular acquisition year) to all never-treated facilities in operation in the year of the merger. I only include target facilities in their acquisition year to avoid comparing contemporaneously-acquired facilities to facilities that have been acquired or will be acquired in the future, which affects their closure likelihoods. The fixed effect γ is a dummy variable for industry sector, year, and first observed year of operation (which is the start of the panel for most of the sample), and σ for county and year.

The outcome Y is either a binary indicator for whether the firm is still open within two years of the acquisition year, or the number of years a firm will continue to operate before closing. This is a descriptive, rather than causal regression, since selection into acquisition is non-random. The estimate for β indicates whether acquired facilities tended to remain open ($\beta > 0$) for longer than facilities whose firms were never the target of an acquisition but were operating at the same time.

The intuition for the potential bias in the effect of mergers on facility closures is not obvious. If acquiring firms target struggling or inefficient facilities that would have shut down had they not been acquired, the comparison to never-acquired facilities leads to an overestimate of the effect of the merger on eventual closure. If acquiring firms select

more successful firms and the control group consists of facilities less likely to survive, this specification underestimates the effect of a merger on facility closure.

3.3.5 Distributional outcomes

Facility-level reductions in pollution may not correspond with a more equitable distribution of pollution. In order to assess whether a merger systematically shifts the distribution of high-risk pollution within acquired firms, I estimate heterogeneity in the effect by neighborhood demographics, and in the aggregate changes in emissions exposure across a target firm's facilities.

Empirically, there are multiple reasonable ways to classify an individual target facility's relative advantage, which could lead to disparate findings and interpretations, so I use two different rankings for facility demographic characteristics. In order to understand the overall effects M&A activity may be having on the distribution of air pollution from toxic facilities among the least-advantaged communities in the United States, I classify facilities across an entire industry into "High", "Medium", or "Low" terciles, based on either Census tract-level median household income and percent of non-Hispanic white residents. Then, in order to understand how an acquiring firm's management of target firms in the years after an acquisition affects the distribution of pollution across its own facilities, I classify plants into terciles within their target firm and industry, so that "high income", for example, is relative to comparable plants over which the parent company may adjust operations. There is less overlap than one might expect between a facility's classification in the first, broader classification and the second, more granular classification (for the income terciles, only 65 percent of the sample has the same sample-wide as within-firm tercile). Because the within-firm terciles require at least 3 facilities within a firm and industry over which to distinguish levels, those results necessarily omit the very smallest target firms.

For the firm level analysis, an observation consists of a treated firm in each state, industry, and year, which I consider to be the relevant unit within which firms reallocate pollution. Though it would be interesting to understand whether the acquiring firm shifts

polluting activity to facilities in the target firm when they are in less-advantaged neighborhoods, and vice versa, it is likely facilities' neighborhood characteristics are correlated with the type of activity at the facility. Instead, I focus on emission shares within the target firm's fixed set of facilities.

To do this, I use a measure of the neighborhood characteristics that reflects the distribution of emissions across sociodemographic groups in each year, which I call the exposure-weighted average. First, an example: suppose a firm consists of just Facility H, located in a high-income neighborhood, and Facility L, located in a low-income neighborhood. They each have some reported emissions in the first period. If in the second period emissions rise at Facility L and fall at Facility H, then the firm's exposure-weighted average income will be lower in the second period; this can be interpreted as a shift towards a greater share of exposure in the lower-income neighborhood. It should not be interpreted as evidence the typical income of households in either neighborhood changed in response to the pollution.

In general, the exposure-weighted average S is:

$$S_{kfst} = \frac{\sum_{i \in kjs} X_i * d_i * E_{it}}{\sum_{i \in kjs} d_i * E_{it}}$$

where X is a tract-level demographic variable (for example, median household income), d is the tract population density, and E is emissions at facility i in a particular state, industry, and year belonging to firm k . Each facility's neighborhood demographics and population density are held at those of the 2010 Census Tracts in which the facility is located, and the set of facilities in each firm, industry, and state are constant—so changes in the exposure-weighted average demographics are only driven by relative changes in emissions across those facilities. If after a merger parent companies shift pollution to poorer neighborhoods, the firm's overall exposure-weighted household income will look less affluent.

I test this with the following, by now familiar, TWFE regression (shown here), event

study regression, and interaction-weighted estimator (which follow analogously as above):

$$S_{kfst} = \beta M_{kfst} + \alpha_k j s + \gamma_{jt} + \sigma_{st} + \epsilon_{kfst}$$

for each firm k comprised of facilities in state s treated by an acquisition. The outcome S is the exposure-weighted demographic descriptor defined above, and α , γ , and σ are firm-industry-state, industry-year and state-year fixed effects, respectively. Larger firms have many more facilities informing the aggregate outcome S , and so in these regressions I weight by the number of facilities. As with all previous regressions, the event study and interaction-weighted designs present an opportunity to confirm, at this aggregate level, the absence of pre-trends.

I also use this design to estimate whether the distribution of emissions within each firm, industry, and state becomes more unequal after an acquisition, using a Gini coefficient (this measure of inequality in emissions has precedent in Millimet and Slottje (2002)). An increase in the Gini coefficient indicates fewer plants within a fixed set of facilities are responsible for a greater share of the total emissions, but does not on its own reveal which plants experienced the relative increase and which the relative decrease. It is possible for the Gini coefficient to rise even as total emissions fall.

Finally, I use the exposure-weighted average and Gini coefficient as an outcome variable in a descriptive regression, relating this distributional measure to changes in the size of acquiring firms. Following the intuition above, I expect as acquiring firms grow, emissions also shift away from plants in more-advantaged neighborhoods. I estimate this with an analogous exposure-weighted average for acquiring firms, again averaging across a fixed set of facilities with a single parent company (facilities not ever the target of an acquisition). I replace the treatment M with the number of facilities held by the parent company in a given year. An increase in the number of held facilities is a proxy for an acquisition, but by no means a perfect one, since firms may add plants for other reasons. As above, changes in the exposure-weighted demographics reflect relative changes in emissions.

3.4 Results

In order to understand how a change in the parent company of a TRI facility affects the pollution it releases to the surrounding neighborhood, I first focus on the facility-specific effects of an acquisition, estimating whether and how much toxic air emissions change at individual plants. I then conduct a firm-level analysis of the distribution of emissions at target and acquiring firms following an acquisition.

I report the inverse-weighted (Sun and Abraham) estimates throughout this section, which average across cohorts and event years as explained above, for the specified years post-treatment. I summarize immediate and longer-term effects with 5- and 10-year windows. In all regression tables, I present estimates from both the TWFE and IW designs, and in the plots I present the event study and the IW estimates; results are often but not always similar, and not consistently different in a way that suggests the source of bias introduced with the TWFE estimator in this setting.

3.4.1 Target facility effects

A change in ownership results in substantially lower emissions at the facility level, in both the pounds of toxic releases and when accounting for the toxicity of releases (Table 3.2). Examining both the total pounds of emissions and the toxicity of the waste stream helps understand both mechanisms—M&A activity may lead to reductions in overall activity, or adjustments to facility operations—and also consequences for health risk in surrounding neighborhoods.

When the sample of treated facilities includes the extensive margin (facilities that close in the years immediately after a merger), an acquisition causes the pounds of plant-level air emissions to decline by about 52 percent in 5 years following, and 62 percent in the 10 years following; toxicity-weighted emissions fall by even more, about 80 percent in the 10 years after the change in ownership. However, much of the effect is driven by facilities closing down entirely (further explored below). Among facilities that continue to operate in the subsequent years, the decline is much smaller in magnitude: about 24 percent

fewer pounds of chemicals are released to the air in the decade following an acquisition at a TRI facility. The total toxicity of that emissions stream falls by an average of 27 percentage points over the decade.

TWFE and IW event study plots both show the decrease is not a one-time shift to a lower level of emissions, but a steady decline over the post-merger period (Figures 3.6 and 3.7). These plots all corroborate the parallel trends assumption, with estimated effects in the years before the acquisition insignificantly different from zero. Among operating facilities, the IW coefficients suggest even as the quantity of pollution falls, the total toxicity decreases in the years following acquisition are not significantly different from zero, suggesting plant-level shifts in the composition of the waste stream.

TRI facility emissions are self-reported, and so it is possible post-acquisition reductions are not representative of real changes in facility operations, but instead changes in the reporting practices or strategy. If acquisitions increase under-reporting of emissions at acquired facilities, the estimated effect of a merger on actual emissions may be biased up in its magnitude; on the other hand, if larger parent companies are less likely to tolerate downward-biased misreporting, the estimated effect of a merger on emissions could be biased downwards (pre-merger emissions reports being too low). I cannot rule out either case, but can evaluate whether the effects presented in Table 3.2 are independently detectable by linking TRI facilities' locations to gridded, annual, satellite-derived estimates of near-surface fine particulates (PM_{2.5}).

Using this outcome variable risks missing the signal for the noise, since many non-TRI sources of pollution contribute to near-surface fine particulates (such as traffic, airports, small facilities, and non-point emitters); the releases from an individual TRI facility may comprise a very small share of the total PM at a given location. Moreover, not all TRI facility releases are particulates, and so are not detected by satellite PM measures (nor PM monitors). Both of these caveats should introduce measurement error that biases towards zero an estimated effect of an acquisition on emissions, and in fact, I fail to find an effect on satellite PM_{2.5} among all facilities (Panels A1 and B1 of Table 3.3).

However, among the top quarter of facilities by release quantities—that is, those

most likely to move the needle of annual, all-source PM_{2.5}—an acquisition shows up as an annual $0.10\mu\text{g}/\text{m}^3$ decrease in local ambient PM_{2.5} (Panels A2 and B2) on average over 10 years.⁵ Panels A3 and B3 further subset facilities, restricting to both the highest-quarter of air emissions and the lowest-quarter of ambient PM; intuitively, in low-overall PM areas, few other sources contribute to ambient PM, making the TRI facility a large share of the detected amount. The imposed restrictions mean a small sample size, but estimates are yet larger ($-0.23\mu\text{g}/\text{m}^3$ on average post-merger). These findings corroborate Table 3.3 that post-merger reported decreases reflect actual changes in air emissions.

I expect a merger to have different effects when the acquiring and target firms operate in the same or different industrial sectors, and this is the case. “Same sector” mergers are not necessarily horizontal, but they identify the set of acquisitions where the new parent company is much more likely to have preexisting experience and capacity; “different sector” mergers are the remainder of M&A activity, where the target and acquiring firms operate in different primary industries. Both types of acquisitions lead to plant closures, but among the set of facilities that continue operating, decreases in the quantity of air emissions are driven almost entirely by same-sector acquisitions (column (3) of Table 3.4). This heterogeneity in outcomes by acquisition type is consistent with several, not mutually exclusive, explanations: acquiring firms already operating in the same sector as the target firm may be more likely to have the technical expertise to improve facility efficiency, and also to be acting anti-competitively and reducing output.

The seeming improvement in local air pollution at operating facilities, particularly for same-sector merger types, does not actually mean a decrease in the health burden on the surrounding neighborhood. While the pounds of point-source pollution falls, the per-unit toxicity, or health risk, of the pollution that is released increases by an average of 36 percent over the next five years (column (3) of Table 3.5). This compositional shift may entirely offset the health benefits from the decrease in quantity: the toxicity-weighted total amount of air pollution hardly falls at operating plants after the acquisition (column (6) of Table 3.4).

⁵For reference, the EPA annual PM_{2.5} standard is $12\mu\text{g}/\text{m}^3$.

Changes in sales and employment also provide insight to operational changes at acquired plants. I find facilities that stay open after a merger increase employment slightly (by about 6 percentage points) and sales do not change (columns (5) and (6) of Table 3.6). Sales would typically be a better proxy for output than labor, but sales are imputed from employment in the YTS data, and no other inputs to production are available. I use both to understand the emissions intensity of production at plants, and estimate whether these ratios change in the post-merger years. I find some evidence emissions intensity falls (columns (7)–(10) of Table 3.6), which is consistent with acquiring parents improving emissions control technology. However, sales and employment data come from a separate source, and are incompletely matched—less than half of the TRI facilities in the sample have a corresponding observation—and estimates for the main specification look quite different among this subset. Columns (2) and (4) of Table 3.6 show facilities with linked employment and sales data decrease emissions by much less than the typical facility in the study. Future research should attempt to link TRI facilities to a more complete survey of establishment-level operational data.

Facilities report more than just emissions to the air, and mergers may also affect those releases. Post-acquisition decreases in output observed across all release media might be more indicative of decreases in output, while medium-specific decreases could be more consistent with operational changes. Table 3.7 shows only air emissions fall after acquisition, while emissions to water and land, and transport off-site, may not change at all. Air emissions have so far been presented as the combination of releases through two channels that the EPA calls “stack” and “fugitive”. Stack emissions leave through controlled outlets, such as smokestacks or chimneys, while fugitive emissions are unintended or unaccounted-for releases that do not pass through an exhaust system (both groups can be particulate or gaseous air emissions). Post-merger fugitive emissions fall by about twice as much (in percentage terms) as stack emissions, which may indicate new parent companies impose better practices to reduce instances of loss and leaks.

I directly estimate the difference in lifespans of facilities that are acquired, compared to those owned by a firm never involved in either side of M&A activity, operating in

the same county and industrial sector, and with the same first operating year in the panel. Target facilities are 2–3 percent more likely to close in the two years following the acquisition, and tend to have shorter lifespans (years from the acquisition year to closure) by about 2 months (Table 3.8).

3.4.2 Distributional effects

It is plausible a change in ownership, which appears to dramatically reduce toxic air pollution at individual facilities, correspondingly lessens broader spatial disparities in exposure to air pollution. This could be the case if plants with the greatest pre-merger inefficiencies in emissions control technology were disproportionately located in disadvantaged communities; widespread descriptive documentation of disproportionate TRI exposure (for example, Brooks and Sethi (1997)) suggests this may be the case. On the other hand, if acquiring firms have discretion over which target facilities they close down and where they prioritize technological and operational changes, air quality improvements may disproportionately accrue to more advantaged neighborhoods; this could happen if firm decisions are in part motivated by community pressure, as Wang et al. (2018) suggest.

To estimate whether the magnitude of the effect of an acquisition at a target facility varies systematically by facility neighborhood characteristics, as described in Section 3.3.5, I classify facilities into terciles based on two demographic metrics (income and racial composition) in two ways: across industries for the entire sample (reflecting overall differences in facility neighborhoods), and within target firms, states, industries (reflecting relative differences within firms' choice sets). In order to assess within-firm heterogeneity in facility neighborhoods, I necessarily focus on target firms that have at least three facilities.

The two classifications produce very similar results for facility income ranks, and suggest the discretionary explanation dominates the efficiency explanation. Facilities whose tracts are in the highest-third of median household income—both sample-wide and within-firm, and particularly for the case of same-sector mergers—have greater decreases in emissions exposure (Table 3.9). There is slightly less robust evidence the

racial composition of facilities' neighborhoods is associated with the degree of emissions reductions.

With the acquired firms as the unit of observation, I then estimate the effect of an acquisition on the emissions-weighted average neighborhood demographics, a statistic that reflects changes in the share of emissions across demographic groups in facilities' neighborhoods, and the Gini coefficient for the inequality in emissions across plants. These findings generally corroborate the facility-level heterogeneity, and allow for event-study estimates. In the decade after acquisition, emissions shift towards facilities in lower-income and less white neighborhoods (Table 3.10 and Figure 3.9), so that the exposure-weighted average neighborhood looks 4 percent poorer and 20 percentage points less white. Because the exposure-weighted average is calculated for a fixed set of facilities (those held by the target firm) with fixed neighborhood characteristics (2010 Census), this can only be explained by relatively greater post-merger decreases in air emissions in higher-income and whiter neighborhoods. (This also underestimates the dynamics of the disparity if neighborhood characteristics are responsive to changes in emissions, and relatively greater decreases in emissions occur in already-higher-income neighborhoods.) The distribution of both the quantity and toxicity of emissions within the target firm also becomes more unequal: the Gini coefficient for the distribution of both the pounds and toxicity of emissions increases by 0.06 and 0.07 units, respectively.

These are not causal estimates, since differences in the types of firms located in neighborhoods with different demographic characteristics may underlie the observed post-merger differences in pollution control. Still, the findings are consistent with acquiring firms having some discretion over where emissions efficiency is improved and where closures happen, and indicates mergers increase exposure disparities across neighborhoods. These findings suggest the greater trend towards consolidation in highly-polluting industries has not lessened environmental justice correlations, and that parent companies—either intentionally or inadvertently—tend to exacerbate exposure disparities by improving operational efficiency at, or closing entirely, plants in wealthier, whiter, and denser neighborhoods.

3.4.3 Acquiring facility effects

After finding emissions decline substantially at target plants after a change in ownership, it is natural to ask whether there are spillover effects among facilities already owned by the parent company. However, most acquiring firms make multiple acquisitions, so there is no clear pre- and post-merger period over which to compare emissions. I estimate the effect of an acquisition on already-held facilities among the small subset of firms that make a single acquisition over the study period, and find mixed evidence the total toxicity of emissions rises at operating plants (Table 3.11), which could mean some decreases in emissions at the target firm's facilities are offset with increases at the acquiring firm's facilities. The absence of strong evidence M&A activity affects the parent firm may mean technology transfer primarily happens from acquiring to target facilities. These single-acquisition firms are clearly a special case that omits large- or fast-growing firms—only 34 percent of acquiring firms make a single acquisition over the study period—so these estimates should not be considered representative of all M&A activity.

A less-precise way to understand how changes in firm size affect the distribution of pollution at acquiring firms uses changes in the number of facilities associated with acquiring firms. Since acquisitions are not the only reasons firms grow, this is an imperfect proxy for acquisitions. As above, I hold fixed the set of facilities over which the exposure-weighted average and the Gini coefficient are calculated, so only relative changes in emissions among an acquiring firm's plants operating in the same industry and state lead to changes in the aggregate statistics. I find as firms grow, average emissions at the original set of plants do not increase or decrease, but they become more unequally distributed across the acquiring firm's plants, meaning some facilities generate a greater share of emissions while others generate less. The quantity of emissions also becomes more concentrated at the plants in less-white neighborhoods as the size of acquiring firm grows.

3.5 Conclusion

This paper documents the effects of a change in ownership at polluting facilities on both plant-level air emissions of toxic chemicals and the firm-level distribution of toxic pollution across plants, using two decades of M&A activity paired with reporting data from large, point-source polluters in the United States. Toxic air emissions have been widely established to threaten public health, and are disproportionately imposed on underserved communities. This empirical evidence suggests concentration does not necessarily lead to better environmental performance, nor a more equitable distribution of harmful environmental exposures, in polluting sectors.

Many facilities close down entirely after acquisition, which implies local loss of employment, may signal more-concentrated markets in affected sectors, and unambiguously improves local air quality by removing a large source of toxic pollutants.

Among facilities that continue to operate, a change in ownership at a highly polluting facility is less obviously beneficial for local air quality. The quantity of annual air emissions only falls when the target and acquiring firms operate in the same sector—highlighting the importance of industry-specific expertise for operational changes at plants—and the emission stream increases in toxicity per pound. There is substantial evidence operational changes at target facilities contribute to observed decreases in air emissions.

Mergers also have troubling consequences for the distribution of toxic environmental exposures across facilities, which is consistent with the acquiring firm having discretion over where to cease or modify operations. After an acquisition, emissions become more unequally distributed across facilities in the target firm, and fall by more at facilities in higher-income and more-white neighborhoods.

This mixed story suggests analyses of mergers should carefully consider both the potential for changes in facility-level environmental performance, and also firm-level shifts in the burden of exposure to toxic air pollution across more- and less-advantaged neighborhoods.

3.6 Figures and tables





Figure 3.1: TRI on-site releases by medium


Overall, on-site TRI releases declined steadily through the first half of the study period, and flattened somewhat after 2010. Land releases consistently comprise the greatest share of on-site emissions (compared with emissions to air and water).

Figure 3.2: Toxics Release Inventory facilities and air emissions over time



Trends in TRI facilities and their air emissions are shown relative to 2001 levels. The number of facilities reporting emissions, as well as the total quantity of those emissions, has declined since the early 2000s. Emissions per facility have also declined, but releases have gotten more harmful to human health. These trends all hold for the subset of facilities in continuous operation during the study period (not shown) and so are not being driven by entry of less-polluting firms and exit of more-polluting firms, suggesting instead changes in facility operation.

Figure 3.3: Counts of TRI facilities by first and final reporting years



Most (63 percent) of the facilities in the sample reported emissions in 2001 (the start of the study period) and nearly half reported emissions in 2018 (the end of the study period). Twenty-nine percent of facilities are in both groups. The remainder started or ended operations during the study period.





Figure 3.4: Counts of TRI facilities acquired in each year, 2001–2018

TRI target facilities with a single matched merger were acquired throughout the study period, but the composition of M&A activity by industrial sector varies from year to year.

Figure 3.5: Air emissions and emission-weighted tract demographics, all TRI facilities



As the total amount of air emissions has fallen over the past two decades, the burden of exposure has also shifted (most recently towards lower-income and less-white neighborhoods, although this was not always the trend). Inequality in the distribution of emissions across TRI facilities remained incredibly high throughout the study period, indicating very few plants are responsible for most of the emissions. For reference, the 2010 U.S. population was about 64 percent non-Hispanic White, 12 percent Black, and 16 percent Hispanic, and the median household income was about 49,500 (in the calculation of the exposure-weighted demographics, neighborhood characteristics are held fixed at 2010 levels).

Table 3.1: Summary statistics

	Never merged (1)	Merged once		Mult. merges (4)
		Pre (2)	Post (3)	
Facilities	27,333	4,560	4,560	1,680
Years open	9.96 (6.71)	14.3 (4.88)	14.3 (4.88)	13.3 (5.8)
Lbs. released to air	52600 (386000)	81300 (490000)	57300 (285000)	125000 (665000)
Lbs. released to water	11700 (225000)	9770 (106000)	12500 (124000)	17100 (176000)
Lbs. on site total	213000 (6740000)	165000 (1210000)	168000 (1760000)	221000 (1030000)
log(Tox. released to air)	20.8 (24.8)	21.6 (25.3)	20.7 (23.1)	20.7 (24)
log(Tox. released to water)	17.5 (21.3)	17.3 (20.4)	17 (20.2)	18.3 (22.9)
log(Tox. on site total)	25.1 (21.3)	27 (20.4)	26.8 (20.2)	22.9 (22.9)
Tract med. hh. inc.	48200 (19900)	47600 (19300)	48600 (20400)	50100 (20800)
Tract pct. White	0.68 (0.29)	0.69 (0.28)	0.69 (0.28)	0.68 (0.29)
log(Tract density)	14.3 (1.76)	14.2 (1.74)	14.1 (1.77)	14.2 (1.74)
Employees*	170 (347)	209 (416)	203 (451)	247 (404)
Sales*	73600 (255000)	96900 (685000)	102000 (365000)	127000 (361000)

*These measures of facility-level sales and employment are not collected in the TRI, but linked from the Your Economy Time Series, and only available for about 60 percent of the facilities in the sample.

Facilities are grouped as follows: (1) never identified as being the target of an acquisition; (2) acquired once, in the years before acquisition; (3) acquired once, in the years after acquisition; and (4) matched to multiple acquisitions during the panel. Summary statistics are calculated across facilities that are open in a given year.

Table 3.2: Average merger effects on TRI releases


A: TWFE Estimates				
	lhs(Lbs air emissions)		lhs(Tox air emissions)	
A1: All facilities	(1)	(2)	(3)	(4)
Post-merger	−0.482*** (0.048)	−0.418*** (0.047)	−0.972*** (0.136)	−0.797*** (0.095)
Observations	65,065	45,900	65,065	45,900
A2: Operating facilities	(1)	(2)	(3)	(4)
Post-merger	−0.081 (0.049)	−0.108*** (0.037)	−0.242** (0.097)	−0.212*** (0.067)
Observations	52,642	37,122	52,642	37,122
Merged within ... years	10	5	10	5
B: IW Estimates				
	lhs(Lbs air emissions)		lhs(Tox air emissions)	
B1: All facilities	(1)	(2)	(3)	(4)
Post-merger	−0.887*** (0.116)	−0.760*** (0.115)	−1.477*** (0.258)	−1.161*** (0.251)
Observations	65,065	45,900	65,065	45,900
B2: Operating facilities	(1)	(2)	(3)	(4)
Post-merger	−0.272*** (0.083)	−0.203*** (0.074)	−0.320** (0.146)	−0.107 (0.092)
Observations	52,642	37,122	52,642	37,122
Merged within ... years	10	5	10	5

Note:


*p<0.1; **p<0.05; ***p<0.01

The interaction-weighted estimator of Sun and Abraham (2020) generally yields larger-magnitude estimates of the effect of an acquisition on emissions at target facilities. For both estimators and both outcomes, effects are larger over the 10-year horizon than over 5 years. Effects are much smaller for the sample restricted to facilities that continue to operate, implying many acquired facilities close down after the acquisition (Table 3.8 explores this further). Regressions in this table, and all future tables unless otherwise specified, include state-by-year, industry-by-year, and facility fixed effects. Standard errors are clustered at the industry level.


Figure 3.6: TWFE event studies - target facility effects




(a) Pounds released, all facilities



(b) Toxicity released, all facilities




(c) Pounds released, operating facilities




(d) Toxicity released, operating facilities

Figure 3.7: IW event studies - target facility effects at all facilities




(a) Pounds released, all facilities




(b) Toxicity released, all facilities

The multicolored point estimates correspond to cohort-specific average treatment effects in each event year ($\hat{\delta}_{el}$), and show there is substantial heterogeneity in treatment magnitudes across acquisition years. The black point estimates are the weighted average across cohorts, the IW equivalent of the coefficients in Figure 3.6.

Figure 3.8: IW event studies - target facility effects at operating facilities



(a) Pounds released, operating facilities



(b) Toxicity released, operating facilities

Among facilities that continue to operate, the emissions-reducing effects of the merger are less clear using the IW estimator—particularly for the total toxicity of air emissions, for which event-study coefficients are often insignificantly different from zero in the post-merger decade.

Table 3.3: Annual satellite-derived PM2.5 at operating TRI facilities

A: TWFE Estimates				
	Satellite PM2.5			
	All facilities		Operating facilities	
	(1)	(2)	(3)	(4)
A1: All facilities				
Post-merger	−0.016*	0.009	−0.020**	−0.007
	(0.008)	(0.022)	(0.010)	(0.008)
Observations	59,326	42,465	48,031	34,346
A2: Only highest emitters	(1)	(2)	(3)	(4)
Post-merger	−0.052***	−0.034*	−0.062***	−0.036*
	(0.019)	(0.018)	(0.021)	(0.019)
Observations	17,452	12,245	15,770	11,079
A3: Only highest emitters & lowest PM	(1)	(2)	(3)	(4)
Post-merger	−0.110**	−0.152**	−0.139***	−0.182***
	(0.045)	(0.065)	(0.043)	(0.063)
Observations	3,963	2,781	3,650	2,580
Merged within ... years	10	5	10	5
B: IW Estimates				
	Satellite PM2.5			
	All facilities		Operating facilities	
	(1)	(2)	(3)	(4)
B1: All facilities				
Post-merger	−0.008	0.010	−0.033	−0.013
	(0.022)	(0.018)	(0.030)	(0.019)
Observations	59,326	42,465	48,031	34,346
B2: Only highest emitters	(1)	(2)	(3)	(4)
Post-merger	−0.103**	−0.054	−0.104*	−0.067*
	(0.051)	(0.036)	(0.053)	(0.040)
Observations	17,452	12,245	15,770	11,079
B3: Only highest emitters & lowest PM	(1)	(2)	(3)	(4)
Post-merger	−0.231**	−0.212**	−0.225*	−0.269***
	(0.112)	(0.093)	(0.125)	(0.103)
Observations	3,963	2,781	3,650	2,580
Merged within ... years	10	5	10	5

Note:

*p<0.1; **p<0.05; ***p<0.01

Large decreases in emissions could show up in annual estimates of ground-level fine particulates, and this is one way to check that reported decreases actually occurred. There is no detected effect among all facilities (Panels A1 and B1), but among subsets where decreases are most likely to be detected (Panels A2, A3, B2, and B3), the magnitudes of estimates corroborate that decreases occur. The middle panel restricts the analysis to facilities in the top 25 percent by air emissions, and the bottom uses facilities in both the top 25 percent by air emissions and also the bottom 25 percent of facilities by typical PM2.5 levels.

Table 3.4: TRI releases to air by acquisition type (different- versus same- sector)

A: TWFE Estimates

	lhs(Lbs air emissions)			lhs(Tox air emissions)		
A1: All facilities	(1)	(2)	(3)	(4)	(5)	(6)
Post-merger	−0.418*** (0.047)	−0.387*** (0.102)	−0.449*** (0.053)	−0.797*** (0.095)	−0.785*** (0.130)	−0.828*** (0.126)
Observations	45,900	19,353	26,547	45,900	19,353	26,547
A2: Operating facilities	(1)	(2)	(3)	(4)	(5)	(6)
Post-merger	−0.108*** (0.037)	−0.005 (0.073)	−0.171*** (0.059)	−0.212*** (0.067)	−0.056 (0.089)	−0.332** (0.145)
Observations	37,122	15,628	21,494	37,122	15,628	21,494
Sector	All	Different	Same	All	Different	Same
Merged within ... years	5	5	5	5	5	5

B: IW Estimates

	lhs(Lbs air emissions)			lhs(Tox air emissions)		
B1: All facilities	(1)	(2)	(3)	(4)	(5)	(6)
Post-merger	−0.760*** (0.115)	−0.881*** (0.233)	−0.703*** (0.104)	−1.161*** (0.251)	−1.758*** (0.356)	−0.926*** (0.252)
Observations	45,900	19,353	26,547	45,900	19,353	26,547
B2: Operating facilities	(1)	(2)	(3)	(4)	(5)	(6)
Post-merger	−0.203*** (0.074)	−0.103 (0.126)	−0.238** (0.096)	−0.107 (0.092)	−0.285 (0.240)	−0.067 (0.141)
Observations	37,122	15,628	21,494	37,122	15,628	21,494
Sector	All	Different	Same	All	Different	Same
Merged within ... years	5	5	5	5	5	5

Note:

*p<0.1; **p<0.05; ***p<0.01

For the sample with all target facilities, including those that close, reductions in air emissions after a merger are substantial across all types of mergers. Among target facilities that continue to operate, the observed reduction in facility-level emissions is driven by mergers in which the acquiring and target firms operate in the same sector. This suggests both types of mergers result in facility closures, while only same-sector mergers result in intensive-margin decreases in emissions.

Table 3.5: Emission toxicity per unit

A: TWFE Estimates

	ihs(Tox per pound of emissions)		
Operating facilities	(1)	(2)	(3)
Post-merger	-0.030 (0.035)	0.028 (0.060)	-0.062 (0.064)
Merger sectors	All	Different	Same
Merged within ... years	5	5	5
Observations	33,481	13,895	19,586

B: IW Estimates

	ihs(Tox per pound of emissions)		
Operating facilities	(1)	(2)	(3)
Post-merger	0.238*** (0.091)	-0.034 (0.179)	0.326*** (0.111)
Merger sectors	All	Different	Same
Merged within ... years	5	5	5
Observations	33,481	13,895	19,586

Note: *p<0.1; **p<0.05; ***p<0.01

The toxicity of each unit of emissions rises for mergers in the same sector, suggesting facilities substitute to more harmful chemicals (or reduce emissions of less-harmful chemicals by more).

Table 3.6: Employment, sales, and emissions intensity

A: TWFE Estimates

	ihs(Pounds)		ihs(Toxicity)		ihs(Emp.)	ihs(Sales)	Lbs. ($\frac{\text{Lbs.}}{\text{Emp.}}$)	ihs($\frac{\text{Lbs.}}{\text{Sales}}$)	Tox. ($\frac{\text{Tox.}}{\text{Emp.}}$)	ihs($\frac{\text{Tox.}}{\text{Sales}}$)
Operating facilities	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Post-merger	-0.108*** (0.037)	-0.031 (0.047)	-0.212*** (0.067)	-0.247** (0.118)	0.032*** (0.010)	0.001 (0.021)	-0.039 (0.027)	-0.007 (0.012)	-0.191** (0.081)	-0.125** (0.061)
Observations	37,122	17,099	37,122	17,099	17,099	15,966	17,099	15,966	17,099	15,966
Only YTS-matched facilities	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes

B: IW Estimates

	ihs(Pounds)		ihs(Toxicity)		ihs(Emp.)	ihs(Sales)	Lbs. ihs($\frac{\text{Lbs.}}{\text{Emp.}}$)	Lbs. ihs($\frac{\text{Lbs.}}{\text{Sales}}$)	Tox. ihs($\frac{\text{Tox.}}{\text{Emp.}}$)	Tox. ihs($\frac{\text{Tox.}}{\text{Sales}}$)
Operating facilities	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Post-merger	-0.203*** (0.074)	-0.096 (0.065)	-0.107 (0.092)	0.040 (0.086)	0.080*** (0.019)	0.015 (0.030)	-0.130** (0.056)	-0.027 (0.026)	-0.033 (0.077)	-0.062 (0.100)
Observations	37,122	17,099	37,122	17,099	17,099	15,966	17,099	15,966	17,099	15,966
Only YTS-matched facilities	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note:

*p<0.1; **p<0.05; ***p<0.01

There is some evidence the number of employees at target facilities that continue to operate rises following an acquisition (which is to say, the average acquisition does not lead to substantial layoffs if the facility remains open) and no evidence sales change. Using employment and sales as a proxy for output, there is mixed evidence the emissions intensity of the facility falls. An important caveat, however, is that many TRI facilities in the sample were not successfully matched in the YTS data, so it is possible columns (5)–(10) reflect a non-representative sample of all TRI facilities; differences in the magnitude of the YTS-only estimates (columns (1) vs. (2) and (3) vs. (4)) indicate this may be likely. Sales are always imputed in the YTS data.

Table 3.7: Releases by medium

A: TWFE Estimates

	ihs(Pounds) <i>released to:</i>					
	Stack	Fugitive	Water	Land	Off site	Total
Operating facilities	(1)	(2)	(3)	(4)	(5)	(6)
Post-merger	-0.048 (0.047)	-0.067** (0.031)	0.011 (0.016)	0.026 (0.022)	0.023 (0.064)	-0.059 (0.039)
Merged within ... years	5	5	5	5	5	5
Observations	37,122	37,122	37,122	37,122	37,122	37,122

B: IW Estimates

	ihs(Pounds) <i>released to:</i>					
	Stack	Fugitive	Water	Land	Off site	Total
Operating facilities	(1)	(2)	(3)	(4)	(5)	(6)
Post-merger	-0.129* (0.074)	-0.218*** (0.067)	-0.005 (0.026)	0.006 (0.034)	-0.121 (0.086)	-0.227** (0.105)
Merged within ... years	5	5	5	5	5	5
Observations	37,122	37,122	37,122	37,122	37,122	37,122

Note:

*p<0.1; **p<0.05; ***p<0.01

Air emissions are reported in two categories: “stack” emissions leave through a smokestack, chimney, or vent, while “fugitive” emissions escape through leaks, loss, accidents, and other channels. Among operating facilities, the overall effect on air emissions is more driven by reductions in fugitive emissions, which tend to be closer to the ground (and so potentially a greater threat to neighborhood health) and less controlled. There is not evidence air emissions are being substituted to other release media (water and land) after a merger, but emissions to these media do not decline, either.

Table 3.8: Facility closures

	Still open in 2 years (0/1)				Years until closure			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Acquired	-0.017*** (0.005)	-0.012** (0.006)	-0.027*** (0.009)	-0.023** (0.009)	-0.151** (0.074)	-0.010 (0.079)	-0.197* (0.113)	-0.061 (0.120)
... × same sector			0.017 (0.010)	0.020* (0.010)			0.082 (0.131)	0.091 (0.131)
Omit acq. firms	No	Yes	No	Yes	No	Yes	No	Yes
Observations	248,347	180,179	248,347	180,179	248,347	180,179	248,347	180,179

Note:

*p<0.1; **p<0.05; ***p<0.01

To (descriptively) understand whether target facilities tend to stay open for more or fewer years than non-acquired facilities, I compare target facilities in their year of acquisition to facilities that were also open in the same year of the acquisition, and whose parent firms were never acquired. Alternating columns omit from the comparison group facilities held by firms that acquired other TRI facilities. The coefficient on “acquired” indicates targeted facilities are 1–2 percentage points more likely to close within the next two years, but don’t have shorter subsequent lifespans after that initial window. The likelihood of immediate closure may be less when the acquiring and target firms operate in the same sector. Facilities that were the target of multiple acquisitions were dropped from all specifications. Acquired and never-acquired facilities may be substantively different in their propensity to close down, and this selection effect is discussed in the main text. Regressions in this table use industry-by-year-by-first panel year and county-by-year fixed effects, and results are robust to alternative choices.

Table 3.9: Effect heterogeneity by facility neighborhood demographics, TWFE estimates

All facilities	lhs(Pounds air emissions)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-merger, Low income	−0.199 (0.156)	−0.323 (0.208)	−0.037 (0.185)	−0.185 (0.284)				
Post-merger, Med. income	0.159 (0.295)	−0.208 (0.334)	−0.121 (0.246)	−0.325 (0.268)				
Post-merger, High income	−0.280** (0.124)	−0.439* (0.229)	−0.171 (0.165)	−0.453* (0.258)				
Post-merger, Low pct. White					−0.071 (0.166)	−0.382 (0.258)	−0.082 (0.131)	−0.270 (0.211)
Post-merger, Med. pct. White					−0.061 (0.141)	−0.244 (0.213)	−0.220 (0.275)	−0.304 (0.418)
Post-merger, High pct. White					−0.223 (0.248)	−0.350 (0.338)	−0.058 (0.138)	−0.394* (0.189)
All facilities	lhs(Toxicity of air emissions)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-merger, Low income	−0.560 (0.362)	−0.983 (0.637)	−0.349 (0.286)	−0.844 (0.677)				
Post-merger, Med. income	−0.643 (0.400)	−1.316* (0.702)	−0.804* (0.438)	−1.218* (0.623)				
Post-merger, High income	−0.729** (0.305)	−1.276* (0.622)	−0.793** (0.361)	−1.505** (0.655)				
Post-merger, Low pct. White					−0.382 (0.314)	−1.162 (0.669)	−0.573* (0.275)	−1.122* (0.581)
Post-merger, Med. pct. White					−0.745** (0.324)	−1.384** (0.570)	−0.607 (0.402)	−0.770 (0.724)
Post-merger, High pct. White					−0.885* (0.506)	−1.046 (0.804)	−0.741* (0.375)	−1.591** (0.653)
Merger sectors	All	Same	All	Same	All	Same	All	Same
Terciles within ...	All	All	Firm	Firm	All	All	Firm	Firm
Acquisition within ... years	10	10	10	10	10	10	10	10
Observations	10,679	6,332	10,679	6,332	10,679	6,332	10,679	6,332

Note:

*p<0.1; **p<0.05; ***p<0.01

The pollution-reducing effects of a merger vary significantly by the demographics of the neighborhood in which a target facility is located—whether or not the relative advantage is industry-wide or within a target firm. Across TRI-reporting industries, facilities in the lowest-third of the income distribution see the smallest decreases in air emissions. Associations between the racial composition of facility neighborhoods and post-merger reductions in pollution are less robust than for income. Interacted coefficients are presented in order to make clear when a group (tercile)’s coefficient is different from zero with significance.

Table 3.10: The pollution burden shifts at target firms after an acquisition, TWFE and IW estimates


	lhs(Lbs per fac.)	Pounds-weighted			Gini (lbs.)
		lhs(HH. income)	Pct. white	Pct. black	
TWFE	(1)	(2)	(3)	(4)	(5)
Post-merger	−0.495*** (0.178)	−0.020** (0.009)	−0.019*** (0.007)	0.016*** (0.006)	0.017* (0.010)
IW	(1)	(2)	(3)	(4)	(5)
Post-merger	−1.420*** (0.173)	−0.039*** (0.013)	−0.017* (0.009)	0.014* (0.008)	0.071*** (0.012)
Acquisition within ... years	10	10	10	10	10
Observations	8,642	8,642	8,642	8,642	7,108
	lhs(Lbs per fac.)	Toxicity-weighted			Gini (tox.)
		lhs(HH. income)	Pct. white	Pct. black	
TWFE	(1)	(2)	(3)	(4)	(5)
Post-merger	−0.849*** (0.301)	−0.016 (0.011)	−0.018** (0.008)	0.016*** (0.006)	0.011 (0.010)
IW	(1)	(2)	(3)	(4)	(5)
Post-merger	−2.305*** (0.363)	−0.016 (0.014)	−0.023** (0.011)	0.016* (0.008)	0.060*** (0.015)
Acquisition within ... years	10	10	10	10	10
Observations	8,642	8,642	8,642	8,642	6,766

Note:

*p<0.1; **p<0.05; ***p<0.01

In the decade after being acquired, firms see a shift in the pollution burden towards lower income, less white, and more black neighborhoods. Here, an observation represents a target firm, rather than individual facilities, operating in a particular state and industry in the years before and after its acquisition. The outcome variables in columns (3)–(8) represent the emissions-weighted average demographics for a firm treated by a merger (explained in detail in the main text), and regressions include weights equal to the number of facilities in each firm. The per-facility decrease in emissions is larger here, because facilities that close down are not dropped from the panel after 3 years (as is the case elsewhere).

Figure 3.9: IW event studies - exposure-weighted demographics



At the target firm-level, air emissions and their total toxicity declines substantially after an acquisition, but they also become more unevenly distributed. Event-year averages (across cohorts) only are presented and cohort-specific estimates are omitted for simplicity. These estimates correspond to the second and fourth rows of Table 3.10. The IW pre-period estimates are often, but not always, insignificantly different from zero. Visual inspection suggests a departure from the pre-period trend in the years after the acquisition for these aggregate (firm-level) measures.

Table 3.11: Releases at facilities held by acquiring firms

A: TWFE Estimates								
	lhs(Pounds air emissions)				lhs(Toxicity of air emissions)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A1: All facilities								
Post-merger	0.030 (0.087)	0.041 (0.071)	0.052 (0.084)	−0.007 (0.076)	0.031 (0.144)	0.072 (0.194)	0.009 (0.153)	−0.043 (0.171)
Observations	31,122	19,163	22,481	13,752	31,122	19,163	22,481	13,752
A2: Operating facilities	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-merger	0.006 (0.060)	0.009 (0.050)	0.046 (0.057)	−0.016 (0.059)	0.056 (0.068)	0.079 (0.122)	0.094 (0.071)	0.042 (0.124)
Observations	26,498	16,190	19,103	11,564	26,498	16,190	19,103	11,564
Merger sectors	All	Same	All	Same	All	Same	All	Same
Acquisition within ... years	10	10	5	5	10	10	5	5
B: IW Estimates								
	lhs(Pounds air emissions)				lhs(Toxicity of air emissions)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
B1: All facilities								
Post-merger	−0.096 (0.110)	0.118 (0.135)	0.063 (0.101)	0.176 (0.127)	−0.004 (0.425)	0.601* (0.335)	0.118 (0.335)	0.437 (0.278)
Observations	31,122	19,163	22,481	13,752	31,122	19,163	22,481	13,752
B2: Operating facilities	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-merger	−0.003 (0.050)	0.012 (0.125)	0.044 (0.058)	0.048 (0.128)	0.358* (0.206)	0.512* (0.296)	0.290* (0.158)	0.288 (0.226)
Observations	26,498	16,190	19,103	11,564	26,498	16,190	19,103	11,564
Merger sectors	All	Same	All	Same	All	Same	All	Same
Acquisition within ... years	10	10	5	5	10	10	5	5

Note:

*p<0.1; **p<0.05; ***p<0.01

I estimate the effect of an acquisition among facilities already held by the acquiring parent, finding no evidence that the quantity of emissions changes post-merger at already-held plants, and mixed evidence the toxicity of emissions rises somewhat. However, this is only possible when the parent company makes a single acquisition during the study period, which is a small subset of facilities held by parent companies that make acquisitions (when parent companies have more than one matched M&A transaction, they do not have a clear pre- and post-period; these tend to be larger and more-active firms). The identifying assumption that acquisitions are unanticipated may also be more tenuous among facilities held by the acquiring firm.

Table 3.12: The pollution burden also shifts as the number of facilities held by acquiring firms changes

	ihs(Lbs per fac.)	Pounds-weighted			Gini (lbs.)
		ihs(HH. income)	Pct. white	Pct. black	
TWFE	(1)	(2)	(3)	(4)	(5)
Post-merger	0.041 (0.402)	−0.019 (0.023)	−0.034** (0.017)	0.012** (0.005)	0.041** (0.020)
Observations	15,567	15,567	15,567	15,567	14,375
	ihs(Lbs per fac.)	Toxicity-weighted			Gini (tox.)
		ihs(HH. income)	Pct. white	Pct. black	
TWFE	(1)	(2)	(3)	(4)	(5)
Post-merger	−0.487 (0.866)	0.018 (0.043)	−0.019 (0.026)	−0.006 (0.020)	0.040** (0.016)
Observations	15,567	15,567	15,567	15,567	13,512

Note:

*p<0.1; **p<0.05; ***p<0.01

As acquiring firms grow in size, the distribution of emissions at the original set of facilities becomes more unequal. The quantity of emissions, but not the total toxicity, shifts towards plants in less-white neighborhoods. Because most acquiring firms—particularly those that are large or fast-growing—make multiple acquisitions, there is no clear pre- and post-merger period. However, I can exploit variation in changes in the number of facilities for each firm, state, and industry, considering the same distributional outcomes as in Table 3.10.

Appendix

3.A Additional figures and tables

Table 3.A1: Changes in the number and type of airborne releases

A: TWFE Estimates

	ihs(Num. airborne chemicals)		Airborne metals (0/1)	
Operating facilities	(1)	(2)	(3)	(4)
Post-merger	−0.009 (0.008)	−0.011 (0.007)	−0.008 (0.006)	−0.008* (0.004)
Observations	52,642	37,122	52,642	37,122
Merged within ... years	10	5	10	5

B: IW Estimates


	ihs(Num. airborne chemicals)		Airborne metals (0/1)	
Operating facilities	(1)	(2)	(3)	(4)
Post-merger	−0.028 (0.025)	0.005 (0.015)	−0.021* (0.011)	−0.010 (0.007)
Observations	52,642	37,122	52,642	37,122
Merged within ... years	10	5	10	5

Note:

*p<0.1; **p<0.05; ***p<0.01

Changes in the types of air releases at target facilities could provide evidence of shifts in the management of toxic chemicals. There is only mixed evidence acquired facilities release fewer airborne metals, which are especially toxic, after acquisition.

Figure 3.A1: Event studies - number and type of airborne chemicals



Event study plots corresponding to Table 3.A1 suggest the number of different chemicals released to the air declines at facilities that continue to operate following an acquisition, as does the likelihood that any metals (which are particularly toxic to human health) are released to air. These changes are consistent with plant-level operational shifts that reduce pollution.

Table 3.A2: Continuously operating facilities

A: TWFE Estimates

	lhs(Lbs air emissions)		lhs(Tox air emissions)	
	(1)	(2)	(3)	(4)
A1: Operating facilities				
Post-merger	−0.081 (0.049)	−0.108*** (0.037)	−0.242** (0.097)	−0.212*** (0.067)
Observations	52,642	37,122	52,642	37,122
A2: Continuously operating facilities	(1)	(2)	(3)	(4)
Post-merger	−0.050 (0.041)	−0.082** (0.034)	−0.196* (0.101)	−0.172** (0.068)
Observations	41,632	29,616	41,632	29,616
Merged within ... years	10	5	10	5

B: IW Estimates

	lhs(Lbs air emissions)		lhs(Tox air emissions)	
	(1)	(2)	(3)	(4)
B1: Operating facilities				
Post-merger	−0.272*** (0.083)	−0.203*** (0.074)	−0.320** (0.146)	−0.107 (0.092)
Observations	52,642	37,122	52,642	37,122
B2: Continuously operating facilities	(1)	(2)	(3)	(4)
Post-merger	−0.219*** (0.070)	−0.170** (0.078)	−0.195 (0.152)	−0.034 (0.117)
Observations	41,632	29,616	41,632	29,616
Merged within ... years	10	5	10	5

Note:

*p<0.1; **p<0.05; ***p<0.01

Small facilities may operate just above and just below the reporting threshold from year to year, in which case their releases in the data would be inappropriately truncated to zero in the years under the threshold, and nonzero otherwise. Zeros reflecting activity below the threshold and “true” zeros are indistinguishable in the data. In order to ensure that this characteristic of the reporting requirements is not driving results, I compare estimates for operating facilities (nonzero emissions in a given year, corresponding to panels A2 and B2 of Table 3.2) and continuously emitting facilities, which never report zero emissions to at least some medium in a year in between two positive-emissions years. Of course, subsetting the data in this way risks throwing out meaningful zeros and estimating effects among a selected sample of facilities.

Table 3.A3: Choice of fixed effects

A: TWFE Estimates

	lhs(Lbs air emissions)				
A1: All facilities	(1)	(2)	(3)	(4)	(5)
Post-merger	−0.509*** (0.054)	−0.512*** (0.055)	−0.491*** (0.045)	−0.482*** (0.048)	−0.465*** (0.033)
Observations	65,065	65,065	65,065	65,065	65,065
A2: Operating facilities	(1)	(2)	(3)	(4)	(5)
Post-merger	−0.083 (0.053)	−0.085 (0.052)	−0.095** (0.045)	−0.081 (0.049)	−0.078** (0.030)
Observations	52,642	52,642	52,642	52,642	52,642
FE	Y+i	S:Y+i	C:Y+i	I:Y+S:Y+i	I:Y+C:Y+i
Merged within ... years	10	10	10	10	10

B: IW Estimates

	lhs(Lbs air emissions)				
B1: All facilities	(1)	(2)	(3)	(4)	(5)
Post-merger	−0.931*** (0.116)	−0.909*** (0.120)	−0.989*** (0.145)	−0.887*** (0.116)	−0.959*** (0.116)
Observations	65,065	65,065	65,065	65,065	65,065
B2: Operating facilities	(1)	(2)	(3)	(4)	(5)
Post-merger	−0.254*** (0.085)	−0.241*** (0.085)	−0.285*** (0.095)	−0.272*** (0.083)	−0.307*** (0.075)
Observations	52,642	52,642	52,642	52,642	52,642
FE	Y+i	S:Y+i	C:Y+i	I:Y+S:Y+i	I:Y+C:Y+i
Merged within ... years	10	10	10	10	10

Note:

*p<0.1; **p<0.05; ***p<0.01

Fixed effects are abbreviated: (Y) year, (i) facility, (S) state, (C) county, and (I) industry. My preferred specification uses a richer set of fixed effects (I:Y+S:Y+i) than just individual and year fixed effects, but results are robust to other choices.

Bibliography

- Akey, Pat and Ian Appel**, “The Limits of Limited Liability: Evidence from Industrial Pollution,” *SSRN Electronic Journal*, 2017, 1 (647).
- Arnold, David**, “Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes,” *Working Paper*, 2020.
- Ash, Michael and James K. Boyce**, “Measuring corporate environmental justice performance,” *Corporate Social Responsibility and Environmental Management*, March 2011, 18 (2), 61–79.
- Banzhaf, Spencer, Lala Ma, and Christopher Timmins**, “Environmental justice: The economics of race, place, and pollution,” *Journal of Economic Perspectives*, 2019, 33 (1), 185–208.
- Becker, Randy A.**, “Local environmental regulation and plant-level productivity,” *Ecological Economics*, October 2011, 70 (12), 2516–2522.
- Berchicci, Luca, Glen Dowell, and Andrew A. King**, “Environmental capabilities and corporate strategy: Exploring acquisitions among US manufacturing firms,” *Strategic Management Journal*, September 2012, 33 (9), 1053–1071.
- Blonigen, Bruce A and Justin R Pierce**, “Evidence for the Effects of Mergers on Market Power and Efficiency,” *NBER Working Paper*, 2016.
- Brooks, Nancy and Rajiv Sethi**, “The Distribution of Pollution: Community Characteristics and Exposure to Air Toxics,” *Journal of Environmental Economics and Management*, 1997, 32 (2), 233–250.
- Campa, Pamela**, “Press and leaks: Do newspapers reduce toxic emissions?,” *Journal of Environmental Economics and Management*, 2018, 91, 184–202.
- Condliffe, Simon and O. Ashton Morgan**, “The effects of air quality regulations on the location decisions of pollution-intensive manufacturing plants,” *Journal of Regulatory Economics*, 2009, 36 (1), 83–93.
- Cui, Jingbo, Harvey Lapan, and Gian Carlo Moschini**, “Productivity, Export, and Environmental Performance: Air Pollutants in the United States,” *American Journal of Agricultural Economics*, 2016, 98 (2), 447–467.
- De Silva, Dakshina G., Timothy P. Hubbard, and Anita R. Schiller**, “Entry and exit patterns of “Toxic” firms,” *American Journal of Agricultural Economics*, 2016, 98 (3), 881–909.

- Eliason, Paul J, Benjamin Heebsh, Ryan C McDevitt, and James W Roberts**, “How Acquisitions Affect Firm Behavior and Performance: Evidence from the Dialysis Industry,” *The Quarterly Journal of Economics*, November 2019, 135 (1), 221–267.
- Fikru, Mahelet G.**, “Merger and Acquisition among Heterogeneous Polluting Firms : Theory and Evidence,” *Working Paper*, 2011.
- **and Luis Gautier**, “Are big mergers welfare enhancing when there is environmental externality?,” *Energy Economics*, March 2020, 87, 104718.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *NBER Working Paper*, 2018.
- Grant, Don and Andrew W. Jones**, “Are subsidiaries more prone to pollute? New evidence from the EPA’s toxics release inventory,” *Social Science Quarterly*, 2003, 84 (1), 162–173.
- Hamilton, James T.**, “Testing for Environmental Racism: Prejudice, Profits, Political Power?,” *Journal of Policy Analysis and Management*, 1995, 14 (1), 107.
- Holladay, J. Scott**, “Exporters and the environment,” *Canadian Journal of Economics*, February 2016, 49 (1), 147–172.
- Kagan, Robert A, Neil Gunningham, and Dorothy Thornton**, “Explaining Corporate Environmental Performance: How Does Regulation Matter?,” *Law & Society Review*, March 2003, 37 (1), 51–90.
- Kim, Incheol, Hong Wan, Bin Wang, and Tina Yang**, “Institutional Investors and Corporate Environmental, Social, and Governance Policies: Evidence from Toxics Release Data,” *Management Science*, 2019, (610), mns.2018.3055.
- Kwon, Ohsung, Sangmin Lim, and Duk Hee Lee**, “Acquiring startups in the energy sector: a study of firm value and environmental policy,” *Business Strategy and the Environment*, December 2018, 27 (8), 1376–1384.
- Li, Xiaoyang and Yue M. Zhou**, “Offshoring Pollution while Offshoring Production?,” *Strategic Management Journal*, November 2017, 38 (11), 2310–2329.
- Millimet, Daniel L. and Daniel Slottje**, “Environmental compliance costs and the distribution of emissions in the U.S,” *Journal of Regional Science*, 2002, 42 (1), 87–105.
- Park, Yongjoon**, “Impact of Airline Mergers on Environmental Externalities,” *Transportation Research Record*, December 2019, 2673 (12), 529–537.
- Powers, Nicholas**, “Measuring the impact of the Toxics Release Inventory: Evidence from Manufacturing Plant Births,” *Working Paper*, 2013.
- Sanders, Nicholas J**, “Toxic Assets : How the Housing Market Responds to Environmental Information Shocks,” *Working Paper*, 2012.
- Shapiro, Joseph S. and Reed Walker**, “Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade,” *American Economic Review*, December 2018, 108 (12), 3814–3854.

- Simon, Daniel H. and Jeffrey T. Prince**, “The effect of competition on toxic pollution releases,” *Journal of Environmental Economics and Management*, 2016, 79 (November), 40–54.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Working Paper*, 2020.
- Wang, Xiao, George Deltas, Madhu Khanna, and Xiang Bi**, “Environmental Disclosure Programs, Community Pressure and the Spatial Redistribution of Pollution: The Relocation of Toxic Releasing Facilities After the TRI,” *SSRN Electronic Journal*, 2018.
- Wolverton, Ann**, “Effects of socio-economic and input-related factors on polluting plants’ location decisions,” *B.E. Journal of Economic Analysis and Policy*, January 2009, 9 (1).