NBER WORKING PAPER SERIES

WHO BENEFITS FROM HAZARDOUS WASTE CLEANUPS? EVIDENCE FROM THE HOUSING MARKET

Alecia W. Cassidy
Elaine L. Hill
Lala Ma

Working Paper 30661
http://www.nber.org/papers/w30661

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2022, Revised October 2025

The authors would like to thank Lint Barrage, David Bernstein, Traviss Cassidy, Susana Ferreira, Dennis Guignet, Robert Hammond, Lauren Lambie-Hanson, Craig Landry, Michael R. Moore, Adam Nowak, Nicolai Kuminoff, Amanda Ross, Steven L. Ross, Joe Shapiro, and Kenny Wunder, as well as seminar participants at the University of Alabama, University of Albany, University of Georgia, AERE Summer Conference, APPAM Fall Conference, ASSA Annual meeting, Indiana University, New York Fed, Rochester Institute of Technology, SEA Annual Conference, SUNY Albany, and West Virginia University for valuable feedback. We would also like to thank three anonymous referees. The authors gratefully acknowledge the EPA Office of Land and Emergency Management (OLEM) for providing program expertise and access to the data. We especially thank David Hockey, Rachel Horton, and Norman Birchfield at the of the OLEM for their partnership. This work was conducted in part while Lala Ma was appointed to an Oak Ridge Institute for Science and Education (ORISE) Research Participation Program at the EPA OLEM. Finally, we thank Grace Sventek, Sydney Pargman, and especially Aniket Patil who did most of the heavy lift, for excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research or of the EPA.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Alecia W. Cassidy, Elaine L. Hill, and Lala Ma. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Who Benefits from Hazardous Waste Cleanups? Evidence from the Housing Market
Alecia W. Cassidy, Elaine L. Hill, and Lala Ma
NBER Working Paper No. 30661
November 2022, Revised October 2025
JEL No. Q51, Q52, Q53, Q58

ABSTRACT

The Resource Conservation and Recovery Act (RCRA) manages cleanup of hazardous waste releases at over 3,500 sites across the US, which covers approximately 17.9 million of all developed land in the country. We use publicly available Census survey data (Decennial and American Communities Survey) to evaluate the national and distributional housing market impacts of these cleanups. We find that cleanups near residential properties yield significant, yet localized, increases in home prices, and that impacts are concentrated in the lower deciles of the price distribution. Importantly, we find fairly weak evidence of sorting along sociodemographic dimensions in response to cleanup. Our findings suggest that cleanup benefits accrue to the residents who are the original 'hosts' of pollution and could correct pre-existing disparities in exposure to land-based contamination.

Alecia W. Cassidy
The University of Alabama
Department of Economics
awcassidy1@cba.ua.edu

Lala Ma
University of Kentucky
Department of Economics
lala.ma@uky.edu

Elaine L. Hill
University of Rochester
Department of Public Health Sciences,
and Department of Economics
and NBER
elaine_hill@urmc.rochester.edu

1 Introduction

The Resource Conservation and Recovery Act (RCRA) aims to "protect human health and the environment from the potential hazards of waste disposal" (U.S. EPA, 2014). The Corrective Action Program, established under RCRA, investigates and cleans releases of hazardous waste at RCRA facilities with the primary goals of controlling human exposure and containing migration of contaminated groundwater. The impacts of this particular program are potentially widespread: As of fiscal year 2011, the RCRA Corrective Action Program tracked 3,747 sites, which spanned 17,946,593 acres (U.S. EPA, 2011). This program alone covers between 15.4 and 17.5% of all developed land in the U.S.[1]

This paper evaluates the benefits of cleanups performed under the RCRA Corrective Action Program (CAP) by estimating the impacts of cleanup on national housing prices and neighborhood change. We combine publicly-available tract-level data from the 1990 and 2000 Decennial Censuses and the 2008–2012 and 2018–2022 5-year American Community Surveys with all cleanups conducted under the program across the continental US, Alaska, and Hawaii. To measure housing prices, we construct deciles of the housing price distribution from Census survey counts. Throughout, we will use "price" to refer to these derived deciles. These deciles allow us to characterize cleanup effects across the distribution of housing prices. Our identification strategy leverages spatial and temporal variation in cleanups, capturing the differential effect for Census tracts whose boundaries are near and far from RCRA sites from before to after cleanup.[2]

A concentration of RCRA sites in disadvantaged neighborhoods means that cleanup efforts could reduce inequitable pollution gaps documented in environmental justice studies (Banzhaf, Ma and Timmins, 2019). The positive distributional effects may not, however, materialize if cleanups trigger re-sorting in response to price changes, altering the composition of those exposed to cleaned sites. Therefore, we follow our housing price analysis with an investigation of the extent to which RCRA cleanups altered neighborhood composition to evaluate to whom cleanup benefits accrue. We first estimate reduced-form regressions of cleanup on 17 different socio-demographic and housing-related outcomes from the Census Summary Files. We then apply a structural sorting model to test whether willingness to pay (WTP) for RCRA cleanup systematically differs by race and income to evaluate the scope for heterogeneous preferences to drive environmental gentrification.[3]

[1] The share of the RCRA program depends on the estimate of the total developed lands in the US. The US EPA (2010) estimates 102.5 million acres of developed land in the US as of 2011, whereas the USDA National Resource Inventory estimates 116 million acres as of 2017 (USDA, 2017).

[2] We empirically determine what constitutes "near" and "far" in Section 5. To preview, we compare tracts containing RCRA sites (0 km away) to those $(0, 10]$ km away in the main specification, where distances are defined as the minimum distance from tract to site.

[3] The additional structure of the empirical sorting model alleviates an identification challenge with using aggregate population changes (i.e., without information on residential origin and destination) to infer preference to avoid pollution (Depro, Timmins and O'Neil, 2015). Moreover, this model allows moving to be costly, an important source of bias in hedonic pricing analyses (Bayer, Keohane and Timmins, 2009).

Our contribution is to provide *nationally representative* evidence of the *distributional* impacts of RCRA cleanups using *publicly available data.* Prior studies have examined the housing price impacts of waste managed under RCRA for specific areas within the country (Smith and Desvousges, 1986; Kinnaman, 2009). Our analysis complements this body of work by extending the scope to the national level. While others have estimated the national housing market impacts of pollution remediation,[4] including Guignet and Nolte (2024) who examine the RCRA program with a national scope,[5] few have tested for sorting in response to these types of remediation activities, which could alter the individuals who ultimately experience the benefits of environmental improvements.[6] Since RCRA facilities are widespread in the US, the scope for cleanup activities to trigger endogenous neighborhood change, or "environmental gentrification" (Banzhaf and McCormick, 2007), is a real concern given its potential to affect both the overall and distribution of cleanup benefits. Finally, that we use publicly available Census data is in contrast with most of the papers in this literature (with a couple of notable exceptions, e.g. Gamper-Rabindran et al. (2011); Gamper-Rabindran and Timmins (2013)). Our techniques can be applied to analyze housing price impacts even when researchers do not have access to individual-transaction data. This is increasingly important since Zillow Ztrax is no longer available to researchers, and other proprietary data sources such as Cotality are prohibitively expensive for the vast majority of economists.

We find that cleanup increases home prices of properties in the same tract as the facility, but does not have a price impact beyond the immediate tract. The housing impacts are higher in percentage terms for properties in the lower deciles of the price distribution, with a 7.6% increase in price for the 1^{st} decile, and no evidence of increase for the 9^{th} decile, though standard errors are not small enough to rule out similar percentage increases between the two deciles. This suggests that

[4]Several studies have estimated the housing market impacts of different nuisances, with the geographic scope ranging from regional to national. Examples include river pollution by wastewater treatment plants (Keiser and Shapiro, 2018), Toxic Release Inventories (Currie, Davis, Greenstone and Walker, 2015; Mastromonaco, 2015), brownfield sites (Linn, 2013; Haninger, Ma and Timmins, 2017; Ma, 2019), Superfund sites (Currie, Greenstone and Moretti, 2011; Greenstone and Gallagher, 2008; Gamper-Rabindran, Mastromonaco and Timmins, 2011; Kohlhase, 1991; Gayer, Hamilton and Viscusi, 2000), industrial chemical accidents (Guignet, Jenkins, Nolte and Belke, 2023), and Areas of Concern in the Great Lakes (Cassidy, Meeks and Moore, 2023). As Banzhaf (2021) illustrates, the capitalization effects found in these papers can be interpreted as a lower bound on welfare effects under assumptions such as a time-invariant hedonic gradient. More work is needed to discern whether the formal results by Banzhaf (2021) characterizing assumptions under which quasi-experiments can reveal a lower bound on welfare extend to the case of quantiles derived from aggregate data.

[5]Our paper complements work by Guignet and Nolte (2024) and differs from theirs in (1) our focus on the distribution of price impacts and post-cleanup neighborhood change, (2) our examination of all sites under the Corrective Action Program as opposed to only Treatment, Storage, and Disposal Facilities (TSDFs), and (3) our use of publicly-available Census data with rich demographic information that enables our investigation into cleanup-related sociodemographic change. Our effect sizes on housing prices are comparable, which is not surprising given evidence from the literature that estimates produced using decile-level census data can detect similar effects to individual-transaction data (Gamper-Rabindran and Timmins, 2013). We provide additional discussion of the differences in Section 5.1.

[6]One example of a study that undertakes a reduced-form investigation of sorting is Greenstone and Gallagher (2008), who find no effect on price and also no sorting. In contrast, we find that there is an effect on price but no sorting. The latter means that cleanup has benefits specifically to those living closest to the sites, whereas the former is a finding of no capitalization of benefits.

cleanups raised housing prices for the least advantaged residents living in tracts near facilities.[7]

Our findings are robust to alternative specifications, exposure definitions, controls, and data sources. First, we find similar effects when we employ Robust High-Dimensional Fixed Effects Poisson estimation to address the skewed nature of our dependent variable (Correia, Guimarães and Zylkin, 2021). Second, although our primary results limit the sample to tracts near one RCRA site, impacts for the lowest deciles are attenuated but still statistically significant at lower deciles of the price distribution when we include tracts near multiple RCRA sites. Third, results are similar when we employ a regression adjustment that accounts for the lack of balance in tract areas between our treatment group (tracts containing a RCRA site) and our comparison group (nearby tracts within 10 km). Fourth, our substantive conclusions (that price impacts are concentrated in the lowest deciles, and mostly insignificant for the highest deciles) are corroborated when we estimate quantile treatment effects using a supplemental dataset containing individual housing transaction data from Ohio and Pennsylvania.

Apart from localized price increases in the lowest deciles, we find limited evidence of re-sorting along socio-demographic dimensions in response to cleanup. Using reduced-form regressions, we find no statistically significant impacts of cleanup on any of the 17 socio-economic and housing-related indicators from the Census. Our structural sorting model also suggests that WTP for cleanup is not economically different between different groups, defined either based on race or income. Furthermore, there is limited evidence of race-based sorting even when we focus on homeowners, who stand to gain from both pollution improvements and asset price appreciation and are disproportionately white.[8]

Taken together, our findings imply that the benefits of cleanup accrued to those living closest to the facilities, who tended to be more disadvantaged compared to those living farther from the facilities. This is particularly important given recent findings in the literature related to environmental equity. Hausman and Stolper (2020) show that, in a framework where people undervalue a clean environment and have partial information, residential sorting on willingness-to-pay leads to an equilibrium where deadweight loss due to pollution is higher for economically disadvantaged segments of the population.[9] Furthermore, Bakkensen and Ma (2020) demonstrate that

[7] The capitalization effects we find *reflect* benefits to homeowners, but caution should be exercised when interpreting them in terms of willingness-to-pay. The median houses in different periods could have very different characteristics. The potential for within-tract sorting further muddies the interpretation. Therefore, we cannot necessarily assume results from Rosen (1974) and Banzhaf (2020) will hold. Furthermore, Banzhaf and Farooque (2013) show that median housing prices are only weakly correlated with variables of interest, including prices from individual transactions, ozone, and income (with correlations of 0.543, -0.425, and 0.284, respectively) using cross-sectional data. If results generalize to our panel setting, low correlations between median housing values and variables of interest would imply downward attenuation of our estimates compared to true willingness-to-pay.

[8] According to the Office of Economic Policy at the US Department of Treasury, the homeownership rate for white households was 75% compared to 45% for Black households and 48% for Hispanic households. These gaps in homeownership have been persistent over time. See https://home.treasury.gov/news/featured-stories/racial-differences-in-economic-security-housing.

[9] Partial information could broadly include factors like inattention to the nuisance or misperception of risks

well-intentioned policies in the housing market can have significant distributive effects, leading the least well-off residents to take on even more exposure to an environmental bad. If RCRA cleanups do not induce sorting along socio-demographic dimensions, they could be corrective of the type of pre-existing disparities that Hausman and Stolper (2020) call attention to, and they are unlikely to exacerbate existing inequities as Bakkensen and Ma (2020) find.

The paper is organized as follows. Section 2 provides a background on the Resource Conservation and Recovery Act and its cleanup program. Section 3 describes our data sources and the data construction process. We describe our empirical models in Section 4 and present the corresponding results in Section 5. Section 6 discusses the implications of the results, and Section 7 concludes.

2 Background

The Resource Conservation and Recovery Act (RCRA) was enacted in 1976 by Congress. The Act consists of ten subtitles, where the two major programs under RCRA are subtitles C and D, which respectively regulate hazardous waste and non-hazardous solid waste. Subtitle C, under which cleanups of hazardous waste are conducted, sets regulations for the handling (i.e., creation, management, and disposal) of hazardous waste and is codified in Title 40 of the Code of Federal Regulations (CFR). There are three main types of RCRA hazardous waste handlers: (1) generators, (2) transporters, and (3) facilities that treat, store, or dispose of waste. As of 2009, there were 460 Treatment, Storage, and Disposal Facilities (TSDF's), 18,000 transporters, and 14,700 generators that handle large quantities of waste.[10] Subtitle C regulations, importantly, grant the EPA the authority to require cleanup for any release of hazardous waste to all environmental media at both RCRA-permitted and non-permitted facilities. The cleanup program, previously known as the RCRA Corrective Action Hazardous Waste Cleanup Program but recently rebranded as the Hazardous Waste Cleanup Program, is the focus of this paper.

The Corrective Action Program (CAP), established under the Hazardous and Solid Waste Amendments to RCRA in 1984, investigates and cleans releases of hazardous waste at a subset of RCRA facilities. Unlike the Superfund program, sites managed under this cleanup program are often in operation.[11] Sites not in operation typically result in clean closure and possible reuse, such as residential developments, shopping plazas, mixed-use commercial, and recreation (parks, wetlands,

(Bakkensen and Barrage, 2021).

[10]Generators are subdivided into three groups based on the amount and type of hazardous waste that is generated – Very Small Quantity Generators (VSQG), Small Quantity Generators (SQG), and Large quantity generators (LQG). VSQG's generate 100 kilograms or less per month of hazardous waste or one kilogram or less per month of acutely hazardous waste; SQG's generate more than 100 kilograms, but less than 1,000 kilograms of hazardous waste per month; LQG's generate 1,000 kilograms per month or more of hazardous waste or more than one kilogram per month of acutely hazardous waste.

[11]Using the Biennial Hazardous Waste Report (BR), we found that of the 2,323 CAP sites that are required to report the BR, 76% reported waste generation, storage, shipping, or receiving after the clean-up. 999 of our sites are exempt from the BR due to being small and very small quantity generators. (U.S. Environmental Protection Agency, 2025a)

and golf courses).[12]

Beginning in 1999, efforts took place to reform the cleanup process and remove bureaucratic hurdles to accelerate the pace of cleanups. The EPA identified RCRA facilities with the potential for unacceptable exposure to pollutants and/or for ground water contamination. Facilities were chosen based on the National Corrective Action Prioritization System (NCAPS), which categorizes facilities as High, Medium, or Low priority.[13] Theoretically, the ranking should have been based on waste type, waste volume, release pathways (ground water, surface water, air, and soil), and the potential for human and ecosystem exposure. However, the scoring of these sites were not systematically based on pollution intensity.[14] In some cases, the ranking can also depend on compliance history or special conditions (e.g. regional initiatives). Most RCRA facilities were ranked by 1993. Since the nature of site selection may not be random, evaluation of causal housing market impacts from changes in pollution at sites would require controlling for confounding factors that prioritize cleanup at some sites but not others. Our empirical strategy (described subsequently) will deal with this selection by ensuring that the variation we use is between tracts in the vicinity the same site, rather than between sites cleaned up at different times.

Relevant to our study, the program set cleanup (or risk reduction) targets based on two environmental indicators (EI), which we use in our analysis to determine the timing of cleanup: (1) Current Human Exposures Under Controls (or the Human Exposure EI), and (2) Migration of Contaminated Groundwater Under Control (or the Groundwater EI).[15] A positive Human Exposure EI determination indicates that there are no "unacceptable" human exposures to contamination that can be reasonably expected under current land- and groundwater-use conditions. A positive Groundwater EI determination indicates that the migration of contaminated groundwater has stabilized, and that monitoring will be conducted to confirm that contaminated groundwater remains within the original area of contamination.[16] With two cleanup targets established, the EPA set goals to control human exposure and migration of contaminated groundwater.

To the extent that housing market participants are aware of these facilities and the corrective actions that have taken place, a portion of the cleanup benefits should be capitalized into housing prices. If, however, households are unaware, then the benefits of corrective actions might not be reflected in the housing market; this does not mean that cleanups have no value, since the public

[12] The EPA highlights economic development in 42 case studies (U.S. Environmental Protection Agency, 2025c). There may be other potential redevelopment options that are not reflected in these case studies. There are no systematic data that we are aware of that capture the nature of economic development following cleanup.

[13] This is somewhat similar to the Hazard Ranking System (HRS) used by Superfund, except requires less detailed input.

[14] This is based on conversations with EPA officers suggest that the initial ranking of sites had a certain degree of randomness. Unfortunately, the scores for the sites are no longer available to verify the nature of the selection by score.

[15] These are identified in our data as Corrective Action (CA) codes 725 and 750 respectively (U.S. Environmental Protection Agency, 2025b).

[16] "Unacceptable" contamination levels refer to contaminant concentrations in excess of appropriate risk-based levels.

may still value these cleanups *had it known* about them (Cassidy, 2023; Gayer et al., 2000; Ma, 2019). We next test for housing impacts with data.

3 Data

3.1 Data Sources

Data come from the following sources: (1) RCRAinfo Corrective Action Program cleanups, (2) Demographic data from the US Census Bureau aggregate tabulations (Summary Files) for the following surveys: 1990 Decennial, 2000 Decennial, 2008-2012 5-year American Community Survey (ACS), and 2018-2022 5-year ACS, and (3) IPUMS USA public-use microdata from the following surveys: 2000 5% Sample, ACS 2008-2012 5-year Sample, and ACS 2018-2022 5-year Sample.

RCRA Corrective Action Program Cleanups Data on Corrective Action Program (CAP) cleanups come from the RCRAinfo database, which is publicly available from the EPA. We begin with all sites in the cleanup program.[17] Each facility is identified by a unique waste handler identifier. Several attributes of the handler are available, including the location of each facility, the primary industry to which it belongs (3-digit NAICS code), whether the facility is a waste generator, transporter, or treatment, storage, or disposal facility (commonly referred to as a TSDF), and the first NCAPS ranking.

We focus on the two Environmental Indicators (EIs) to define our cleanup event. The data entry system for the EIs worked as follows: the government official would review data associated with the site, and enter the corrective action into the system with a status code of "NO" or "IN" if the objectives had not yet been achieved. The government official would enter the status code of "YE" if the objectives had been achieved. This means that for some facilities, there are multiple dates associated with the corrective action. We take the date of cleanup to be the date on which the last "YE" entry was made for either of the two EIs, whichever comes later. We take the date at which cleanup began to be the date of the first entry associated with either of the two EIs, whichever comes earlier. The time between when cleanup began and when it was finalized is categorized as "In Progress"; this way, we are not capturing what happens during cleanup in the measurement of our post-cleanup effect.

In Figure 1, we visualize the spatial distribution of CAP sites in our dataset. We start with

[17] These are the sites designated in the 2005, 2008, and 2020 CA baselines. "Cleanup program" is used interchangeably with the Corrective Action Program throughout recent EPA materials; we follow this convention throughout our paper. While the two earlier baseline years (2005 and 2008) were prioritized lists of sites needing corrective action, the baseline sites should be seen as constituting all sites in the program. According to the EPA's RCRA Orientation Manual, "The RCRA cleanup baseline has expanded to include all 3,779 facilities expected to need corrective action. Because EPA has set ambitious goals for 2020 that relate to these facilities, the group is called the 2020 Corrective Action Universe. The goals for 2020 apply to the full corrective action universe and are to have human exposures controlled at 65% of facilities, the migration of contaminated groundwater controlled at 55% of these facilities, and final remedies constructed at 32% of these facilities. " (Available at: https://www.epa.gov/sites/default/files/2015-07/documents/rom.pdf) We note, however, that additional sites may have corrective actions take place, that are not part of the cleanup program. These are not part of the scope of our paper.

3,873 sites in the program from the RCRAinfo database. Of those, 3,492 have non-duplicate valid coordinates for US locations, excluding Guam, Puerto Rico and the U.S. Virgin, and are the nearest site to at least one tract. Panel (a) shows all 3,492 sites that could be geolocated to Census tracts. Panel (b) shows the 1,557 sites when we restrict to tracts within 10 km of a single site, as we do in our main specification.

Census Summary Files Data From the Census, we collect census tract-level statistics from publicly available Summary Files for the following surveys: 1990 Decennial, 2000 Decennial, 2008-2012 5-year ACS, and 2018-2022 5-year ACS (see Data Appendix C for details).[18] Census tracts typically contain between 1,200 and 8,000 people, and their borders are defined using permanent, salient and natural geographic and manmade features that could delineate neighborhoods, such as rivers and highways. This makes them an ideal geographic unit for studying demographic changes that reflect neighborhood-scale dynamics, such as local housing markets, gentrification and socioeconomic shifts.[19]

We collect data on the value of owner-occupied housing (reported in various price bins), counts of houses located in each price bin, and neighborhood demographic characteristics related to, e.g., race, income, and education.[20] Using the counts of the number of houses in each price bin, we construct deciles of the price distribution for each census tract following Gamper-Rabindran et al. (2011).[21] We then inflation-adjust prices to 2023 dollars (see Appendix C.2 for details).

Tracts may expand or condense over time; this necessitates a method to compare tract-level information over time (i.e., interpolation). We do this with National Historical Geographic Information System (NHGIS) crosswalks with provides interpolation weights (Manson, Schroeder, Riper, Knowles, Kugler, Roberts and Ruggles, 2024).[22] Each interpolation weight indicates the proportion of a source zone's characteristics to allocate to a specific target zone. Using these NHGIS crosswalks, we construct a panel data set of 2010 census tracts over four decennial census years: 1990, 2000, 2010, 2020 (See Appendix Section C.1 for full details).

Our main specification uses owner-occupied weights for interpolation (and renter-occupied weights for rental prices), given our focus on housing values of owner-occupied housing; when

[18]Connecticut changed their county-equivalents from 8 to 9 counties for the ACS 2022. There are no interpolation crosswalks for this change, therefore, we use 2017-2021 ACS for CT.

[19]For more, see the Geographic Areas Reference Manual U.S. Dept. of Commerce, Economics and Statistics Administration, Bureau of the Census (1994).

[20]The underlying questions for these variables of interest change over time; we describe changes in owner-occupied housing in Appendix Section C.4, calculation of percentiles in Section C.3, and changes in race/ethnicity in Section C.5.

[21]In a study examining the housing impacts of Superfund cleanups, Gamper-Rabindran et al. (2011) show their decile approach can detect similar magnitudes of benefits from cleanup as approaches using repeat-sales data. We use the Decennial censuses when long form was available (1990, 2000) and the ACS 5-year estimates when long form is not (2010, 2020) because they contain similar content. We use the 2008-2012 ACS because 2010 is the midpoint and similarly, 2018-2022 ACS for 2020 as the midpoint. This approach is commonly used in place of the long form Census (Logan, Xu and Stults, 2014).

[22]NHGIS produces geographically-standardized time-series tables, however, our housing variables of interest are not contained in these files.

Figure 1: Spatial Visualization of Sites

(a) All RCRA Corrective Action Program (CAP) Sites

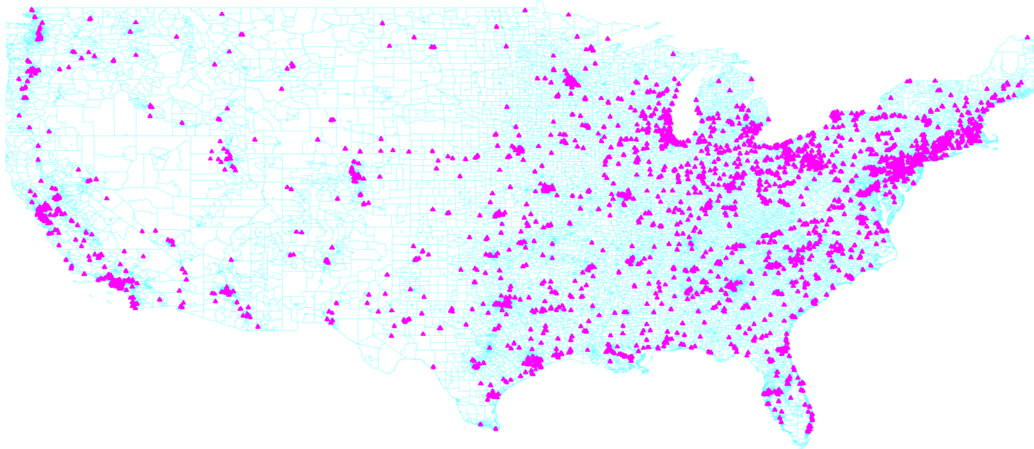

(b) RCRA CAP Sites when Restricting to Tracts within 10 km of a Single Site

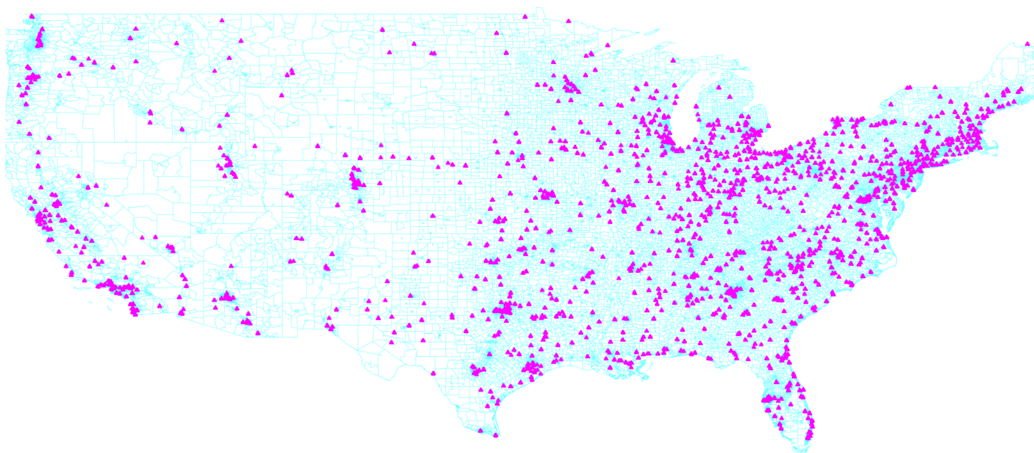

Notes: Panel (a) shows all RCRA Corrective Action Program sites in our dataset in the Continental USA (purple triangles). There are 3,492 sites (only counting non-duplicate sites not in Puerto Rico, Guam or the United States Virgin Islands with valid coordinates that are the nearest site to at least 1 tract) in our dataset. Of those, 1,557 are represented when we restrict to tracts within 10 km of a site (Panel b). Note that Alaska and Hawaii are not pictured, but do contain sites present in our dataset. Lines denote 2010 Tract boundaries. Tracts are only included if they are in the NHGIS 2010 tract-level shapefile, which removes entirely water-based tracts.

owner-occupied weights are missing, we use proportion-area weights as recommended by NHGIS (Manson et al., 2024). We show robustness checks for alternative NHGIS interpolation weights and interpolation using the Longitudinal Tract Data Base (LTDB) (Logan et al., 2014), which also interpolates census summary statistics from different decennial censuses into estimates based on 2000 or 2010 tract boundaries. The two primary differences between the NHGIS and LTDB harmonization techniques are as follows. First, the NHGIS crosswalks allow building up from the block or block part level (depending on the variables available) which increases accuracy. Second, whereas LTDB only blocks out water bodies as uninhabited land, NHGIS incorporates road buffers and land cover like forests into the uninhabited space. However, the techniques should only substantially differ for a minority of tracts that had large changes in boundaries. We discuss differences in these approaches, details about the interpolation weights, and summarize the interpolation steps in Appendix Section C.1. Table C.1 summarizes specifics about the weights used and source information for each robustness check. We show trends for all harmonized Census variables used in the paper in Appendix Section C.9.

Aside from interpolation, it is worth understanding how housing price data collected from Census summary files might compare with individual transaction data. The first consideration here is the use of survey data in this context. In collecting Census data, homeowners are asked, "About how much do you think this house and lot, apartment, or mobile home (and lot, if owned) would sell for if it were for sale?" The disadvantage of using survey data is that homeowners might have an inaccurate view of how much their home is worth.[23] While recent work has found that individual level survey-based estimates of home values are highly correlated with contemporaneous transaction values (Banzhaf and Farooque, 2013) and proprietary value estimates (Bakkensen and Barrage, 2021),[24] homeowners might not correctly factor in information about local amenities (Banzhaf and Farooque, 2013) and their valuation of local amenities might be disproportionate to the size of the associated environmental benefits (Diamond and Hausman, 1994; Hausman, 2012).[25] Taken together, it is possible that our estimates are biased upward. We explore this possibility and discuss implications for our study in detail in Appendix B.1, where we compare our Census data to individual-level transaction data from Ohio and Pennsylvania. The general finding is that homeowners do appear to overvalue their houses, in line with findings from Kiel and Zabel (1999); however, we do not find that the overvaluation varies by price percentile. See Figure B.1 in Appendix

[23]The literature that investigates the differences between the two types of data and the determinants of that gap generally finds that owners overestimate the market value of their homes (Banzhaf and Farooque, 2013; Diamond and Hausman, 1994; Hausman, 2012; Goodman Jr and Ittner, 1992; Kiel and Zabel, 1999). In terms of correlated characteristics, the consensus is less clear: While some find that overvaluation increases with income and housing tenure and decreases college degree and home value (Dreesen and Damen, 2023; Tur-Sinai, Fleishman and Romanov, 2020), others find no consistent correlation with market value and characteristics (Goodman Jr and Ittner, 1992; Kiel and Zabel, 1999).

[24]Bakkensen and Barrage (2021) find a correlation coefficient of 0.89 and Banzhaf and Farooque (2013) find the correlation to be approximately 95–97%

[25]A related concern is that households in less expensive homes overvalue the amenities, but recent literature seems to indicate that the opposite is more likely (Hausman and Stolper, 2020).

B that plots the average deciles for tract-level Zillow and Census prices. The degree of potential overvaluation does not systematically vary by price percentile (i.e., it is fairly parallel with the 45 degree line). The advantage of using survey data over individual transaction data is that not every house is sold every year, and so frequently transacted houses tend to be over-represented in transaction data. Those also tend to be more likely to be "flipped."[26]

The second consideration is geographic data coverage- twelve states in the United States are non-disclosure states, where transaction data does not represent sale prices. Instead, the recorded transaction values are an opaque combination of mortgage values and whatever the buyer and seller want to portray the price as.[27] These states are: Alaska, Idaho, Kansas, Missouri (some counties), Mississippi, Louisiana, Wyoming, Utah, Texas, North Dakota, New Mexico, and Montana. These non-disclosure states account for approximately 18% of observations in our sample, and 14% of sites in the cleanup program.[28]

It is also worth discussing our use of deciles of the housing price distribution (see Appendix Section C.3 for details on how these are calculated). These are the most appropriate dependent variable because our research question is about distributional impacts. An additional advantage of using percentiles (or any moment of the distribution) of the housing price instead of housing price data itself is that percentiles have a less skewed distribution. Researchers often use the logarithm of price as the dependent variable in regressions with transaction data to alleviate concerns with outliers posing a threat to small-sample inference when the distribution of the dependent variable is skewed (Bishop, Kuminoff, Banzhaf, Boyle, von Gravenitz, Pope and Timmins, 2020). We show histograms of these price percentiles in Figures A.1 in the appendix.[29]

Despite observing skew in our price variables, we do not follow the convention of using the log of price as the dependent variable. Instead, to address small-sample bias that could arise due to the fact that our dependent variable is skewed, we provide a robustness check using Poisson regression. We prefer this because, as described by Gould (2011), robust Poisson regression should be used in situations where researchers would be inclined to take the logarithm of the dependent variable to report a semi-elasticity. When logarithms are used instead, the estimates can be biased and inconsistent (Manning and Mullahy, 2001; Santos Silva and Tenreyo, 2006). Correia et al. (2021) prove that the high-dimensional fixed effects Poisson estimator we employ yields unbiased and consistent estimates.

[26] For comparison purposes, we present robustness checks using individual transaction-level data in Appendix Section B.

[27] Concrete examples are documented in Cassidy (2023).

[28] An alternative source of data that could be helpful in the case of non-disclosure states is assessed values, but there are concerns with using these to proxy housing prices as outlined in Appendix Section B.1.

[29] To get a sense of the how the skewness compares between these percentile variables and raw transaction price data, one can compare these histograms to appendix Figure B.5, which presents the distribution of raw transaction prices from Ohio and Pennsylvania for the treatment and comparison groups.

Census Microdata We supplement the aggregate summary files from the Census with publicly available census microdata from IPUMS USA to estimate our sorting model (Ruggles, Flood, Sobek, Backman, Cooper, Drew, Richards, Rodgers, Schroeder and Williams, 2025). This is used to construct the share of households that stay in their tract over time for individuals of different groups, categorized based on race, income, or homeownership. Specifically, the survey asks in what year each person (for 1960-1970) or the householder (for 1980 on) moved into the dwelling unit (apartment, house, or mobile home). IPUMS then recodes the responses as the number of years ago that the householder moved into the housing unit. Based on the year of the survey, we create a binary variable for whether a person moved into the dwelling by the previous decade (e.g., for the year 2000, this is an indicator equal to 1 if the person moved in 11 to 20 years ago). We then calculate group-specific averages of the probability that the individual stayed in that dwelling between 1990 and 2000. We weight the group-specific stay shares by the person weights provided by IPUMS. We download data for each year using the following samples: 2000 5% sample, 5-year ACS 2008-2012 sample, and 5-year ACS 2018-2022 sample. Since the ACS 2008-2012 and 2018-2022 samples are 5-year averages, we use the midpoint of the 5-year period to calculate when a person moved into the dwelling.

3.2 Data Construction & Summary Statistics

We construct a national dataset linking our outcomes of interest (housing and sociodemographic variables) to RCRA cleanups. For the housing data, we first identify all census tracts for which any portion of the tract's boundary is within a 10-kilometer (km) buffer of a RCRA Corrective Action baseline facility. Using this spatial relationship, we then create a census tract-by-year level data set that describes the deciles of the census tract's housing price distribution, the deciles of the Census tract's rental price distribution, and sociodemographic characteristics of the tract in that particular year, all matched to cleanup timing.

We limit the sample to tracts within 10 km (minimum distance) of RCRA sites in order to avoid comparing neighborhoods that are very different (removing 24,378 tracts that are not within 10 km of any site). For example, Table 1 provides summary statistics of census tract characteristics by whether a tract is within 10 km of any RCRA facility. Areas with facilities have lower housing prices in the higher deciles of the price distribution, have lower income, and are more diverse. They are slightly less likely to be on public assistance, or below the poverty line. That these observable characteristics are correlated with RCRA site location suggests that other correlated, unobserved factors may exist. Our initial sample limitation thus removes some of these potential unobserved confounders, assuming that the composition of tracts is relatively constant around 10 km away from the RCRA sites. Figure 2 visualizes tracts near sites. Panel (a) shows whether a tract contains at least one site. Panel (b) shows how many sites are within 10 km by tract.

We further limit the sample to areas within 10 km of a single RCRA facility so that cleanup timing is well-defined (See Appendix Figure A.2a for a map of tracts within 10 km and Figure A.2b

Figure 2: Spatial Visualization of Tracts

(a) Tracts Containing Sites

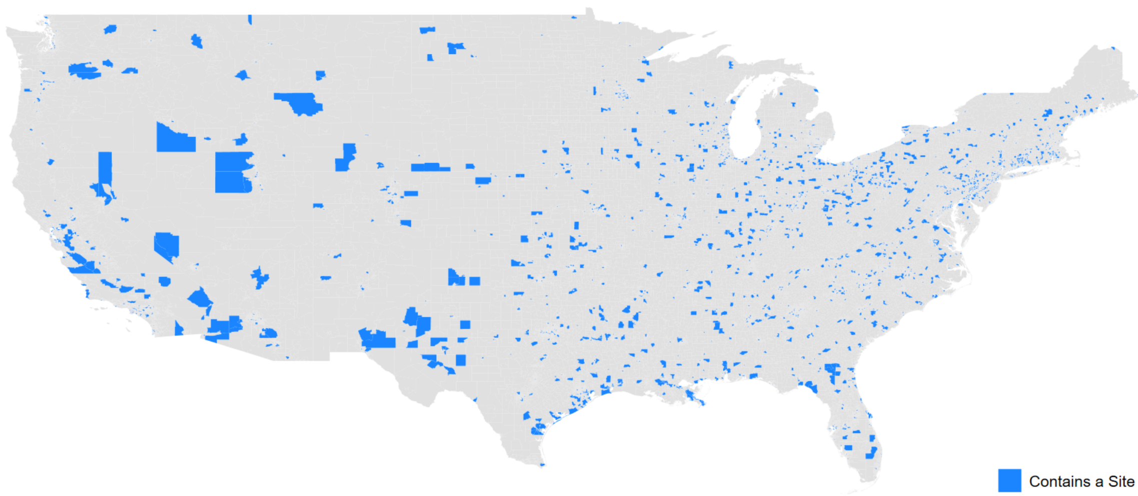

(b) Tracts within 10 km of Sites

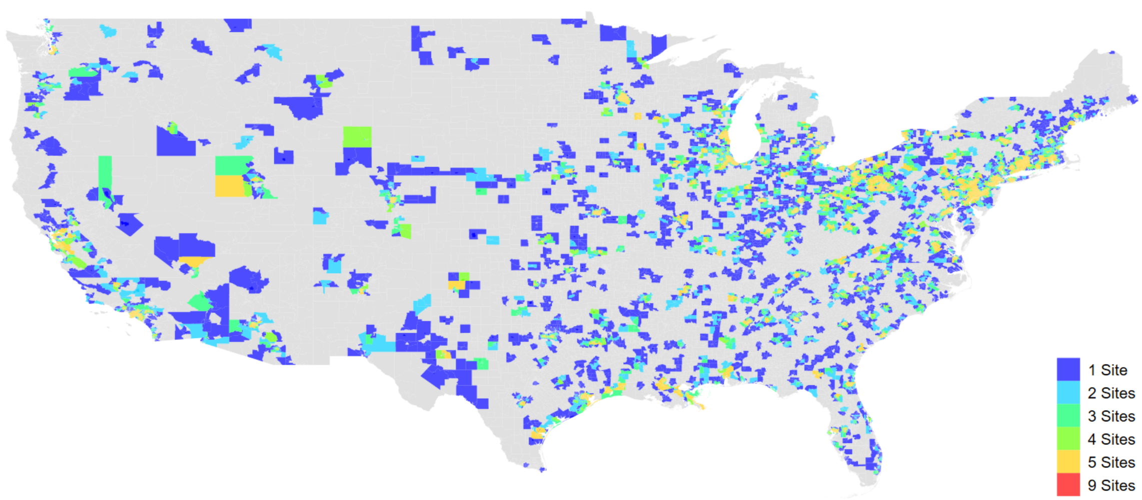

Notes: Panel (a) shows 2010 tracts (shaded blue) containing at least one RCRA site in the Continental USA in our dataset. Panel (b) shows the number of sites in 10 km by tract in the Continental USA. Tracts are only included if they are in the NHGIS 2010 tract-level shapefile, which removes entirely water-based tracts.

for tracts within 10 km of one site). Based on these limitations, we are left with a final sample of Census Tracts near 1,557 RCRA sites (13,604 tracts), out of a total of 3,492 sites that had Census data within 10 km and within the continental US. Table A.1 examines whether these areas are different than the tracts within 10 km of multiple RCRA facilities. Generally, tracts near a single facility seem to be more well-off than those near multiple RCRA facilities.[30] A concern is that this sample restriction reduces the external validity of the results, since doing so removes more than 50 percent of the tracts. Therefore, we include a second more generalizable sample, restricting to tracts that have more than one site within 10 km, but all sites have to either a) be within the tract or b) be within $(0, 10]$ km. This approach incorporates 3,385 of the 3,492 potential sites (and 46,173 of the potential 48,387 tracts within 10 km), excluding only 1,869 tracts that have sites within the tract AND are $(0, 10]$ km away from RCRA sites, making them neither pure treatment nor control cases. We show a map of this "generalized sample" in Figure A.2c.

4 Empirical Models

4.1 Housing Price Impacts

We begin with the following difference-in-differences (DID) strategy to estimate the impact of RCRA cleanups on housing prices:

$$\begin{aligned} Y_{i,t}^{k} = {} & \beta_1 Near_i^d \times Post_{r(i),t} + \beta_2 Near_i^d \times InProg_{r(i),t} + \beta_3 Post_{r(i),t} + \beta_4 InProg_{r(i),t} \\ & + \beta_5 Near_i^d + \lambda_{r(i),t} + \delta_{s(i),t} + \gamma_{b(i),t} + \theta_i + \epsilon_{i,t} \end{aligned} \tag{1}$$

In (1), $Y_{i,t}^{k}$ is the k^{th} decile of the house price (inflation-adjusted to 2023 dollars) in tract i in Census year t, where t =1990, 2000, 2010, or 2020. $Post_{r(i),t}$ is an indicator variable that takes value 1 if the site nearest to the tract (RCRA site $r(i)$) has been cleaned up by year t and 0 otherwise. $InProg_{r(i),t}$ is an indicator for cleanup in progress. It takes value one if the census year overlaps with the time between the first EI entry and the last EI entry recorded for the nearest RCRA facility; see Section 3 for more details. $Near_i^d$ is an indicator variable that takes value one if tract i is within distance $d < D$ km of the RCRA site, where D is the cutoff distance from sites beyond which we do not use observations.[31] Since the data are at the Census tract-level, we calculate distance from each RCRA site to the nearest border of the corresponding census tract. Our main results use $D = 10$, and we only use sites within 10 km of at most one RCRA site. $\lambda_{r(i),t}$ is a set of RCRA site by year effects, and $\delta_{s(i),t}$ is a set of state by year fixed effects. $\gamma_{b(i),t}$ is a set of distance bin by year fixed

[30] Specifically, tracts near a single facility have higher housing values at higher deciles of the price distribution, higher average household income, lower unemployment rates, are less likely to be college-educated, and have lower shares of Hispanic and Black population than those excluded from our sample. In Appendix Figures C.2 through C.7 of the appendix, we provide trends over time in variables of interest, separately for tracts 0 km from a RCRA site, tracts $(0, 10]$ km from a RCRA site, and all tracts in the United States.

[31] In equation (1), we include the variables $Near_i^d$, $InProg_{r(i),t}$, and $Post_{r(i),t}$ for purposes of exposition. However, the corresponding parameters β_3, β_4, and β_5 cannot be separately identified with the inclusion of site by year ($\lambda_{r(i),t}$) and tract fixed effects (θ_i).

effects,[32] θ_i is a set of tract fixed effects, and $\epsilon_{i,t}$ is a (hopefully idiosyncratic) error term.

The coefficient on the interaction between $Near_i^d$ and $Post_{r(i),t}$, β_1, estimates the change in price $Y_{i,t}^k$ after RCRA cleanup for units near the site relative to the same change for those far from the site. This parameter represents the causal impact of cleanup on price under the assumption that the changes in prices of tracts far from (but still within a vicinity of D km of) a RCRA site represent a valid counterfactual for what would have happened to tracts near the site if the nearby RCRA site was not cleaned.

The baseline regressions include tract, bin by year, state by year, and RCRA site by year fixed effects. The tract fixed effects account for idiosyncratic time-invariant features of the tract and net out unobservables that might be correlated with being near a RCRA site. The staggered treatment timing in our context allows us to use bin-by-year fixed effects to net out time-varying unobservables affecting all homes in each distance bin. These fixed effects address the concern that homes closer to and farther away from RCRA sites might not be on parallel trajectories over time; for example, awareness of the harms associated with living near a RCRA site could grow over time nation-wide and could mean that price growth in the near bin lags price growth in the far bin. State-by-year fixed effects allow for time-varying trends at the state-level that coincide with cleanup and affect housing price.

RCRA site by year fixed effects ensure that our effects are estimated off of variation between houses near and far from the same site, rather than relying on variation in when sites are cleaned up. These fixed effects address concerns expressed by the recent staggered adoption DD literature (Gardner (2022); de Chaisemartin and D'Haultfœuille (2022); Wooldridge (2021); Callaway and Sant'Anna (2021); Goodman-Bacon (2021)). This literature is focused on the problem of "bad comparisons"– situations in which already treated units are being used as controls for treated units. Our fixed effects structure guarantees we do not make bad comparisons: within a given site and year, the comparison is between the treated (near) units and the control (far) units, which are never treated. Our setup differs from the setup inherent to the staggered adoption DD, in that instead of having a single set of pure controls that are used for every cohort-specific treatment effect estimate, we effectively have many control groups, each of which is a pure set of controls specific to the treated unit since it is near the same site.[33]

We estimate cleanup impacts at 1 km distance bins to empirically determine the point at which exposure to RCRA sites no longer matter, to inform our choice of control group. Because our measure of distance from a tract is the minimum distance from the nearest facility to the *boundary* of the tract, we break up the $[0, 1)$ km bin into a 0 km bin (for tracts on which the facility resides) and a bin that contains tracts whose boundary is $\in (0, 1]$ km. We specify seven distance bins,

[32] In this specification, there are two distance bins- Near and Far. The number of distance bins is expanded in equation (2).

[33] Guignet and Nolte (2024) and Guignet et al. (2023) also note this. They argue that conceptually, a near-far design with staggered roll-out resembles multiple 2x2 DD estimators and thus could be likened to a stacked DD setup.

indexed by $d = 0, ..., 6$, the last of which captures distances from 5-10 km. From (1), we substitute $Near_i^d$ with a 0 km bin (indicating the facility is on the tract) and 1 km distance bin indicators up to 5 km ($Dist_i^{(d-1,d]}$ for $d = 1, \ldots, 5$):

$$Y_{it}^k = \alpha_1^0 Dist_i^0 \cdot Post_{r(i),t} + \sum_{d=1}^{5} \alpha_1^d \left(Dist_i^{(d-1,d]} \cdot Post_{r(i),t} \right) + \alpha_2^0 Dist_i^0 \cdot InProg_{r(i),t}$$
$$+ \sum_{d=1}^{5} \alpha_2^d \left(Dist_i^{(d-1,d]} \cdot InProg_{r(i),t} \right) + \alpha_3 Post_{r(i),t} + \alpha_4 InProg_{r(i),t} + \lambda_{r(i),t} + \delta_{s(i),t} + \gamma_{b(i),t} + \theta_i + \epsilon_{i,t} \quad (2)$$

Since we exclude the $d = 6$ distance bin indicator from the summation in (2), all effects of cleanup α_1^d are relative to this distance bin, and α_3 can be seen to capture effects of cleanup for this bin. In other words, the coefficients α_1^d for $d = 0, \ldots, 5$, would return the impact of cleaning up a RCRA site located in bin d relative to the impact of cleanup between 5 and 10 kilometers away. Analogous interpretations apply to $InProg_{r(i),t}$ and its interactions.[34]

We define distance from tract to site as the minimum distance, calculated from the point representing the RCRA site to the nearest border of the corresponding census tract. This distance definition is employed since our unit of analysis is a tract. Tracts can be greater than 1 square km in area, so our 0 km bin will naturally contain houses close but not next to the site, and analogously, each of our tract-based distance bins will be nominally smaller than the actual distance from house to site. In standalone appendix section B, we provide summary statistics describing how tract-based distance bins translate to actual distance to site for transaction data from Ohio and Pennsylvania, as well as a thorough discussion of how distances compare. The average actual distance is approximately 2.5 km for our 0 km bin. Distance from tract to site serves as a reasonable proxy for distance from house to site, at least relatively speaking. In appendix D, we justify our choice of distance.

After discussion of the results from the main housing price model, we additionally present robustness checks and alternative price specifications, including an event study specification. The methodology used for each is described alongside corresponding results in Section 5. As an additional robustness check, standalone appendix section B uses individual transaction data from Ohio and Pennsylvania to estimate quantile treatment effects using a recentered influence function regression.

[34] In equation (2), we include the variables $InProg_{r(i),t}$, and $Post_{r(i),t}$ for purposes of exposition. However, the corresponding parameters α_3 and α_4 cannot be separately identified with the inclusion of site by year ($\lambda_{r(i),t}$) fixed effects.

4.2 Neighborhood Composition and Sorting

The preceding property value hedonic model investigates how RCRA cleanups have impacted housing prices at different points in the price distribution. If the price effects vary across the distribution and since pollution is often located in the less desirable neighborhoods within a locality, then remediation has the potential to reverse exposure to such nuisances and decrease gaps in pollution exposure based on socioeconomic status. Of course, the positive distributional impacts may be completely undone by re-sorting in response to cleanup. This would be made more likely if cleanup effects are large enough to trigger endogenous neighborhood change, further altering the composition of a neighborhood. Relatedly, the distributional impacts of cleanup could significantly differ depending on a household's status as an owner or a renter due to capital price appreciation in response to cleanup.

We will undertake two tests for neighborhood change- a reduced-form approach described in Section 4.2.1, and a structural approach described in Section 4.2.2.

4.2.1 Reduced-Form Investigation of Neighborhood Change

We assess the potential for re-sorting using the Census Summary Files data. We first check whether cleanups yielded changes in the composition of residents by changing the dependent variable in our main specification (equation 1) to be one of the following 17 outcomes from the Census:

- Income/Education: average household income, percent below poverty, percent college educated, percent on public assistance, and percent of the population that is unemployed
- Demographic: percent of the population that is Black, percent of homes with a female head of household, percent of the population that is Hispanic, population density, percent of the population that is white, and percent of the population that is under 18 years old
- Housing: percent of homes with four or more bedrooms, percent of homes built in the last ten years, percent that are mobile homes, percent of households that moved in the last ten years, percent of homes that are owner occupied, and percent of homes that are vacant

An advantage of checking for changes in neighborhood composition in this manner is that we can estimate cleanup's impacts at the exact same geographic scale and with the same power as our house price analysis, making the tests comparable.

4.2.2 Structural Sorting Model

A limitation of our reduced form checks on demographic changes is that without knowing the characteristics of the origin and destination of a mover, it is difficult to determine whether a person is actually moving away from or towards pollution. This identification concern was raised by Depro et al. (2015) on using aggregate data to test for sorting behavior.[35] Thus, sorting behavior may be present even without aggregate changes in neighborhood composition. Depro et al. (2015) show that the identification concern can be overcome if one is willing to make additional assumptions on the structure on the problem using an equilibrium sorting model.

Equilibrium sorting models of the housing market infer behavioral parameters (e.g., preferences for amenities) from decisions about where to live (Kuminoff, Smith and Timmins, 2013), taking prices and amenities as given. By explicitly modeling the sorting decisions underlying the hedonic equilibrium, sorting models allow one to make theoretical predictions about aggregate moving behavior that can be matched to data. In this case, the assumed structure of the problem provides additional moment conditions that facilitate identification of sorting behavior. An added benefit is that our sorting model can accommodate stickiness in the moving decision by allowing for moving costs. We apply a structural sorting model to the data to overcome this identification problem, as proposed in Depro et al. (2015). We modify their approach to accommodate our strategy to control for unobserved heterogeneity correlated with pollution.

To test for gentrification, we allow preferences for amenities and the disutility of moving costs to vary by group characteristics that reflect socioeconomic advantage. Any differences in preferences for cleanup would suggest that such an environmental improvement might trigger changes in neighborhood composition, which could then reverse improvements in exposure to RCRA sites. We consider groupings based on two dimensions: race/ethnicity and income. A comprehensive literature on environmental justice, spanning multiple disciplines, investigates disproportionate exposure to pollution based on race and class. We thus follow the extant literature to test whether demand for cleanup varies along these characteristics.

In addition, understanding the benefits of cleanup involves quantifying both changes in exposure as well as price effects. The latter is complicated by differences in home ownership. Renters may become worse off if the increase in rental prices is higher than what they are willing to pay for the environmental improvement. Owners may additionally benefit from asset price appreciation or potentially lose from higher taxes. Property taxes, in particular, are regressive, disproportionately burdening occupants of lower-valued housing (Berry, 2021), which tend to be low-income and minority (Derenoncourt, Kim, Kuhn and Schularick, 2024). Recent work also finds that Black and Hispanic owners pay a higher effective housing tax for the same bundle of public services (Avenancio-León and Howard, 2022). Taken together, this literature suggests that re-sorting in response to RCRA cleanups, whose price effects are concentrated in the lower quantiles of the price distribution, could potentially incentivize owners to leave, without reaping much of the ben-

[35] Sorting behavior is characterized by the tendency to stay or move between locations, i.e. "transition probabilities". For example, if there are 2 locations, then there are 4 values that characterize movement (including the decision to stay in a particular location). Aggregate data by location, however, only reveal how each location's population changed. The identification issue boils down to trying to identify more variables (i.e., the 4 transition probabilities governing movement) with insufficient information (i.e., overall population changes at 2 locations). See Depro et al. (2015) for specific examples of the identification problem.

efits from capital gains.[36] While our model follows the existing literature and effectively treats everyone as renters (Kuminoff et al., 2013), the data include both owners and renters. To clarify the welfare impact of cleanup, we thus separately estimate racial preferences for cleanup for the sample of homeowners only, a group that does not have to contend with higher rental prices due to cleanup. Next, we build a simple model of how people sort into neighborhoods. The baseline model presented below focuses on preference heterogeneity by race. After presenting the baseline, we discuss modifications for the alternative models as well as limitations.

Suppose that an individual (belonging to a racial/ethnic group R), at time period t, observes the characteristics and prices of all locations in that period, and decides whether to move to a different location by time period $t + 1$. Specifically, she chooses whether to live in one of J neighborhoods (characterized by census tracts) within a state, to move out of the state ($J + 1$), or to stay in her current location. For the moment, we suppress individuals' group index R to simplify notation. The individual's preference for tract i follows:

$$U_{i,t} = \delta_{i,t} + \epsilon_{i,t} \tag{3}$$

where $\delta_{i,t}$ represents the average utility that all residents receive from living in tract i at time t; $\epsilon i, t$ is the idiosyncratic utility that the individual receives from locating in i, which is assumed to be distributed Type I Extreme Value. The mean utility, which captures the attractiveness of location i at time t, can be thought of as a quality of life index (e.g., Blomquist, Berger and Hoehn (1988), Kahn (1995), Albouy (2016)) that is determined by the location's attributes ($X_{i,t}$, $\xi_{i,t}$) and the costs of living there ($p_{i,t}$):

$$\delta_{i,t} = \alpha p_{i,t} + X_{i,t}\beta + \xi_{i,t} \tag{4}$$

The location's amenities include ones that are observed by the analyst ($X_{i,t}$), such as exposure to cleaned RCRA sites, and those that are unobserved ($\xi_{i,t}$). The coefficient β on a particular X, e.g., RCRA cleanup, represents the preference for cleanup, where $\beta > 0$ ($\beta < 0$) means that the individual derives positive (negative) utility from cleanup and would sort towards (away from) cleaned locations. Moreover, differences in β by socioeconomic status would reveal differential sorting behavior.[37]

If the individual chooses to move to tract i during a period t (from, e.g., tract k), then she incurs a financial moving cost $MC_{i,k,t}$. We characterize the moving cost as 3 percent of the average housing value of the origin location (tract k) plus 3 percent of the average house value of the

[36] Berry (2021) finds that low-price homes are assessed at higher value compared to their actual sale price.

[37] The amenity vector $X_{i,t}$ may include endogenous amenities for which individuals have preferences, such as congestion (Timmins and Murdock, 2007) or peer effects (Bayer, McMillan and Rueben, 2004). In practice, rather than estimating preferences for such amenities, we include a neighborhood (or tract) fixed effect in estimation (to be described next). To the extent that groups (whether by race or income) have the same disutility of these attributes, the tract fixed effects in the mean utility decomposition should account for such effects while still allowing us to recover *differential* preferences for RCRA cleanup.

destination location (tract i) during time t, following Depro et al. (2015).[38] Since realtor fees are typically applicable to homebuyers rather than renters, we weight moving cost by the share of the tract that is owner-occupied in the baseline model. Individuals who do not move or stay in their tract are assumed to have no moving costs, i.e., $MC_{k,k,t} = 0$. The utility that an individual receives from moving from k to i is:

$$\Delta U_{i,k,t} = \left(\delta_{i,t} - \delta_{k,t}\right) + \left(\epsilon_{i,t} - \epsilon_{k,t}\right) - \mu MC_{i,k,t} \tag{5}$$

where the parameter on the moving cost μ is the disutility of a dollar in moving costs and represents the marginal utility of income. We use this parameter to monetize preference parameters later on (e.g., β/μ).[39]

The individual will choose the location that maximizes her utility. Assuming that the idiosyncratic error term ϵ follows a Type I Extreme Value distribution, the share of people that moves from k to i in the population during time t is characterized by the following logit probability:

$$s_{i,k,t} = \frac{e^{\delta_{i,t} - \delta_{k,t} - \mu MC_{i,k,t}}}{\sum_{\ell} e^{\delta_{\ell,t} - \delta_{k,t} - \mu MC_{\ell,k,t}}} \tag{6}$$

Similarly, the share of people staying in tract k in time t is given by:

$$s_{k,k,t} = \frac{1}{\sum_{\ell} e^{\delta_{\ell,t} - \delta_{k,t} - \mu MC_{\ell,k,t}}} \tag{7}$$

By definition, the population in tract i in $t+1$ is the sum of all people who move to i from each of the $J+1$ neighborhoods during time t. We can therefore use the above shares to relate population counts across time periods t and $t+1$ in the following manner:

$$pop_i^{t+1} = \sum_{k=1}^{J+1} s_{i,k,t} pop_k^t \tag{8}$$

To estimate the preference parameters governing moving decisions, we then use equations 6 through 8 to predict two quantitative measures that are available in the aggregated Census Summary Files and the Census microdata: (1) the total and share of the population in each tract (for each group), and (2) the share of the population that stayed in the current residence (for each group). For consistency, we use the census years $(t, t+1) = (1990, 2000)$, $(t, t+1) = (2000, 2010)$, and $(t, t+1) = (2010, 2020)$, similar to our hedonic model. The following describes the prediction

[38] This assumption reflects the 6% realtor fee convention in the United States over the period we study, with equal incidence between buyer and seller. For more, see: https://www.washingtonpost.com/opinions/2024/03/20/settelement-realtors-commission-six-percent/. By assuming this, we ignore other monetary and non-monetary costs associated with moving, and thus is likely an under-estimate of true moving costs.

[39] Following Depro et al. (2015) and Bayer, McMillan, Murphy and Timmins (2016), this also allows us to circumvent finding an instrument to estimate the coefficient on price α, since price is endogenous.

of these shares:

1. We obtain the total and share of a particular group R, e.g., non-Hispanic Black, for each tract in different census years. We denote the population of group R in year t for tract i as $pop_i^{R,t}$. Dividing both sides of equation 8 by the total population in the region and using group-specific movement shares (i.e., $s_{i,k,t}^R$), we can predict the share of group R living in tract i at time $t+1$ using the time t population shares:

$$\sigma_{R,i}^{t+1} = \sum_{k=1}^{J+1} s_{i,k,t}^R \sigma_{R,k}^t \tag{9}$$

   Here, $\sigma_{R,k}^t$ and $\sigma_{R,i}^{t+1}$ are the share of group R, respectively, living in tract k at time t and tract i at $t+1$.

2. We also obtain the share of the population that stayed in the current residence from Census microdata.[40] Our model predicts the share of people who chose to stay in their time t locations at time $t+1$ using aggregate population counts in each location:

$$\%Stay_R^{t+1} = \frac{\sum_{k=1}^{J+1} s_{k,k,t} pop_{R,k}^t}{totpop_t} \tag{10}$$

With the moving share predictions (equation 9), the stay share predictions (equation 10), and the corresponding estimates from the data, we estimate our parameters of interest using the following two-step procedure:

Step 1 We first solve for the moving cost parameter μ and the vector of mean utilities $\delta_{i,t}$ using a bisection method that nests a Berry (1994) contraction mapping. Specifically, given a guess of μ, we use equation 9 and a guess of the mean utilities at time t, $\delta_{i,t}^{(old)}$, to predict the population shares at $t+1$. We update the vector of mean utilities to be $\delta_{i,t}^{(new)}$ according to the following rule until the vector of mean utilities has converged:

$$\delta_{i,t}^{(new)} = \delta_{i,t}^{(old)} + \log \sigma_i^{t+1} - \log \tilde{\sigma}_i^{t+1} \tag{11}$$

Recall that σ_i is our prediction of population shares (based on a guess of the parameters); we add a "~" to indicate the corresponding shares from the data. We next combine the converged vector $\delta_{i,t}$ and the initial guess of the moving cost parameter μ to predict the share of stayers using equation 10. We then update our guess of μ using a bisection method, solving for the vector of mean utilities

[40] The survey asks in what year the householder (for 1980 on) moved into the dwelling unit (apartment, house, or mobile home). IPUMS then recodes the responses as the number of years ago that the householder moved into the housing unit. Thus, we interpret a household staying in their unit as staying in their tract. In our model, people who move within a tract will not experience a change in mean utility nor will they incur a moving cost $MC_{k,k,t} = 0$.

($\delta_{i,t}$) at each guess of μ. We repeat this process separately for each group R and for each time period, $(t, t+1) = (1990, 2000)$, $(t, t+1) = (2000, 2010)$, and $(t, t+1) = (2010, 2020)$ resulting in group- and time- specific μ and $\delta_{i,t}$ estimates. Lastly, since we limit focus on within-state relocation decisions (but allow individuals to leave the state as a catchall decision), we estimate the model separately for each state.

Step 2 We next stack the mean utility estimates for each group R in each state and each time period, and decompose the mean utility with respect to tract characteristics (interacted with race indicators) to recover preferences. Before doing so, we make the mean utilities comparable across groups, time, and location by dividing the mean utility estimates for a particular group/time/location by the corresponding moving parameter estimate, $\widehat{\delta}^R_{i,t} = \delta^R_{i,t} \,/\, \mu^R_t$.[41] In the final step, we estimate the following specification, similar to our reduced-form model:

$$\begin{aligned}
\widehat{\delta}^R_{i,t} =& \beta_0 + \beta_1 Near^0_i \cdot InProg_{r(i),t} + \beta_2 1[R = B] + \beta_3 1[R = H] + \beta_4 Near^0_i \cdot Post_{r(i),t} && (12)\\
&+ \beta_5 Near^0_i \cdot Post_{r(i),t} \cdot 1[R = B] + \beta_6 Near^0_i \cdot Post_{r(i),t} \cdot 1[R = H] \\
&+ \lambda_{r(i),t} + \delta_{s(i),t} + \gamma_{b(i),t} + \theta_i + \xi_{i,t} && (13)
\end{aligned}$$

where $1[R = \cdot]$ is a group indicator and $Near^0_i$, $Post_{r(i),t}$, and $InProg_{r(i),t}$ are as previously defined in the reduced form model, and $\lambda_{r(i),t}$, $\delta_{s(i),t}$,$\gamma_{b(i),t}$ and θ_i are respectively site-by-year, state-by-year, bin-by-year, and tract fixed effects. The groups we examine in the baseline model are non-Hispanic white ($R = W$), non-Hispanic Black ($R = B$), and Hispanic ($R = H$). The coefficients β_5 and β_6 are, respectively, the Black-white gap and Hispanic-white gap in willingness to pay to live in location i, before and after RCRA cleanup. In this stacked regression, each observation is a particular race in a particular tract, in either 2000, 2010, and 2020.

The mean utility decomposition is set up to be as similar as possible to our reduced form model. Still, important differences exist between our structural and the reduced form models that limit their comparability. The sorting model is best suited to reveal differences in preferences between groups, rather than overall willingness to pay. Furthermore, the reduced form model assumes free mobility and a national housing market, whereas the structural model builds in moving costs and focuses on within-state location decisions.

We repeat the above estimation procedure for the sample of owners. In the baseline model, moving costs are weighted by the share of a tract that are owners. We do not apply these weights when re-estimating the model for the sample of owners. In addition to differential preferences by race, we re-estimate the model to allow for heterogeneity in preferences by income. We categorize

[41] While the moving cost parameter is positive in most instances (i.e., people dislike higher moving costs), a small fraction (< 1%) of the MC parameters is negative. In these cases, we use the absolute value of the μ^R_t in order to preserve the signs of the mean utilities.

individuals as high versus low income (relative to the individual's state median income) based on income as observed in the year that individuals are surveyed in the public-use Census microdata.

The following limitations are important to keep in mind for interpretation. First, at least part of the heterogeneity by race that we estimate stems from differences in income. Variation in preferences may also reflect housing market dynamics that differentially affect racial minorities. Recent literature has found evidence that minorities are steered towards less desirable neighborhoods and that ignoring these differential constraints will attenuate demand estimates for these group (Christensen and Timmins, 2018; Christensen, Sarmiento-Barbieri and Timmins, 2020). In our context, this may be accounted for to some extent by allowing differential switching costs and group-specific indicators in the mean utility decomposition. But what may seem like 'low preference' for environmental quality for racial minorities may partially be driven by more choice constraints. Second, for the models that examine preferences separately for owners, we abstract from the rent/own decision since we do not observe individuals over time in the data and, thus, any changes in homeownership status. Therefore, caution should be exercised in interpreting these results since we have to assume that those who own stay as owners. Similarly, for our results by income, we must assume individual income types are fixed even though income is endogenous, a limitation that is again due to our inability to track individuals over time. Finally, for all models, we focus on within-state location decisions and abstract from labor market considerations in residential choice (Roback, 1982).[42]

5 Results

We organize our results as follows. In Section 5.1, we discuss the results of our main specification. In Section 5.2, we present event studies to check for pre-trends driving our main results. In Section 5.3, we present and discuss results from our reduced-form and structural investigations of sorting.

5.1 Impacts on Housing Prices

We test the specifications proposed in equations (1) and (2) on all nine deciles of the price distribution and from within 10 km away from a RCRA facility. To start, we employ a flexible exposure buffer specification (2) to see how far the treatment effects might extend. We are interested in the coefficients on the interaction effects between distance bins and the post-cleanup indicator. We excluded the 5–10 km bin, so all of our estimates can be interpreted as the differential effect of cleanup on homes in a particular distance bin and homes in the 5–10 km bin. Figure 3 presents the point estimates for the interaction between various distance bins and the post-cleanup indicator. We find that the 0 km bin stands out as having large and significant impacts for most price

[42] In the empirical model, we allow for a 'catch-all' location choice that is associated with moving outside of a state. Practically, we do not infer preferences from this choice and use it to avoid dealing with an open migration system. For details, see Depro et al. (2015).

Figure 3: Housing Price - Distance Gradient

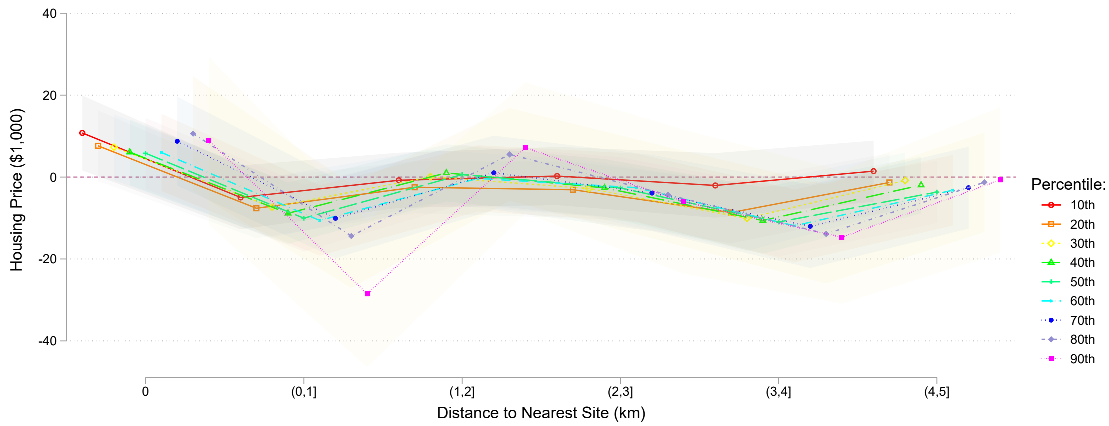

Notes: This figure plots the set of Difference-in-Differences (DD) coefficients on the interaction $Dist_i^0 \cdot Post_{r(i),t}$ and $Dist_i^{(d-1,d]} \cdot Post_{r(i),t}$ from 2 for each percentile from the 10^{th} through the 90^{th}. All specifications include fixed effects for tract, bin by year, site by year, and state by year. The sample is the set of tracts containing exactly one RCRA site within 10km in 1990, 2000, 2010 (ACS 2008-2012), and 2020 (ACS 2018-2022).

deciles. After 0 km, the effect on housing prices dips below 0 and is insignificant for most of the price percentiles, indicating that this is where our treatment effect ends. Therefore, we designate 0 km as our definition of the 'near' (treated) bin in equation (1), and use $(0, 10]$ km as the control group.[43]

Summary statistics by whether the tract was in the 0 km bin can be found in Table A.2 in the Appendix. Housing prices and income tend to be higher in the 0 km bin, but those living in the 0 km bin are more likely to be Black, less likely to be college graduates, and are more likely to be below the poverty line or on public assistance.

Results from a more parsimonious specification, following equation (1), are presented in Table 2, where we have grouped all homes (0–10] km away from the facility into one group so that we have just two distance bins. We also plot the change from before to after cleanup in the 0 km distance bin in Figure 4. The impacts on the price percentiles range from \$7,828 to \$10,881 depending on the percentile, and are statistically significant at the 10% level, except for the impact on the 90^{th} percentile of the price distribution.

Although the magnitudes appear to be relatively similar in levels, in percentage terms, the impacts are stronger for the lower percentiles of the price distribution. In the bottom panel of Table 2, we present the percentage impacts by dividing the estimates by their sample mean. [44] We see that

[43]We note that negative estimates in the $(0, 1]$ km bin could reflect SUTVA violations; these are particularly prominent for the higher deciles of the price distribution. Therefore, we also provide a robustness check in Table A.8 that uses 5–10 km as the control group.

[44]Another way to present percentage effects would be to compute the average of $\frac{\hat{\beta}}{\hat{P}^k}$, where $\hat{P}^k$ is the predicted price percentile from each regression. Typically, this mean is estimated after the regression is run, and standard errors are

Figure 4: Housing Price Impacts at 0 km bin by decile from Near-Far Comparison

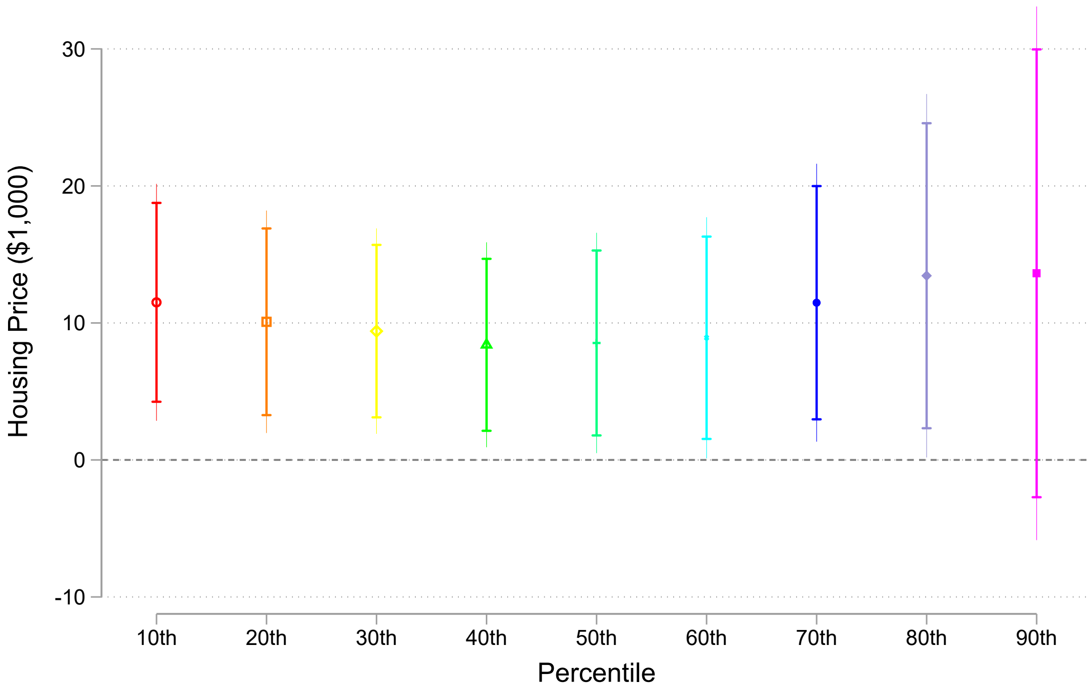

Notes: This figure plots the DD coefficients by decile of housing price impact from Table 2, which correspond to the coefficient on the interaction $Near_i^d \cdot Post_{r(i),t}$ from (1), where $Near_i^d$ is the indicator for the 0 km bin. All specifications include fixed effects for tract, bin by year, site by year, and state by year. The sample is the set of tracts containing exactly one RCRA site within 10km in 1990, 2000, 2010 (ACS 2008-2012), and 2020 (ACS 2018-2022). Standard errors are clustered at the tract level. Whiskers marked with horizontal lines and vertical protruding segments indicate 95% and 99% CIs respectively.

the effects are much larger for the lower deciles of the price distribution. For example, the relative impact on the 10^{th} percentile price for the 0 km bin versus the (0–10] km bin is around 10.1%, but for the 90^{th} percentile price specification, the percentage difference is only approximately 3.0%. For the most part, the percentage impacts are not different in a statistically significant sense. Table A.3 presents p-values from the null hypothesis that each decile has the same impact as the others in percentage terms. The p-values are produced using stacked regression. The 10th percentile of price is higher than the 40th through 70th percentile impacts in percentage terms, but not the the 80th and 90th percentiles. This may be because the 80th and 90th deciles are so imprecisely

obtained using the Delta Method. We found this method to be slightly fragile for a handful of robustness checks where we added linear controls interacted with the high dimensional fixed effects. In particular, when the terms including linear controls interacted with fixed effects were collinear with fixed effects, some nuisance parameters were exactly identified. These parameters are not important to us, and were partialled out when estimating our primary effects of interest, but were required to calculate the Delta Method standard errors, resulting in no standard error reported. Therefore, we elected to stick with a simple transformation of dividing the estimate by its sample mean ($\frac{\hat{\beta}}{\bar{P}}$). The difference between the two methods was not apparent to the third decimal place.

estimated that we cannot reject the null of similar effect sizes.

We do not take a stand a priori on which price deciles will experience effects of cleanup. However, we could see a hypothesis that price is impacted by cleanup as actually constituting a family of 9 hypotheses. The probability of rejecting at least one true null hypothesis in a family of hypotheses is known as the familywise error rate (FWER), and grows with the number of hypotheses in the family (Hochberg and Tamhane, 1987). To address this concern, we show *p*-values for a variety of correction methods in Table E.1 of the appendix for our main estimates, with corresponding discussion in Appendix Section E. The p-value for the 10th percentile is less than 0.1 for all multiple hypothesis testing (MHT) correction methods, and is less than 0.05 for both of the resampling-based methods, which gives us confidence that our results are not merely the product of testing nine hypotheses.

It is ideal to cluster at the level of treatment (Bertrand, Duflo and Mullainathan, 2004; Abadie, Athey, Imbens and Wooldridge, 2022), but the level of treatment is ambiguous in our context– census tracts are varying distances away from RCRA sites, and thus the impacts from before to after could theoretically differ by these distances. Hence, in our main specifications, we cluster on the tract level. However, because RCRA sites are cleaned up at different times, it could be argued that site or site by year is the level of treatment. Furthermore, spatial correlation between tracts in the same county could also be important. Given these considerations, we provide results with alternative clustering levels in Appendix Figure A.3. The primary finding that the 10th percentile of price is statistically significant is preserved across alternative clustering levels.

We present the high-dimensional fixed effects Poisson estimator (Correia et al., 2021) in Table 3 that directly recovers effects as semi-elasticities.[45] The clustered standard errors are fully robust to mis-specifications of functional form, and do not make any assumption that the mean equals the variance, as with traditional Poisson regression. The top panel of the table presents the Poisson parameters, which are the semi-elasticities. The middle panel of the table presents the corresponding marginal effects in levels. Comparing to our main estimates, we see that the estimated marginal effects are similar.

The primary analysis requires both distinct treatment and control groups and well-defined cleanup timing for the event study framework. This results in only about half of sites being used in our analysis. As a check that our results are likely to generalize, we also present estimates from a sample maintaining only the treatment-control distinction. That is, we define treatment tracts as those containing at least one RCRA site with all sites within 10 km contained in the tract, and control tracts as those with no contained sites but at least one site within 10 km (See Appendix

[45]While the estimates in our Poisson model represent the log difference in expected housing prices between houses in the 0 km and $(0, 10]$ km bins (which approximates a percentage change for small values), we will refer to these coefficients as semi-elasticities, though technically this term is more applicable to marginal changes in continuous variables.

Figure A.2c).[46] The results are shown in Table 4. The magnitudes of the coefficients are smaller than those found in Table 2, which is to be expected: effects should be attenuated since the timing used to categorize tracts into "Post", "In Prog" or "Pre" is only the timing for the nearest site, but the tract might be in the vicinity of multiple sites that had different timing.

One concern with our approach is that we are using data that are surveyed in spatial units that are changed each decennial year and require interpolation. In our main result, we use NHGIS to interpolate to 2010 boundaries using owner-occupied weights. In Appendix C, we show that our results are robust to three other interpolation weights and to interpolating to 2000 boundaries (with distance to site calculated also using 2000 boundaries). See Figure C.1 for summary of these results and Table C.1 for details.

One disadvantage of using minimum distance from any point in the tract to the nearest RCRA site is that larger tracts are more likely to contain RCRA sites and hence are more likely to fall in the "0 km" bin; see Appendix Section D for more. This motivates a robustness check employing a regression adjustment using tract size. To construct our regression adjustment specification, we center the log of tract area around the Near group mean, and then we add to the regression the interaction between the centered log of tract area with fixed effects.[47] The result is a fully saturated model, as recommended by Wooldridge (2010). Table 5 presents the coefficients on 0 km $\times$ Post, 0 km $\times$ In Prog and each of those interacted with centered log of area. The coefficients on 0 km $\times$ Post and 0 km $\times$ In Prog are very similar to our main effects (Table 2). The coefficients on the interactions between centered log of tract area and the main variables of interest are all statistically insignificant. We also show that our results are robust to using a population-weighted centroid measure of distance in Appendix Section D.

We briefly discuss other checks on robustness. Results are presented in Appendix A.1. First, standard hedonic models often include house and neighborhood characteristics (Bishop et al., 2020), so we include various house and neighborhood characteristics as additional controls in Tables A.4 through A.7. Results are consistent with those from our main specification in magnitude and significance. Next, our primary control group uses $(0, 10]$ km include as many sites as possible. To ensure that SUTVA violations do not bias our findings, we additionally provide a robustness check where the control group is $(5, 10]$ km in Table A.8. Results are robust.

Finally, we discuss how Guignet and Nolte (2024), which complements our work, differs. There are three important differences. First, our research questions are different: their study focuses on average welfare impacts of the cleanups, whereas we characterize effects across the price distribution and thoroughly investigate post-cleanup neighborhood change and sorting. Second, they focus only on housing price responses to cleanup of the most hazardous sites (known as Treat-

[46] Not all sites will be used by the regression to produce our estimates because of singletons (Correia, 2016); however, this is not the result of deliberate sample selection, as our regression is run on the entire sample.

[47] The exception is the tract FE, which already absorb tract-level area, and hence would be collinear.

ment, Storage, and Disposal Facilities (TSDFs)), whereas we also study cleanups at non-TSDFs.[48] Third, Guignet and Nolte (2024) employ Zillow ZTrax data on individual transactions collected from county assessors, whereas we use publicly available Census data to take advantage of abundant demographic information that is more suitable to our research questions.

Given these differences, it is still worth addressing how our results compare and explaining how our different data sources and identification strategies might affect the precision of our estimates. Guignet and Nolte (2024) highlight the advantages of individual transaction data over aggregated data. However, their sample size drops from 28,312 identifying transactions to just 829 in their house fixed effects sample, indicating that variation in their study is not coming from repeated sales of individual houses, but rather from sales of *different* houses in the same tract in different years. So, even if their lefthand-side variable differs at the individual sale level, the identifying variation still is within tract over time- identical to ours.

Moreover, our main results are based off of aggregated data, which means we lack the fine spatial and temporal resolution of individual sales data that Guignet and Nolte (2024) exploit in their study. There are often two concerns with aggregated data: aggregation bias and loss of precision. Aggregation bias would attenuate our estimates towards zero, and we would expect large confidence intervals due to loss of precision. However, we still detect economically and statistically significant effects comparable to those of Guignet and Nolte (2024). While the comparison is hard to interpret since we study deciles of the price distribution and Guignet and Nolte (2024) study average effects, we are doubtful that aggregation bias is a major problem in our context. For a concrete comparison, we find housing price impacts at the median of about 3.2%, whereas their cleanup effects are 6-7%; that our results are lower is unsurprising given their focus on TSDFs (a subset of RCRA sites expected to experience larger impacts).[49]

Considering that (a) their primary source of variation is at the tract level, and (b) we find economically and statistically significant impacts similar to theirs, it does not appear that the use of aggregated data is a major drawback of our study. Furthermore, an even broader point applies- with Zillow no longer available to researchers, it is useful to show that results like those of Guignet and Nolte (2024) replicate using publicly-available data. That our study finds similar effects to

[48] Guignet and Nolte (2024) study 2,389 TSDF facilities, 689 of which have a corrective action and 195 of which have completed the corrective action. Our dataset, by contrast, starts with 3,873 sites in the cleanup program (there are a total of 3,983 sites in the program; however, this list contains sites in Puerto Rico, Guam, and the Virgin Islands, as well as sites with invalid or duplicate coordinates which were removed). After geocoding to the geographic areas covered by the NHGIS shapefile of 2010 Census tracts, we have a total of 3,492 sites represented in our dataset. Of those, 2,770 are TSDFs. The TSDF sites in Guignet and Nolte (2024) may not be a strict subset of those we study, especially in our main specification, where we restrict to tracts within 10 km of exactly one site to ensure that we have a clean definition of cleanup timing to allow for diagnostics using an event study framework.

[49] The estimates in Guignet and Nolte (2024) are very similar to our finding for the 10th percentile of a 7.6% rise in housing prices post-cleanup. This is potentially explained by the fact that Guignet and Nolte (2024) study houses in closer proximity to sites (within 750 meters), and show that home values are much lower near these RCRA sites, even when conditioning on tract fixed effects, indicating that these housing values are low compared to other places within the same tract.

Guignet and Nolte (2024) is further evidence that using publicly available data from Census summary files can detect similar effects as individual-level transaction data, a finding documented in Gamper-Rabindran and Timmins (2013). It has also been shown that effects at the median may not capitalize the impact of neighborhood improvements if the effects are local and sources of change are concentrated in certain types of areas within a tract (e.g., the less desirable neighborhoods), a point shown by Gamper-Rabindran et al. (2011) in the context of Superfund housing market impacts.

We additionally include a standalone supplementary appendix (Appendix B) which shows robustness of our conclusions to using individual-level transaction data from Ohio and Pennsylvania. Section B.1 of that appendix, previously discussed with our data source descriptions, investigates the extent to which people over-estimate the value of their house more at different portions of the price distribution. Section B.2 puts our effects into context by describing how Census tract-based treatment translates to distance from a given house to a RCRA site and compares effects using both measures of exposure. We reproduce our main results for the sample of Ohio and Pennsylvania Tracts in section B.3. Section B.4.1 describes our quantile treatment effects model applied to transactions data and assumptions. Section B.4.2 then shows the quantile treatment effect results. The broad conclusion is that our main results are robust for the lowest deciles of the price distribution, but not the higher deciles. We find the strongest evidence of a positive cleanup effect for the first decile, both in economic and statistical terms.

Overall, we document robust evidence of capitalization of RCRA cleanups into housing prices, especially at the lower deciles of the price distribution. This could indicate either that citizens are aware of and directly value the cleanups, or that they value the redevelopments and other aspects of area revitalization that are sometimes bundled with the cleanups. No matter which is the case, the impacts we document here are noteworthy given the vast scope and expense of the RCRA cleanup program– the program has provided $97.3 million in federal grant funding to state governments (EPA, 2025).

5.2 Event Study of Impacts on Housing Prices

One potential threat to identification is differential pre-trends between the houses closest to the RCRA sites and those further away. As suggestive evidence that differential pre-trends do not drive our results, we produce an event-study graph that depicts treatment effects over time. That is, we graph the coefficients for the 10^{th} percentile price from the following regression, for the 0 km bin, treating the (0, 10] km distance bin as a control group:

$$Y_{i,t}^{k} = \sum_{\tau \neq -2} \beta_{1\tau} Near_i^0 \times \mathbb{1}\left\{t \in [\tau, \tau+2)\right\} + \sum_{\tau} \beta_{2\tau} \mathbb{1}\left\{t \in [\tau, \tau+2)\right\} + \lambda_{r(i),t} + \delta_{s(i),t} + \gamma_{b(i),t} + \theta_i + \epsilon_{i,t} \quad (14)$$

We focus mainly on the 10th percentile since that was the decile for which we found the strongest effect, in percentage terms. In the above, Y_{it}^k is the house price (kth percentile), $\mathbb{1}\{t \in [\tau, \tau+2)\}$ is a dummy variable that takes value 1 if the Census year t is between τ and $\tau+2$ years relative to the cleanup period of the nearest site and 0 otherwise, and $Near_i$ is a dummy variable that takes value 1 if the facility is on tract i. In a slight abuse of notation, we treat the in progress period from (1) as an event time occurring between time -1 and time 0, which may not be 2 years in length, and we omit it from the above equation for brevity. Excluding one event time in the second summation scales the treatment effect in the two years just prior to cleanup to 0 for ease of interpretation. Note that while equation 14 includes event time dummies, these will be absorbed by our set of fixed effects in estimation.

The interpretations of parameters in our event study differ slightly from the standard event study because the far bin is able to serve as a control group for the near bin in every event time due to the inclusion of a full set of event time indicators $\mathbb{1}\{t \in [\tau, \tau+2)\}$. As such, $\beta_{2\tau}$ is interpreted as the housing price for homes in the far bin during event time $[\tau, \tau+2)$, net of bin-by-census year, state-by-census year, and tract-level averages. $\beta_{1\tau}$ is interpreted as the difference in the housing price for homes in the near and far bins during event time $[\tau, \tau+2)$, net of bin-by-census year, state-by-census year, and tract-level averages.

It is worth emphasizing that the setup of our event study takes advantage of variation in when sites were cleaned up relative to census years. This is similar to the setup in Asker, Brunner and Ross (2022). We could alternatively measure event time in decades, but that approach fails to leverage the rich data we have on cleanup timing. Our setup can be seen as an unbalanced panel in event time– we only have 4 observations for each site. In Figure A.9 of the appendix, we plot the number of sites represented in each event time in our regressions.

Figure 5 is our event study graph and plots $\beta_{1\tau}$ over time. The figure shows that home prices for the first decile of the distribution increase immediately following cleanup, peaking between 2 and 4 years after cleanup, and subsequently decrease but do not reach their pre-cleanup levels within 10 years. Confidence intervals exclude zero for all post-period event times, confirming the statistically significant result in the first column of Table 2. We also plot $\beta_{1\tau}$ over time for other percentiles of the housing price distribution in Figures A.4 in the Appendix and observe a similar pattern through the 40th percentile (although we note that confidence intervals include zero in a minority of post-periods for the 20th-40th percentiles). Effects in the 50th through 90th percentiles appear to be potentially driven by a pre-trend. We thus expect their coefficients to be over-estimated.

Figure 5: Event Study for the 10^{th} Percentile of the Housing Price Distribution

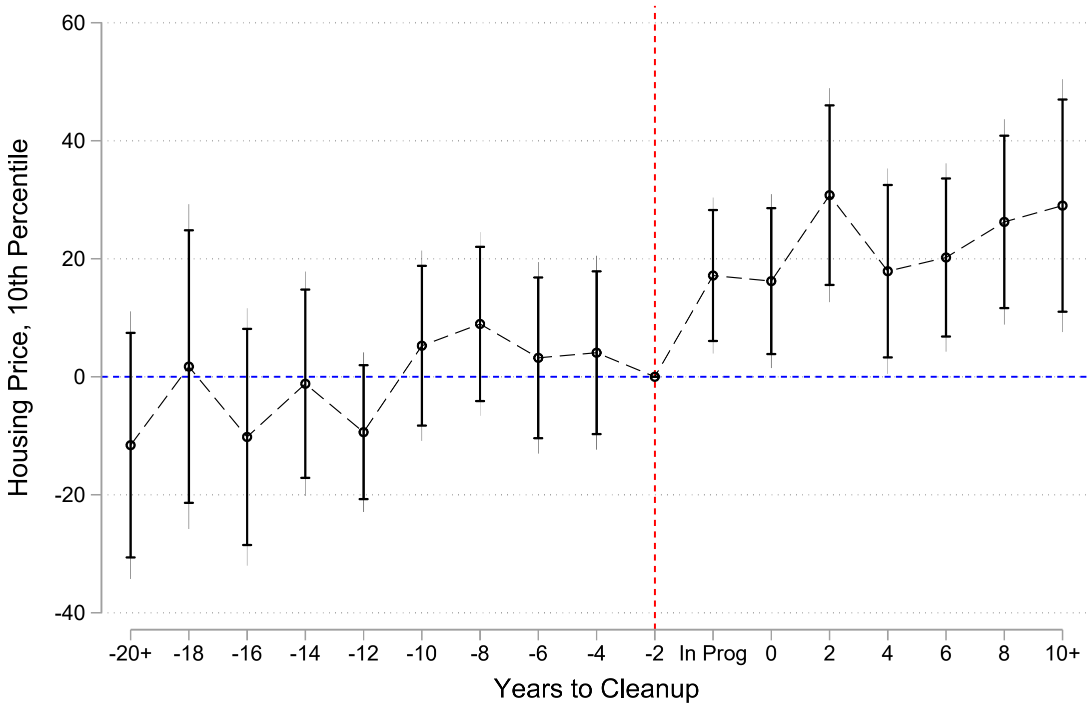

Notes: This figure shows the coefficient representing the difference in the near (0 km) and far ((0, 10] km) bins over time from the event study specification in equation 14. We use the same fixed effects and sample as in the main regression. The coefficient for the two years just prior to the cleanup (at position -2) is normalized to 0 by excluding the dummy on Near× Event time=-2 from the regression. Data from the during-cleanup phase is represented by "In Prog," and "0" represents the two-year immediately following cleanup completion. Whiskers marked with horizontal lines and vertical protruding segments indicate 95% and 99% CIs respectively, clustering at the tract level. Prices are denominated in thousands of dollars.

5.3 Impacts on Neighborhood Composition, Rental Price, and Sorting

Next, we test whether RCRA cleanups impacted various socio-economic and housing-related indicators from the Census in a reduced-form framework. Specifically, we examine 17 other outcomes, and find only weakly statistically significant evidence for one of them ($p < .10$, for percent college educated; 3% increase). In Table 6a, we explore impacts on five income and education-related outcomes: average household income, percent below poverty, percent college educated, percent on public assistance, and percent of the population that is unemployed. The impacts are neither statistically nor economically significant. In Table 6b, we explore impacts on six demographic outcomes: percent of the population that is Black, percent of homes with a female head of household, percent of the population that is Hispanic, population density, percent of the population that is White, and percent of the population that is under 18 years old. No impacts are statistically significant, with none reaching even 1% of baseline. In Table 6c, we explore impacts on six housing-related outcomes: percent of homes with four or more bedrooms, percent of homes built in the last ten years, percent that are mobile homes, percent of households that moved in the last ten years, percent of homes that are owner occupied, and percent of homes that are vacant. The impacts are not statistically or economically significant. This is particularly surprising given that we would expect to find statistically significant effects at the 10 percent level for one out of every 10 outcomes studied, even if no true impacts existed.

In Appendix Tables C.2-C.5, we show that these results are robust to four different interpolation approaches, including interpolating to 2000 tract boundaries.[50]

One might be concerned that socio-economic and housing-related indicators would not respond immediately to RCRA cleanups. Households might sort on distance to RCRA facilities with a lag because of moving frictions, even if prices adjust immediately, which would bias us towards finding no sorting even when sorting was indeed happening. To visualize the timing of potential impacts, we also produce event-study graphs similar to those we make for housing impacts in Figures A.5 through A.7. We see no clear evidence of lagged effects. While we detect price effects from cleanup, the data suggest that they do not seem large enough to trigger moving.

One trend of note is the upward trend in Average HH Income (and possibly % college educated), which at first glance might seem to point to a gentrification story. On this point, we first note that in column 1 of Table 6a, we cannot reject the null hypothesis of no effect, so we do not find any evidence of gentrification. Even more importantly, the magnitude of the point estimate is only \$92, when the average income is \$101,218 (2023\$). This is an economically insignificant magnitude, so we do not find evidence that gentrification is driving our results. And as described above, the increase in % college educated is a small 3% increase that is not robust across interpolation approaches.

Next, we investigate whether rental prices are affected by clean-ups. In Table A.9, we present equation (1) with deciles of rental prices and find no statistically or economically meaningful impacts on rental prices following RCRA clean-up. In Figure A.8, we show similar event studies for each decile and again find no evidence that clean-up impacts rental prices.

Despite the reduced form exploration of neighborhood change showing no evidence of gentrification or sorting, we are sympathetic to the view that a reduced form approach might not uncover these effects. Thus, in light of the identification concerns raised in Depro et al. (2015), we also test for differential sorting by race using a structural sorting model. Table 7 presents our differential willingness to pay (WTP) estimates from the mean utility decomposition in equation 12, where the outcome variable is the monetized value of the mean utility of a tract for each group (race or income) for a given year (2000, 2010, 2020). We estimate the decomposition to recover race-

[50]Table C.2 shows a marginally statistically significant increase in mobile homes; Table C.3 finds a similar result for college education as Table 6 and a statistically significant reduction in % under 18 (3%); Tables C.4-C.5 show no statistically significant changes. Table C.5 uses universe defined weights that are specific to each variable universe and likely provides the most appropriate weights for this exercise.

specific cleanup preferences for all individuals (column 1) and for owners (column 2), and also by income group (column 3). We limit to tracts containing exactly one RCRA site within 10 km, like the main analysis. Since the sorting model estimates preferences using changes across decades, the specifications recover preferences for years 2000 (from 1990 to 2000 moves), 2010 (from moves between 2000 and 2010), and 2020 (from moves between 2010 and 2020). All specifications include site-by-year, state-by-year, bin-by-year, and tract fixed effects, similar to our reduced form specifications.

For the baseline model for all individuals regardless of home ownership status, the point estimates suggest that the WTP for cleanup of Black and Hispanic households are, respectively, \$116 and \$120 lower than that of white households, but the differences are not statistically significant. We next examine the WTP differential by race for homeowners. Column 2 of Table 7 finds stronger evidence that Black and Hispanic owners are willing to pay less than white owners, with a WTP that is \$302 lower for Black owners and \$295 lower for Hispanic owners. That race-based sorting is augmented for the owner group is consistent with large differences in US housing values by race (Derenoncourt et al., 2024). Considering the reduced form evidence showing small and statistically insignificant changes in race shares in response to cleanup (Table 6b), this suggests that the magnitudes of the differences in preferences are not large enough to trigger re-sorting.

The average WTP may be underscored by significant heterogeneity across geography. We investigate the WTP differentials for owners by state and plot the estimates against the state's median income in year 2000. Figures 6a and 6b present these estimates for Black and Hispanic groups, respectively. While in some states, minority WTP is significantly higher than their white counterparts, in other cases the WTP is lower. In addition, it does not seem to be the case that income is the underlying driver of heterogeneity in WTP across geography: The slope coefficient of a fitted line for Black households is 0.0011 and for Hispanic households is -0.0145. To test for income-based sorting directly, we allow preferences for cleanup to vary by income (Column 3 of Table 7). While higher income individuals are willing to pay \$37 more for cleanup than lower income individuals, the WTP differential is not statistically significant.

Taken together, our results suggest limited evidence of income-based sorting for cleanup. There is some evidence that minorities have lower demand for cleanup, when focusing on owners. In light of the limited evidence of income-based sorting, this heterogeneity may reflect non-market, race-specific constraints rather than preference. For all individuals, however, the model still finds limited evidence of differential sorting by race. These results are broadly consistent with our previous tests that cleanup would not change the racial composition of a neighborhood.

6 Discussion

Our empirical findings yield two broad conclusions. First, we found housing price impacts that were limited to the tracts nearest RCRA sites and strongest in relative terms for the lower

Figure 6: Heterogeneity in WTP by State Median Income for Owners

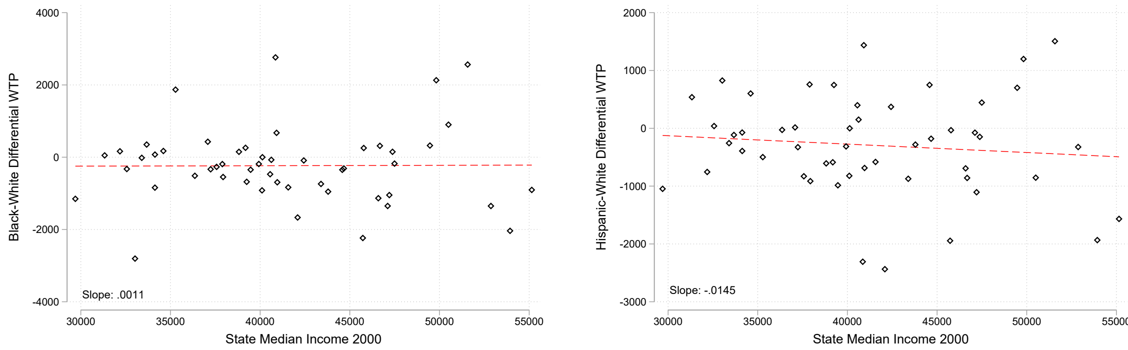

(a) Black-White Gap (b) Hispanic-White Gap

Notes: This figure presents state-level estimates of the (a) Black-white and (b) Hispanic-white differential in WTP from our structural sorting model for a sample of homeowners living in tracts containing exactly 1 RCRA site within 10km. We re-estimate equation 12 to include interactions between state indicators and the main interaction of interest $Near \times Post \times Black$ or $Near \times Post \times Hispanic$ and plot the coefficients on the price differential for each state against the respective state's median income in the year 2000. The dashed, red line indicate a best-fit line from a linear regression of the point estimate on state median income, with the corresponding slope shown in the bottom left corner.

deciles of the price distribution. This finding that home prices increased after cleanup is most prominent and robust for the 10th percentile. It is not driven by a pre-trend. It is also robust to using transaction-level data to estimate quantile treatment effects.

Several mechanisms could underlie these price impacts, such as improved neighborhood reputation (Messer, Schulze, Hackett, Cameron and McClelland, 2006) and endogenous changes in non-targeted amenities (Banzhaf and Walsh, 2013). The welfare implications also depend on whether individuals are owners or renters. The price effects we document could cause renters to become worse off if such price increases are larger than what renters are willingness to pay for cleanup. Unlike housing prices, however, we find minimal impacts on rental prices.

Second, we found that RCRA cleanups were unlikely to cause residents to re-sort. Two approaches were used- we found no evidence of neighborhood change using a reduced-form model, and we found no evidence of a differential tendency to sort by income and limited evidence of sorting by race using a structural sorting model. Evidence of race-based sorting is predominantly from the group of homeowners. Given that we find limited evidence of income-based sorting, sorting by race may possibly be driven by differential gains in homeownership between racial groups (Derenoncourt et al., 2024; Avenancio-León and Howard, 2019) or due to differential market frictions faced by minority households (Christensen and Timmins, 2018; Christensen et al., 2020). Combined with the absence of neighborhood change from the reduced-form models, the magnitude of the race-based sorting we find is unlikely to drive neighborhood turnover.

Taken together, our findings imply that the benefits of cleanup accrued to residents who already lived near the sites, and helped the poorer segment of the housing market more than the richer segment. Our findings are inconsistent with a gentrification story in which more well-off citizens move closer to the sites after cleanup, changing the composition of the population of the tracts on which the sites are located. One possible explanation for the finding of no sorting along these socio-demographic dimensions is that the RCRA cleanups were not a large enough shock, relative to moving costs, to induce increased moving (Palmquist, 1992). This is corroborated by the finding of no effect on the percent of households who moved in the last 10 years in our reduced-form investigation of neighborhood change.

The limited evidence on sorting is a hopeful one in light of recent work. Hausman and Stolper (2020) show that when there is partial information in the housing market, people undervalue a clean environment, and households sort according to their willingness to pay for a clean environment on this partial information, and deadweight loss due to pollution is higher for low-income households. If people do not sort after RCRA cleanups, then these cleanups theoretically could mitigate any pre-existing exposure disparities between the rich and poor that stem from the channels that Hausman and Stolper (2020) pinpoint (partial information and under-valuation of a clean environment). Bakkensen and Ma (2020) present the case where well-meaning policies can cause significant sorting that exacerbates pre-existing disparities in exposure to an environmental bad between advantaged and disadvantaged groups. Our finding that there is no evidence of sorting shows that this need not always be the case.

7 Conclusion

This paper evaluates the housing market impacts of cleanups conducted under the Resource Conservation and Recovery Act (RCRA), an expansive hazardous waste cleanup program. We find that the positive environmental impacts from RCRA cleanups are reflected in the housing market. The price increases that we find are driven by cleanups concentrated among the lowest price deciles of the census tract in which the RCRA facility is located: prices increase by between 7.6-9.8% for the 1^{st} decile of the price distribution, and we detect no evidence of a price increase for the 9^{th} decile. This indicates cleanups raise housing values for those occupying the lowest-value homes, that tend to be poorer and minority (Derenoncourt et al., 2024), which are likely to face other disadvantageous circumstances in life and are typically more vulnerable to the deleterious effects of pollution (see, e.g. Apelberg, Buckley and White, 2005).

Furthermore, we find that the benefits of cleanups accrued to those living closest to the sites and, notably, do not find that cleanups induced re-sorting. This is consistent with the localized price impacts that we find, but somewhat surprising given how expansive RCRA cleanups were and the recent literature that has highlighted the potential for policies to worsen underlying inequities (Hausman and Stolper, 2020; Bakkensen and Ma, 2020). Our results point to a hopeful conclusion-cleanup can help those living nearest to nuisances.

A fruitful direction for future work would be to further explore the attributes of both environmental policies and housing market conditions that explain such neighborhood dynamics. Further research could also examine the relationships between policy scope, information provision, and compositional changes in housing markets to determine the distribution of benefits across various groups from different types of policies.

8 Tables

Table 1: Attributes by Whether There is a RCRA Site within 10 km

	≥ 1 Site within 10 km		No Sites within 10 km			
Attribute	Mean	St. Dev.	Mean	St. Dev.	Δ Mean	*t*-statistic
Housing Price, 10th Percentile	156.703	104.196	124.918	103.179	31.784	41.202
Housing Price, 20th Percentile	187.289	120.381	160.670	120.141	26.619	29.657
Housing Price, 30th Percentile	210.204	133.129	187.951	133.632	22.252	22.395
Housing Price, 40th Percentile	231.064	145.266	213.141	146.633	17.923	16.544
Housing Price, 50th Percentile	252.183	157.901	239.142	160.408	13.040	11.088
Housing Price, 60th Percentile	275.823	172.644	268.516	176.463	7.307	5.698
Housing Price, 70th Percentile	303.854	190.521	304.761	196.742	-0.907	-0.641
Housing Price, 80th Percentile	341.640	215.175	355.619	226.283	-13.979	-8.736
Housing Price, 90th Percentile	407.385	259.766	448.148	283.325	-40.764	-20.955
Rent, 10th Percentile	647.190	303.611	566.929	334.942	80.261	34.394
Rent, 20th Percentile	783.625	316.479	695.242	357.324	88.383	35.472
Rent, 30th Percentile	880.810	329.767	791.056	378.205	89.754	34.143
Rent, 40th Percentile	963.778	345.153	873.877	400.697	89.901	32.427
Rent, 50th Percentile	1,043.379	364.277	953.811	425.830	89.568	30.505
Rent, 60th Percentile	1,127.480	388.424	1,037.931	455.456	89.549	28.604
Rent, 70th Percentile	1,223.033	419.299	1,133.867	490.917	89.166	26.514
Rent, 80th Percentile	1,346.096	463.863	1,256.968	537.923	89.128	24.260
Rent, 90th Percentile	1,546.165	542.104	1,455.534	615.360	90.631	21.696
% Below Poverty	13.332	12.609	13.720	10.351	-0.388	-4.624
% on Public Assistance	5.960	7.430	3.892	4.526	2.068	56.901
% White	68.594	31.499	75.293	25.479	-6.699	-30.934
% Black	14.564	24.627	8.256	16.194	6.308	40.249
% Hispanic	11.546	19.222	11.144	18.443	0.402	2.740
Population Density	2,773.502	5,775.713	767.319	2,039.192	2,006.183	60.149
Avg HH Income	98.020	50.391	93.019	43.321	5.000	13.993
% College Educated	23.494	17.521	22.798	15.660	0.697	5.571
% Owner Occupied	62.336	24.968	71.538	18.664	-9.202	-55.022
% Female Head of Household	19.584	14.026	15.861	10.171	3.723	42.275
% Unemployment	5.605	5.018	6.081	4.707	-0.476	-16.790
% Vacant	7.788	7.417	13.627	12.963	-5.839	-69.102
% Moved in Last 10 Years	63.824	15.032	61.192	14.332	2.632	25.522
% Mobile Home	5.106	36.366	11.574	40.697	-6.469	-33.038
% Built in Last 10 Years	17.107	19.946	18.641	18.640	-1.534	-12.888
% Under 18	24.883	7.336	24.126	6.928	0.757	14.899
% 4+ Bedrooms	16.993	14.280	19.112	13.898	-2.118	-19.966

Notes: This table provides summary statistics on characteristics at the Census tract-level (interpolated to 2010 boundaries) for the years 1990 (Decennial), 2000, (Decennial), 2010 (2008-2012 ACS), and 2020 (2018-2022 ACS). It compares the average and standard deviation for characteristics of tracts whose boundaries are within 10 km from a RCRA site (columns 1 and 2, the sample restriction we use) to tracts that are not within 10 km of any RCRA sites (columns 3 and 4). It also provides the t-statistic associated with whether the mean difference between the samples are different from 0 (columns 5 and 6). Housing prices and average household (HH) income are in thousands of dollars, rental prices are in dollars, population density is number of people per km^2, and the remaining variables are in percentages.

Table 2: Price Impacts of Cleanup by Decile, Near-Far Comparison

Dep. var: Price^{kth}	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	11.504***	10.085**	9.403**	8.405**	8.539**	8.918**	11.477**	13.447**	13.622
	(4.411)	(4.142)	(3.826)	(3.815)	(4.103)	(4.488)	(5.174)	(6.766)	(9.936)
0 km × In Prog	10.099***	8.484**	7.070**	5.894**	5.483*	6.506*	7.878*	10.711**	13.341
	(3.801)	(3.371)	(2.953)	(2.962)	(3.186)	(3.608)	(4.192)	(5.306)	(8.278)
% Impact:									
0 km × Post	0.076***	0.053**	0.043**	0.035**	0.032**	0.030**	0.034**	0.035**	0.029
	(0.029)	(0.022)	(0.018)	(0.016)	(0.015)	(0.015)	(0.015)	(0.018)	(0.021)
0 km × In Prog	0.066***	0.045**	0.032**	0.024**	0.020*	0.022*	0.024*	0.028**	0.028
	(0.025)	(0.018)	(0.014)	(0.012)	(0.012)	(0.012)	(0.013)	(0.014)	(0.018)
Avg Price	151.876	189.696	217.879	243.520	269.621	298.921	334.088	382.495	468.803
Adj R^2	0.823	0.876	0.897	0.907	0.911	0.910	0.908	0.900	0.874
Clusters	12,770	12,770	12,770	12,770	12,770	12,770	12,770	12,770	12,770
Obs	48,585	48,585	48,585	48,585	48,585	48,585	48,585	48,585	48,585

Notes: This table presents the price impacts of cleanup by housing price decile from equation 1, using the 0 km bin as the $Near_{it}$ distance indicator. *Post* is the post-cleanup indicator, and *InProg* is an indicator for if cleanup is in progress but not complete. The top panel reports the point estimates (in 2023 thousands of dollars), the middle panel reports the price impacts as a percentage of the dependent variable mean. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(0, 10]$ km away from a facility. All standard errors are clustered on census tract.

Table 3: Semi-elasticity Price Impacts by Decile with Poisson Regression

Dep. var: Price^{kth}	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	0.098^{**}	0.061^{**}	0.045^{**}	0.031^{*}	0.026	0.023	0.029	0.032	0.025
	(0.040)	(0.028)	(0.021)	(0.019)	(0.018)	(0.018)	(0.019)	(0.021)	(0.024)
0 km × In Prog	0.089^{**}	0.048^{**}	0.029	0.017	0.012	0.013	0.015	0.022	0.020
	(0.036)	(0.024)	(0.018)	(0.016)	(0.015)	(0.016)	(0.016)	(0.018)	(0.021)
0 km × Post	14.905^{**}	11.528^{**}	9.773^{**}	7.646^{*}	7.020	7.011	9.636	12.413	11.753
	(6.101)	(5.261)	(4.666)	(4.569)	(4.911)	(5.358)	(6.199)	(8.142)	(11.389)
0 km × In Prog	13.566^{**}	9.152^{**}	6.264	4.096	3.110	3.914	4.918	8.245	9.607
	(5.510)	(4.546)	(3.935)	(3.863)	(4.166)	(4.742)	(5.510)	(6.878)	(9.760)
Avg Price	151.876	189.696	217.879	243.520	269.621	298.921	334.088	382.495	468.803
Pseudo R^2	0.826	0.860	0.874	0.883	0.889	0.893	0.896	0.898	0.894
Clusters	12,770	12,770	12,770	12,770	12,770	12,770	12,770	12,770	12,770
Obs	48,585	48,585	48,585	48,585	48,585	48,585	48,585	48,585	48,585

Notes: This table presents the price impacts of cleanup by housing price decile using the high-dimensional fixed effects Poisson estimator from Correia, Guimarães and Zylkin (2020). The top panel presents the Poisson parameters, where each is interpreted as a semi-elasticity. The middle panel presents the marginal effects in levels (thousands of 2023 dollars), with standard errors calculated using the Delta method. As before, 0 km is an indicator for whether a tract contains a RCRA site, *Post* is the post-cleanup indicator, and *InProg* is an indicator for if cleanup is in progress but not complete. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(0, 10]$ km away from a facility. All standard errors are clustered on census tract.

Table 4: Price Impacts of Cleanup by Decile, Near-Far Comparison, Allowing Multiple Sites

Dep. var: Price^{kth}	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	9.327**	7.935**	7.352**	7.138**	7.707**	7.301*	8.833*	9.273	8.774
	(4.021)	(3.767)	(3.497)	(3.468)	(3.714)	(4.125)	(4.758)	(6.251)	(9.062)
0 km × In Prog	8.069**	6.594**	5.624**	4.674	4.021	4.611	6.011	7.611	9.475
	(3.539)	(3.093)	(2.828)	(2.892)	(3.091)	(3.436)	(3.947)	(4.982)	(7.628)
% Impact:									
0 km × Post	0.056**	0.039**	0.031**	0.027**	0.027**	0.023*	0.025*	0.023	0.018
	(0.024)	(0.018)	(0.015)	(0.013)	(0.013)	(0.013)	(0.014)	(0.016)	(0.019)
0 km × In Prog	0.049**	0.032**	0.024**	0.018	0.014	0.015	0.017	0.019	0.020
	(0.021)	(0.015)	(0.012)	(0.011)	(0.011)	(0.011)	(0.011)	(0.012)	(0.016)
Avg Price	166.204	205.525	234.621	260.932	287.493	316.864	351.900	399.412	482.753
Adj R^2	0.813	0.871	0.893	0.904	0.910	0.911	0.910	0.905	0.884
Clusters	43,234	43,234	43,234	43,234	43,234	43,234	43,234	43,234	43,234
Obs	164,566	164,566	164,566	164,566	164,566	164,566	164,566	164,566	164,566

Notes: This table presents the price impacts of cleanup by housing price decile using the same specification as that in Table 2 but a different sample that includes tracts with *at least one* RCRA site within 10 km. Treatment tracts are those containing at least one RCRA site with all sites within 10 km contained in the tract, and control tracts are those with no contained sites but at least one site within 10 km. The variables Post and In Prog are as previously defined, using the timing of the nearest site to a tract. All regressions include fixed effects for tract, bin by year, nearest-site by year, and state by year. The excluded category is tracts with at least one site $(0, 10]$ km away from a facility. All standard errors are clustered on census tract.

Table 5: Price Impacts of Cleanup by Decile, Near-Far Comparison with Regression Adjustment

Dep. var: Price^{kth}	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	13.115***	12.207***	12.373***	9.351**	8.562*	9.214*	11.684**	9.774	2.364
	(4.766)	(4.376)	(4.202)	(4.219)	(4.494)	(4.909)	(5.631)	(7.561)	(11.372)
× Ln(Tract Area (km^2))	0.499	-0.315	-0.631	-2.609	-3.694	-4.079	-3.748	-6.140	-9.704
	(3.181)	(2.930)	(2.398)	(2.364)	(2.504)	(2.767)	(3.203)	(4.699)	(6.993)
0 km × In Prog	11.029***	9.905***	9.213***	7.772**	7.279**	8.174**	8.860**	10.516*	13.149
	(3.657)	(3.331)	(3.057)	(3.039)	(3.252)	(3.603)	(4.215)	(5.701)	(9.559)
× Ln(Tract Area (km^2))	-0.250	-0.143	1.101	0.817	0.556	-0.322	-0.533	0.596	1.275
	(3.406)	(3.466)	(3.105)	(3.231)	(3.464)	(3.716)	(3.976)	(4.582)	(6.222)
Avg Price	151.876	189.696	217.879	243.520	269.621	298.921	334.088	382.495	468.803
Adj R^2	0.808	0.868	0.890	0.902	0.908	0.909	0.907	0.899	0.873
Clusters	12,770	12,770	12,770	12,770	12,770	12,770	12,770	12,770	12,770
Obs	48,585	48,585	48,585	48,585	48,585	48,585	48,585	48,585	48,585

Notes: This table presents the regression adjustment version of our main specification to address the fact that tracts in the 0 km bin have different areas than those $(0, 10]$ km from a RCRA site. The first and third row are the un-interacted 0 km × Post and 0 km × In Prog, whereas the 2nd and 4th rows are each interacted with the log of tract area centered about the treatment group mean. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(0, 10]$ km away from a facility. All standard errors are clustered on census tract.

Table 6: Impacts on Sociodemographic and Housing Variables, Near-Far Comparison

(a) Income and Education-related Variables

Dep. var:	Avg HH Income	% Below Poverty	% College Educated	% on Public Assistance	% Unemployment
0 km × Post	0.092	-0.421	0.893*	0.287	0.531
	(1.318)	(0.498)	(0.479)	(0.284)	(0.360)
0 km × In Prog	-0.874	-0.467	0.614*	0.152	0.190
	(1.040)	(0.440)	(0.358)	(0.221)	(0.317)
Avg Outcome	101.218	12.701	26.467	3.660	6.021
Adj R^2	0.871	0.787	0.896	0.586	0.512
Clusters	13,137	13,140	13,150	13,137	13,146
Obs	52,387	52,436	52,513	52,422	52,478

(b) Impacts on Demographic Variables, Near-Far Comparison

Dep. var:	% Black	% Female Head of Household	% Hispanic	Population Density	% White	% Under 18
0 km × Post	-0.046	0.262	-0.101	-2.097	-0.249	-0.285
	(0.363)	(0.565)	(0.376)	(10.974)	(0.515)	(0.308)
0 km × In Prog	0.277	0.453	-0.491	8.589	0.054	-0.143
	(0.259)	(0.423)	(0.309)	(8.314)	(0.417)	(0.256)
Avg Outcome	11.088	17.728	12.781	1,506.599	69.678	24.026
Adj R^2	0.928	0.757	0.934	0.971	0.939	0.733
Clusters	13,147	13,129	13,147	13,155	13,147	13,147
Obs	52,482	52,370	52,482	52,607	52,482	52,482

(c) Impacts on Housing-Related Variables, Near-Far Comparison

Dep. var:	% 4+ Bedrooms	% Built in Last 10 Years	% Mobile Home	% Moved in Last 10 Years	% Owner Occupied	% Vacant
0 km × Post	-0.420	-0.117	0.910	-0.004	0.336	0.285
	(0.468)	(0.864)	(0.595)	(0.613)	(0.575)	(0.478)
0 km × In Prog	-0.238	0.712	1.418	0.747	0.002	-0.060
	(0.393)	(0.738)	(0.905)	(0.535)	(0.457)	(0.358)
Avg Outcome	20.049	17.568	7.519	63.097	67.339	9.632
Adj R^2	0.846	0.610	0.070	0.714	0.907	0.763
Clusters	13,138	13,138	13,137	13,137	13,130	13,134
Obs	52,431	52,431	52,422	52,422	52,382	52,400

Notes: This table tests the impacts of RCRA cleanups on 17 socioeconomic and housing variables, comparing changes in outcomes for tracts containing a RCRA site (0 km) over time. The regression model is the same as the main model in Table 2 except with one of the 17 socioeconomic and housing variables as the dependent variable. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(0, 10]$ km away from a facility. All standard errors are clustered on census tract. Avg. HH Income is in $1,000.

Table 7: Mean Utility Decomposition (In Dollars)

	By Race		By Income
	All	Owners	All
Near × Post × Black	-115.940	-302.149***	
	(95.417)	(67.104)	
Near × Post × Hispanic	-120.432	-295.400***	
	(413.072)	(76.723)	
Near × Post × High Income			37.353
			(23.828)
Clusters	13,433	13,433	12,708
Obs	120,897	120,897	55,516

Notes: This table presents estimates of differential willingness-to-pay between groups from the second stage of the sorting model, where mean utilities for each group and year are stacked and regressed on a full set of interactions between group, 0km bin (or 'near'), and post cleanup for the sample of tracts exposed to at most 1 facility within 10km. Groups are based on 1) race (non-Hispanic white, non-Hispanic Black, and Hispanic), 2) race for the sample of homeowners, and 3) income (below versus above state median income). The base group for comparison is white residents (by race) and low income (by income). All specifications include tract, distance bin-by-year, site-id-by-year, and state-by-year fixed effects. All standard errors are clustered on census tract.

References

Abadie, Alberto, Susan Athey, Guido Imbens, and Jeffrey Wooldridge, "When Should You Adjust Standard Errors for Clustering?," *Quarterly Journal of Economics*, 2022.

Albouy, David, "What are cities worth? Land rents, local productivity, and the total value of amenities," *Review of Economics and Statistics*, 2016, *98* (3), 477–487.

Amornsiripanitch, Natee, "Why are residential property tax rates regressive?," *Available at SSRN 3729072*, 2020.

Apelberg, Benjamin J., Timothy J. Buckley, and Ronald H. White, "Socioeconomic and Racial Disparities in Cancer Risk from Air Toxics in Maryland," *Environmental Health Perspectives*, 2005, *113* (6), 693–699.

Arias, Elizabeth, Carolyn A. Liebler, Marc A. Garcia, and Rogelio Sáenz, "Data Impacts of Changes in U.S. Census Bureau Procedures for Race and Ethnicity Data," *SSM - Population Health*, 2025, *29*, 101742.

Asker, Erdal, Eric Brunner, and Stephen Ross, "The Impact of School Spending on Civic Engagement: Evidence from School Finance Reforms," Technical Report, National Bureau of Economic Research 2022.

Avenancio-León, Carlos F. and Troup Howard, "The Assessment Gap: Racial Inequalities in Property Taxation," *PSN: Politics of Race (Topic)*, 2019.

Avenancio-León, Carlos F and Troup Howard, "The assessment gap: Racial inequalities in property taxation," *The Quarterly Journal of Economics*, 2022, *137* (3), 1383–1434.

Bakkensen, Laura A and Lint Barrage, "Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics," *The Review of Financial Studies*, 11 2021, *35* (8), 3666–3709.

Bakkensen, Laura and Lala Ma, "Sorting Over Flood Risk and Implications for Policy Reform," *Journal of Environmental Economics and Management*, November 2020, *104* (102362).

Banzhaf, H Spencer, "Panel Data Hedonics: Rosen'S First Stage as a "Sufficient Statistic"," *International Economic Review*, 2020, *61* (2), 973–1000.

— , "Difference-in-Differences Hedonics," *Journal of Political Economy*, 2021, *129* (8), 2385–2414.

— **and Omar Farooque**, "Interjurisdictional housing prices and spatial amenities" Which measures of housing prices reflect local public goods?," *Regional Science and Urban Economics*, 2013, *43* (4), 635–648.

— **and Randall P Walsh**, "Segregation and Tiebout sorting: The link between place-based investments and neighborhood tipping," *Journal of Urban Economics*, 2013, *74*, 83–98.

— **, Lala Ma, and Christopher Timmins**, "Environmental justice: Establishing causal relationships," *Annual Review of Resource Economics*, 2019, *11*, 377–398.

Banzhaf, Spencer and Eleanor McCormick, "Moving Beyond Cleanup: Identifying the Crucibles of Environmental Gentrification," *Working Paper*, 2007.

Bayer, Patrick, Nathaniel Keohane, and Christopher Timmins, "Migration and hedonic valuation: The case of air quality," *Journal of Environmental Economics and Management*, 2009, *58* (1), 1–14.

—, Robert McMillan, Alvin Murphy, and Christopher Timmins, "A Dynamic Model of Demand for Houses and Neighborhoods," *Econometrica*, 2016, *84* (3), 893–942.

—, —, and Kim Rueben, "An Equilibrium Model of Sorting in an Urban Housing Market," 2004. Working Paper No. w10865.

Berrens, Robert P. and Michael McKee, "What Price Nondisclosure? The Effects of Nondisclosure of Real Estate Sales Prices," *Social Science Quarterly*, 2004, *85* (2), 509–520.

Berry, Christopher R, "Reassessing the property tax," *Available at SSRN 3800536*, 2021.

Berry, Steven T, "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 1994, *25* (2), 242–262.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-In-Differences Estimates?*," *The Quarterly Journal of Economics*, 02 2004, *119* (1), 249–275.

Bischof, Daniel, "BLINDSCHEMES: Stata module to provide graph schemes sensitive to color vision deficiency," Statistical Software Components, Boston College Department of Economics Oct 2016.

Bishop, Kelly C., Nicolai V. Kuminoff, H. Spencer Banzhaf, Kevin J. Boyle, Klaus von Gravenitz, Jaren C. Pope, and Christopher Timmins, "Best practices in using hedonic property value models for welfare measurement," *Review of Environmental Economics and Policy*, 2020, *43.*

Blomquist, Glenn C, Mark C Berger, and John P Hoehn, "New Estimates of Quality of Life in Urban Areas," *American Economic Review*, 1988, pp. 89–107.

Bollum, Tanner, "Public revenue leakage from real estate non-disclosure laws." PhD dissertation, Montana State University-Bozeman, College of Agriculture 2021.

Brumfield, Cara, Casey Goldvale, and Christopher Brown, "Race and Origin Questions in Context," Technical Report, Georgetown Center on Poverty and Inequality, Economic Security and Opportunity Initiative June 2019. Accessed: 2025-03-20.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230. Themed Issue: Treatment Effect 1.

Cassidy, Alecia, "How Does Mandatory Energy Efficiency Disclosure Affect Housing Prices?," *Journal of the Association of Environment and Resource Economists*, 2023, *10* (3).

—, Robyn C. Meeks, and Michael R. Moore, "Cleaning up the Great Lakes: Housing market impacts of removing legacy pollutants," *Journal of Public Economics*, 2023, *226*, 104979.

Christensen, Peter and Christopher Timmins, "Sorting or steering: Experimental evidence on the economic effects of housing discrimination," Technical Report w24826, National Bureau of Economic Research 2018.

— , **Ignacio Sarmiento-Barbieri, and Christopher Timmins**, "Housing Discrimination and Pollution Exposures in the United States," Technical Report w26805, National Bureau of Economic Research 2020.

Clarke, Damian, Joseph P. Romano, and Michael Wolf, "The Romano–Wolf multiple-hypothesis correction in Stata," *The Stata Journal*, 2020, *20* (4), 812–843.

Correia, Sergio, "Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator," Technical Report 2016. Working Paper.

— , **Paulo Guimarães, and Thomas Zylkin**, "Verifying the existence of maximum likelihood estimates for generalized linear models," 2021.

— , — , **and Tom Zylkin**, "Fast Poisson estimation with high-dimensional fixed effects," *The Stata Journal: Promoting communications on statistics and Stata*, March 2020, *20* (1), 95–115.

Currie, Janet, Lucas Davis, Michael Greenstone, and Reed Walker, "Environmental Health Risks and Housing Values: Evidence from 1,600 Toxic Plant Openings and Closings," *American Economic Review*, 2015, *105* (2), 678–709.

— , **Michael Greenstone, and Enrico Moretti**, "Superfund cleanups and infant health," *American Economic Review*, 2011, *101* (3), 435–441.

Davenport, Lauren, "The Fluidity of Racial Classifications," *Annual Review of Political Science*, 2020, *23*, 221–240.

Depro, Brooks, Christopher Timmins, and Maggie O'Neil, "White Flight and Coming to the Nuisance: Can Residential Mobility Explain Environmental Injustice?," *Journal of the Association of Environmental and Resource Economists*, 2015, *2* (3), 439–468.

Derenoncourt, Ellora, Chi Hyun Kim, Moritz Kuhn, and Moritz Schularick, "Wealth of two nations: The US racial wealth gap, 1860–2020," *The Quarterly Journal of Economics*, 2024, *139* (2), 693–750.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey," *The Econometrics Journal*, 06 2022, *26* (3), C1–C30.

Diamond, Peter A. and Jerry A. Hausman, "Contingent valuation: Is some number better than no number?," *Journal of Economic Perspectives*, 1994, *8* (4), 45–64.

Dornfest, Alan S, Jennifer Rearich, T Douglas Brydon III, and Richard Almy, "State and provincial property tax policies and administrative practices (PTAPP): 2017 findings and report," *Journal of Property Tax Assessment & Administration*, 2019, *16* (1), 43–130.

Dreesen, Stijn and Sven Damen, "The accuracy of homeowners' valuations in the twenty-first century," *Empirical Economics*, 2023, *65* (1), 513–566.

Du Preez, Mario and Michael Sale, "Municipal assessments versus actual sales prices in hedonic price studies," *Journal of Economic and Financial Sciences*, 2015, *8* (1), 35–46.

Dube, Arindrajit, "Minimum Wages and the Distribution of Family Incomes," *American Economic Journal: Applied Economics*, October 2019, *11* (4), 268–304.

Economic Policy Institute, "Racial and Ethnic Disparities in the United States: An Interactive Chartbook," https://www.epi.org/publication/anti-racist-policy-research/ June 2022. Updated November 2024.

EPA, US, "Resource Conservation and Recovery Act (RCRA) Overview," March 2025. Available from https://www.epa.gov/rcra/resource-conservation-and-recovery-act-rcra-overview.

Firpo, Sergio, "Efficient Semiparametric Estimation of Quantile Treatment Effects," *Econometrica*, 2007, *75* (1), 259–276.

__ , Nicole M. Fortin, and Thomas Lemieux, "Unconditional Quantile Regressions," *Econometrica*, 2009, *77* (3), 953–973.

Gamper-Rabindran, Shanti and Christopher Timmins, "Does cleanup of hazardous waste sites raise housing values? Evidence of spatially localized benefits," *Journal of Environmental Economics and Management*, 2013, *65* (3), 345–360.

__ , Ralph Mastromonaco, and Christopher Timmins, "Valuing the Benefits of Superfund Site Remediation: Three Approaches to Measuring Localized Externalities," *National Bureau of Economic Research No. w16655*, 2011.

Gardner, John, "Two-stage differences in differences," *arXiv preprint arXiv:2207.05943*, 2022.

Gayer, T., J.T. Hamilton, and W.K. Viscusi, "Private Values of Risk Tradeoffs at Superfund Sites: Housing Market Evidence on Learning about Risk," *Review of Economics and Statistics*, 2000, *82* (3), 439–451.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277. Themed Issue: Treatment Effect 1.

Goodman Jr, John L and John B Ittner, "The accuracy of home owners' estimates of house value," *Journal of housing economics*, 1992, *2* (4), 339–357.

Gould, William, "Use Poisson rather than regress; tell a friend," *Stata Blog*, 2011. Accessed: 2024-02-04.

Greenstone, Michael and Justin Gallagher, "Does hazardous waste matter? Evidence from the housing market and the superfund program," *Quarterly Journal of Economics*, 2008, *123* (3), 951–1003.

Guignet, Dennis and Christopher Nolte, "Hazardous Waste and Home Values: An Analysis of Treatment and Disposal Sites in the United States," *Journal of the Association of Environmental and Resource Economists*, 2024, *11* (2), 487–521.

__ , Robin R. Jenkins, Christoph Nolte, and James Belke, "The External Costs of Industrial Chemical Accidents: A Nationwide Property Value Study," *Journal of Housing Economics*, 2023, *62*, 101954.

Haninger, Kevin, Lala Ma, and Christopher Timmins, "The Value of Brownfield Remediation," *Journal of the Association of Environmental and Resource Economists*, 2017, *4* (1), 197–241.

Harleman, Max, "Who bears the cost of renewable power transmission lines? Evidence from housing values," *Energy Policy*, 2024, *191*, 114179.

Hausman, Catherine and Samuel Stolper, "Inequality, Information Failures, and Air Pollution," *National Bureau of Economic Research No. w26682*, 2020, pp. 1–75.

Hausman, Jerry, "Contingent valuation: from dubious to hopeless," *Journal of Economic Perspectives*, 2012, pp. 43–56.

Havnes, Tarjei and Magne Mogstad, "Is universal child care leveling the playing field?," *Journal of Public Economics*, 2015, *127*, 100–114. The Nordic Model.

Hochberg, Yosef and Ajit C Tamhane, *Multiple comparison procedures*, John Wiley & Sons, Inc., 1987.

Holm, Sture, "A simple sequentially rejective multiple test procedure," *Scandinavian journal of statistics*, 1979, pp. 65–70.

Jann, Ben, "GEOPLOT: Stata module to draw maps," Statistical Software Components, Boston College Department of Economics Jun 2023.

Jensen, Eric, Nicholas Jones, Megan Rabe, Beverly Pratt, Lauren Medina, Kimberly Orozco, and Lindsay Spell, "2020 U.S. Population More Racially and Ethnically Diverse Than Measured in 2010," https://www.census.gov/library/stories/2021/08/2020-united-states-population-more-racially-ethnically-diverse-than-2010.html August 2021. America Counts: Stories, U.S. Census Bureau. Accessed: 2025-03-20.

Jones, Damon, David Molitor, and Julian Reif, "What do Workplace Wellness Programs do? Evidence from the Illinois Workplace Wellness Study*," *The Quarterly Journal of Economics*, 08 2019, *134* (4), 1747–1791.

Jones, Nicholas A. and Jungmiwha Bullock, "The Two or More Races Population: 2010," Technical Report C2010BR-13, U.S. Census Bureau September 2012. 2010 Census Brief. Accessed: 2025-03-20.

Jones, Nicholas, Rachel Marks, Roberto Ramirez, and Merarys Rios-Vargas, "Census Illuminates Racial and Ethnic Composition of the Country," https://www.census.gov/library/stories/2020/10/2020-census-data-products-illuminate-racial-and-ethnic-composition-of-the-country.html 2020. Accessed: 2022-07-13.

Kahn, Matthew E, "A Revealed Preference Approach to Ranking City Quality of Life," *Journal of Urban Economics*, 1995, *38* (2), 221–235.

Keiser, David A and Joseph S Shapiro, "Consequences of the Clean Water Act and the demand for water quality," *The Quarterly Journal of Economics*, 2018, *134* (1), 349–396.

Kiel, Katherine A. and Jeffrey E. Zabel, "The accuracy of owner-provided house values: The 1978–1991 American Housing Survey," *Real Estate Economics*, 1999, *27* (2), 263–298.

Kinnaman, Thomas C, "A landfill closure and housing values," *Contemporary Economic Policy*, 2009, *27* (3), 380–389.

Kohlhase, Janet E, "The Impact of Toxic Waste Sites on Housing Values," *Journal of Urban Economics*, 1991, *30* (1), 1–26.

Kuminoff, Nicolai V, V Kerry Smith, and Christopher Timmins, "The new economics of equilibrium sorting and policy evaluation using housing markets," *Journal of Economic Literature*, 2013, *51* (4), 1007–1062.

Liebler, Carolyn A., "On the Boundaries of Race: Identification of Mixed-heritage Children in the United States, 1960 to 2010," *Sociology of Race and Ethnicity*, 2016, *2* (4), 548–568.

_ , Sonya R. Porter, Leticia E. Fernandez, James M. Noon, and Sharon R. Ennis, "America's Churning Races: Race and Ethnicity Response Changes Between Census 2000 and the 2010 Census," *Demography*, 01 2017, *54* (1), 259–284.

Linn, Joshua, "The Effect of Voluntary Brownfields Programs on Nearby Property Balues: Evidence from Illinois," *Journal of Urban Economics*, 2013, *78*, 1–18.

Logan, John R., Brian J. Stults, and Zengwang Xu and, "Validating Population Estimates for Harmonized Census Tract Data, 2000–2010," *Annals of the American Association of Geographers*, 2016, *106* (5), 1013–1029.

_ , Wenquan Zhang, Brian J. Stults, and Todd Gardner, "Improving Estimates of Neighborhood Change with Constant Tract Boundaries," *Applied Geography*, 2021, *132*, 102476.

Logan, John R, Zengwang Xu, and Brian J Stults, "Interpolating US decennial census tract data from as early as 1970 to 2010: A longitudinal tract database," *The Professional Geographer*, 2014, *66* (3), 412–420.

Ma, Lala, "Learning in a Hedonic Framework: Valuing Brownfield Remediation," *International Economic Review*, 2019, *60* (3), 1355–1387.

Manning, Willard G and John Mullahy, "Estimating log models: to transform or not to transform?," *Journal of Health Economics*, 2001, *20* (4), 461–494.

Manson, Steven, Jonathan Schroeder, David Van Riper, Katherine Knowles, Tracy Kugler, Finn Roberts, and Steven Ruggles, "IPUMS National Historical Geographic Information System: Version 19.0 [dataset]," https://doi.org/10.18128/D050.V19.0 2024. Minneapolis, MN: IPUMS.

Marks, Rachel and Merarys Rios-Vargas, "Improvements to the 2020 Census Race and Hispanic Origin Question Designs, Data Processing, and Coding Procedures," August 2021. United States Census Bureau. Accessed: 2025-03-20.

Mastromonaco, Ralph, "Do Environmental Right-to-Know Laws Affect Markets? Capitalization of Information in the Toxic Release Inventory," *Journal of Environmental Economics and Management*, 2015, *71*, 54–70.

Messer, Kent D, William D Schulze, Katherine F Hackett, Trudy A Cameron, and Gary H McClelland, "Can Stigma Explain Large Property Value Losses? The Psychology and Economics of Superfund," *Environmental and Resource Economics*, 2006, *33* (3), 299–324.

Nolte, Christoph, Kevin J. Boyle, Anita Chaudhry, Christopher Clapp, Dennis Guignet, Hannah Hennighausen, Ido Kushner, Yanjun Liao, Saleh Mamun, Adam Pollack, Jesse Richardson, Shelby Sundquist, Kristen Swedberg, and Johannes H. Uhl, "Data practices for studying the impacts of environmental amenities and hazards with nationwide property data," *Land Economics*, 2024, *100* (1), 200–221.

Paglin, Morton and Michael Fogarty, "Equity and the property tax: A new conceptual focus," *National Tax Journal*, 1972, *25* (4), 557–565.

Palmquist, Raymond B, "A note on transaction costs, moving costs, and benefit measurement," *Journal of Urban Economics*, 1992, *32* (1), 40–44.

Pew Research Center, "Multiracial in America: Proud, Diverse and Growing in Numbers," Technical Report, Pew Research Center June 2015. Accessed: 2025-03-20.

Reif, Julian, "WYOUNG: Stata module to perform multiple testing corrections," Statistical Software Components, Boston College Department of Economics December 2017.

Roback, Jennifer, "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 1982, *90* (6), 1257–1278.

Romano, Joseph P and Michael Wolf, "Exact and Approximate Stepdown Methods for Multiple Hypothesis Testing," *Journal of the American Statistical Association*, 2005, *100* (469), 94–108.

Romano, Joseph P. and Michael Wolf, "Stepwise Multiple Testing as Formalized Data Snooping," *Econometrica*, 2005, *73* (4), 1237–1282.

Romano, Joseph P and Michael Wolf, "Efficient computation of adjusted p-values for resampling-based stepdown multiple testing," *Statistics & Probability Letters*, 2016, *113*, 38–40.

_ , Azeem M Shaikh, and Michael Wolf, "Formalized data snooping based on generalized error rates," *Econometric Theory*, 2008, *24* (2), 404–447.

Rosen, S., "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, 1974, *82* (1), 34–55.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, Renae Rodgers, Jonathan Schroeder, and Kari C.W. Williams, "IPUMS USA: Version 16.0 [dataset]," 2025.

Schroeder, Jonathan P., "Target-Density Weighting Interpolation and Uncertainty Evaluation for Temporal Analysis of Census Data," *Geographical Analysis*, 2007, *39* (3), 311–335.

Silva, J. M. C. Santos and Silvana Tenreyo, "The Log of Gravity," *The Review of Economics and Statistics*, 2006, *88* (4), 641–658.

Smith, V Kerry and William H Desvousges, "The value of avoiding a LULU: hazardous waste disposal sites," *The Review of Economics and Statistics*, 1986, pp. 293–299.

Starr, Paul and Christina Pao, "The Multiracial Complication: The 2020 Census and the Fictitious Multiracial Boom," *Sociological Science*, 2024, *11* (40), 1107–1123.

Strmic-Pawl, Hephzibah V., Brandon A. Jackson, and Steve Garner, "Race Counts: Racial and Ethnic Data on the U.S. Census and the Implications for Tracking Inequality," *Sociology of Race and Ethnicity*, 2018, *4* (1), 1–13.

Timmins, Christopher and Jennifer Murdock, "A Revealed Preference Approach to the Measurement of Congestion in Travel Cost Models," *Journal of Environmental Economics and Management*, 2007, *53* (2), 230–249.

Tur-Sinai, Aviad, Larisa Fleishman, and Dmitri Romanov, "The accuracy of self-reported dwelling valuation," *Journal of Housing Economics*, 2020, *48*, 101660.

United States Census Bureau, "Home Value Estimates in Current Dollars," *US Census Bureau website*, September 2023, *https://www.census.gov/programs-surveys/acs/technical-documentation/user-notes/2023-07.html*.

U.S. Bureau of Labor Statistics, "Consumer Price Index for All Urban Consumers: All Items (CPIAUCSL)," https://fred.stlouisfed.org/series/CPIAUCSL 2025. Retrieved from FRED, Federal Reserve Bank of St. Louis.

U.S. Census Bureau, "2010 Census Shows Multiple-Race Population Grew Faster Than Single-Race Population," https://www.census.gov/newsroom/releases/archives/2010_census/cb12-182.html September 2012. Accessed: 2025-03-20.

—, "Comparing 2010-2012 ACS 3-Year Estimates to Other ACS Estimates," https://www.census.gov/programs-surveys/acs/guidance/comparing-acs-data/2012/3-year-comparison.html 2012. Accessed: 2025-03-20.

—, "Improvements to the Race Question," https://www.census.gov/programs-surveys/acs/technical-documentation/user-notes/2021-11.html November 2021. Accessed: 2025-03-20.

U.S. Dept. of Commerce, Economics and Statistics Administration, Bureau of the Census, "Geographic areas reference manual," *US Census Bureau website*, November 1994, *1* (Ch 10), 10–1–10–17.

U.S. Environmental Protection Agency, ""Biennial Hazardous Waste Report"," https://www.epa.gov/hwgenerators/biennial-hazardous-waste-report 2025. Retrieved from U.S. Environmental Protection Agency Website.

— , ""Nationally Defined Values for Corrective Action Event Code"," https://19january2017snapshot.epa.gov/sites/production/files/2016-01/documents/ca-evcds.pdf 2025. Retrieved from U.S. Environmental Protection Agency Website.

— , "Redevelopment Economics at Hazardous Waste Cleanup Facilities," https://www.epa.gov/hw/redevelopment-economics-hazardous-waste-cleanup-facilities 2025. Retrieved from U.S. Environmental Protection Agency Website.

US EPA, "Report on the Environment, Chapter 4 Exhibit 4-2," 2010.

U.S. EPA, "OSWER Cross-Program Revitalization Measures," 2011.

— , "Resource Conservation and Recovery Act (RCRA) Orientation Manual," 2014. Available online at: $https://www.epa.gov/sites/production/files/2015-07/documents/rom.pdf$.

USDA, "The 2017 NRI: Changes in Land Cover/Use," 2017.

Weber, Rachel N. and Daniel P. McMillen, "Ask and Ye Shall Receive? Predicting the Successful Appeal of Property Tax Assessments," *Public Finance Review*, 2010, *38* (1), 74–101.

Westfall, Peter H and James F Troendle, "Multiple testing with minimal assumptions," *Biometrical Journal: Journal of Mathematical Methods in Biosciences*, 2008, *50* (5), 745–755.

— **and S Stanley Young**, *Resampling-based multiple testing: Examples and methods for p-value adjustment*, Vol. 279, John Wiley & Sons, 1993.

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data*, 2nd ed., Cambridge, Massachusetts: MIT Press, 2010.

Wooldridge, Jeffrey M, "Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators," *Available at SSRN 3906345*, 2021.

Appendices to *Who Benefits from Hazardous Waste Cleanups? Evidence from the Housing Market*

September 5, 2025

Alecia Cassidy[51]
University of Alabama

Elaine L. Hill[52]
University of Rochester
and NBER

Lala Ma[53]
University of Kentucky
and Resources for the Future

[51]Cassidy: Department of Economics, Finance, & Legal Studies, University of Alabama, 246 Alston Hall, Tuscaloosa, AL 35487, awcassidy1@ua.edu.

[52]Hill: Department of Public Health Sciences & Department of Economics, University of Rochester & NBER, 265 Crittenden Blvd., Box 420644, Rochester, NY 14642, elaine_hill@urmc.rochester.edu.

[53]Ma: Department of Economics, University of Kentucky & Resources for the Future, Business & Economics Building, Lexington, KY 40506, lala.ma@uky.edu.

A Appendix: Additional Tables and Figures

A.1 Appendix Tables

Table A.1: Attributes by 1 RCRA Site vs. Multiple Sites, within 10 km

	1 RCRA Site		Multiple RCRA Sites			
Attribute	Mean	St. Dev.	Mean	St. Dev.	Δ Mean	t-statistic
Housing Price, 10th Percentile	143.801	94.548	161.516	107.177	-17.715	-17.296
Housing Price, 20th Percentile	172.490	108.300	192.810	124.138	-20.320	-17.153
Housing Price, 30th Percentile	194.273	119.268	216.147	137.473	-21.873	-16.719
Housing Price, 40th Percentile	214.380	129.797	237.288	150.157	-22.908	-16.083
Housing Price, 50th Percentile	234.872	140.693	258.641	163.391	-23.768	-15.396
Housing Price, 60th Percentile	258.193	153.889	282.401	178.694	-24.208	-14.368
Housing Price, 70th Percentile	286.187	170.141	310.445	197.183	-24.258	-13.061
Housing Price, 80th Percentile	324.600	193.734	347.997	222.313	-23.397	-11.139
Housing Price, 90th Percentile	391.815	237.397	413.193	267.403	-21.379	-8.396
Rent, 10th Percentile	623.082	307.369	655.912	301.769	-32.830	-10.371
Rent, 20th Percentile	754.758	321.937	794.068	313.830	-39.310	-11.862
Rent, 30th Percentile	848.692	336.828	892.429	326.401	-43.737	-12.641
Rent, 40th Percentile	928.858	352.879	976.410	341.440	-47.552	-13.152
Rent, 50th Percentile	1,006.538	371.774	1,056.707	360.604	-50.169	-13.191
Rent, 60th Percentile	1,089.345	395.512	1,141.275	384.902	-51.930	-12.854
Rent, 70th Percentile	1,182.470	424.652	1,237.707	416.377	-55.237	-12.777
Rent, 80th Percentile	1,302.164	467.233	1,361.988	461.613	-59.824	-12.616
Rent, 90th Percentile	1,497.929	542.500	1,563.615	540.906	-65.686	-12.013
% Below Poverty	11.882	10.654	13.869	13.221	-1.987	-16.563
% on Public Assistance	4.844	5.298	6.374	8.039	-1.530	-27.225
% White	75.914	26.677	65.883	32.693	10.031	32.626
% Black	10.374	19.274	16.116	26.165	-5.741	-24.643
% Hispanic	9.318	16.874	12.370	19.959	-3.052	-15.787
Population Density	1,511.732	4,221.320	3,240.518	6,188.972	-1,728.785	-29.448
Avg HH Income	97.330	46.910	98.275	51.619	-0.945	-1.829
% College Educated	22.587	15.979	23.830	18.047	-1.243	-6.933
% Owner Occupied	68.145	22.178	60.185	25.593	7.960	31.623
% Female Head of Household	15.851	10.754	20.968	14.823	-5.117	-40.049
% Unemployment	4.810	3.926	5.899	5.336	-1.089	-27.595
% Vacant	8.909	8.871	7.373	6.752	1.536	18.309
% Moved in Last 10 Years	65.189	15.010	63.318	15.009	1.871	12.244
% Mobile Home	8.667	40.217	3.786	34.738	4.881	15.836
% Built in Last 10 Years	23.381	21.931	14.783	18.630	8.598	42.968
% Under 18	25.456	6.917	24.671	7.474	0.785	10.799
% 4+ Bedrooms	17.992	14.405	16.624	14.216	1.368	8.978
TSDF	0.804	0.397	0.803	0.398	0.001	0.175
Number of waste types	64.856	125.613	72.317	137.266	-7.460	-5.161
High NCAPS score	0.340	0.474	0.276	0.447	0.064	13.140
Medium NCAPS score	0.289	0.453	0.247	0.431	0.041	8.198
Low NCAPS score	0.220	0.414	0.303	0.459	-0.083	-17.185

Notes: This table provides summary statistics on characteristics at the Census tract-level (interpolated to 2010 boundaries) for the years 1990 (Decennial), 2000 (Decennial), 2010 (2008-2012 ACS), and 2020 (2018-2022 ACS). It compares the average and standard deviation for characteristics of tracts whose boundaries are < 10 km from a single RCRA site (columns 1 and 2, the sample restriction we use) to tracts within 10 km from multiple RCRA sites (columns 3 and 4). It also provides the t-statistic associated with whether the mean difference between the samples are different from 0 (columns 5 and 6). Housing prices and average household (HH) income are in thousands of year dollars, rental prices are in dollars, population density is number of people per km^2, and the remaining variables are in percentages. Acronyms are as follows: HH=household, TSDF=Treatment, Storage, and Disposal Facility, NCAPS=National Corrective Action Prioritization System

Table A.2: Attributes by Near vs. Far from 1 RCRA Site, within 10 km

	0 km		(0, 10] km			
Attribute	Mean	St. Dev.	Mean	St. Dev.	Δ Mean	t-statistic
Housing Price, 10th Percentile	94.071	63.751	146.323	95.157	-52.253	-19.665
Housing Price, 20th Percentile	117.413	72.010	175.284	109.084	-57.871	-19.161
Housing Price, 30th Percentile	135.119	78.544	197.274	120.197	-62.155	-18.848
Housing Price, 40th Percentile	151.320	84.934	217.579	130.860	-66.259	-18.621
Housing Price, 50th Percentile	167.997	92.126	238.265	141.881	-70.268	-18.275
Housing Price, 60th Percentile	187.360	103.161	261.786	155.166	-74.426	-17.448
Housing Price, 70th Percentile	210.741	114.355	290.014	171.611	-79.273	-16.899
Housing Price, 80th Percentile	243.011	130.824	328.739	195.485	-85.728	-16.107
Housing Price, 90th Percentile	301.612	166.062	396.390	239.552	-94.778	-14.219
Rent, 10th Percentile	462.683	226.518	631.086	308.682	-168.403	-18.157
Rent, 20th Percentile	585.736	240.222	763.192	323.175	-177.456	-17.852
Rent, 30th Percentile	673.294	249.293	857.445	338.238	-184.150	-17.860
Rent, 40th Percentile	746.693	260.748	937.948	354.415	-191.255	-17.726
Rent, 50th Percentile	814.480	275.029	1,016.121	373.377	-201.642	-17.780
Rent, 60th Percentile	884.566	289.852	1,099.563	397.304	-214.997	-18.012
Rent, 70th Percentile	963.934	310.867	1,193.375	426.625	-229.441	-17.855
Rent, 80th Percentile	1,066.988	344.025	1,313.900	469.471	-246.912	-17.430
Rent, 90th Percentile	1,230.191	405.349	1,511.289	545.020	-281.098	-16.944
% Below Poverty	14.328	9.735	11.750	10.685	2.578	6.410
% on Public Assistance	5.762	5.158	4.795	5.301	0.967	5.904
% White	81.304	21.455	75.626	26.898	5.679	6.176
% Black	9.256	16.572	10.434	19.406	-1.178	-1.650
% Hispanic	6.720	13.390	9.458	17.030	-2.738	-4.839
Population Density	263.328	479.974	1,578.643	4,321.378	-1,315.315	-25.876
Avg HH Income	77.875	27.527	98.374	47.504	-20.499	-16.748
% College Educated	15.439	10.727	22.971	16.123	-7.532	-16.556
% Owner Occupied	67.361	19.535	68.187	22.310	-0.826	-1.024
% Female Head of Household	15.139	9.106	15.889	10.834	-0.750	-1.956
% Unemployment	5.069	3.736	4.796	3.935	0.273	2.171
% Vacant	11.164	9.722	8.788	8.808	2.375	6.149
% Moved in Last 10 Years	62.992	13.385	65.307	15.083	-2.315	-4.351
% Mobile Home	13.297	13.172	8.419	41.154	4.878	8.252
% Built in Last 10 Years	19.945	14.574	23.565	22.242	-3.621	-6.594
% Under 18	26.094	5.759	25.422	6.972	0.673	2.923
% 4+ Bedrooms	14.816	8.900	18.162	14.622	-3.346	-8.950
TSDF	0.789	0.408	0.805	0.397	-0.016	-0.873
Number of waste types	65.644	140.262	64.814	124.769	0.830	0.137
High NCAPS score	0.399	0.490	0.337	0.473	0.062	3.082
Medium NCAPS score	0.296	0.457	0.288	0.453	0.008	0.416
Low NCAPS score	0.201	0.401	0.221	0.415	-0.020	-1.106

Notes: This table provides summary statistics on characteristics at the Census tract-level (interpolated to 2010 boundaries) for the years 1990 (Decennial), 2000 (Decennial), 2010 (2008-2012 ACS), and 2020 (2018-2022 ACS). It compares the average and standard deviation for characteristics of tracts that contain a single (or is within 0 km of a) RCRA site (columns 1 and 2) to tracts that are (0, 10] km from a single RCRA site (columns 3 and 4). It also provides the t-statistic associated with whether the mean difference between the samples are different from 0 (columns 5 and 6). Housing prices and average household (HH) income are in thousands of year dollars, rental prices are in dollars, population density is number of people per km^2, and the remaining variables are in percentages. Acronyms are as follows: HH=household, TSDF=Treatment, Storage, and Disposal Facility, NCAPS=National Corrective Action Prioritization System

Table A.3: p-values from Tests that Percentile Impacts are Equal in Percentage Terms

	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
10^{th}									
20^{th}	0.104								
30^{th}	0.090	0.295							
40^{th}	0.061	0.135	0.099						
50^{th}	0.060	0.131	0.124	0.424					
60^{th}	0.061	0.138	0.176	0.493	0.698				
70^{th}	0.103	0.273	0.457	0.986	0.736	0.419			
80^{th}	0.158	0.401	0.631	0.964	0.786	0.642	0.928		
90^{th}	0.154	0.348	0.514	0.778	0.886	0.964	0.731	0.566	

Notes: This table presents the p-values corresponding to the null hypothesis that each pair of housing price percentile impacts from our main specification is statistically different from one another in percentage terms. We present p-values from testing that percentage impacts are the same for each pair of housing price outcomes. These are estimated via stacked regression.

Table A.4: Robustness Check: House Characteristic Controls

Dep. var: Price^{kth}									
	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	11.932***	10.616**	9.998***	9.039**	9.245**	9.710**	12.367**	14.460**	14.817
	(4.474)	(4.145)	(3.792)	(3.739)	(3.994)	(4.347)	(4.994)	(6.610)	(9.697)
0 km × In Prog	10.287***	8.796***	7.443**	6.315**	5.978*	7.059**	8.504**	11.503**	14.337*
	(3.832)	(3.368)	(2.912)	(2.887)	(3.089)	(3.491)	(4.022)	(5.138)	(8.024)
Avg Price	151.873	189.693	217.877	243.518	269.620	298.920	334.087	382.495	468.804
Adj R^2	0.827	0.881	0.902	0.912	0.916	0.916	0.913	0.905	0.879
Clusters	12,770	12,770	12,770	12,770	12,770	12,770	12,770	12,770	12,770
Obs	48,584	48,584	48,584	48,584	48,584	48,584	48,584	48,584	48,584

Notes: This table presents the results from re-estimating our main effects of interest on housing price from Table 2 but adding house characteristic controls (% 1 bedrooms, % 2 bedrooms, ... % 4 bedrooms, %5+ bedrooms). All sample and controls are the same as in Table 2.

Table A.5: Robustness Check: Neighborhood-level Controls

Dep. var: Price^{kth}	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	11.150**	9.796**	9.134**	8.097**	8.280**	8.681**	11.233**	13.303**	13.649
	(4.389)	(4.096)	(3.765)	(3.746)	(4.011)	(4.379)	(5.080)	(6.731)	(9.977)
0 km × In Prog	9.913***	8.331**	6.883**	5.647*	5.258*	6.289*	7.644*	10.531**	13.251
	(3.798)	(3.350)	(2.929)	(2.944)	(3.164)	(3.577)	(4.164)	(5.299)	(8.324)
Avg Price	151.846	189.670	217.857	243.500	269.602	298.901	334.071	382.485	468.806
Adj R^2	0.825	0.879	0.899	0.909	0.913	0.913	0.910	0.902	0.877
Clusters	12,768	12,768	12,768	12,768	12,768	12,768	12,768	12,768	12,768
Obs	48,569	48,569	48,569	48,569	48,569	48,569	48,569	48,569	48,569

Notes: This table presents the results from re-estimating our main effects of interest on housing price from Table 2 but adding neighborhood controls (house count, pop density, % white, % black and % Hispanic). All sample and controls are the same as in Table 2.

Table A.6: Robustness Check: Land Use Controls

Dep. var: Price^{kth}	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	11.150**	9.796**	9.134**	8.097**	8.280**	8.681**	11.233**	13.303**	13.649
	(4.389)	(4.096)	(3.765)	(3.746)	(4.011)	(4.379)	(5.080)	(6.731)	(9.977)
0 km × In Prog	9.913***	8.331**	6.883**	5.647*	5.258*	6.289*	7.644*	10.531**	13.251
	(3.798)	(3.350)	(2.929)	(2.944)	(3.164)	(3.577)	(4.164)	(5.299)	(8.324)
Avg Price	151.846	189.670	217.857	243.500	269.602	298.901	334.071	382.485	468.806
Adj R^2	0.825	0.879	0.899	0.909	0.913	0.913	0.910	0.902	0.877
Clusters	12,768	12,768	12,768	12,768	12,768	12,768	12,768	12,768	12,768
Obs	48,569	48,569	48,569	48,569	48,569	48,569	48,569	48,569	48,569

Notes: This table presents the results from re-estimating our main effects of interest on housing price from Table 2 but adding land-use controls at the 1990 county level, each interacted with year, to the main specification. Information on lot size is unavailable, but land use is available, however only at the county level for 1990, where it is tabulated in two levels: "In housing units on properties of less than 1 acre" and "In housing units on properties of 1 acre or more." We construct the variable on rural/urban land use from NHGIS and merged each into our final data set based on county. We then convert each variable into a percent, by dividing by total people represented. See Appendix Section C.11 for more details on variable construction. All sample and controls are the same as in Table 2.

Table A.7: Robustness Check: House Characteristic, Neighborhood-level, and Land Use Controls

Dep. var: $Price^{kth}$	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	10.262**	9.005**	8.385**	7.328*	7.434*	7.848*	10.439**	12.361*	11.866
	(4.344)	(4.073)	(3.762)	(3.743)	(4.034)	(4.423)	(5.097)	(6.686)	(9.846)
0 km × In Prog	9.594**	7.991**	6.643**	5.434*	5.032*	5.962*	7.522*	10.517**	13.323
	(3.729)	(3.243)	(2.844)	(2.852)	(3.057)	(3.477)	(4.046)	(5.164)	(8.156)
Avg Price	151.575	189.358	217.526	243.150	269.236	298.520	333.670	382.060	468.348
Adj R^2	0.829	0.883	0.902	0.912	0.915	0.915	0.911	0.903	0.877
Clusters	12,733	12,733	12,733	12,733	12,733	12,733	12,733	12,733	12,733
Obs	48,471	48,471	48,471	48,471	48,471	48,471	48,471	48,471	48,471

Notes: This table presents the results from re-estimating our main effects of interest on housing price from Table 2 but adding the following controls to our primary specification: year by 1990 bedrooms, house count, population density, racial, and rural/urban land use/acreage variables. See Appendix Section C.11 for more details on variable construction. All sample and controls are the same as in Table 2.

Table A.8: Robustness Check: Price Impacts of Cleanup by Decile, Near-Far Comparison, $(5, 10]$ km as Comparison Group

Dep. var: $Price^{kth}$	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	17.676***	14.287***	13.470***	11.631**	10.918**	11.183**	13.517**	16.290**	19.818*
	(5.267)	(5.090)	(4.708)	(4.712)	(5.069)	(5.487)	(6.287)	(7.864)	(11.455)
0 km × In Prog	10.601**	8.390**	6.826*	5.217	4.749	6.257	7.441	12.505**	20.714**
	(4.331)	(4.129)	(3.638)	(3.645)	(3.863)	(4.237)	(4.896)	(6.162)	(9.684)
% Impact:									
0 km × Post	0.110***	0.072***	0.059***	0.046**	0.039**	0.036**	0.039**	0.041**	0.040*
	(0.033)	(0.025)	(0.021)	(0.018)	(0.018)	(0.018)	(0.018)	(0.020)	(0.023)
0 km × In Prog	0.066**	0.042**	0.030*	0.020	0.017	0.020	0.021	0.031**	0.042**
	(0.027)	(0.021)	(0.016)	(0.014)	(0.014)	(0.014)	(0.014)	(0.015)	(0.020)
Avg Price	160.490	199.722	228.930	255.413	282.537	312.995	349.569	399.935	489.504
Adj R^2	0.812	0.870	0.892	0.903	0.907	0.907	0.905	0.897	0.872
Clusters	7,655	7,655	7,655	7,655	7,655	7,655	7,655	7,655	7,655
Obs	29,092	29,092	29,092	29,092	29,092	29,092	29,092	29,092	29,092

Notes: This table re-estimates the price impacts from Table 2, except compares tracts 0 km away from a RCRA site to those $(5, 10]$ km away from the site. The purpose of this exercise is to address SUTVA concerns. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(5, 10]$ km away from a facility. All standard errors are clustered on census tract.

Table A.9: Rent Impacts

Dep. var: Rent^{kth}	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	2.630	4.748	6.930	5.075	-8.344	-14.918	-4.506	0.990	8.164
	(15.488)	(14.620)	(14.304)	(14.218)	(14.866)	(16.536)	(20.248)	(24.470)	(31.977)
0 km × In Prog	-3.659	-6.203	-5.896	-9.938	-12.972	-15.561	-4.219	-4.588	-0.432
	(13.325)	(12.505)	(12.230)	(12.389)	(12.782)	(14.620)	(16.123)	(18.732)	(24.797)
Avg Rent	650.526	789.466	890.090	976.961	1,061.540	1,151.139	1,252.790	1,384.002	1,594.734
Adj R^2	0.715	0.769	0.796	0.808	0.819	0.825	0.824	0.814	0.780
Clusters	12,458	12,458	12,458	12,458	12,458	12,458	12,458	12,458	12,458
Obs	44,790	44,790	44,790	44,790	44,790	44,790	44,790	44,790	44,790

Notes: This table re-estimates the price impacts from Table 2 except uses deciles of the rental price as the dependent variable, where we apply our decile calculation procedure to rental variables instead of home prices. Whereas housing prices are in $1,000 units throughout, rent is in dollars. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(5, 10]$ km away from a facility. All standard errors are clustered on census tract.

A.2 Appendix Figures

Figure A.1: Histograms: 10^{th}–90^{th} Percentiles of Housing Price

(a) 10^{th} Percentile

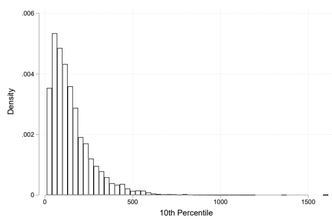

(b) 20^{th} Percentile

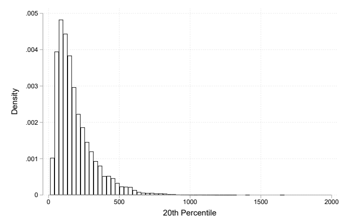

(c) 30^{th} Percentile

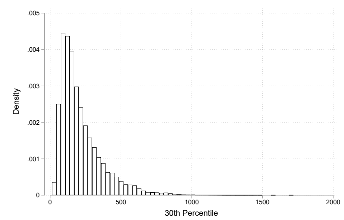

(d) 40^{th} Percentile

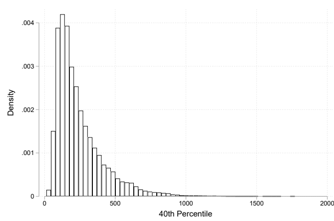

(e) 50^{th} Percentile

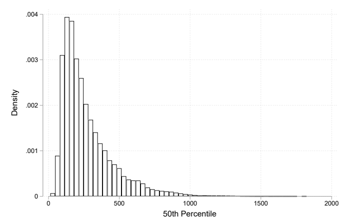

(f) 60^{th} Percentile

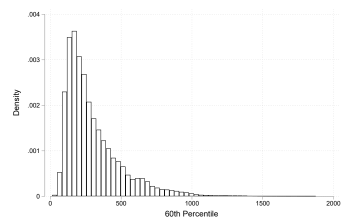

(g) 70^{th} Percentile

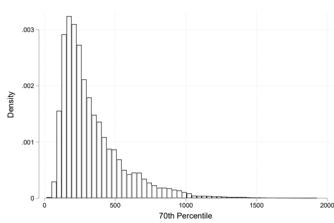

(h) 80^{th} Percentile

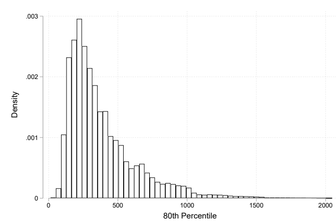

(i) 90^{th} Percentile

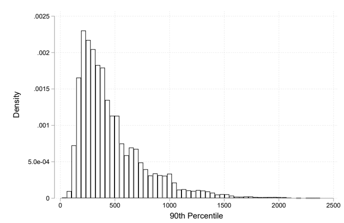

Notes: This figure shows histograms of prices from each percentile of prices from the 10^{th} to the 90^{th} in order to gauge the skewness of prices across percentiles. The sample includes all tracts in the NHGIS shapefile.

Figure A.2: Spatial Visualization of Tracts Near CAP Sites

(a) All Possible Tracts within 10 km of a Site

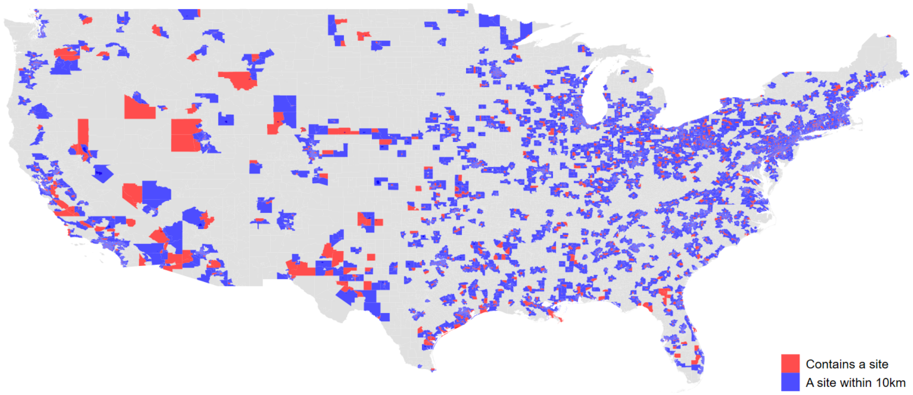

(b) Tracts with 1 Site within 10 km (main sample)

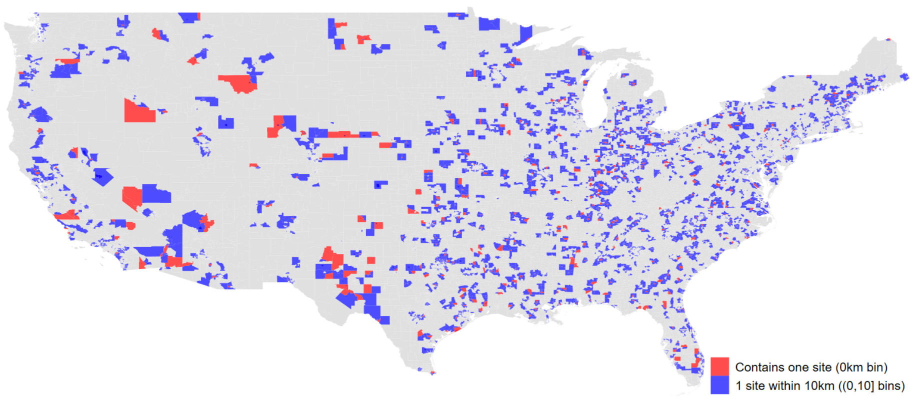

(c) Tracts with Multiple Sites within 10 km (generalized sample)

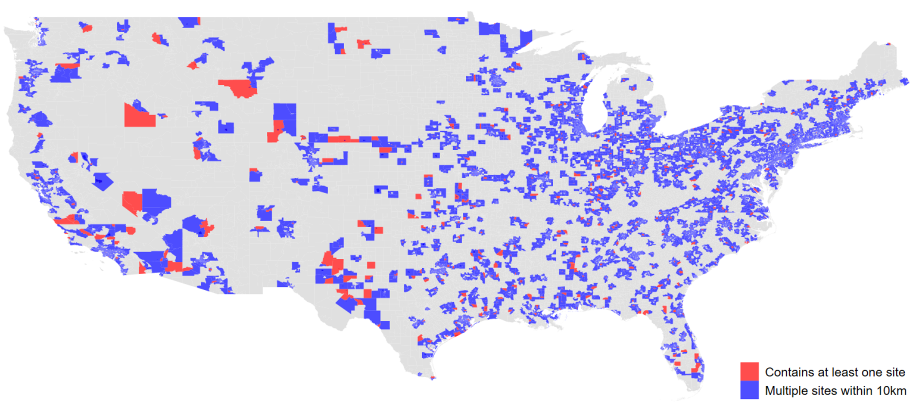

Notes: Panel (a) shows tracts restricting to within 1 km of a RCRA cleanup. Panel (b) is our main sample of tracts restricting to a single site within 10 km. Panel (c) is a generalized sample that restricts to tracts with at least one site with all sites within 10 km contained in the tract and control tracts with no contained sites but at least one site within 10 km. Note that Alaska and Hawaii are not pictured, but do contain sites present in our dataset. Tracts are only included if they are in the NHGIS 2010 tract-level shapefile, which removes entirely water-based tracts.

Figure A.3: Main Results with Alternative Clustering Levels

(a) Site

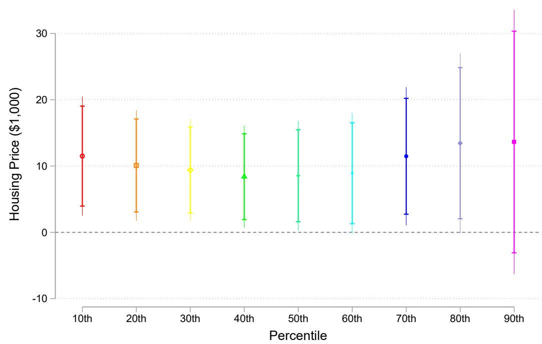

(b) County

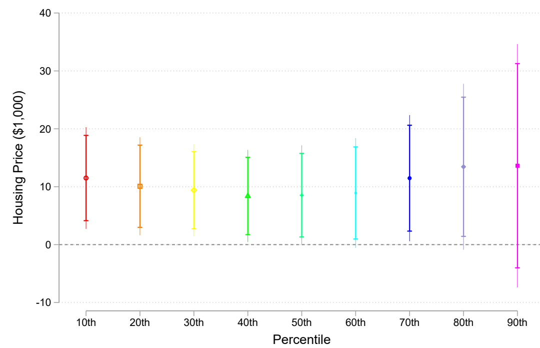

(c) Site × Year

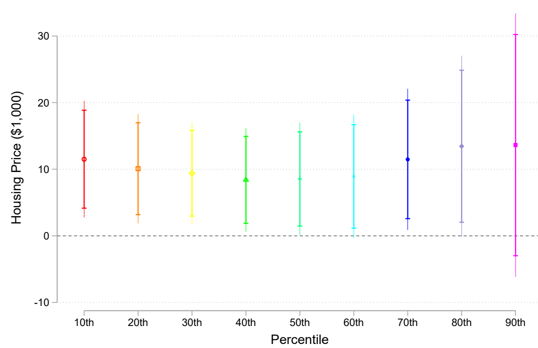

Notes: This figure plots the DD coefficients by decile of housing price impact from varying the clustering levels in our main results (Table 2). Each effect depicted is a coefficient on the interaction $Near_i^d \cdot Post_{r(i),t}$ from (1), where $Near_i^d$ is the indicator for the 0 km bin. All specifications include fixed effects for tract, bin by year, site by year, and state by year. The sample is the set of tracts containing exactly one RCRA site within 10km in 1990, 2000, 2010 (ACS 2008-2012), and 2020 (ACS 2018-2022). Standard errors are clustered at the site level (Panel a), the county level (Panel b) and the site by year level (Panel c). Whiskers marked with horizontal lines and vertical protruding segments indicate 95% and 99% CIs respectively.

Figure A.4: 10^{th}–90^{th} Percentiles of Housing Price Over Time

(a) 10^{th} Percentile

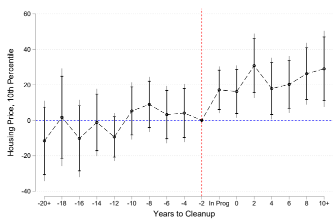

(b) 20^{th} Percentile

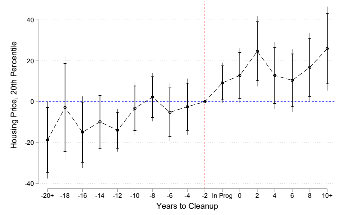

(c) 30^{th} Percentile

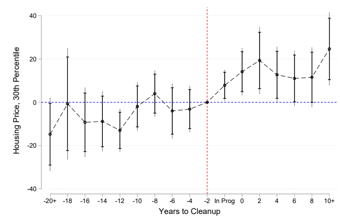

(d) 40^{th} Percentile

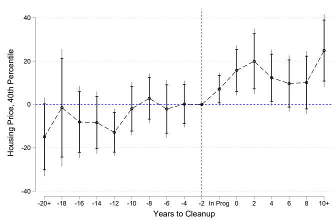

(e) 50^{th} Percentile

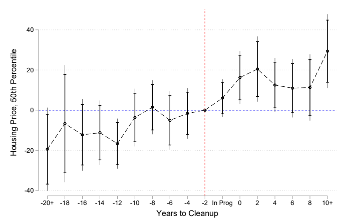

(f) 60^{th} Percentile

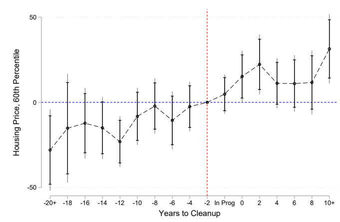

(g) 70^{th} Percentile

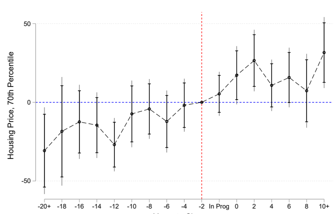

(h) 80^{th} Percentile

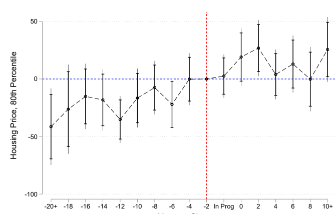

(i) 90^{th} Percentile

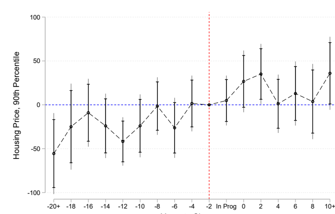

Notes: This figure re-estimates the price impacts at each percentile of the price distribution (as denoted in the subcaption) (similar to Figure 5 in the main text). It shows the coefficient representing the difference in the near (0 km) and far ((0, 10] km) bins over time from the event study specification in equation 14. We use the same fixed effects and sample as in the main regression. The coefficient for the two years just prior to the cleanup (at position -2) is normalized to 0 by excluding the dummy on Near× Event time=-2 from the regression. Data from the during-cleanup phase is represented by “In Prog,” and “0” represents the two-year immediately following cleanup completion. Whiskers marked with horizontal lines and vertical protruding segments indicate 95% and 99% CIs respectively, clustering at the tract level. Prices are denominated in thousands of dollars.

Figure A.5: Income and Education-related Variables Over Time

(a) Avg HH Income

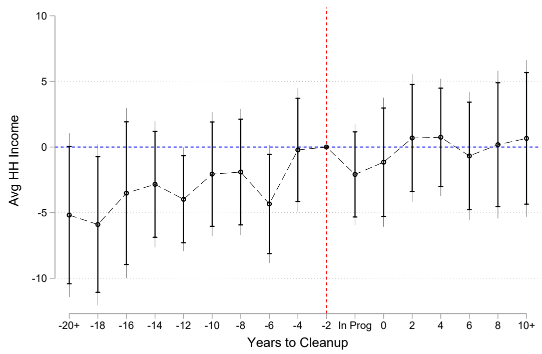

(b) % Below Poverty Line

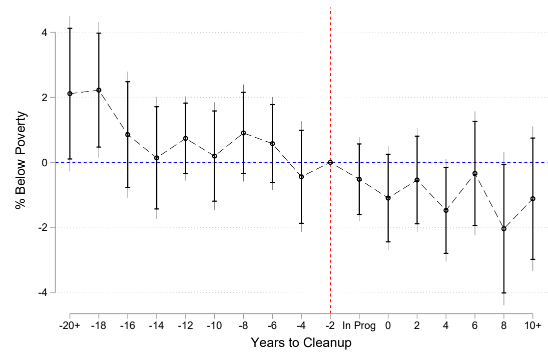

(c) % Public Assistance

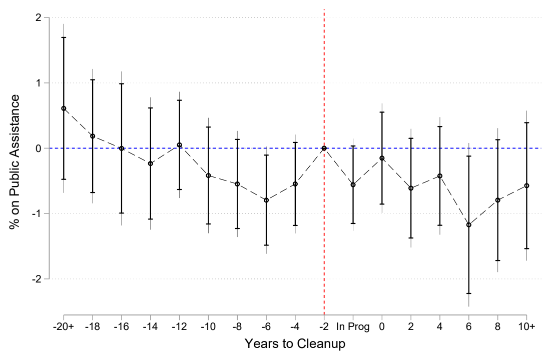

(d) % Unemployment

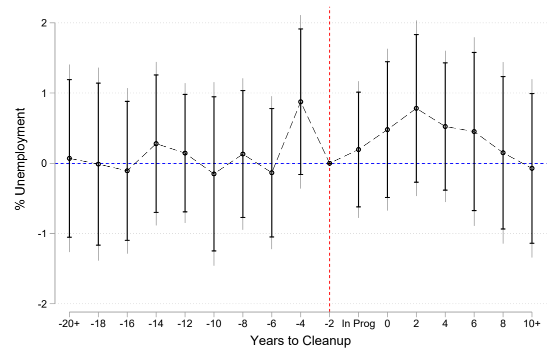

(e) % College Educated

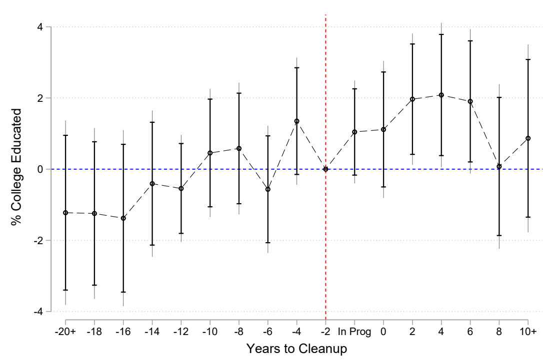

Notes: This figure shows the coefficient representing the difference in the near (0 km) and far ((0, 10] km) bins over event time for each outcome indicated in the sub-caption. We use the same fixed effects as in the main regression. The coefficient for the two years just prior to the cleanup (at position -2) is normalized to 0 by excluding the dummy on Near× Event time=-2 from the regression. Data from the during-cleanup phase is represented by "In Prog," and "0" represents the two-year immediately following cleanup completion.

Figure A.6: Demographic Variables Over Time

(a) % Black

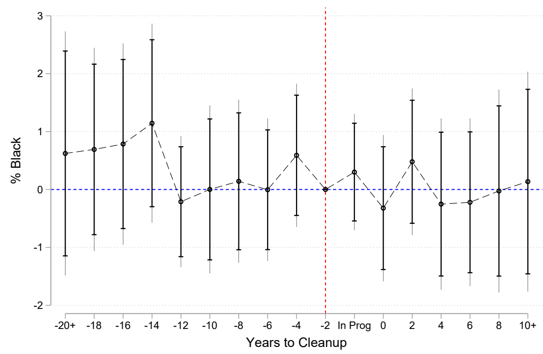

(b) % Hispanic

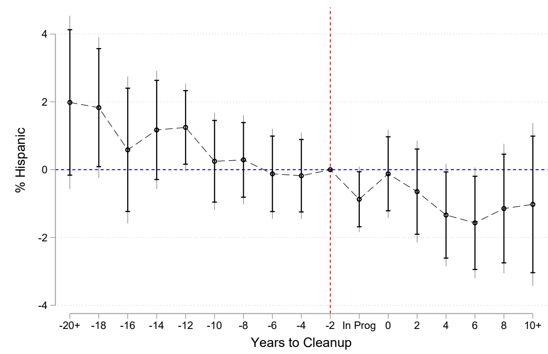

(c) % White

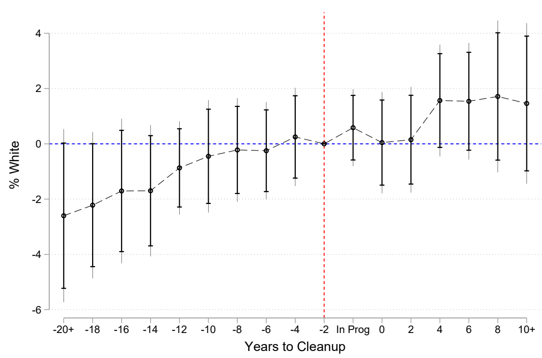

(d) % Female Head of Household

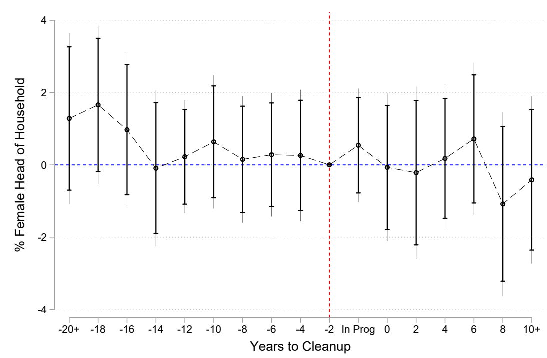

(e) % Under 18

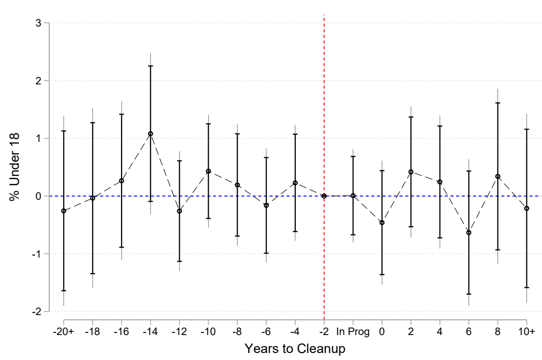

(f) Population Density

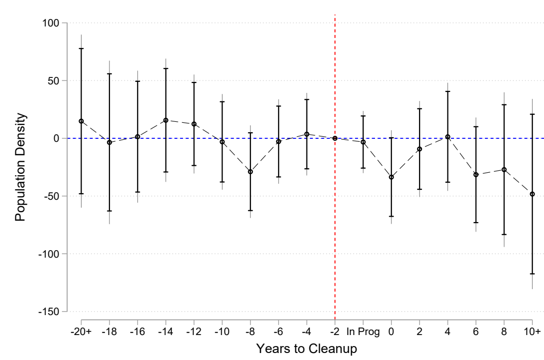

Notes: This figure shows the coefficient representing the difference in the near (0 km) and far ((0, 10] km) bins over event time for each outcome indicated in the sub-caption. We use the same fixed effects as in the main regression. The coefficient for the two years just prior to the cleanup (at position -2) is normalized to 0 by excluding the dummy on Near× Event time=-2 from the regression. Data from the during-cleanup phase is represented by "In Prog," and "0" represents the two-year immediately following cleanup completion.

Figure A.7: Housing-related Variables Over Time

(a) % 4+ Beds

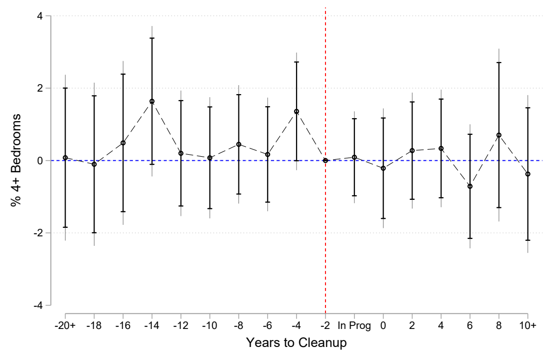

(b) % Mobile Home

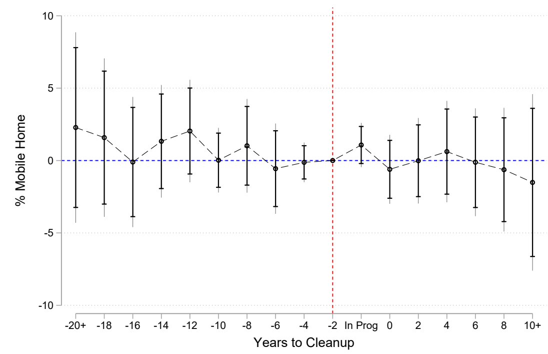

(c) % Owner Occupied

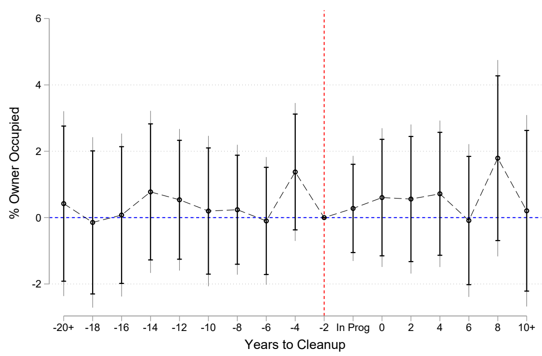

(d) % Vacant

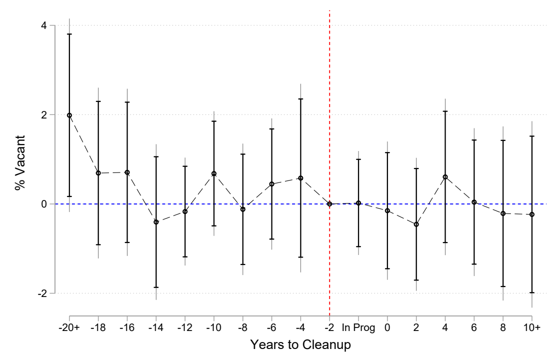

(e) % Built in Last 10 years

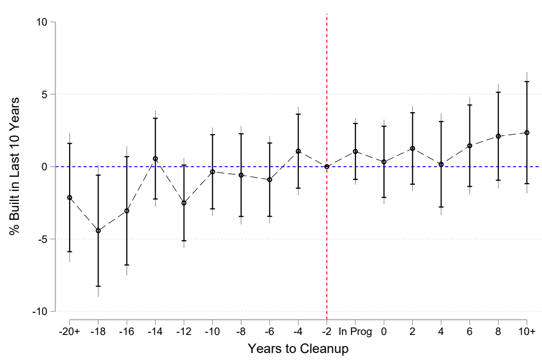

(f) % Moved in Last 10 Years

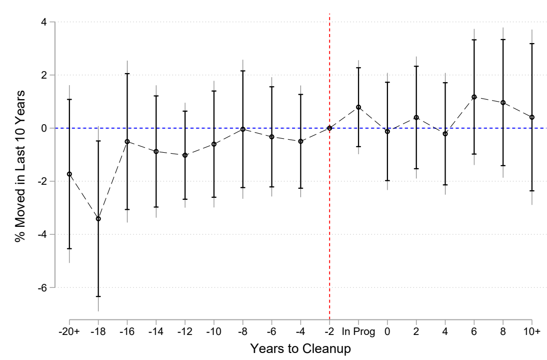

Notes: This figure shows the coefficient representing the difference in the near (0 km) and far ((0, 10] km) bins over event time for each outcome indicated in the sub-caption. We use the same fixed effects as in the main regression. The coefficient for the two years just prior to the cleanup (at position -2) is normalized to 0 by excluding the dummy on Near× Event time=-2 from the regression. Data from the during-cleanup phase is represented by "In Prog," and "0" represents the two-year immediately following cleanup completion.

Figure A.8: 10^{th}–90^{th} Percentiles of Rental Price Over Event Time

(a) 10^{th} Percentile

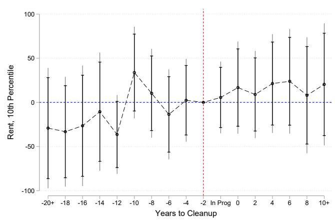

(b) 20^{th} Percentile

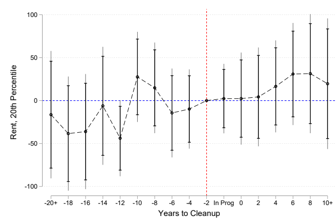

(c) 30^{th} Percentile

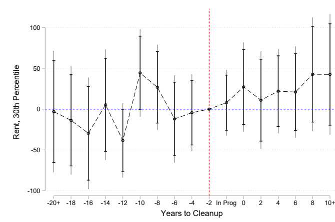

(d) 40^{th} Percentile

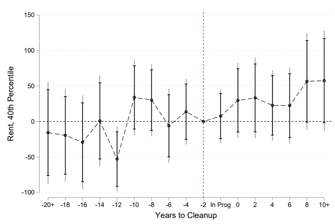

(e) 50^{th} Percentile

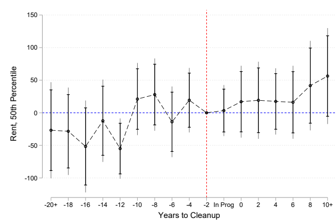

(f) 60^{th} Percentile

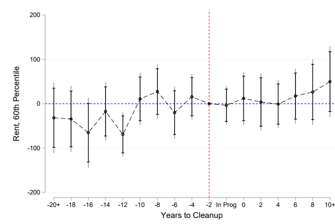

(g) 70^{th} Percentile

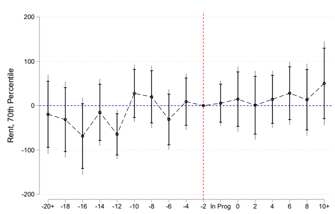

(h) 80^{th} Percentile

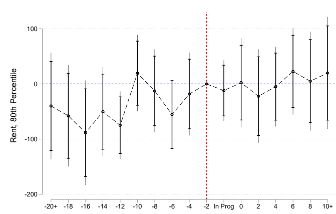

(i) 90^{th} Percentile

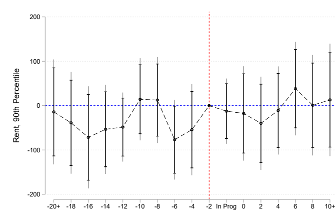

Notes: This figure re-estimates the rental price impacts at each percentile of the price distribution (as denoted in the subcaption). It shows the coefficient representing the difference in the near (0 km) and far ((0, 10] km) bins over time from the event study specification in equation 14. We use the same fixed effects and sample as in the main regression. The coefficient for the two years just prior to the cleanup (at position -2) is normalized to 0 by excluding the dummy on Near× Event time=-2 from the regression. Data from the during-cleanup phase is represented by "In Prog," and "0" represents the two-year immediately following cleanup completion. Whiskers marked with horizontal lines and vertical protruding segments indicate 95% and 99% CIs respectively, clustering at the tract level. Whereas housing prices are in \$1,000 units throughout, rent is in dollars.

Figure A.9: Sites Represented Over Event Time

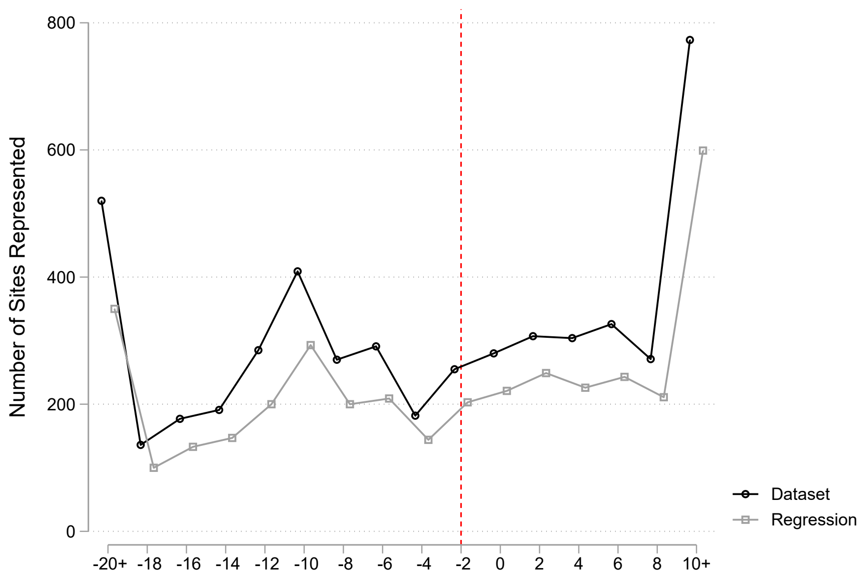

Notes: This figure shows the number of RCRA sites represented in each event time, where event times are defined as the difference in years between the survey year and the cleanup period, in two-year intervals. The blue line depicts the count of sites for the entire dataset, whereas the red line depicts the same count, only for sites that appear in our main regression.

B Supplementary Appendix on Individual Transaction Data

We obtained Zillow Ztrax data from Ohio and Pennsylvania for comparison purposes. This data contains individual housing transactions spanning the years 1994-2019,[1] collected from publicly available sources such as county offices.

- Section B.1 compares our survey-based housing price data to other sources of housing price data, including Zillow
- Section B.2 compares exposure to RCRA based on tract-based distance versus property-based distance
- Section B.3 reproduces our main results using Census data with OH-PA tracts only (the Zillow ZTRAX sample)
- Section B.4 estimates our main models with transactions level data
- Section B.5 provides a description of the Zillow dataset that we use (used in Harleman (2024)).

[1] There are only a negligible number of observations from before 1994 (about 0.04%).

B.1 Survey Based House Values vs. Other Price Data Sources

In this appendix section, we compare our survey-based housing price data from the Census to data from other sources. One key question is the extent to which people over-estimate the value of their house more at different portions of the price distribution. We first compare our Census data to Zillow data for PA and OH. Figure B.1 shows the average deciles of housing price for Zillow and the Census at the tract level. Like previous work, we find that survey-based responses from the Census are systematically higher than Zillow estimates. However, the figure also shows that the degree of potential overvaluation does not systematically vary by price percentile (i.e., it is fairly parallel with the 45 degree line). In addition, we plot within-percentile comparisons of price in Figure B.2. Each subfigure depicts a binned scatterplot of house values comparing the Census percentiles to tract-level deciles from Zillow. We find a similar story in terms of the magnitude of the bias across the range of home values in each quantile.

Figure B.1: Average Price Percentiles from Zillow vs. Census Data

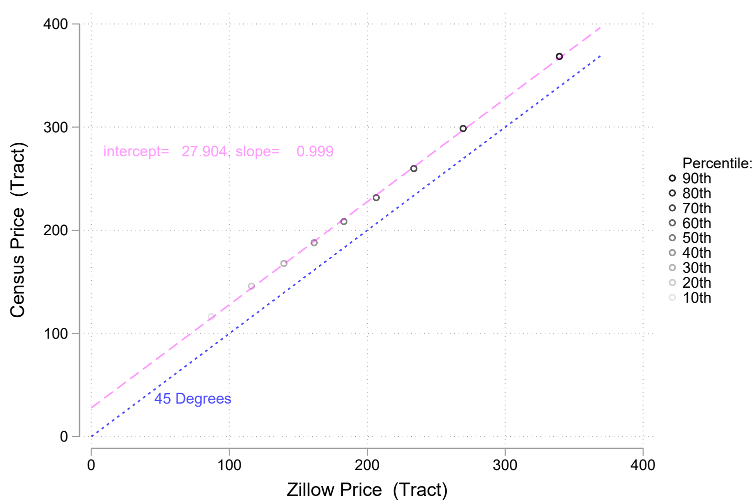

Notes: This figure plots the average deciles for tract-level Zillow and Census prices. The pink line depicts the line of best fit from the nine points, and the pink text depicts the intercept and slope of said line. A slope of 1 would indicate that the degree of overvaluation is constant across price percentiles.

Figure B.2: Census vs. Zillow Price By Decile

(a) 10^{th} Percentile

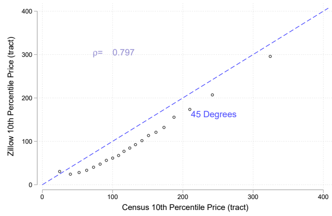

(b) 20^{th} Percentile

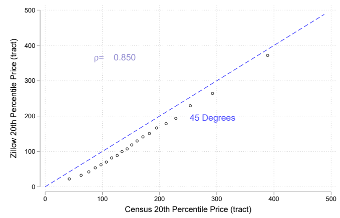

(c) 30^{th} Percentile

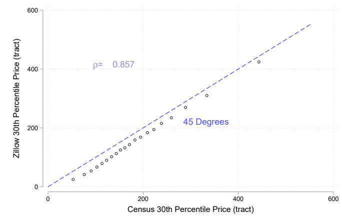

(d) 40^{th} Percentile

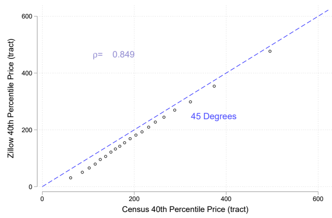

(e) 50^{th} Percentile

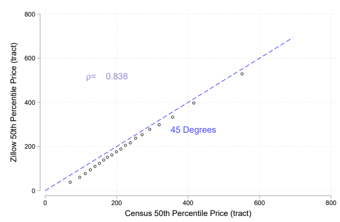

(f) 60^{th} Percentile

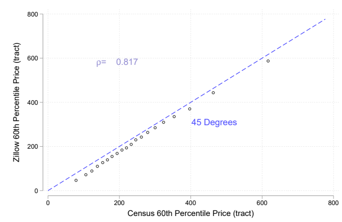

(g) 70^{th} Percentile

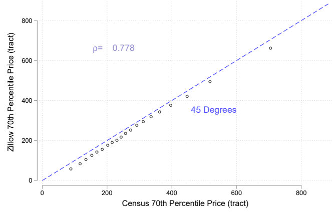

(h) 80^{th} Percentile

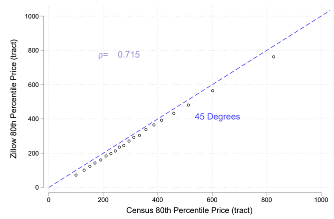

(i) 90^{th} Percentile

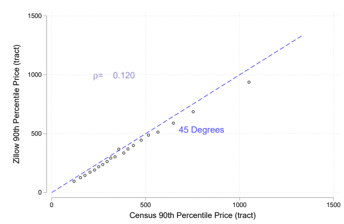

Notes: This figure compares Zillow to Census tract-level deciles of price, using a binned scatter plot with 20 bins. The 45 degree line is indicated, as well as the correlation coefficient between the variables on the axes. In the sample of Zillow houses used to construct tract-level deciles, we did not use tracts where there were less than 9 observations.

We also proceed with another point of comparison using Redfin's "All Residential Sales" data, which is available starting in 2012. Figure B.3 shows state-level median prices from Zillow Ztrax compared to Redfin for the states of PA and OH. We see that Redfin data has a much higher median than Zillow Ztrax data for both states. We overlay the median from our Census tract-level dataset. Our Census data is closer to Redfin's median in both cases. The right-hand panel presents the transaction volume for Zillow vs. Redfin. Zillow transaction volume is close to Redfin for Ohio, but not for Pennsylvania. Our Zillow data ends in the first quarter of 2020, leading to a drop for that year. Other than 2020, the graph shows that Zillow's transaction volume is higher than Redfin's for each year, and closest in 2019. This reflects the possibility that some Zillow records are re-recordings, mortgages, and other non-sale-related documents in years prior to 2019, even though we have taken steps to clean the data according to best practices.

Lastly, we want to point out that an alternative source of data that could be considered is assessed values. However, the literature on assessed values suggests that they are inaccurate rep-

Figure B.3: Comparing OH and PA State-Level Statistics

(a) Median Price

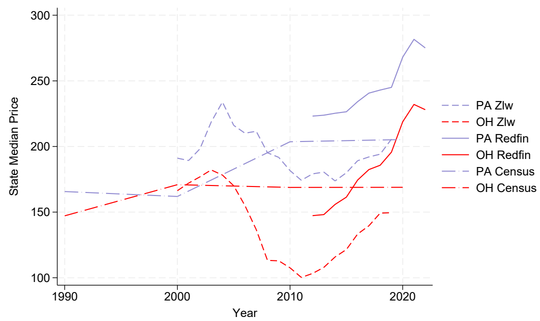

(b) Transaction Volume

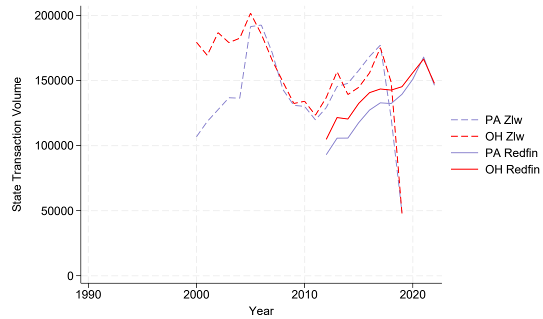

Notes: This figure compares prices and volume of transactions for OH and PA across three data sources: Zillow ZTRAX, Redfin, and Census. Panel (a) presents state level median price computed from Ztrax data, state-level median price from Redfin transactions, and the median of our 50th percentile price (where median is taken across tracts). Panel (b) shows transaction volumes computed by Zillow Ztrax versus that reported by Redfin.

resentations of actual values (Du Preez and Sale, 2015), as they tend to systematically over-estimate house values because of skewed incentives (Amornsiripanitch, 2020), and thus suffer from their own distributional biases (Avenancio-León and Howard, 2019), making them particularly unsuitable to a study of differential impacts at different percentiles of the price distribution. They also might not reflect the change in transaction prices from before to after cleanup, as they are updated infrequently (Paglin and Fogarty, 1972) and capped in the magnitude of change that is possible over time (Dornfest, Rearich, Brydon III and Almy, 2019). When it comes to non-disclosure states in particular, tax assessors work off of worse information in these states because sale prices are not available (Berrens and McKee, 2004). Furthermore, tax assessment data is likely to be more biased downward in non-disclosure states given homeowner incentives to appeal only if they exceed the sale price (Berrens and McKee, 2004), leading to significant downward bias due to non-disclosure of sale prices (Bollum, 2021). Given that information has been shown to play a vital role in how accurate assessments are (Weber and McMillen, 2010), any systematic biases in assessments might be compounded in non-disclosure states. That said, like assessors, homeowners probably also lack information to properly value their houses when they do not have access to sales information for neighboring or comparable houses. However, to the extent that some of the above issues are driven by asymmetric information and differential incentives (e.g. selective appealing), it is less likely that these issues would to extend to homeowner reporting of valuations in the Census.

B.2 Understanding Census Tract-based Treatment and Comparison Groups

We first use the transaction data to understand our treatment effects in the context of actual distance from a house to a RCRA site. Because we employ Census tract-level data in this paper, we define distance to a RCRA site as the minimum distance from the Census tract to the site. This means that we may include in our treatment group homes relatively far from a site but still on the same tract. This is a form of measurement error that could cause attenuation of our estimates.

To investigate the extent of this issue, we provide information on actual distance to site in Figure B.4. In panel B.4a, we show summary statistics by bin, where the bins are defined as the minimum distance from tract to site. The median, 5th, and 95th percentiles are all increasing in census distance bin. Panel B.4b of the same figure shows the kernel densities of actual distance to site for both the treatment (same tract as RCRA site) and the comparison (Census tract (0-10] km away) groups. The distributions demonstrate a small amount of overlap- for the most part, houses in our census-tract based treatment group are near the sites, and those in our comparison group are far from RCRA sites. In summary, actual distance to site is well-proxied by Census tract to site, at least relatively speaking.

Individual transaction data also allows us to examine the price distribution. In Figure B.5, we show the kernel density plots of price for both the treatment and comparison groups, measured prior to cleanup. The distributions are similar.

Figure B.4: Actual Distance to Site vs. Census Tract Distance

(a) Actual Distance: Summary Statistics by Census tract-based Distance Bin

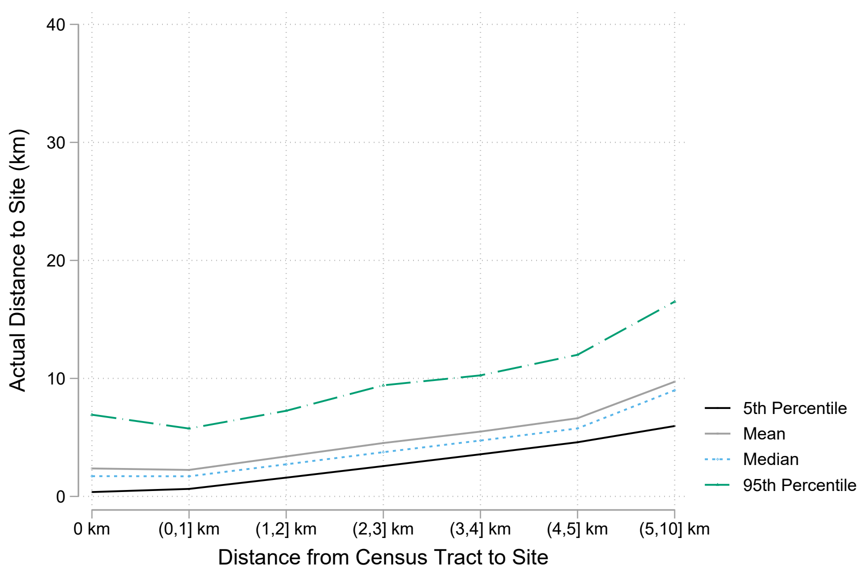

(b) Distance to Site Kernel Density, Census Tract-based Treatment and Comparison Groups

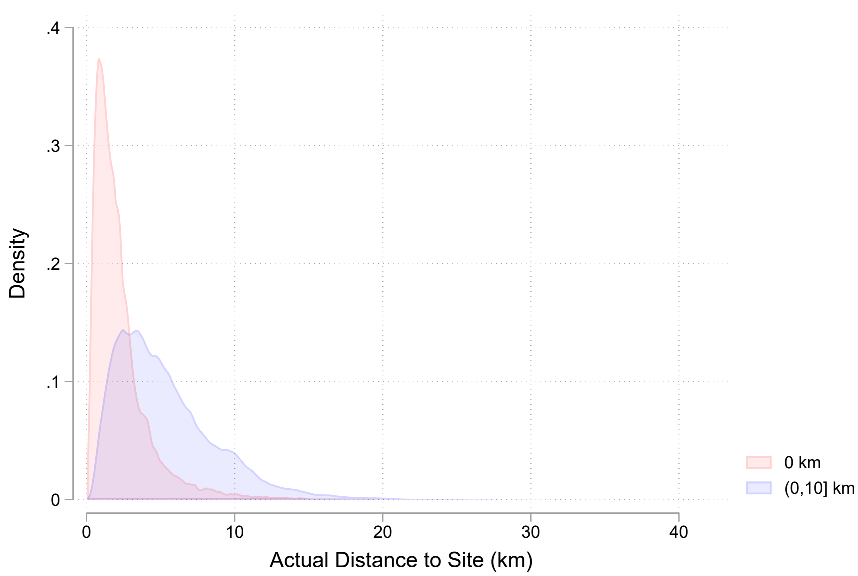

Notes: This figure depicts summary statistics and a kernel density of actual distances to site for houses in Ohio and Pennsylvania from the Zillow Ztrax database, for the treatment and comparison groups in our main analysis (which are defined by distance from Census tract to site).

Figure B.5: Price Kernel Density, Census Tract-based Treatment and Comparison Groups

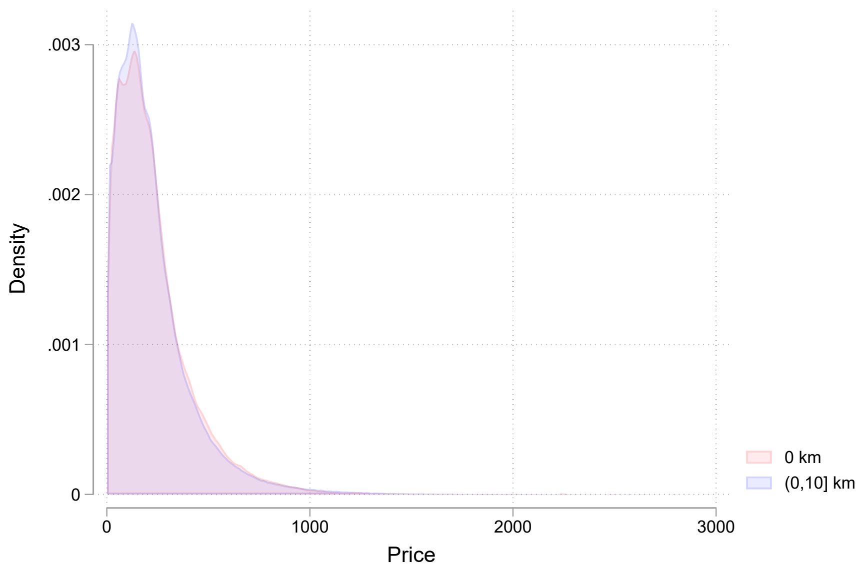

Notes: This figure depicts a kernel density of pre-cleanup transaction prices for houses in Ohio and Pennsylvania from the Zillow Ztrax database, separately for the treatment and comparison groups defined in our main analysis. Note that the figure only plots prices less than 3 million dollars (3000 in thousand dollar units) for readability. These discarded prices account for less than 1% of observations. Prices are adjusted to 2023 dollars.

B.3 Results Using Census Data With OH-PA Subset of Tracts

Here, we re-estimate our main specification using the sample of tracts in OH and PA. Price effects generally follow our main sample results in sign, with larger magnitudes, but are less precise (likely due to the restricted sample size).

Table B.1: Main Spec, OH-PA Sample

Dep. var: Price^{kth}	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	22.940*	12.968*	8.392	9.050	10.409	13.796	19.231	16.378	21.231
	(13.615)	(6.871)	(5.856)	(5.931)	(6.897)	(9.904)	(12.477)	(15.942)	(19.489)
0 km × In Prog	11.037	3.424	3.362	6.997*	7.803*	8.814	12.155*	17.896**	16.176
	(10.638)	(4.534)	(3.426)	(3.697)	(4.425)	(5.800)	(6.698)	(8.225)	(13.073)
Avg Price	95.966	123.455	143.634	162.031	181.147	202.627	228.495	264.652	328.496
Adj R^2	0.827	0.896	0.926	0.931	0.929	0.922	0.911	0.889	0.861
Clusters	721	721	721	721	721	721	721	721	721
Obs	2,774	2,774	2,774	2,774	2,774	2,774	2,774	2,774	2,774

Notes: This table presents estimates from our main specification where the sample consists only of tracts in Ohio and Pennsylvania.

B.4 Quantile Treatment Effects Using Individual Transaction Data

In this sub-section, we use Zillow data to show that our main results are robust to use of individual transaction-level data.

B.4.1 Quantile Treatment Effect Approach

We employ a recentered influence function approach (Firpo, 2007; Firpo, Fortin and Lemieux, 2009; Havnes and Mogstad, 2015; Dube, 2019) to estimate quantile treatment effects of cleanup on housing prices using individual housing transaction data.

The recentered influence function, or RIF, is a transformation that represents a distributional statistic.[2] For the τ^{th} quantile q_τ, the RIF is:

$$RIF(p, q_\tau) = q_\tau + \frac{\tau - \mathbb{1}\{p \leq q_\tau\})}{f(q_\tau)} \tag{B.1}$$

In the above, p is the outcome variable and $f(q_\tau)$ is the population distribution function for the outcome variable.

Estimation is a two-step process: first, obtain the RIF for each value of the outcome variable (price). The only element of the expression that varies across observations is $\mathbb{1}\{p \leq q_\tau\}$, which is estimated as a linear probability model. Second, regress the RIF against $Near_i^0$, $Post_{r(i),t}$, and $Near_i^0 \times Post_{r(i),t}$:

$$RIF_i(q_\tau) = \beta_1 Near_i^0 \times Post_{r(i),t} + \beta_2 Near_i^0 \times InProg_{r(i),t} + \beta_3 Post_{r(i),t} + \beta_4 InProg_{r(i),t} + \beta_5 Near_i^0 + \epsilon_{i,t} \tag{B.2}$$

The coefficient β_1 is interpreted as an unconditional partial effect (UPE) of cleanup on the τ^{th} quantile of the price distribution. The identifying assumption is that, in the absence of treatment, the change in shares or frequency of houses from before to after treatment around a given level of the housing price would be the same in the treatment group as in the comparison group.[3]

The estimator is invariant to monotonic transformations of the outcome variable. Therefore, it allows a common trend in the outcome variable, both in absolute and relative terms.[4] As Firpo et al. (2009) recommend, we bootstrap the standard errors (clustering at the tract level to parallel other parts of this paper).

[2]It amounts to re-centering an influence function. In (B.1), the second term is the influence function.

[3]Following the notation in Havnes and Mogstad (2015), assume that $F_t(p)$ is the counterfactual price distribution for houses on the same tract as the site, and $G_t(p)$ is the counterfactual price distribution for houses in the comparison group (houses on a Census tract 5–10 km away from the site.) We will consider two periods for simplicity, so $t = 0$ indicates pre-cleanup and $t = 1$ indicates post-cleanup. The counterfactual price distribution for post-cleanup houses is taken to be:

$$\mathcal{F}_1(p) = F_0(p) + G_1(p) - G_0(p) \tag{B.3}$$

So, for a given price level p, the estimated quantile treatment effect is the difference between the actual and counterfactual price distribution at that price level, net of the same difference for the comparison group.

[4]Therefore, taking a logarithmic transformation of price would not change our conclusions.

We see this investigation as a robustness check of our main effects. Recall that our main effects are estimated at the Census tract level.

To ensure that results are directly comparable, we use the exact same definition of distance here. That is, we compare houses on the same tract as a RCRA site ("0 km") to houses on a Census tract that is $(0, 10]$ km away from the nearest site. Instead of using distance from house to nearest site, we employ a Census-tract based measure of distance to be consistent with the design used in the rest of the paper. If the percentile regressions using aggregate data in the main body of our paper yield estimates similar to the true quantile treatment effects, we will consider our primary specification to be validated. In order to ensure we compare apples to apples in this exercise, we need to keep the group of houses the same, so we define "Near" and "Far" using minimum Census distance.

B.4.2 Quantile Treatment Effect Results

We plot the estimates of β_1 from (B.2) in Figure B.6; the full results are presented in Table B.2. The results show that only the lower quantile effects are similar to those in Figure 4. The only estimated coefficient that is statistically significant at any reasonable level of significance is the 10th percentile quantile effect, which is statistically significant at the 10% confidence level. The magnitudes of the coefficients are positive and comparable in magnitude to the main results for the lower deciles, but are very small for the 80th percentile and become negative and very large at the 90th percentile (although confidence intervals are so large that we cannot rule out positive effects). The substantive conclusion of our main analysis– that there is a positive cleanup effect for the lower deciles of the price distribution– is corroborated by this robustness check.

Figure B.6: Quantile Treatment Effects, 0 km vs. (0, 10] km

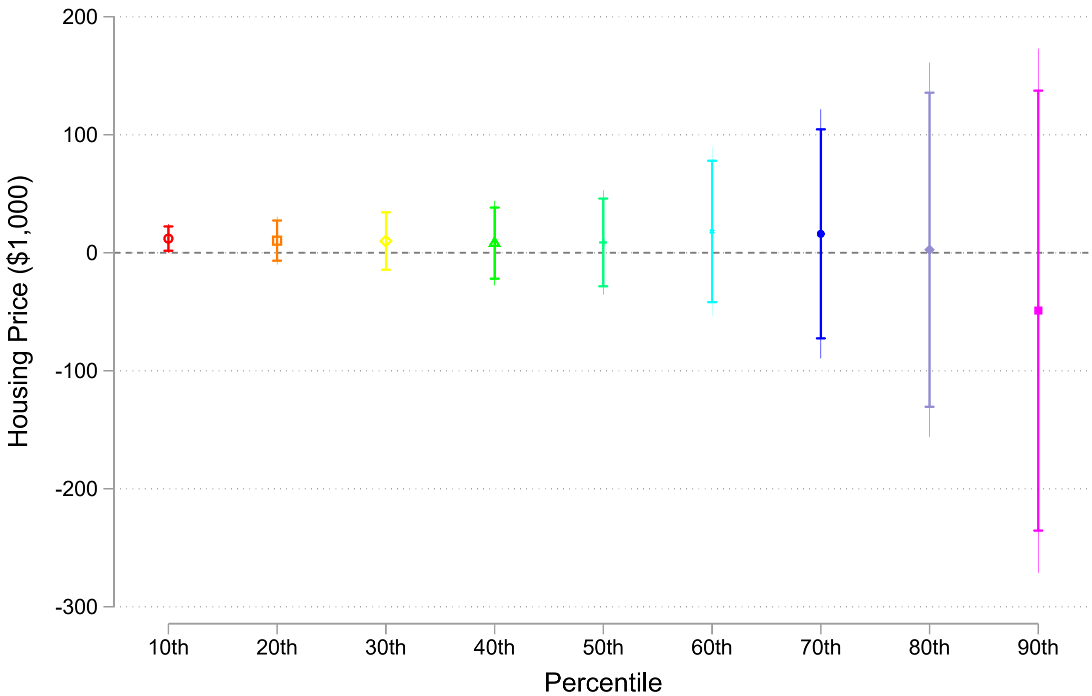

Notes: This figure depicts the estimated Quantile Treatment Effects using the OH-PA Zillow dataset. These are the 0 km ×Post estimates from Table B.2.

Table B.2: Quantile Treatment Effects, OH-PA Transaction Data

Dep. var: Price^{kth}	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	11.988*	10.268	9.915	8.140	8.730	18.046	16.016	2.568	-49.053
	(6.258)	(10.351)	(14.770)	(18.315)	(22.608)	(36.433)	(53.844)	(80.903)	(113.328)
0 km × In Prog	31.681	39.676	48.083*	57.015*	66.819**	72.843	58.718	12.524	-39.086
	(23.256)	(26.189)	(27.162)	(29.413)	(28.910)	(46.065)	(65.310)	(83.760)	(98.780)
Post	-23.121***	-29.336***	-34.427***	-36.994***	-38.113***	-42.673***	-51.858***	-59.706***	-77.140***
	(2.105)	(3.209)	(3.880)	(4.993)	(5.697)	(6.695)	(8.463)	(11.969)	(17.624)
In Prog	-44.780***	-54.253***	-62.150***	-66.737***	-69.483***	-70.108***	-75.756***	-75.264***	-68.604**
	(6.673)	(8.298)	(9.167)	(9.959)	(10.603)	(11.427)	(14.110)	(17.977)	(29.150)
0 km	-20.837***	-28.170**	-34.092**	-37.487*	-44.426*	-51.434	-51.405	-45.827	-13.727
	(7.601)	(11.843)	(15.675)	(20.040)	(23.963)	(38.689)	(54.854)	(82.157)	(113.271)
Avg Price	44.752	77.216	108.843	139.335	172.803	208.528	253.184	318.009	434.216
Clusters	721	721	721	721	721	721	721	721	721
Obs	714,200	714,200	714,200	714,200	714,200	714,200	714,200	714,200	714,200

Notes: This table presents Quantile Treatment Effect estimates using individual housing transaction data from Zillow Ztrax spanning the states of Ohio and Pennsylvania. We use transaction data from all tracts within 10 km of at most one RCRA facility in these regressions. Standard errors are bootstrapped (999 bootstrap replications), clustering on Census Tract to parallel other sections of this manuscript.

B.5 Construction of the Zillow Data Set for OH and PA

In this appendix section, we detail our Zillow data cleaning. We obtained pre-cleaned data as well as code used to produce it (written by researcher Max Harlemann and used for Harleman (2024)), and the source files have been deleted as per our contract with Zillow (Ztrax was discontinued Sept 30, 2023). The code provided by Max filtered for arms-length transactions by:

- Keeping only deed transfers (DataClassStndCode "D", "H", "U", "J"). "U" and "J" apply only to Hawaii so they will not affect our analysis. "D" and "H" are characterized by Nolte, Boyle, Chaudhry, Clapp, Guignet, Hennighausen, Kushner, Liao, Mamun, Pollack, Richardson, Sundquist, Swedberg and Uhl (2024) as highly probable indicators of sale.
- Dropping refinancing (LoanTypeStndCode "RE")
- Dropping intra-family transfers (IntraFamilyTransferFlag "Y" or other value that's not missing), "INTR"
- Dropping transactions exempt from transfer tax (TransferTaxExemptFlag)
- Dropping partial interest transactions (PartialInterestTransferStndCode)

In terms of identifying property types,

- The code uses the AssessmentLandUseStndCode to categorize property types as recommended by Nolte et al. (2024).
- The code drops coop, government, exempt, and recreational properties. Specific dropped indicators are: GV EI HI RC RI EX CP MS PP RC TR AP CP EX IM RC MX PD.

Lastly, the code drops observations with missing or zero values for sales price and lot size.

After receiving the data and code, we enacted three further restrictions to bring the sample in line with best practices due to Nolte et al. (2024):

- We drop sales with price < $1,000.
- We drop records with missing coordinate information, and discard any observation where the lat/long value does not match to the county and state recorded by Zillow.
- We filter on "RR" (Residential) use code plus a building indicator, which is 0 whenever building area is 0 or missing.

C Census Data Appendix

Our Census and American Community Survey data comes from the National Historical GIS (NHGIS) Website (Manson et al., 2024). In this paper, our main specifications combine the long-form decennial Census summary files for 1990 and 2000 and American Community Survey (ACS) 5-year estimates for 2008-2012 and 2018-2022 (and replace Connecticut with 2017-2021 due to changes in geography). Due to changes in how race is recorded over time, we use short form decennial Census summary files for 2010 and 2020 to include variables on owner-occupied and renter-occupied tenure by race in our sorting models (Section C.10). In this Appendix, we describe how we interpolate these files to consistent geographic boundaries (i.e., 2010 tracts) and show robustness checks around various interpolation methods (i.e., NHGIS versus LTDB). We also show robustness to interpolation back to 2000 tracts. These robustness checks are described in section C.1.1.

There has been a shift in the racial categorization in the Census and ACS over time as well as changes in how respondents self-identify their race and ethnicity (described in detail in Section C.5 and Section C.7). Therefore, how any respondent is recorded over time in the Census or ACS may be influenced by measurement changes as well as changes in identity. For the purposes of this paper, these changes have led to challenges in harmonizing racial categories over time for 1990, 2000, 2010 and 2020 especially involving multiracial and Hispanic individuals. We describe trends in our harmonized data (main approach) in Section C.9.

The following describe each subsection in this appendix.

- Section C.1 provides details regarding interpolation
  - Section C.1.1 presents robustness checks for different interpolation methods and tract boundaries
  - Section C.1.2 presents the crosswalks and weighting mechanisms
  - Section C.1.3 highlights the method used to interpolate Census data to 2010 tract boundaries
  - Section C.1.4 explains further interpolation to the 2000 tract boundaries
  - Section C.1.5 presents primary specification and robustness check details
- Section C.2 describes general sample selection decisions
- Section C.3 describes how we calculate percentiles
- Section C.4 discusses changes in owner occupied housing values
- Section C.5 describes in detail how race/ethnicity measurement has changed

- Section C.5.1 provides background on the 1977 OMB Directive on Race
- Section C.5.2 provides background on the single-race categorization in the 1990 Census
- Section C.5.3 provides background on the shift to multiple race categorization on the 2000 Census
- Section C.5.4 highlights similarities between the racial categorization questions in the 2010 Census and the 2008 - 2012 ACS
- Section C.5.5 addresses the changes in the racial categorization questions in the 2020 Census

· Section C.6 describes how changes in coding race for Hispanics led to large amounts of Hispanics coded as "some other race" in 2020

· Section C.7 discusses literature suggesting growth in survey respondents' identification as multiracial

· Section C.8 describes NHGIS approach to coding race/ethnicity, which we follow

· Section C.9 shows trends in Census and ACS variables over our sample

· Section C.10 describes data specifics for the tenure by race variables used in the sorting model (i.e., racial/ethnic breakdown of owners)

· Section C.11 describes the acreage and land use data we include for Table 5

C.1 Overview of Interpolation & Data Harmonization

Table C.1: Crosswalk and Weighting Methods

Specification	Boundary	Weight	Crosswalk
Main	2010	wt_ownhu; wt_renthu[1]	NHGIS[2]
Robustness Estimate 1	2010	Area and population weights by NHGIS (PAREA)	NHGIS[2]
Robustness Estimate 2	2010	Area and population weights by LTDB (weight)[3]	LTDB[4]
Robustness Estimate 3	2000	Area and population weights by LTDB (weight)[5]	"Backwards LTDB"[6]
Robustness Estimate 4	2010	Universe Defined Weights (UDW)[7]	NHGIS[8]

Notes: [1] When `wt_ownhu` or `wt_renthu` is missing, PAREA is used as a backup weight.
[2] NHGIS crosswalk available at nhgis.org/geographic-crosswalks.
[3] For Robustness Estimate 2, the LTDB crosswalk for 2020 did not work and so we swapped it with the NHGIS tract-level crosswalk for 2022 ACS.
[4] LTDB crosswalk available at https://s4.ad.brown.edu/projects/diversity/researcher/LTBDDload/DataList.aspx.
[5] For Robustness Estimate 3, the Robustness Estimate 2 data are interpolated to 2000 via the "Backwards LTDB" Crosswalk.
[6] Backwards LTDB crosswalk available at https://s4.ad.brown.edu/projects/diversity/researcher/LTBDDload/DataList.aspx
[7] NHGIS crosswalks details available at nhgis.org/geographic-crosswalks.
[8] NHGIS crosswalks available at nhgis.org/geographic-crosswalks.

As this paper spans data across 30 years, with Census and ACS reported at different geographic levels (e.g., blocks, block group parts, block groups), we aim to standardize and interpolate this data to 2010 tract level boundaries. We use Stata code provided by the Longitudinal Tract Data Base (LTDB) that allows users to provide user-defined source data, crosswalks, and choice of counts and medians (LTDB stata code). We modified it to interpolate using both NHGIS crosswalks and LTDB crosswalks (described in detail below). See Table C.1 for a summary of our approaches in our main specification and robustness checks. We have tested robustness to different spatial boundaries (2010 tracts versus 2000) and to different interpolation weights.

To analyze spatial and demographic changes consistently across Census years, it is crucial to harmonize historical Census data to a standardized geographic framework. The key methodological challenge arises because Census geographic units, such as blocks, block groups, and Census tracts, are frequently redefined due to population growth, urban expansion, and redistricting. Census blocks are the smallest geographic unit, block groups are aggregations of blocks, and Census tracts are larger units consisting of multiple block groups. To address these boundary changes, our methodology relies on interpolation techniques that leverage geographic crosswalks. At a conceptual level, interpolation involves redistributing demographic and housing characteristics from historical geographic units (source units) into more recent, standardized units (target units). This approach ensures comparability and consistency across multiple Census periods despite geographic shifts.

The interpolation methodology proceeds in three primary stages. First, we begin with the smallest available geographic units, typically blocks for Census short-form (full-count) data and block group parts or block groups for Census long-form (sample-based) and ACS data. Using smaller units reduces estimation error because these units typically nest more precisely into subsequent geographic definitions. Second, we apply geographic crosswalks obtained from the National Historical Geographic Information System (NHGIS). These crosswalks establish the precise geographic correspondence between historical source units and target units by assigning interpolation weights that reflect the expected proportion of a source unit's characteristics residing within a target unit. For the most granular units, such as Census blocks, NHGIS provides a single interpolation weight. This weight, calculated using target-density weighting methods, assumes a homogeneous distribution of population and housing characteristics within each source block, allowing proportional allocation to target units (Schroeder, 2007).

For larger geographic units, such as block group parts, block groups, and Census tracts, NHGIS provides multiple universe-specific interpolation weights. These distinct weights address the heterogeneous spatial distribution of various demographic and housing characteristics. For example, separate weights are available for total population, adult population, households, families, total housing units, owner-occupied housing units, and renter-occupied housing units. By using these universe-specific weights, we ensure that each demographic group or housing characteristic is allocated to the target geography according to its unique spatial distribution, thereby improving the accuracy of our interpolated estimates.

In situations where universe specific weights are unavailable or unreliable, we resort to an area-based weighting measure known as PAREA provided by NHGIS in their crosswalks. The PAREA weight represents the proportion of land area from the source geographic unit that falls within each target unit. While less precise than population-based measures, this fallback ensures data continuity and provides a robust alternative when other weighting information is limited or missing. Finally, after allocating historical data according to these interpolation weights, we aggregate redistributed characteristics across all source units within each target unit. This step ensures that the final estimates accurately reflect the demographic and spatial composition of the standardized geographic boundaries. This approach was adapted from code provided by the LTDB.[5]

We perform a few robustness checks (Section C.1.1) using alternative NHGIS weights, geographies and using the LTDB (Table C.1 summarizes all robustness checks). The LTDB interpolation method primarily relies on a combination of area-based and population-based interpolation to estimate past Census tract data within 2010 boundaries (Logan et al., 2014). When tract boundaries have changed over time, LTDB redistributes populations based on block-level data, assuming spatial uniformity or stationarity — that demographic characteristics are evenly distributed within source tracts. However, this assumption introduces errors, particularly in tracts characterized by demographic clustering or uneven distribution (Logan, Stults and and, 2016; Logan, Zhang, Stults

[5] LTDB crosswalk available at https://s4.ad.brown.edu/projects/diversity/researcher/LTBDDload/DataList.aspx.

and Gardner, 2021).

To improve on these limitations, LTDB developed a more refined approach incorporating ancillary data such as water boundaries and using block-level population counts to minimize misallocations into uninhabited areas, though this method still treats all non-water land as uniformly habitable (Logan et al., 2014, 2016). While this method considerably improves estimates over basic area weighting, it still faces limitations when internal populations of census blocks are unevenly distributed or when block boundaries themselves are split across new census tracts (Logan et al., 2016)

The two primary methodological differences between the NHGIS and LTDB harmonization techniques are as follows. First, NHGIS crosswalks build up from the smallest available geographic units (blocks for short-form Census data or block-group parts for long-form and ACS data), improving interpolation accuracy (Manson et al., 2024). In contrast, the LTDB approach initially interpolated primarily at the tract level (for data from 1970–1990), relying only on aggregate tract-to-tract area interpolation, and subsequently incorporated block-level population weighting from 2000 onward (Logan et al., 2014). Second, whereas the LTDB interpolation method primarily omits water bodies and treats all other land equally as potentially inhabited (Logan et al., 2014), for 2000 Blocks to 2010 geographies NHGIS incorporates a refined dasymetric interpolation approach that excludes not only water bodies but also non-residential spaces such as large forests, industrial zones, highways, and open fields, utilizing residential road buffers and detailed land cover data (Schroeder, 2007).

C.1.1 Robustness Checks: Alternative Interpolation Approaches & Various Tract Boundaries

As a set of robustness checks, we employ the LTDB interpolation approach for 2010 tracts and also use their reverse interpolation crosswalks for 2000 tracts (summarized in Table C.1). All interpolation details are described in Section C.1.2.

Our main approach uses NHGIS owner-occupied weights to harmonize the data to 2010 tract boundaries. As described in the previous section, when owner-occupied weights are missing, we use parea weights; this gives us the same number of tracts across weights. To make sure our results are not sensitive to different interpolation weights, we employ 4 robustness checks: (1) NHGIS parea weights with 2010 boundaries, (2) LTDB weights with 2010 boundaries, (3) LTDB weights with 2000 boundaries, (4) NHGIS universe defined weights (i.e., separate weights for each universe of the underlying data, e.g., population, households, housing units, owner-occupied, rental occupied) with 2010 boundaries. Our robustness checks (see C.1) support limited differences in estimates across the various crosswalks and geographic boundaries. In Tables C.2-C.4 we show similar robustness checks for Table 6. As with our housing price results, we find minimal differences across weights and tract boundaries for our reduced form results for changing demographics around clean-ups.

Figure C.1: Robustness Check: Varying Census Tract Boundaries and Data Interpolation Methods

(a) 2010 Tracts & owner-occupied weights

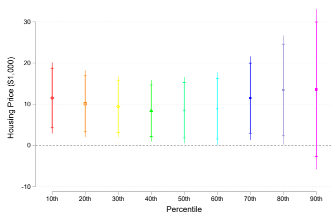

(b) 2010 Tracts & parea weights

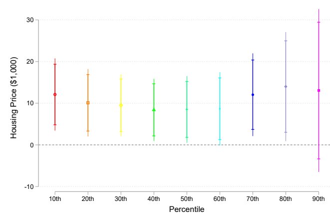

(c) 2010 Tracts & LTDB weights

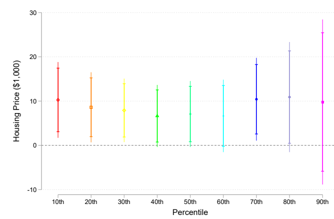

(d) 2000 Tracts & LTDB weights

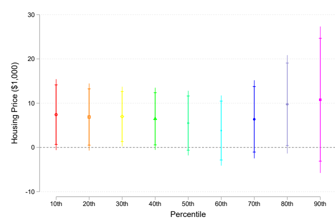

(e) 2010 Tracts & UDW weights

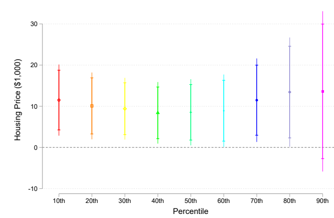

Notes: This figure shows five combinations of tract boundaries and interpolation weights. See Section C.1 for crosswalk sources and details. These robustness checks are summarized in Table C.1. Owner-occupied, parea, UDW = universe defined weights are all NHGIS weights available in their crosswalks. This figure is similar to Figure 4 in the main text; Panel a is exactly the same. We use the same fixed effects and sample as in the main regression. Whiskers marked with horizontal lines and vertical protruding segments indicate 95% and 99% CIs respectively, clustering at the tract level. Prices are denominated in thousands of dollars.

While methodological differences between NHGIS and LTDB interpolation procedures enhance NHGIS's accuracy, particularly in complex tract boundary changes, these differences are minor for most tracts. Logan, Stults, and Xu (2016) indicate that both NHGIS and LTDB yield highly comparable estimates for tracts with stable or minimally adjusted boundaries, such as unchanged or merged tracts. However, for tracts involving more complex boundary changes—especially those where census blocks were divided—NHGIS's incorporation of additional ancillary data (such as land cover and road networks) slightly improves accuracy relative to LTDB, reducing estimation errors. Despite this improvement, the overall difference between NHGIS and LTDB estimates remains small for most cases (Logan et al., 2016). Our robustness checks reflect this.

C.1.2 Crosswalks and Weighing Mechanisms

NHGIS crosswalks provide interpolation weights that estimate how a source unit's population and housing characteristics should be redistributed across target units. These weights account for geographic splits and allow users to proportionally allocate data. The types of Weights in NHGIS Crosswalks are listed below.

Block-Level Weights (Simplest case)

- Single `WEIGHT` field

Table C.2: Robustness Check 1: Impacts on Sociodemographic and Housing Variables, Near-Far Comparison

(a) Income and Education-related Variables

Dep. var:	Avg HH Income	% Below Poverty	% College Educated	% on Public Assistance	% Unemployment
0 km × Post	0.152	-0.515	0.768	0.223	0.520
	(1.319)	(0.512)	(0.469)	(0.282)	(0.361)
0 km × In Prog	-0.807	-0.553	0.552	0.077	0.168
	(1.042)	(0.450)	(0.358)	(0.224)	(0.314)
Avg Outcome	101.199	12.706	26.469	3.661	6.020
Adj R^2	0.876	0.789	0.899	0.590	0.517
Clusters	13,137	13,140	13,150	13,137	13,146
Obs	52,387	52,436	52,513	52,422	52,478

(b) Impacts on Demographic Variables, Near-Far Comparison

Dep. var:	% Black	% Female Head of Household	% Hispanic	Population Density	% White	% Under 18
0 km × Post	-0.050	0.264	-0.108	-3.587	-0.250	-0.308
	(0.363)	(0.568)	(0.376)	(11.118)	(0.515)	(0.308)
0 km × In Prog	0.271	0.433	-0.494	7.949	0.059	-0.141
	(0.259)	(0.421)	(0.309)	(8.378)	(0.417)	(0.256)
Avg Outcome	11.089	17.725	12.781	1,506.397	69.677	24.027
Adj R^2	0.928	0.758	0.935	0.971	0.939	0.734
Clusters	13,147	13,129	13,147	13,155	13,147	13,147
Obs	52,482	52,370	52,482	52,607	52,482	52,482

(c) Impacts on Housing-Related Variables, Near-Far Comparison

Dep. var:	% 4+ Bedrooms	% Built in Last 10 Years	% Mobile Home	% Moved in Last 10 Years	% Owner Occupied	% Vacant
0 km × Post	-0.163	0.068	1.078*	-0.126	0.394	0.292
	(0.493)	(0.835)	(0.642)	(0.617)	(0.578)	(0.478)
0 km × In Prog	-0.032	0.952	1.402	0.775	-0.013	-0.073
	(0.396)	(0.739)	(0.902)	(0.536)	(0.457)	(0.357)
Avg Outcome	20.046	17.599	7.492	63.123	67.337	9.630
Adj R^2	0.847	0.621	0.264	0.721	0.907	0.765
Clusters	13,138	13,138	13,137	13,137	13,130	13,134
Obs	52,431	52,431	52,422	52,422	52,382	52,400

Notes: The data in this table is harmonized using the NHGIS parea weight and harmonized to 2010 tract boundaries (Robustness Check 1 in Table C.1). This table tests the impacts of RCRA cleanups on 17 socioeconomic and housing variables, comparing changes in outcomes for tracts containing a RCRA site (0 km) over time. The regression model is the same as the main model in Table 2 except with one of the 17 socioeconomic and housing variables as the dependent variable. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(0, 10]$ km away from a facility. All standard errors are clustered on census tract.

Table C.3: Robustness Check 2: Impacts on Sociodemographic and Housing Variables, Near-Far Comparison

(a) Income and Education-related Variables

Dep. var:	Avg HH Income	% Below Poverty	% College Educated	% on Public Assistance	% Unemployment
0 km × Post	-0.283	-0.496	0.810*	0.206	0.344
	(1.324)	(0.487)	(0.467)	(0.276)	(0.348)
0 km × In Prog	-0.812	-0.554	0.732**	0.004	-0.036
	(1.016)	(0.456)	(0.353)	(0.221)	(0.308)
Avg Outcome	101.025	12.711	26.468	3.668	6.024
Adj R^2	0.876	0.789	0.900	0.613	0.524
Clusters	13,140	13,141	13,150	13,140	13,146
Obs	52,405	52,451	52,522	52,440	52,488

(b) Impacts on Demographic Variables, Near-Far Comparison

Dep. var:	% Black	% Female Head of Household	% Hispanic	Population Density	% White	% Under 18
0 km × Post	-0.164	0.035	-0.215	-1.916	-0.197	-0.692**
	(0.380)	(0.535)	(0.382)	(11.795)	(0.537)	(0.342)
0 km × In Prog	0.085	0.190	-0.547*	7.865	0.273	-0.358
	(0.274)	(0.410)	(0.322)	(9.118)	(0.431)	(0.260)
Avg Outcome	11.119	17.751	12.804	1,497.131	69.625	24.083
Adj R^2	0.926	0.758	0.936	0.977	0.938	0.749
Clusters	13,148	13,134	13,148	13,155	13,148	13,148
Obs	52,503	52,395	52,503	52,619	52,503	52,503

(c) Impacts on Housing-Related Variables, Near-Far Comparison

Dep. var:	% 4+ Bedrooms	% Built in Last 10 Years	% Mobile Home	% Moved in Last 10 Years	% Owner Occupied	% Vacant
0 km × Post	-0.223	-0.061	0.823	-0.228	0.527	0.348
	(0.506)	(0.850)	(0.512)	(0.625)	(0.670)	(0.544)
0 km × In Prog	0.059	0.889	0.498	0.565	0.053	0.069
	(0.389)	(0.726)	(0.410)	(0.512)	(0.521)	(0.364)
Avg Outcome	19.978	17.614	7.301	63.185	67.261	9.631
Adj R^2	0.833	0.654	0.878	0.736	0.890	0.784
Clusters	13,141	13,141	13,140	13,140	13,135	13,137
Obs	52,448	52,448	52,445	52,440	52,414	52,428

Notes: The data in this table is harmonized using the LTDB weights and harmonized to 2010 tract boundaries (Robustness Check 2 in Table C.1). This table tests the impacts of RCRA cleanups on 17 socioeconomic and housing variables, comparing changes in outcomes for tracts containing a RCRA site (0 km) over time. The regression model is the same as the main model in Table 2 except with one of the 17 socioeconomic and housing variables as the dependent variable. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(0, 10]$ km away from a facility. All standard errors are clustered on census tract.

Table C.4: Robustness Check 3: Impacts on Sociodemographic and Housing Variables, Near-Far Comparison

(a) Income and Education-related Variables

Dep. var:	Avg HH Income	% Below Poverty	% College Educated	% on Public Assistance	% Unemployment
0 km × Post	-0.968	-0.430	0.405	0.390	0.295
	(1.288)	(0.455)	(0.441)	(0.245)	(0.322)
0 km × In Prog	-0.789	-0.608	0.553	0.103	0.098
	(0.931)	(0.425)	(0.347)	(0.193)	(0.299)
Avg Outcome	98.888	13.274	25.579	3.860	6.194
Adj R^2	0.903	0.805	0.913	0.622	0.531
Clusters	11,298	11,301	11,311	11,298	11,305
Obs	45,139	45,164	45,214	45,152	45,181

(b) Impacts on Demographic Variables, Near-Far Comparison

Dep. var:	% Black	% Female Head of Household	% Hispanic	Population Density	% White	% Under 18
0 km × Post	-0.176	0.101	-0.071	-7.018	-0.187	-0.185
	(0.391)	(0.518)	(0.366)	(10.353)	(0.541)	(0.313)
0 km × In Prog	0.213	0.078	-0.305	5.428	-0.069	-0.127
	(0.282)	(0.409)	(0.303)	(8.655)	(0.427)	(0.279)
Avg Outcome	11.651	18.334	12.520	1,480.461	69.567	23.977
Adj R^2	0.934	0.783	0.942	0.978	0.944	0.762
Clusters	11,311	11,299	11,311	11,322	11,311	11,311
Obs	45,219	45,151	45,219	45,276	45,219	45,219

(c) Impacts on Housing-Related Variables, Near-Far Comparison

Dep. var:	% 4+ Bedrooms	% Built in Last 10 Years	% Mobile Home	% Moved in Last 10 Years	% Owner Occupied	% Vacant
0 km × Post	-0.652	-0.405	0.709	-0.151	-0.141	0.244
	(0.436)	(1.011)	(0.475)	(0.642)	(0.649)	(0.468)
0 km × In Prog	-0.407	-0.298	0.391	0.490	-0.640	0.504
	(0.370)	(0.770)	(0.414)	(0.508)	(0.578)	(0.404)
Avg Outcome	19.391	15.556	7.161	61.914	66.834	9.736
Adj R^2	0.866	0.652	0.900	0.760	0.917	0.810
Clusters	11,299	11,299	11,299	11,298	11,300	11,303
Obs	45,157	45,157	45,157	45,152	45,168	45,176

Notes: The data in this table is harmonized using the LTDB weights and harmonized to 2000 tract boundaries (Robustness Check 3 in Table C.1). This table tests the impacts of RCRA cleanups on 17 socioeconomic and housing variables, comparing changes in outcomes for tracts containing a RCRA site (0 km) over time. The regression model is the same as the main model in Table 2 except with one of the 17 socioeconomic and housing variables as the dependent variable. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(0, 10]$ km away from a facility. All standard errors are clustered on census tract.

Table C.5: Robustness Check 4: Impacts on Sociodemographic and Housing Variables, Near-Far Comparison

(a) Income and Education-related Variables

Dep. var:	Avg HH Income	% Below Poverty	% College Educated	% on Public Assistance	% Unemployment
0 km × Post	0.035	-0.470	0.763	0.297	0.508
	(1.313)	(0.499)	(0.474)	(0.284)	(0.359)
0 km × In Prog	-0.907	-0.509	0.508	0.157	0.176
	(1.039)	(0.441)	(0.358)	(0.221)	(0.317)
Avg Outcome	101.222	12.703	26.467	3.660	6.020
Adj R^2	0.872	0.788	0.896	0.587	0.516
Clusters	13,137	13,140	13,150	13,137	13,146
Obs	52,387	52,436	52,513	52,422	52,478

(b) Impacts on Demographic Variables, Near-Far Comparison

Dep. var:	% Black	% Female Head of Household	% Hispanic	Population Density	% White	% Under 18
0 km × Post	-0.046	0.256	-0.102	-2.864	-0.251	-0.284
	(0.363)	(0.565)	(0.375)	(11.052)	(0.515)	(0.308)
0 km × In Prog	0.278	0.447	-0.488	8.053	0.049	-0.143
	(0.259)	(0.422)	(0.309)	(8.383)	(0.417)	(0.256)
Avg Outcome	11.088	17.725	12.781	1,506.876	69.679	24.026
Adj R^2	0.928	0.757	0.934	0.971	0.939	0.734
Clusters	13,147	13,129	13,147	13,155	13,147	13,147
Obs	52,482	52,370	52,482	52,607	52,482	52,482

(c) Impacts on Housing-Related Variables, Near-Far Comparison

Dep. var:	% 4+ Bedrooms	% Built in Last 10 Years	% Mobile Home	% Moved in Last 10 Years	% Owner Occupied	% Vacant
0 km × Post	-0.392	-0.223	0.830	-0.083	0.327	0.286
	(0.469)	(0.858)	(0.561)	(0.613)	(0.575)	(0.478)
0 km × In Prog	-0.224	0.639	1.351	0.697	-0.013	-0.057
	(0.390)	(0.739)	(0.887)	(0.535)	(0.457)	(0.358)
Avg Outcome	20.053	17.559	7.468	63.091	67.339	9.632
Adj R^2	0.847	0.609	0.194	0.715	0.907	0.763
Clusters	13,138	13,138	13,137	13,137	13,130	13,134
Obs	52,431	52,431	52,422	52,422	52,382	52,400

Notes: The data in this table is harmonized using the NHGIS universe defined weights (UDW) weights and harmonized to 2010 tract boundaries (Robustness Check 4 in Table C.1). This table tests the impacts of RCRA cleanups on 17 socioeconomic and housing variables, comparing changes in outcomes for tracts containing a RCRA site (0 km) over time. The regression model is the same as the main model in Table 2 except with one of the 17 socioeconomic and housing variables as the dependent variable. We use all tracts within 10 km of one RCRA facility in this regression. All regressions include fixed effects for tract, bin by year, site by year, and state by year. The excluded category is tracts $(0, 10]$ km away from a facility. All standard errors are clustered on census tract.

- Represents the proportion of a source block's population and housing units that fall within a target tract.
- Applied uniformly to all characteristics.

Block Group Parts and Higher Geographic Levels (More complex cases)

- Provide universe-specific weights to improve accuracy. These include:
  - wt_pop (total population)
  - wt_adult (adult population)
  - wt_fam (families)
  - wt_hh (households)
  - wt_hu (total housing units)
  - wt_ownhu (owner-occupied housing units)
  - wt_renthu (renter-occupied housing units)
- These weights account for spatial clustering of demographic groups and prevent errors associated with spatial stationarity assumptions.

PAREA (Area Proportion Weights) (Used for robustness checks)

- Represents the proportion of a source unit's land area that falls within a target unit.
- Used as a fallback when other population or housing weights are unavailable.

LTDB Weights

- Single interpolation weight per record
  - Primarily based on total population proportions derived from census block-level data.
  - Applied uniformly across all demographic and housing characteristics. However, recent iterations include trait-based adjustments to improve accuracy for specific full-count variables.

C.1.3 Interpolation to 2010 Census Tract Boundaries

To harmonize historical Census data to 2010 tracts, we follow a bottom-up interpolation approach, utilizing NHGIS-provided crosswalks. Each step depends on the smallest available geographic level for each Census year.

Selecting the Source and Target Geographies

For each Census year, we determine the source geographic level (i.e., where the data originates) and the target level (i.e., where it is being mapped).

- 1990 Census:
  - Short-form data (full count) → STF1 (Block-level) interpolated to 2010 tracts using NHGIS block-to-tract crosswalk.
  - Long-form data (sample-based) → STF3 (Block group parts) interpolated separately to 2010 tracts using NHGIS block group parts-to-tract crosswalk.
  - Final step: Merging block-based and block group parts-based estimates to produce 1990 data in 2010 tract boundaries.
- 2000 Census:
  - Short-form data → SF1b (Block-level) interpolated to 2010 tracts using NHGIS block-to-tract crosswalk.
  - Long-form data → SF3b (Block group parts) interpolated separately to 2010 tracts using NHGIS block group parts-to-tract crosswalk.
  - Final step: Merging these two estimates.
- 2020 ACS Data:
  - ACS 2018-2022 and ACS 2017-2021 (for Connecticut correction) are provided at block group level.
  - Interpolated to 2010 tracts using NHGIS block group-to-tract crosswalk.

Interpolation via LTDB

- As part of our robustness checks, we also employ the LTDB crosswalk for interpolating 1990 and 2000 data to 2010 boundaries and further interpolating data from 2010 to 2000 tract boundaries. However, due to data quality issues encountered when applying the LTDB crosswalk to interpolate 2020 data back to 2010 tract boundaries, we rely exclusively on the NHGIS crosswalk for the 2020-to-2010 interpolation, in particular when interpolating the 2020_DHCa Tract File, ACS 2021 Tract file and the ACS 2022 Tract File to 2010 Tract Boundaries.

Once all data is mapped to 2010 tracts, we proceed to reverse-interpolate it to 2000 Census tract boundaries using "Backwards LTDB" Crosswalk.

C.1.4 Interpolation to 2000 Census Tract Boundaries

After standardizing all data to 2010 tract boundaries, we further harmonize it to 2000 tract boundaries using LTDB's 2010-to-2000 crosswalk ("Backwards LTDB" Crosswalk).

- Unlike NHGIS, LTDB provides a single weight per record, meaning all characteristics are allocated uniformly when shifting from 2010 tracts to 2000 tracts.
- Since 2010 tracts were formed by modifying 2000 tracts (e.g., splitting or merging tracts), the crosswalk ensures that past data remains comparable.

C.1.5 Implementation in This Study

Primary Specification

- For block-to-tract interpolation, a single `WEIGHT` field is applied.
- Uses wt_ownhu (for owner-occupied housing) and wt_renthu (for renter-occupied housing) for block group parts and higher level interpolation.
  - When wt_ownhu and wt_renthu are unavailable, PAREA weights are substituted following NHGIS best practices.

Robustness Checks

- PAREA-based interpolation to assess sensitivity of results.
- LTDB weights
- UDW weights
  - The 2020 crosswalk was not usable so we created 2022 estimates using the NHGIS PAREA weight and used LTDB crosswalks for other years (1990, 2000).

C.2 General Sample Selection & Data Cleaning

- We adjusted prices to 2023 dollars using (U.S. Bureau of Labor Statistics, 2025) . The inflation adjustment was made as follows:
  - For the 2008-2012 ACS, we treated the values as if they were from the year 2010 (the midpoint of the year range) when adjusting for inflation. This is because the Census Bureau did not make any adjustments when producing aggregate statistics.

- For the 2018–2022 ACS, there was a change in how the Census Bureau aggregated the housing values across years- they undertook an inflation adjustment to get all values in 2022 dollars already. So, for these ACS observations, we treated values as if they were denominated in 2022 dollars. For details on the change and how it impacts interpretation, please see (United States Census Bureau, 2023).
- For the ACS (2017–2021), which was used as the 2020 value for the state of Connecticut, we treated the values as if they were denominated in 2019 dollars.
- The universe-specific weights sometimes add to 0 for a particular 2010 geographic target unit in years other than 2010, and this can happen even when population and/or housing is nonzero in that target unit & year. For example, Census Tract 228, Placer County, California (2010 tract id 06061022800) had a population of 2,230 in 2010. In 2000, using area weighting, the tract has a non-missing count of both population and housing, but universe-specific weights show counts of 0 and thus variables are missing. In line with practices of NHGIS in construction of their time series data, if a variable is missing in the housing or universe-specific weighted data but not in the area weight data, we use the value from the area-weight interpolated data to fill in the missing value. [6]
- We only include tracts near sites that had either date started or date completed for the two EIs we study (coded in our data as Corrective Action Codes 725 and 750).

C.3 How We Calculate Percentiles

We construct a cumulative distribution from binned price data by allocating counts of data to construct an empirical distribution by linear interpolation.

Table C.6 shows a numerical example based on the count of data in each price bin for a hypothetical Census tract.[7] The first column lists the bins. The second column lists the count of observations in that bin for the Census tract, and the third column shows the percent of observations calculated to be strictly below the upper bound price value of the bin.

Table C.7 shows the implied deciles based on the data in Table C.6 produced by our procedure.

To ensure that all quantiles calculated by the procedure are valid, we do not allow for situations where any of the computed quantiles fall in the bottom-most or top-most category of prices. For example, in 2010, if the 10th percentile computed from this procedure fell in the <10 category, we would designate the entire geographic unit as having invalid quantiles. This restriction

[6] For an example, see their revision history page at https://www.nhgis.org/revision-history, where they state, "For example, in the crosswalk from 2010 block groups to 2020 tracts, if a 2010 block group contains no families, it is not possible to compute a valid wt_fam using family counts. We compute wt_fam instead by using the sum of population and housing units, or, if that sum is also zero for the 2010 block group, we use the proportion of the 2010 block group's area in the 2020 tract."

[7] To make the exercise realistic, the hypothetical tract data was constructed by rounding mean price counts in each bin for the year 2010.

Table C.6: Numerical Example Data

	Count	Count Percent < Upper Bound
< 10	14	1.350
[10, 15)	10	2.314
[15, 20)	9	3.182
[20, 25)	9	4.050
[25, 30)	9	4.918
[30, 35)	10	5.882
[35, 40)	7	6.557
[40, 50)	20	8.486
[50, 60)	24	10.800
[60, 70)	28	13.500
[70, 80)	32	16.586
[80, 90)	39	20.347
[90, 100)	35	23.722
[100, 125)	87	32.112
[125, 150)	76	39.441
[150, 175)	92	48.312
[175, 200)	64	54.484
[200, 250)	109	64.995
[250, 300)	83	72.999
[300, 400)	108	83.414
[400, 500)	59	89.103
[500, 750)	66	95.468
[750, 1000)	24	97.782
≥ 1000	23	100.000
Total	1,037	100.000

Notes: This table depicts count data for a hypothetical census tract based on 2010 price bins. The first column presents the bins. The second column is the count in each bin for the tract. The third column is the implied percent of houses below the upper bound value of the bin.

accomplishes two goals. First, it ensures that the quantiles are relatively accurate and not skewed by outlier values. Second, it avoids the need for a researcher-chosen "top" and "bottom" value that houses could take in each year.

Table C.7: Numerical Example Computed Deciles

	Price Decile
10	56.542
20	89.077
30	118.707
40	151.576
50	181.836
60	226.238
70	281.265
80	367.222
90	535.227

Notes: This table presents the computed deciles produced by our linear interpolation procedure when we input the counts in Table C.6.

C.4 Specified-Owner Occupied Housing Values

Our study uses data from four major surveys to analyze owner-occupied housing values: the 1990 Census, the 2000 Census, the 2008-2012 American Community Survey (ACS) and the 2018-2022 ACS. Our main variable of interest is owner-occupied housing values across these datasets. In Table C.8, we summarize how owner-occupied housing has changed definitions over time and show the definition of owner-occupied home values in our main specification and our robustness check. In particular, the definition shifted from specified owner occupied (SOO) in 1990 and 2000 to owner-occupied (OO) in 2000, 2012 ACS, 2022 ACS.

Before 1990, much of the owner-occupied housing inventory comprised of single-family homes, either detached or attached. Therefore, earlier census data provided financial housing characteristics for the specified owner-occupied unit universe. However, the housing market began to change during the 1990s as an increasing number of units in multi-unit structures were constructed and sold as condominiums and mobile homes became a more common option for lower-income home-owners. As a result of these changes, the ACS abandoned the concept of the specified owner-occupied universe to ensure housing data was provided for all owner-occupied units, including mobile homes and units in multi-unit structures. This shift is reflected in the preambles to the sections on owner-occupied housing values.

Owner-occupied housing is defined in the 1990 Census as "the owner or co-owner lives in the unit even if it is mortgaged or not fully paid for."[8] The 2000 Census[9] and 2008-2012 ACS[10] have

[8] 1990 Decennial Census of Population and Housing Technical Documentation
[9] Census 2000 Summary File 1; prepared by the U.S. Census Bureau, 2001
[10] 2010 Census Summary File 1; prepared by the U.S. Census Bureau, 2011

similar definitions. Where they differ is described in detail below.

1990 Census In the 1990 Census, respondents were asked to "Answer only if you or someone in this household OWNS OR IS BUYING this house or apartment" regarding the question on declaring the value of their owner-occupied home, which asked, "What is the value of this property; that is, how much do you think this house and lot or condominium unit would sell for if it were for sale?"[11]

About tabulating owner-occupied housing values, the 1990 census states: "Value is tabulated separately for all owner-occupied and vacant-for-sale-only housing units, owner-occupied and vacant-for-sale mobile homes or trailers, and specified owner-occupied and specified vacant-for-sale-only housing units. Specified owner-occupied and specified vacant-for-sale-only housing units include only one-family houses on less than 10 acres without a business or medical office on the property. The data for 'specified' units exclude mobile homes, houses with a business or medical office, houses on 10 or more acres, and housing units in multi-unit buildings."[12]

2000 Census In the 2000 Census, the question was expanded to: "Answer questions 47a—53 if you or someone in this household owns or is buying this house, apartment, or mobile home; otherwise, skip to questions for Person 2." The question asked, "What is the value of this property; that is, how much do you think this house and lot, apartment, or mobile home and lot would sell for if it were for sale?"[13]

In the available 2000 Census long-form summary file, both specified-owner occupied and owner-occupied housing values are available for analysis.

2008-2012 ACS The 2012 ACS survey states: "Please answer the following questions about the house, apartment, or mobile home at the address on the mailing label." The survey then asks, "About how much do you think this house and lot, apartment, or mobile home (and lot, if owned) would sell for if it were for sale?" Respondents fill in a free response number regarding the value of their owner-occupied housing.[14]

2018-2022 ACS In the 2022 ACS survey respondents report the estimated value of their property through an open-ended question: "About how much do you think this house and lot, apartment, or mobile home (and lot, if owned) would sell for if it were for sale?" This ensures consistency in data collection across different housing structures and ownership.[15]

Summary As shown above, the composition of houses included in owner-occupied housing values in these surveys has changed. The 2000 Census, 2008-2012 ACS and 2018-2022 ACS include mobile homes, whereas the 1990 Census primarily focused on single-family homes when asking to

[11] 1990 Decennial Census of Population and Housing Technical Documentation
[12] 1990 Decennial Census of Population and Housing Technical Documentation
[13] Census 2000 Summary File 1 prepared; by the U.S. Census Bureau, 2001
[14] 2010 Census Summary File 1; prepared by the U.S. Census Bureau, 2011
[15] American Community Survey and Puerto Rico Community Survey 2022 Subject Definitions

declare owner-occupied housing values. The 2000 Census provides financial housing characteristics for both all owner-occupied units and the more restricted universe of specified owner-occupied units. The ACS, however, only publishes financial housing characteristics for all units. To ensure compatibility and consistency over time, specified owner-occupied housing values can be used for 1990 to 2000. As the data for specified units exclude mobile homes, houses with a business or medical office, houses on 10 or more acres, and housing units in multi-unit buildings, we chose to use specified units for 1990 and 2000, but can only use owner-occupied for 2008-2012 ACS and 2018-2022 ACS due to this change in survey. We perform robustness checks instead using owner-occupied in 2000, 2008-2012 ACS and 2018-2022 ACS with 1990 as specified only. See Table C.8 for a summary of these definitions and how they are used in this paper.

Table C.8: Housing Types Included in Census Housing Value Surveys

Census Year	**Specified Owner Occupied (SOO)**	**Owner-Occupied (OO)**	**Main Spec**	**Robustness**
1990	Y		SOO	SOO
2000	Y	Y	SOO	OO
2010		Y	OO	OO
2020		Y	OO	OO

Definitions

Specified Owner Occupied (SOO):	Single-family houses on less than 10 acres
Owner-Occupied (OO):	Multi-family, condos, mobile homes, trailers, and single-family

Notes: This table summarizes changes in the types of houses surveyed for our housing value variables over different years. 1990 and 2000 refer to Decennial Censes, whereas 2010 and 2020 refer to the 2008–2012 and 2018–2022 American Communities Surveys respectively.

We present a robustness check of our main results when substituting the owner occupied housing values in Table C.9.

Table C.9: Robustness Check: Including Data from Owner Occupied Housing Units in 2000

Dep. var: Price^{kth}									
	10^{th}	20^{th}	30^{th}	40^{th}	50^{th}	60^{th}	70^{th}	80^{th}	90^{th}
0 km × Post	8.178**	8.081**	6.629*	5.690	6.306	4.970	7.284	7.552	3.651
	(3.772)	(3.597)	(3.649)	(3.761)	(3.983)	(4.211)	(4.671)	(5.801)	(8.778)
0 km × In Prog	7.843**	6.197**	4.479	3.664	3.543	2.665	4.254	4.368	5.027
	(3.285)	(3.075)	(3.082)	(3.348)	(3.266)	(3.440)	(3.810)	(5.014)	(8.114)
Avg Price	152.342	190.306	218.593	244.312	270.512	299.897	335.210	383.801	470.379
Adj R^2	0.840	0.889	0.906	0.915	0.918	0.918	0.914	0.903	0.874
Clusters	12,699	12,699	12,699	12,699	12,699	12,699	12,699	12,699	12,699
Obs	48,224	48,224	48,224	48,224	48,224	48,224	48,224	48,224	48,224

Notes: This table presents the results from a version of the dependent variable using housing values from owner occupied housing units from the year 2000. This is in contrast to our main results which use specified owner occupied values in 2000.

C.5 Change In Race/Ethnicity Questions Over Time

The U.S. Census has changed the way that race and ethnicity are collected and recorded during the period relevant to this analysis, creating measurement issues when harmonizing data across geography and time, with particular challenges for defining multiracial and Hispanic populations. We summarize these changes in Table C.10 and describe them in detail below. The main changes occur from 1990 to 2000, when race categorization changes from choosing a single race (1990) to allowing for multiple racial identities (2000) and another major change occurred from 2010 to 2020, when the survey allowed for write-ins for all major racial groups (requesting origins), and expanded the length of write-in boxes from 30 characters to 200 characters. There were also changes in coding any Hispanic categorization as "some other race" leading to increases in multiracial categorization. These changes are described in more detail below.

Table C.10: Summary of Census and ACS Race/Ethnicity Questions & Coding

Data	**Form**	**Coding**	**Categories Used in Paper**
1990 Census	**Race:** Fill in one race: · White · Black · Indian [American] (print tribe) · Eskimo · Aleut · Asian or Pacific Islander (API) – Chinese – Japanese – Filipino – Asian Indian – Hawaiian – Samoan – Korean – Guamanian – Vietnamese – Other API · Other Race **Hispanic:** Is this person of Spanish/Hispanic origin? Fill ONE circle for each person. · No (not Spanish/Hispanic) · Yes, Mexican, Mexican-Am., Chicano · Yes, Puerto Rican · Yes, Cuban · Yes, other Spanish/Hispanic	**One Major Race Group:** · White · Black · American Indian · Chinese · Japanese · Other API · Other race	**Hispanic and non-Hispanic Racial Breakdown:** · Non-Hispanic White · Non-Hispanic Black · Hispanic

Table C.10 – continued from previous page

Data	Form	Coding	Categories Used in Paper
2000 Census	**Race:** Fill in one or more race: · White · Black · American Indian or Alaskan Native (print tribe) · Asian Indian · Chinese · Filipino · Japanese · Korean · Vietnamese · Other Asian · Native Hawaiian · Guamanian or Chamorro · Samoan · Other PI (print) · Some other race (print) **Hispanic:** Is this person Spanish/Hispanic/Latino? Mark [X] the "No" box if not Spanish/Hispanic/Latino. · No, not Spanish/Hispanic/Latino · Yes, Mexican, Mexican Am., Chicano · Yes, Puerto Rican · Yes, Cuban · Yes, other Spanish/Hispanic/Latino — Print group:	**One Major Race Group:** · White · Black · American Indian or Alaskan Native · Chinese · Japanese · Other API · Other race · Two or More Major Races · Three or More Major Races	**Hispanic and non-Hispanic Breakdown:** · Non-Hispanic White · Non-Hispanic Black · Non-Hispanic Two or More · Hispanic

Table C.10 – continued from previous page

Data	Form	Coding	Categories Used in Paper
2010 Census/ 2012 ACS	**Race:** Mark (X) one or more boxes: · White · Black, African Am., or Negro · American Indian or Alaska Native – Print name of enrolled or principal tribe · Asian Indian · Japanese · Chinese · Korean · Filipino · Vietnamese · Other Asian – Print race, for example, Hmong, Laotian, Thai, Pakistani, Cambodian and so on · Native Hawaiian · Guamanian or Chamorro · Samoan · Other Pacific Islander – Print race, for example, Fijian, Tongan, and so on · Some other race – Print race. **Hispanic:** Is Person 1 of Hispanic, Latino, or Spanish origin? · No, not of Hispanic, Latino, or Spanish origin · Yes, Mexican, Mexican Am., Chicano · Yes, Puerto Rican · Yes, Cuban · Yes, another Hispanic, Latino, or Spanish origin — Print origin, for example, Argentinean, Colombian, Dominican, Nicaraguan, Salvadoran, Spaniard, and so on.	**One Major Race Group:** · White · Black · American Indian or Alaska Native · Chinese · Japanese · Other API · Other race · Two or More Major Races · Three or More Major Races	**Hispanic and non-Hispanic Breakdown:** · Non-Hispanic White · Non-Hispanic Black · Non-Hispanic Two or More · Hispanic

Table C.10 – continued from previous page

Data	Form	Coding	Categories Used in Paper
2020 Census/ 2022 ACS	**Race:** What is Person X's race? Mark (X) one or more boxes AND print origins. · White — Print, for example, German, Irish, English, Italian, Lebanese, Egyptian, etc. · Black or African Am. — Print, for example, African American, Jamaican, Haitian, Nigerian, Ethiopian, Somali, etc. · American Indian or Alaska Native — Print name of enrolled or principal tribe(s), for example, Navajo Nation, Blackfeet Tribe, ..., etc. · Chinese · Vietnamese · Native Hawaiian · Filipino · Korean · Samoan · Asian Indian · Japanese · Chamorro · Other Asian — Print, for example, Pakistani, Cambodian, Hmong, etc. · Other Pacific Islander — Print, for example, Tongan, Fijian, Marshallese, etc. · Some other race — Print race or origin.	**One Major Race Group:** · White · Black · American Indian or Alaska Native · Chinese · Japanese · Other API · Other race · Two or More Major Races · Three or More Major Races	**Hispanic and non-Hispanic Breakdown:** · Non-Hispanic White · Non-Hispanic Black · Non-Hispanic Two or More · Hispanic

Notes: This table summarizes the evolution of race and ethnicity questions across U.S. Census and American Community Survey data collection periods from 1990 to 2020. The "Form" column describes the specific question wording and response options available to respondents in each survey year. The "Coding" column shows how responses were processed and categorized by the Census Bureau. The "Categories Used in Paper" column indicates the final racial/ethnic classifications employed in this study's analysis. Key changes include: (1) the transition from single-race selection (1990) to multiple-race selection (2000 onward), (2) expanded write-in options and longer character limits in 2020, and (3) modifications to coding procedures that significantly affected multiracial classification rates, particularly for Hispanic respondents.

C.5.1 OMB Directive on Race from 1977

In 1977, the Office of Management and Budget (OMB) established a clear taxonomy for racial categorization, called Directive 15, identifying five primary groups, American Indians and Alaska Natives, Asians and Pacific Islanders, Non-Hispanic Blacks, Non-Hispanic Whites, and Hispanics of any race. This directive emphasized that these categories were not to be viewed as scientific or anthropological, signaling a deliberate departure from biological notions of race. (Strmic-Pawl, Jackson and Garner, 2018)

C.5.2 1990 Census Allowed Single Race Categorization

Starting from the 1980 Census, this standardized "ethnoracial pentagon" (White, Black, Indigenous, Asian, and Hispanic) was established and was used for the 1990 Census. In the 1990 Census, respondents could select only one racial category. If a respondent marked multiple racial categories, the Census Bureau recoded the individual into a single racial group, generally using the mother's race, if available. Additionally, the 1990 form asked separately about Hispanic origin (ethnicity), continuing the two-question format initiated in 1980, asking first whether the respondent was of "Spanish/Hispanic origin" and, if so, which national origin (e.g., Mexican, Puerto Rican, Cuban). (Pew Research Center, 2015)

C.5.3 2000 Census Changed to Multiple Race Categorization

A major shift came in the 2000 Census, when respondents, for the first time, could select more than one racial category to self-identify (see Table C.10). The racial identification question now explicitly instructed respondents to "mark one or more races," providing options such as "White," "Black or African American," "American Indian or Alaska Native," various Asian national origins (e.g., "Chinese," "Japanese," "Filipino"), "Native Hawaiian or Pacific Islander," and "Some other race." This resulted in 2.4% of Americans identifying with two or more races in 2000, significantly altering the measurement compared to 1990 (Pew Research Center, 2015). Though an explicit "multiracial" category was not adopted, this allowed for significantly greater complexity.

C.5.4 From 2010–2012 the Questions Were Consistent but Allowed for More Detailed Categories

In the 2010 Census, the format for race and Hispanic ethnicity questions remained consistent with 2000, maintaining the option for multiple racial selections (see Table C.10). However, the detailed categories expanded further, resulting in 63 possible race categories (six single-race options plus 57 possible combinations), and up to 126 combinations when including Hispanic ethnicity options (Strmic-Pawl et al., 2018). This resulted in greater specificity and complexity of self-reported racial data. By 2010, the percentage of Americans selecting two or more races increased slightly to approximately 2.9%, underscoring the continued evolution of multiracial identification in the Census data (Pew Research Center, 2015).

The 2008–2012 ACS closely mirrored the 2010 Census in its race and Hispanic origin questions, following the OMB's 1997 standards for collecting race and ethnicity separately. The Hispanic origin question was updated in 2008 to match the 2010 Census wording, asking, "Is this person of Hispanic, Latino, or Spanish origin?" with specific checkboxes for Mexican, Puerto Rican, and Cuban, plus a write-in option for other origins (U.S. Census Bureau, 2012b) . The race question allowed multiple selections and included 15 categories, with separate checkboxes for major Asian and Pacific Islander groups and a write-in option for "Other Asian" and "Other Pacific Islander" (U.S. Census Bureau, 2012b) .

C.5.5 2020 Census and ACS Changes in Race and Hispanic Origin Questions

- **Key Changes Introduced in the 2020 Census:**
  - Addition of write-in fields under the "White" and "Black or African American" categories, allowing respondents to specify detailed ethnic origins (e.g., "German," "Jamaican," "Nigerian").
  - Revision of Hispanic origin examples in write-in instructions to better represent major U.S. Hispanic subgroups, including the addition of "Ecuadorian" and "Guatemalan."
  - Modification of the "Some Other Race" (SOR) instructions from "Print race" to "Print race or origin," encouraging respondents to provide more detailed ethnic identification.
  - Expansion of coding capacity from 30 characters and two codes (in 2010) to 200 characters and six codes (in 2020).
  - Consolidation of coding lists: In 2020, race and Hispanic-origin write-ins were coded using a single combined list, rather than separate lists for ethnicity and race.

While these updates maintained consistency with OMB standards—separating race and Hispanic ethnicity and classifying individuals as either "Hispanic or Latino" or "Not Hispanic"—the new procedures significantly altered racial and ethnic classification (Marks and Rios-Vargas, 2021; Brumfield, Goldvale and Brown, 2019). Specifically, the combined coding list and expanded character limit led to dramatic increases in multiracial classifications. These procedural changes, rather than demographic shifts, largely explain the 275.5% growth in Americans classified as multiracial from 2010 to 2020 (Jones, Marks, Ramirez and Rios-Vargas, 2020).

As Arias, Liebler, Garcia and Sáenz (2025) highlight, the revised coding rules disproportionately impacted Hispanic and American Indian and Alaska Native (AIAN) populations. For instance, the Hispanic "White alone" population dropped by 75%, while the "White and SOR" category increased by 2,427%. Similarly, due to coding adjustments, a respondent marking "Black" and writing "Cuban" could be classified as multiracial ("Black and SOR"), causing Hispanic multiracial numbers to spike by 576.7%. The AIAN population in combination with another race grew by 160%, adding approximately 5.9 million individuals—a change likely driven by new write-in coding rather than genuine demographic trends (Arias et al., 2025).

By 2020, the ACS fully incorporated these 2020 Census changes, adopting identical write-in fields, coding procedures, and removing outdated terms such as "Negro" (U.S. Census Bureau, 2021). This alignment increased data detail and accuracy but introduced substantial discontinuities in historical trend comparisons, necessitating bridging methods for accurate longitudinal analyses (Arias et al., 2025).

C.6 Change in Coding Race for Hispanics Leading to Large #s of "Some Other Race"

From 1990 to 2020, the U.S. Census consistently asked respondents about Hispanic ethnicity separately from race. However, significant ambiguity persisted, as many Hispanic respondents felt their racial identity was inadequately captured by the available racial categories, frequently selecting "Some Other Race" (SOR) and providing Hispanic-origin write-in responses such as "Mexican" or "Latino." This issue was substantial: in 2010, for example, 37% of Hispanic respondents chose "Some Other Race," underscoring persistent measurement difficulties that complicate the comparability of racial demographic data over time (Pew Research Center, 2015).

The transition from the 2010 to the 2020 Census significantly exacerbated these measurement issues, particularly among Hispanic and multiracial populations, primarily due to methodological rather than demographic shifts. Between 2010 and 2020, the proportion of the U.S. population identifying as "Two or More Races" rose sharply from 2.9% to 10.2%, while the "White alone" category dropped from 72.4% to 61.6% (Starr and Pao, 2024). This change is largely attributed to Census Bureau revisions in question design and coding practices, rather than genuine shifts in racial self-identification. In 2020, approximately 94% of respondents categorized as "Some Other Race alone" were Hispanic individuals providing Hispanic-origin write-in responses (Jensen, Jones, Rabe, Pratt, Medina, Orozco and Spell, 2021). Furthermore, Hispanic respondents who selected both a conventional racial category and provided a Hispanic-origin write-in were often reclassified as multiracial (one conventional race plus SOR), whereas in 2010, similar responses typically remained single-race. Consequently, among Hispanics specifically, identification as "White alone" fell dramatically from 53.0% in 2010 to 20.3% in 2020, while the proportion classified as multiracial rose from 6% to 32.7% (Arias et al., 2025).

This reclassification practice, for example, recoded respondents who selected only "Black" but wrote "Cuban" in the race question as "Black plus Some Other Race," inflating multiracial counts even if respondents did not perceive themselves as multiracial (Starr and Pao, 2024). Similar reclassification widely affected respondents of Hispanic, Middle Eastern, or North African origin who had initially identified solely as "White" or "Black." Such coding practices significantly inflated multiracial counts and substantially reduced the "White alone" category.

These methodological changes created discontinuities not only in the 2020 Census but also in subsequent American Community Survey (ACS) data, beginning in 2021. Similar to the Census, the ACS adopted identical question design and coding updates, leading to parallel discontinuities between the 2019 ACS and the 2021–2022 ACS (Arias et al., 2025). While exact magnitudes differ slightly from the Census, the ACS experienced comparable shifts—decreases in "White alone" and substantial increases in multiracial identification among Hispanic respondents.

Recognizing these measurement challenges, the Census Bureau has employed bridging estimates to align recent data with historical racial definitions. For instance, official bridged population estimates indicated that the non-Hispanic White share of the population declined steadily from 69.1% in 2000 to 63.7% in 2010 and 57.8% in 2022 (Economic Policy Institute, 2022), reflecting more gradual demographic trends once methodological discontinuities are accounted for. **In the ACS data we use, we observe the share of the population identifying as non-Hispanic**

White declined from 63.16% in 2012 to 58.97% in 2022, which is very similar to the bridging estimates performed by the Census Bureau. Our understanding is that by using mutually exclusive categories (non-Hispanic White, non-Hispanic Black, non-Hispanic two or more, and Hispanic, see Table C.11), our data is less affected by these coding changes.

Table C.11: Racial Composition for Selected ACS Years

	ACS 2008–2012	ACS 2018–2022
% Non-Hispanic White	63.160	58.966
% Non-Hispanic Black	13.244	12.825
% Non-Hispanic Two or More	0.196	3.460
% Hispanic	16.239	18.386

Notes: This table reports racial and ethnic composition from the American Community Survey (ACS) 5-year estimates. Only the focal racial/ethnic categories examined in this analysis are presented; additional classifications in the NHGIS data are omitted, resulting in percentages that do not sum to 100%. The data indicate a 4.19 percentage point decline in the Non-Hispanic White population and increases of 2.15 and 3.26 percentage points in the Hispanic and Non-Hispanic Two or More Race populations, respectively.

C.7 Growth in Multiracial Identification of Survey Respondents

For the purposes of our project, we have two issues: (1) changes in Census questions and coding of racial/ethnic categories over time and (2) changes in how individuals identify their racial and ethnic identities. This section describes in detail the research on the stability of racial identities over time and the fact that there is a sociological change that can influence our sorting results along racial lines.

Since the U.S. Census Bureau began allowing individuals to select more than one racial category in Census 2000, the reported multiracial population has grown significantly. Between 2000 and 2010, the number of individuals identifying as multiracial increased by 32%, reflecting both demographic shifts and evolving social norms around racial identification (U.S. Census Bureau, 2012a). In the 2000 Census, approximately 6.8 million individuals, or 2.4% of the U.S. population, reported two or more races. By 2010, this number had risen to approximately 9 million, comprising 2.9% of the population (Jones and Bullock, 2012).

Between 2000 and 2010, the multiracial population in the United States grew substantially, reflecting shifting social norms and evolving patterns of racial self-identification. The number of individuals identifying as both White and Black or African American more than doubled, while those identifying as White and Asian increased by 87% (Pew Research Center, 2015). More broadly, the share of the U.S. population with two-race ancestry nearly doubled, rising from 2.2% in 1980 to 4.3% in the 2010–2012 period (Pew Research Center, 2015).

Racial self-identification is also highly fluid, with about 30% of adults with multiracial backgrounds reporting that they have changed how they describe their racial identity over their lifetimes (Davenport, 2020). This fluidity is especially pronounced among mixed-race adolescents, who are significantly more likely to revise their racial identification over time. The intersection of racial and ethnic identity further complicates self-identification, particularly among Hispanic individuals. While the Census Bureau classifies Hispanic origin as an ethnicity rather than a race, two-thirds of Hispanics perceive their Hispanic identity as part of their racial background (Davenport, 2020). As a result, changes in racial classification may reflect not only individual identity shifts but also the ambiguity in how Hispanic individuals navigate race within official classification systems.

Analyzing linked individual-level Census data from 2000 and 2010, Liebler, Porter, Fernandez, Noon and Ennis (2017) found that 6.1% of individuals (9.8 million people) changed their reported race and/or Hispanic origin between the two Censuses. However, this figure conceals significant differences across racial groups. Racial identification remained relatively stable among non-Hispanic Whites, Blacks, and Asians, with at least 90% consistency over time. In contrast, American Indians, Pacific Islanders, Hispanics, and multiracial individuals exhibited significantly higher levels of response change. Notably, only one-third of individuals who identified as American Indian in 2000 retained the same racial identification in 2010 (Liebler, 2016).

Measurement inconsistencies may also contribute to observed shifts in racial identification. The 2010 Census expanded racial categorization from six single-race options to 63 possible selections, introducing greater specificity but also potential confusion for respondents. This change likely increased measurement error, particularly for multiracial individuals and those identifying as "Some Other Race." Among Hispanic respondents, racial classification was especially unstable, with 13% reporting a different Hispanic or non-Hispanic designation in 2010. Additional factors such as age, geographic location, and survey enumeration method also influenced response changes, with younger individuals and those interviewed in person showing greater variation over time. These findings underscore the challenges of harmonizing racial and ethnic data across time periods, as changes in survey design may introduce variation that does not necessarily reflect underlying demographic or social shifts.

The expansion of racial classification options in Census 2000 and beyond reflects and reinforces an increased recognition of racial fluidity. As attitudes toward racial boundaries have shifted, particularly among younger generations, individuals are increasingly open to identifying with multiple racial backgrounds (Davenport, 2020).

C.8 Best Practices as per IPUMS NHGIS for Addressing Changes in Race-Ethnicity Questions and Coding Over Time

NHGIS addresses challenges posed by changes in racial identification over time by harmonizing multi-race responses in a consistent manner, particularly in its standardized-to-2010-tract time series data. Prior to the 2000 Census, respondents could select only one racial category, making it impossible to determine which individuals might have identified as multiracial in earlier Censuses. To address this limitation, NHGIS aligns pre-2000 race categories with the corresponding "one race alone" categories introduced in 2000 and used in subsequent Censuses. For instance, the pre-2000 category of “White” is aligned with “White alone” from 2000 onward. Individuals reporting “Two or More Races” in 2000 and later Censuses are treated separately, with no direct linkage to pre-2000 data. As a result, some of the observed shifts in single-race categories between 1990 and 2000 reflect this reclassification of multiracial individuals rather than actual demographic change.

While NHGIS primarily focuses on harmonizing pre-2000 racial classifications to align with post-2000 Census categories, it continues to rely on the racial categories provided by the Census Bureau and the American Community Survey (ACS) for post-2000 data, including the 2020 Census and 2020 and beyond ACS. Despite changes in how race is coded and processed, particularly in the 2020 Census, which modified race question instructions and expanded the detail captured in write-in responses, the racial categories themselves remain consistent from 2000 onward. NHGIS therefore adopts the Census and ACS classifications for post-2000 racial data without modification while focusing its harmonization efforts on pre-1990 classifications.

We consider NHGIS’s methodology to represent best practices for harmonizing racial categorization in time-series analyses. Accordingly, our analysis adopts this same approach, aligning pre-2000 race data with the single-race categories from 2000 onward and separately accounting for those reporting two or more races starting in 2000. This approach allows for a consistent examination of racial trends over time, ensuring that shifts in Census racial classifications do not introduce artificial discontinuities in the data.

C.9 Exploration of Variables from Census Summary Files

In Figures C.2, we show trends over time for all the variables of interest in our study, with separate lines depicting the mean for tracts 0 km from a RCRA site, tracts $(0, 10]$ km from a RCRA site, and all tracts in the United States.

Figure C.2: 10^{th}–90^{th} Percentiles of Housing Price Over Time

(a) 10^{th} Percentile

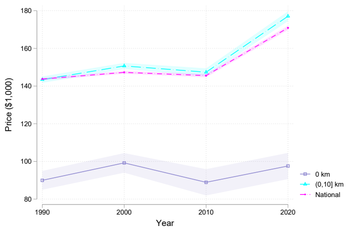

(b) 20^{th} Percentile

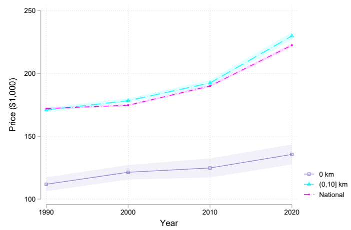

(c) 30^{th} Percentile

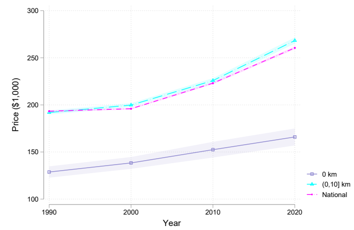

(d) 40^{th} Percentile

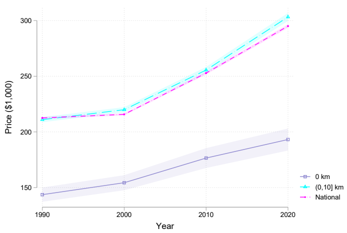

(e) 50^{th} Percentile

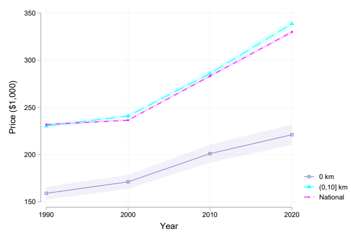

(f) 60^{th} Percentile

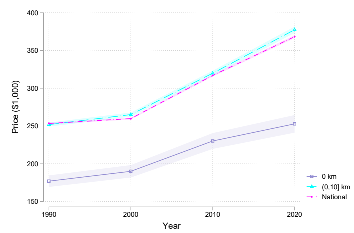

(g) 70^{th} Percentile

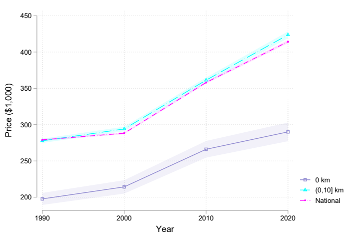

(h) 80^{th} Percentile

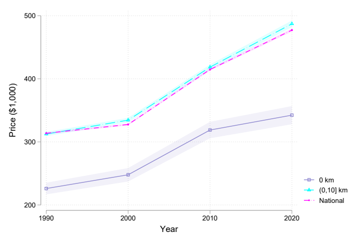

(i) 90^{th} Percentile

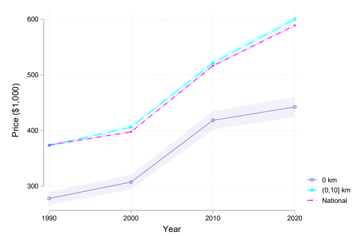

Notes: This figure depicts trends in housing percentiles over time. Lavender (solid) lines depict the mean for tracts containing a RCRA site, cyan (long dash) lines depict the mean for tracts $(0, 10]$ km away from the nearest site, and magenta (short dash) lines show the national mean. Shaded regions are constructed by adding/subtracting from the mean 1.96 times the standard error of the mean (defined as the standard deviation over the square root of the number count of observations).

Figure C.3: Income and Education-related Variables Over Time

(a) Avg HH Income

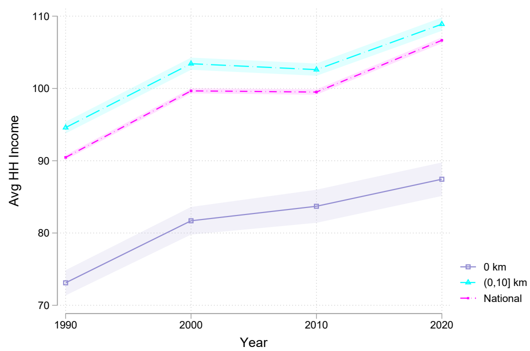

(b) % Below Poverty Line

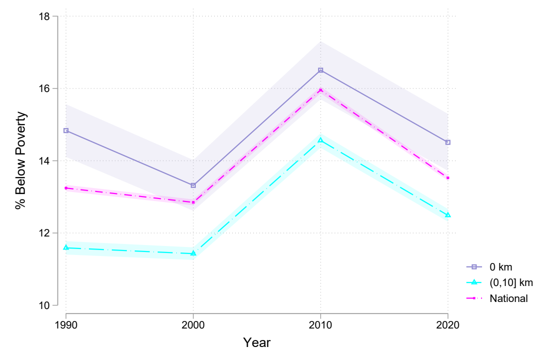

(c) % Public Assistance

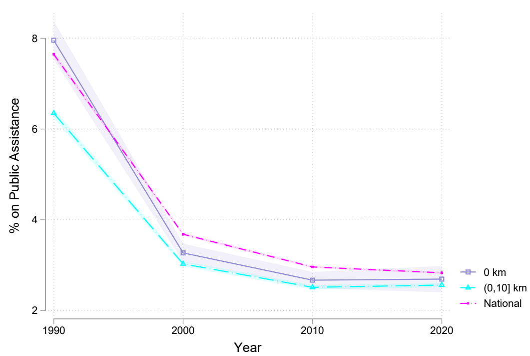

(d) % Unemployment

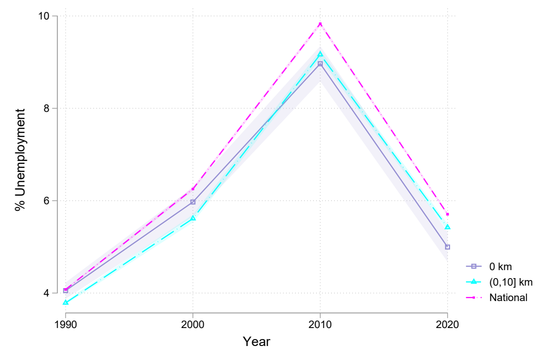

(e) % College Educated

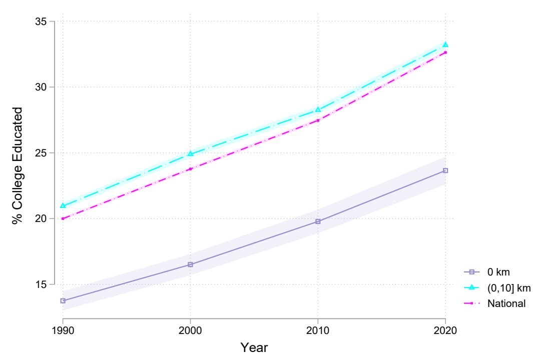

Notes: This figure depicts trends in income and education-related variables over time. Lavender (solid) lines depict the mean for tracts containing a RCRA site, cyan (long dash) lines depict the mean for tracts (0, 10] km away from the nearest site, and magenta (short dash) lines show the national mean. Shaded regions are constructed by adding/subtracting from the mean 1.96 times the standard error of the mean (defined as the standard deviation over the square root of the number count of observations).

Figure C.4: Demographic Variables Over Time

(a) % Black

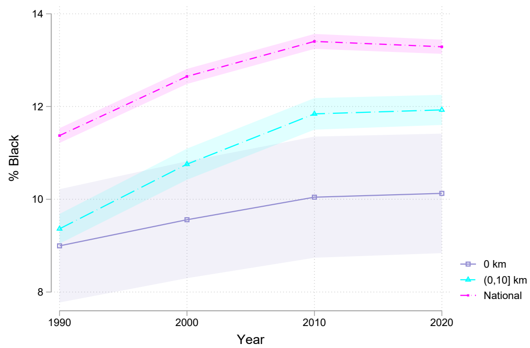

(b) % Hispanic

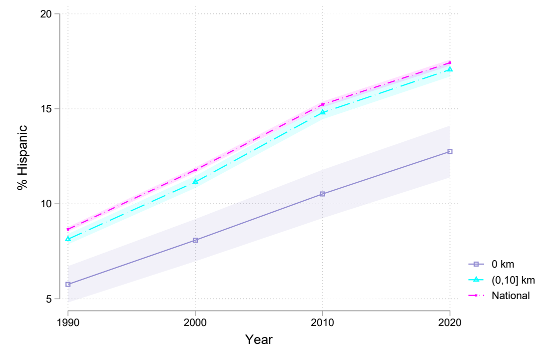

(c) % White

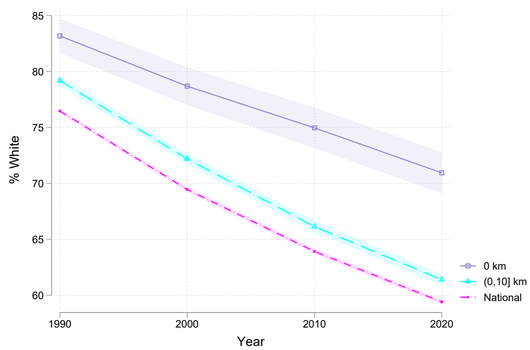

(d) % Female Head of Household

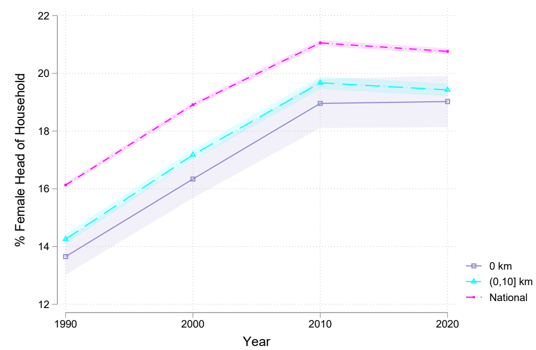

(e) % Under 18

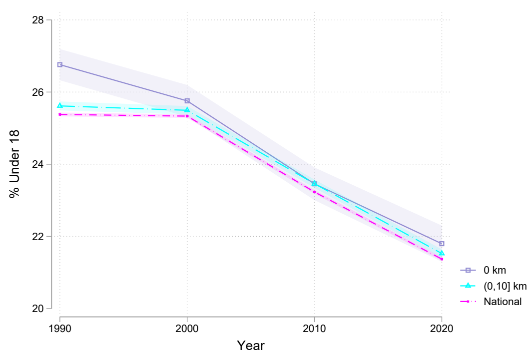

(f) Population Density

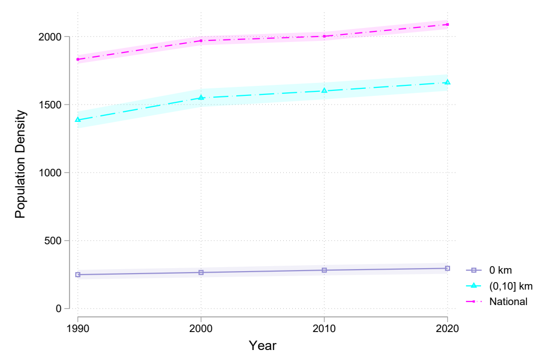

Notes: This figure depicts trends in demographic variables over time. All major racial categories (White, Black, Two or More) are non-Hispanic. Lavender (solid) lines depict the mean for tracts containing a RCRA site, cyan (long dash) lines depict the mean for tracts $(0, 10]$ km away from the nearest site, and magenta (short dash) lines show the national mean. Shaded regions are constructed by adding/subtracting from the mean 1.96 times the standard error of the mean (defined as the standard deviation over the square root of the number count of observations).

Figure C.5: Tenure by Race Variables Over Time

(a) % Black Owner

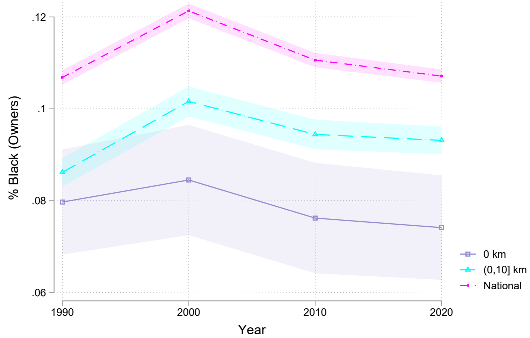

(b) % Hispanic Owner

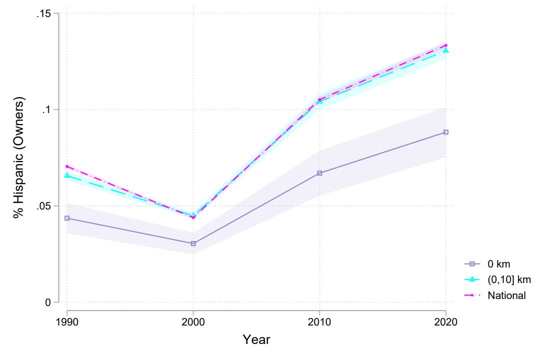

(c) % White Owner

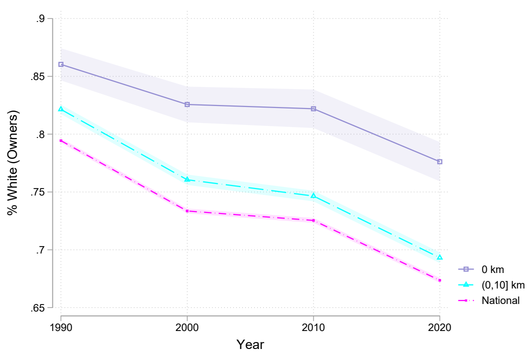

(d) % Two-or-more Race Owner

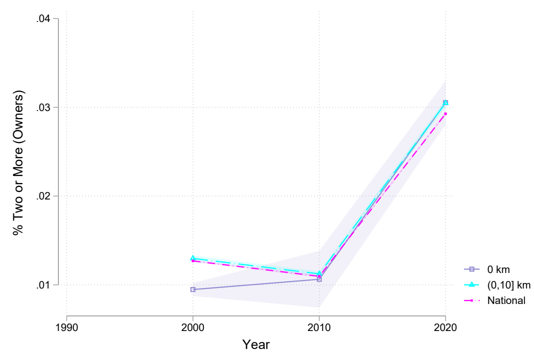

Notes: This figure depicts trends in demographic variables over time. All major racial categories (White, Black, Two or More) are non-Hispanic. Lavender (solid) lines depict the mean for tracts containing a RCRA site, cyan (long dash) lines depict the mean for tracts $(0, 10]$ km away from the nearest site, and magenta (short dash) lines show the national mean. Shaded regions are constructed by adding/subtracting from the mean 1.96 times the standard error of the mean (defined as the standard deviation over the square root of the number count of observations).

Figure C.6: Housing-related Variables Over Time

(a) % 4+ Beds

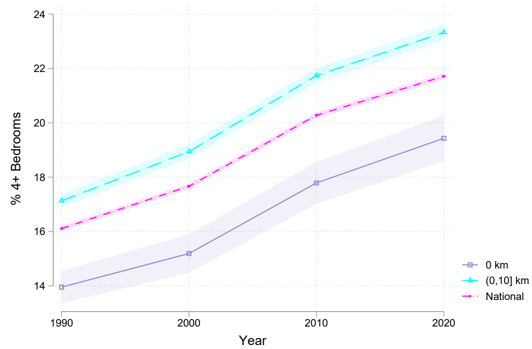

(b) % Mobile Home

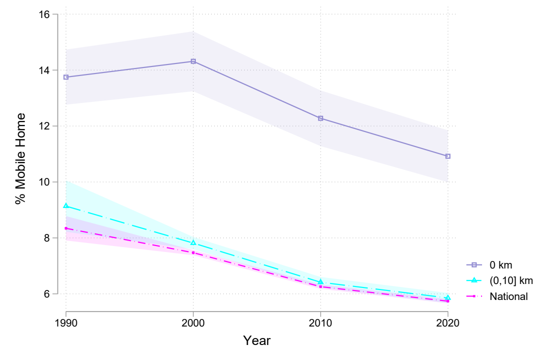

(c) % Owner Occupied

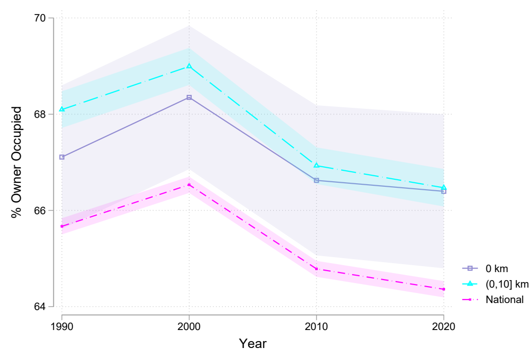

(d) % Vacant

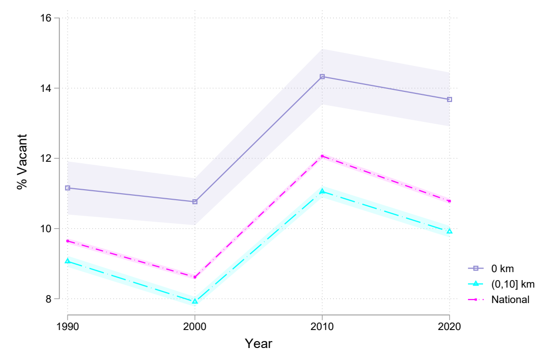

(e) % Built in Last 10 years

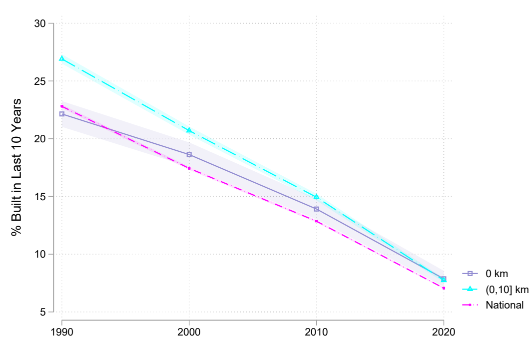

(f) % Moved in Last 10 Years

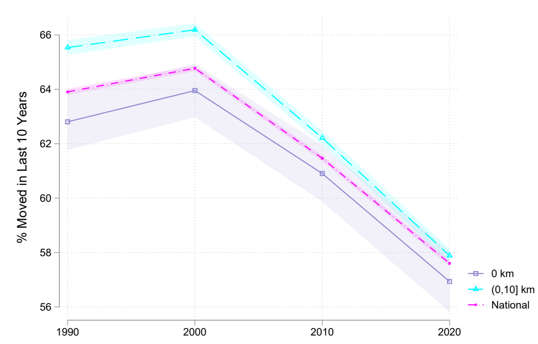

Notes: This figure depicts trends in housing-related variables over time. Lavender (solid) lines depict the mean for tracts containing a RCRA site, cyan (long dash) lines depict the mean for tracts $(0, 10]$ km away from the nearest site, and magenta (short dash) lines show the national mean. Shaded regions are constructed by adding/subtracting from the mean 1.96 times the standard error of the mean (defined as the standard deviation over the square root of the number count of observations).

Figure C.7: 10^{th}–90^{th} Percentiles of Rent Over Time

(a) 10^{th} Percentile

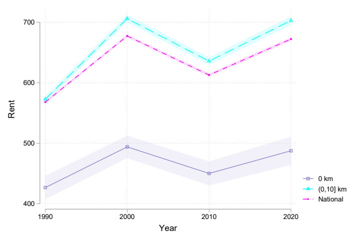

(b) 20^{th} Percentile

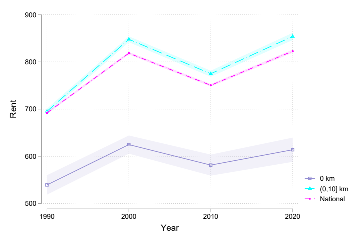

(c) 30^{th} Percentile

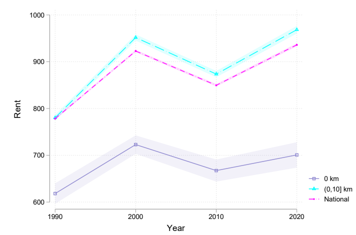

(d) 40^{th} Percentile

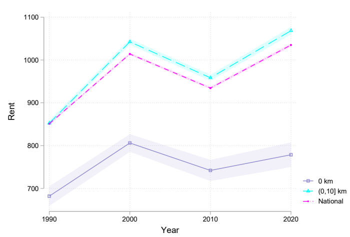

(e) 50^{th} Percentile

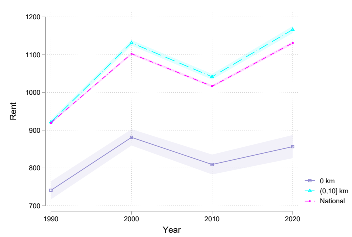

(f) 60^{th} Percentile

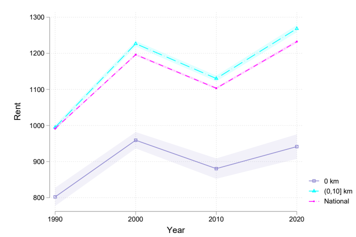

(g) 70^{th} Percentile

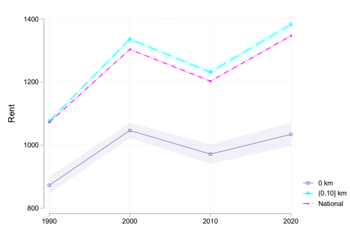

(h) 80^{th} Percentile

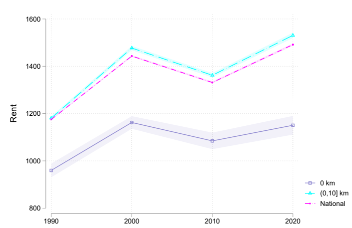

(i) 90^{th} Percentile

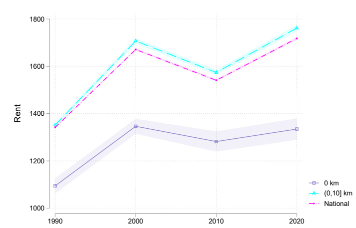

Notes: This figure depicts trends in rent percentiles over time. Lavender (solid) lines depict the mean for tracts containing a RCRA site, cyan (long dash) lines depict the mean for tracts (0, 10] km away from the nearest site, and magenta (short dash) lines show the national mean. Shaded regions are constructed by adding/subtracting from the mean 1.96 times the standard error of the mean (defined as the standard deviation over the square root of the number count of observations).

C.10 Tenure by Race

Table C.12: Tenure by Race Variable Sources and Definitions by Year

Year	Dataset Used by Type	ME and NME Vars
1990	**ME:** 1990 STF1 **NME:** 1990 STF1	**ME Owner:** owner_nonhisp_white, owner_nonhisp_afamer, base_owner_byrace **ME Renter:** renter_nonhisp_white, renter_nonhisp_afamer, base_renter_byrace **NME Owner:** owner_white_alone, owner_afamer_alone, owner_hispanic, base_owner_nme **NME Renter:** renter_white_alone, renter_afamer_alone, renter_hispanic, base_renter_nme
2000	**ME:** 2000 SF1b; 2000 SF2 **NME:** 2000 SF1b; 2000 SF2	**ME Owner:** owner_nonhisp_white, owner_nonhisp_afamer, owner_nonhisp_two_or_more, base_owner_byrace **NME Owner:** owner_white_alone, owner_afamer_alone, owner_two_alone, owner_hispanic, base_owner_nme **NME Renter:** renter_white_alone, renter_afamer_alone, renter_two_alone, renter_hispanic, base_renter_nme
2010	**ME:** 2010 SF1b **NME:** 2008–2012 ACS; 2010 SF1a	**ME Owner:** owner_nonhisp_white, owner_nonhisp_afamer, owner_nonhisp_two_or_more, base_owner_byrace **ME Renter:** renter_nonhisp_white, renter_nonhisp_afamer, renter_nonhisp_two_or_more, base_renter_byrace **NME Owner:** owner_white_alone, owner_afamer_alone, owner_two_alone, owner_hispanic, base_owner_nme **NME Renter:** renter_white_alone, renter_afamer_alone, renter_two_alone, renter_hispanic, base_renter_nme
2022	**ME:** 2020_DHCa **NME:** 2020_DHCa; ACS 2018–2022	**ME Owner:** owner_nonhisp_white, owner_nonhisp_afamer, owner_nonhisp_two_or_more, base_owner_byrace **ME Renter:** renter_nonhisp_white, renter_nonhisp_afamer, renter_nonhisp_two_or_more, base_renter_byrace **NME Owner:** owner_white_alone, owner_afamer_alone, owner_two_alone, owner_hispanic, base_owner_nme **NME Renter:** renter_white_alone, renter_afamer_alone, renter_two_alone, renter_hispanic, base_renter_nme

Notes: ME = Mutually Exclusive; NME = Non-Mutually Exclusive. For each survey year, we identify the specific datasets used and classify tenure data by race. ME variables exclude Hispanics, while NME variables allow overlap with Hispanic origin.

As part of our analysis, we utilize NHGIS data on housing tenure, classifying units as either "Owner Occupied" or "Renter Occupied," stratified by race. We define racial categories using both mutually exclusive(ME) and non-mutually exclusive(NME) variables.. The mutually exclusive racial categories we include to match our main approach in the paper are "White Non-Hispanic," "African American Non-Hispanic," "Two or More Non-Hispanic," and "Hispanic." Due to changes in the summary files over time, this requires us to use the Decennial Summary Files for 2010 and 2020. Furthermore, 2000 does not have tenure-by-race at a lower geography than tract and so we are interpolating tract to tract for 2000 data. Figure C.5 show trends in the variables used in our paper over time. Table C.12 provides a detailed overview of the specific datasets and racial categorizations used by survey year.

The ACS 5-year estimates used in our paper only provide non-mutually exclusive racial categories, such as “White Alone,” “African American Alone,” and “Two or More Alone”. The “Alone” classification allows for both Hispanic and Non-Hispanic individuals to be included hence, being labeled as Non-Mutually Exclusive. We do not use these non-mutually exclusive categorizations in any descriptive summaries as tenure-by-race data are only used in Table 7.

C.11 Acreage and Land Use

Lot size information was needed for robustness checks in Tables A.6 and A.7, but was not publicly available at the tract level (or any geographic level below tract) for 1990, 2000, or 2010 (including both Decennial Census and ACS). The data is only available at the county level for 1990, and is limited to just two categories: “In housing units on properties of less than 1 acre” and “In housing units on properties of 1 acre or more.” The county-level variables downloaded represent counts of persons (not households or housing units) occupying housing units across various categories:

- Urban: Inside urbanized area
  - In households on properties of less than 1 acre
  - In households on properties of 1 acre or more
  - In group quarters
- Urban: Outside urbanized area
  - In households on properties of less than 1 acre
  - In households on properties of 1 acre or more
  - In group quarters
- Rural
  - In households on properties of less than 1 acre
  - In households on properties of 1 acre or more: Farm
  - In households on properties of 1 acre or more: Nonfarm
  - In group quarters

These county-level data were merged under the assumption that county boundaries remain stable over time (and only experience nominal changes, e.g. Dade County vs. Miami-Dade). Each count variable was then converted to a percentage by dividing by the total represented population.

D Choice of Distance Measure

In this appendix section, we describe the reasoning for our choice to use minimum distance to Census tract boundary as our distance measure. We additionally provide robustness checks using an alternative definition of distance- distance to population-weighted centroid.

One disadvantage of using minimum distance from any point in the tract to the nearest RCRA site is that larger tracts are more likely to contain RCRA sites and hence are more likely to fall in the "0 km" bin. In panel (a) of Figure D.1, we plot the 0 km distance bin as well as the other bins within 10 km. We see that the treatment and control groups are quite different when it comes to tract area. We think of treatment as possibly affecting an entire tract, since a tract represents a neighborhood and percentiles of the housing price would depend on all the houses in the neighborhood. Therefore, we see this issue as natural to our context, and thus primarily a challenge for estimation and inference, rather than a problem with how we are measuring distance. As such, we address these challenges using a regression adjustment in Table 5.

An alternative distance measure that could be employed in our main specification is the distance from the centroid to the site. Below, we discuss why we chose the minimum distance measure rather than measuring the distance from centroid to tract.

The first reason we do not prefer the centroid-based measures is that the centroid-based measures resulted in the sample lacking larger tracts. In the case of a non-population-weighted centroid, the issue is largely mechanical. For a census tract containing a site, the likelihood that the centroid is, say, $[0, 1]$ km from the site depends on the total area of the tract. This effect is partially mitigated by using population-weighted centroids, though, since these need not be centrally located.

To illustrate the problem, we produce kernel densities in Figure D.1 and summary statistics in Table D.1. The very largest geographic units nationally are less represented among the treatment or control groups defined according to the centroid-based distance measure than the minimum-distance based one. The population-weighted centroid situation is an in-between case. However, as a whole, the 5-10 km control group for the minimum distance measure is more similar to the overall national area distribution. In fact, the entire $[0, 10]$ km radius around a RCRA site in minimum distance is more similar to the national area distribution than the $[0, 10]$ km radius around a RCRA site in centroid distance for either centroid-based measure.

Another reason why the minimum-distance based bins are helpful in our context is they ensure that control units are not "close" to sites. To illustrate the issue, consider a comparison with the $(5, 10]$ km distance bin (Figure D.2). With minimum distance, we know for sure that tracts in the $(5, 10]$ km minimum distance bin cannot contain any houses within 5 km of a RCRA site. This is not true of the centroid-based approach- the $(5, 10]$ km minimum distance bin defined according to centroid based measures could contain houses even within 1 km of a RCRA site. This could happen in two ways. First, the site might be on an adjacent tract and close to some of the houses

Figure D.1: Kernel Densities of Area by Distance Measure

(a) Min Distance

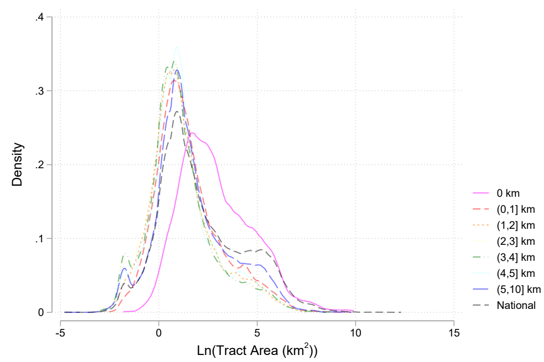

(b) Centroid Distance

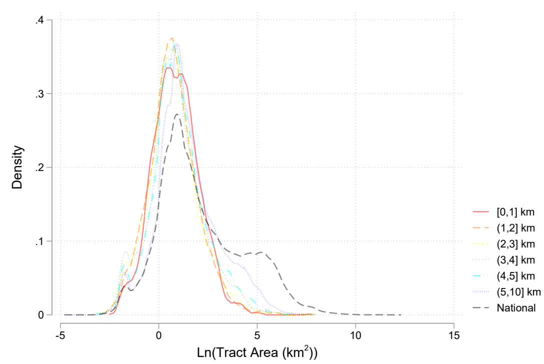

(c) Population-weighted Centroid Distance

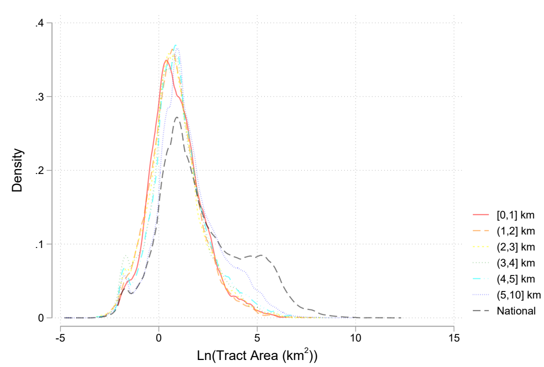

Notes: This figure compares the distribution of tract areas in square kilometers (measured in logarithms for readability) for three different candidate measures of distance to RCRA site. Colored lines represent binned distances, and the dashed dark gray line depicts the overall distribution of tract areas for all tracts in the United States.

in the tract. Second, the site might be on the tract, but the tract is large enough that the centroid is far from the tract. In either situation, the control group is "treated" to some extent, which is undesirable for inference.

Table D.1: Area by Distance Measure

	Mean	St. Dev.	Percentiles			Min	Max
			10^{th}	50^{th}	90^{th}		
Minimum							
0 km	159.401	923.311	2.213	13.712	238.378	0.164	20,042.633
(0, 1] km	39.253	326.306	0.722	3.307	62.220	0.079	17,168.701
(1, 2] km	26.906	152.233	0.539	2.548	40.462	0.041	4,461.864
(2, 3] km	26.377	185.274	0.444	2.400	33.579	0.047	10,829.257
(3, 4] km	25.016	140.990	0.463	2.408	32.871	0.051	5,987.342
(4, 5] km	34.585	277.451	0.574	2.677	47.998	0.041	9,353.257
(5, 10] km	46.722	278.971	0.689	3.595	98.000	0.008	18,028.428
[0, 10] km	43.467	330.743	0.616	3.112	69.755	0.008	20,042.633
Centroid							
[0, 1] km	6.340	56.770	0.567	2.322	10.911	0.079	2,321.146
(1, 2] km	6.622	49.364	0.439	1.958	10.032	0.056	2,579.057
(2, 3] km	6.763	46.036	0.423	1.993	11.454	0.041	3,063.702
(3, 4] km	8.425	45.889	0.410	2.170	13.979	0.051	2,561.966
(4, 5] km	9.515	32.319	0.535	2.460	17.447	0.041	889.067
(5, 10] km	18.670	67.082	0.708	3.311	40.795	0.008	2,801.969
[0, 10] km	12.553	55.925	0.539	2.595	21.661	0.008	3,063.702
Pop-Weighted Centroid							
[0, 1] km	7.901	27.614	0.551	2.180	13.283	0.125	645.208
(1, 2] km	8.099	49.856	0.465	2.079	12.647	0.056	2,429.864
(2, 3] km	9.670	69.809	0.438	2.086	14.501	0.041	4,313.722
(3, 4] km	12.259	116.354	0.424	2.249	15.930	0.051	7,275.870
(4, 5] km	12.541	49.297	0.533	2.479	19.391	0.041	1,020.624
(5, 10] km	25.931	242.063	0.709	3.309	44.552	0.008	19,679.541
[0, 10] km	17.213	169.360	0.545	2.617	24.454	0.008	19,679.541
All Tracts							
National	127.965	1,428.722	0.815	5.037	218.217	0.008	222,977.203

Notes: This table compares summary statistics on tract area for the three candidate measures of distance in this paper. The distance bins are each defined according to the distance measures directly above them in italics. The first group of distance bins is defined according to the minimum-distance from tract to site measure. The second is defined as the distance from the geometric centroid to the nearest site. The third is defined as the distance from the population-weighted centroid to the nearest site. At the bottom of the table, we present statistics for all tracts in the United States for comparison purposes.

Despite the potential drawbacks of a centroid-based distance measure, it is valuable as a check on our primary estimates. For this activity, we choose to use the population-weighted centroid since it is less likely than the geometric centroid to suffer from the aforementioned issues.

It is unclear a priori what the equivalent of our "0 km" distance measure is in centroid distance (especially given that some homes in "0 km" minimum distance are even 10 km away from the site). Therefore, we determine the distance empirically using the analogue of Figure 3, as well as

Figure D.2: Kernel Densities of Actual Distance by Distance Measure

(a) Min Distance

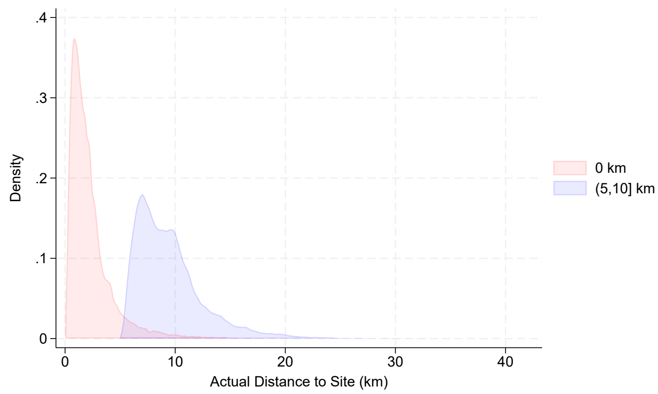

(b) Centroid Distance

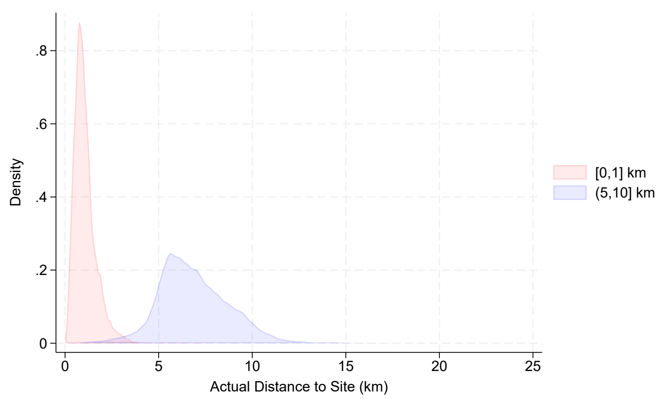

(c) Population-weighted Centroid Distance

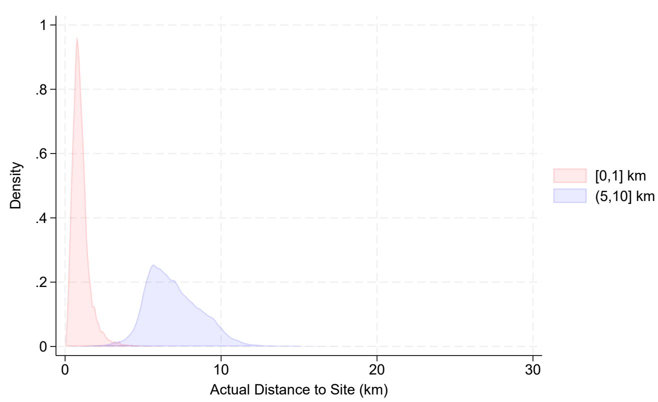

Notes: In this diagram, we employ Zillow Ztrax data from Ohio and Pennsylvania to show the kernel densities of actual distance from house to nearest RCRA site for different measures of distance that could be employed in the main analysis. In panel (a), the peach region is the density for tracts containing a RCRA site ("0 km") and the lavender region is the density for tracts with minimum distance between 5 and 10 km away from the site. In panel (b) the peach region is the kernel density for tracts whose geometric centroid is between 0 and 1 km away from the nearest RCRA site, whereas the lavender region is those whose geometric centroid is 5–10 km away from the nearest RCRA site. Panel (c) is the same as panel (b), but applying the population-weighted centroid instead.

Figure D.3: Housing Price - Distance Gradient Using Population-Weighted Centroid Distance

(a) Excluded Distance Bin: $(5, 10]$ km

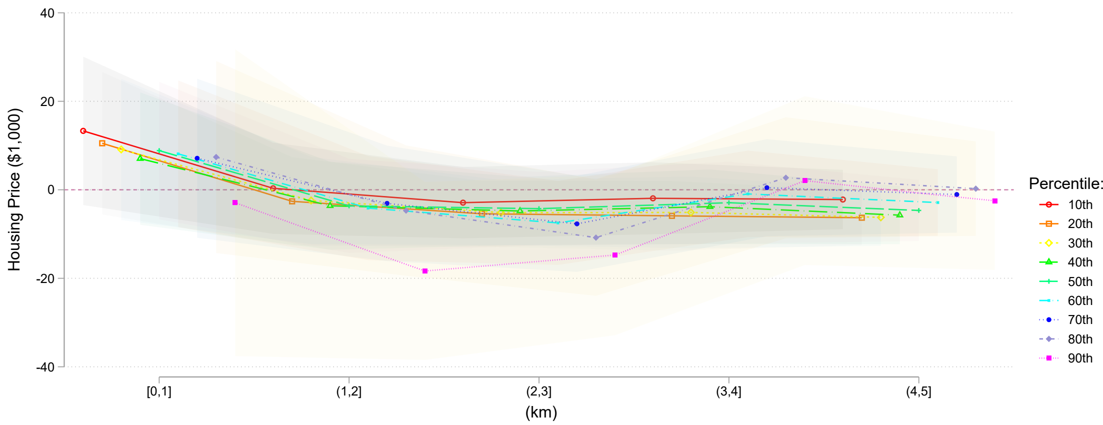

(b) Excluded Distance Bin: $(10, 20]$ km

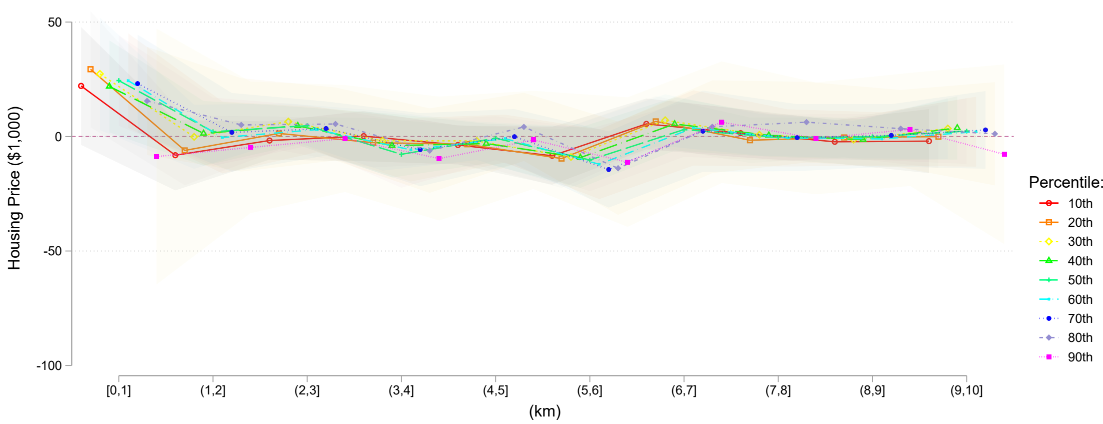

Notes: This figure plots the set of Difference-in-Differences (DD) coefficients on the interaction of each 1-km distance bin, defined according to population-weighted centroid distance, for each percentile from the 10^{th} through the 90^{th}. All specifications include fixed effects for tract, bin by year, site by year, and state by year. The sample is the set of tracts containing exactly one RCRA site within 10km (Panel a) or 20km (Panel b) in 1990, 2000, 2010 (ACS 2008-2012), and 2020 (ACS 2018-2022).

a version of the same figure that takes the bins out to 10 km and uses $(10, 20]$ km as the excluded distance. The results are found in Figure D.3. Both indicate that the impacts are concentrated in the $[0, 1]$ km centroid-distance bin. Therefore, we will designate $[0, 1]$ km to be the "Near" bin for this exercise.

Next, we show the point estimates from the population-weighted centroid regressions in Figure D.4. For panel (a), where we designate $[0, 1]$ km as the "Near" bin and $(1, 10]$ as the excluded bin, we see that magnitudes of lower percentiles are higher than in our primary regressions, whereas magnitudes approach 0 at the highest percentiles. The figure also confirms that the highest sta-

Figure D.4: Robustness Check: Population-Weighted Centroid Distance

(a) Excluded Distance Bin: $(1, 10]$ km

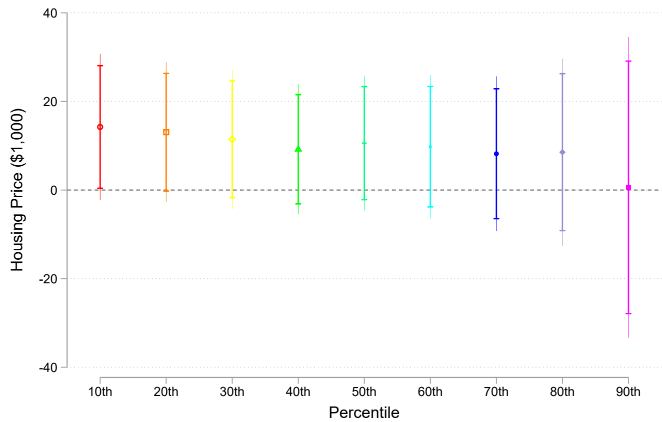

(b) Excluded Distance Bin: $(1, 20]$ km

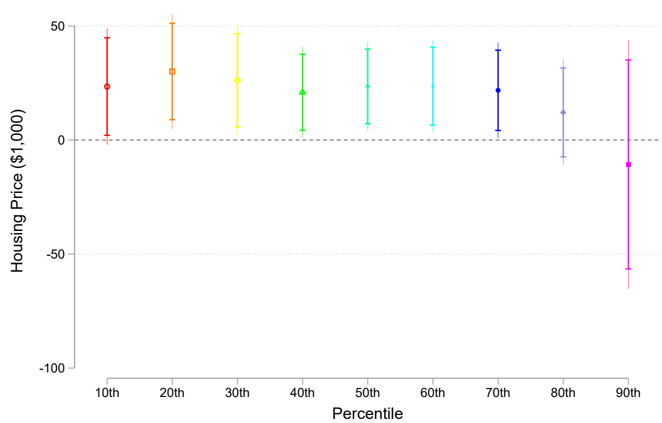

Notes: This figure plots the set of Difference-in-Differences (DD) coefficients from the interaction of Post $\times$ $[0, 1]$ km, where the distance is calculated from population-weighted centroid to site, for each percentile from the 10^{th} through the 90^{th}. Panel (a) employs an excluded distance bin of $(1, 10]$ km, whereas panel (b) uses an excluded distance bin of $(1, 20]$ km. All specifications include fixed effects for tract, bin by year, site by year, and state by year. The sample is the set of tracts containing exactly one RCRA site within (Panel a) or 20km (Panel b) in 1990, 2000, 2010 (ACS 2008-2012), and 2020 (ACS 2018-2022).

tistical significance is seen for the lower percentiles of the price distribution, as with our primary results. For panel (b), where we instead use $(1, 20]$ as the excluded distance bin, the magnitudes of the impacts are even higher at lower deciles of the price distribution, and the 90th percentile is very imprecisely estimated. The overall takeaway is that the impacts we detect at the lowest price deciles are also present using this alternative distance measure.

E Multiple Hypothesis Testing

In this appendix section, we present and discuss Table E.1, which addresses multiple hypothesis testing for our main results.

In the table, the second column shows the unadjusted conventional cluster-robust standard errors from our main regressions, whereas the remaining columns feature different methods of adjusting for multiple hypothesis testing. The third column shows the Bonferroni-corrected p-values, which are the most conservative, but have a documented tendency to over-reject the null hypothesis. The fourth column shows p-values corrected using the Holm method (Holm, 1979). This is a step-down method that allows for arbitrary dependence between hypothesis tests.[16] The method fixes the FWER rather than merely fixing the significance level of each individual hypothesis test. The sixth column shows the Westfall-Young corrected p-values (Westfall and Young, 1993). This is a bootstrap resampling-based correction that also fixes the FWER, but takes into account the fact that the estimators are not necessarily independent. The resampling procedure ensures that the FWER will still be controlled under dependence across outcomes.

Table E.1: Adjustment of p-values for Multiple Hypothesis Testing

Dep. var: Price^{kth}	Conventional	Bonferroni	Holm	Westfall-Young	Romano-Wolf
10^{th}	0.009	0.082	0.079	0.038	0.026
20^{th}	0.015	0.112	0.107	0.054	0.031
30^{th}	0.014	0.112	0.107	0.053	0.030
40^{th}	0.028	0.159	0.149	0.068	0.046
50^{th}	0.037	0.159	0.149	0.087	0.057
60^{th}	0.047	0.159	0.149	0.092	0.058
70^{th}	0.027	0.159	0.149	0.068	0.046
80^{th}	0.047	0.159	0.149	0.092	0.058
90^{th}	0.170	0.170	0.170	0.158	0.122

Notes: This table presents p-values adjusted for multiple hypothesis testing, for the hypothesis test that Near × Post=0. For comparison, the first column reproduces unadjusted p-values from our main specification (Table 2) - these are constructed from clustered standard errors. The second column features Bonferroni corrected p-values. The third column shows Holm adjusted p-values. The fourth and fifth columns show Westfall-Young and Romano-Wolf p-values, which are produced using resampling with 999 clustered bootstrap replications.

The Westfall-Young procedure requires the assumption of subset pivotality, which entails imposing the restriction that the constraints of all null hypotheses jointly hold. Romano, Shaikh and Wolf (2008) assert that this assumption is overly stringent, while Westfall and Troendle (2008) argue that it is not. In our case, the Westfall-Young p-values are mostly between the Holm and Romano-Wolf p-values described below. While subset pivotality might not be necessary in our

[16] Starting with the individual p-values, we iteratively inflate them such that the smallest p-value is inflated the most, and the largest p-value is not inflated at all (unless one of the smaller p-values is inflated above the largest p-value, then the largest p-value will be inflated to preserve the order of p-values).

context, we see the Westfall-Young p-values as complimentary to the Romano-Wolf p-values we present, providing additional evidence on multiple hypothesis testing.

The seventh column shows the the Romano-Wolf corrected p-values (Romano and Wolf, 2005a,b), computed using the algorithm in Romano and Wolf (2016). The Romano-Wolf method accounts for the fact that there is dependence among the estimators. The intuition is as follows: consider two extreme cases where hypothesis tests for two estimators are conducted. In the first case, the two estimators are perfectly dependent (say, they come from the exact same regression using the exact same data, and produce an identical coefficient, for example). In that case, we should not penalize testing the two hypotheses at all. In the second case, the two estimators are independent. In that case, we should heavily penalize testing the two hypotheses.

For both bootstrap-based adjustments, we use 999 replications and cluster the bootstrap at the Census tract level.[17] Holm-Sidak, Westfall-Young, and Romano Wolf all fix the family-wise error rate (Romano and Wolf, 2005a).

The overall conclusion is that the two resampling-based multiple hypothesis correction techniques indicate that the evidence for the 10th percentile is statistically significant at the 5% level even after accounting for MHT. The evidence is weaker for the higher deciles, though not always monotonically so.

[17] Jones, Molitor and Reif (2019) show in their online appendix the Westfall-Young algorithm performs poorly with clustered standard errors but a non-clustered bootstrap.

F Software and Packages Used

- We geocoded using ArcPy (ArcGIS Pro 10 Python 3).
- Stata Packages used:
  - wyoung package (Reif, 2017)
  - rwolf2 package (Clarke, Romano and Wolf, 2020)
  - reghdfe package (Correia, 2016)
  - ppmlhdfe package (Correia et al., 2020)
  - blindschemes package Bischof (2016)
  - geoplot package (Jann, 2023)