Title
Essays on Energy and Public Economics

Author
Lyubich, Eva

Publication Date
2022

Peer reviewed|Thesis/dissertation

Essays on Energy and Public Economics

by

Eva Lyubich

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Reed Walker, Chair
Professor Patrick Kline
Professor Emmanuel Saez
Professor Joseph Shapiro

Spring 2022

Essays on Energy and Public Economics

Abstract

Essays on Energy and Public Economics

by

Eva Lyubich

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Reed Walker, Chair

Energy is an essential input into the lives of individuals and the production of firms. While energy use has numerous private benefits, it imposes an external social cost as the combustion of fossil fuels increases the concentration of greenhouse gases and accelerates the climate crisis. A central role of the public sector is to reduce such externalities. Governments can intervene by creating financial incentives such as taxes or subsidies, by investing into public goods such as research and development of green technologies or infrastructure that decreases aggregate fossil fuel energy demand, or by directly regulating emissions through caps or bans. The aggregate and distributional impact of such policies depends on baseline energy use patterns and available alternatives. My dissertation examines heterogeneity in energy use and carbon emissions — documenting it, exploring its drivers, and discussing its implications for the impact of different public sector interventions.

In Chapter 1, I examine place-based differences in individual energy use and carbon emissions. There is substantial spatial heterogeneity in household carbon emissions across the US, and a strong association between emissions and local amenities such as density, transportation infrastructure, and housing characteristics. I estimate what share of heterogeneity in carbon emissions is attributable to places themselves, and what share reflects individual characteristics and sorting. To do this, I construct a longitudinal panel of residential energy use and commute characteristics for over a million individuals from two decades of administrative Decennial Census and American Community Survey data. I use movers in my data to estimate place effects – the amount by which carbon emissions change for the same individual living in different places – for almost 1,000 cities and roughly 60,500 neighborhoods across the US. I find that place effects explain more than half of differences between places, and about 15-25% of overall variation in carbon emissions. My estimates suggest that decreasing neighborhood-level place effects from one standard deviation above the mean to one standard deviation below the mean would decrease household carbon emissions from residential energy use and commuting by about 40%.

In Chapter 2, I examine racial differences in individual energy expenditures. Using publicly available data from the American Community Survey from 2010-2017, I show that Black households have higher residential energy expenditures than white households in the US. This residential energy expenditure gap persists after controlling for income, household size, home-owner status, and city of residence. It decreased but did not disappear between 2010 and 2017, and it is fairly stable in levels across the income distribution, except at the top. Controlling for home type or vintage does not eliminate the gap, but survey evidence on housing characteristics and available appliances is consistent with the gap being driven at least in part by differences in housing stock and related energy efficiency investments.

In Chapter 3, which is co-authored with Joe Shapiro and Reed Walker, we examine firm-level variation in carbon emissions. We provide the first estimates of within-industry heterogeneity in energy and CO_2 productivity for the entire U.S. manufacturing sector. We measure energy and CO_2 productivity as output per dollar energy input or per ton CO_2 emitted. Three findings emerge. First, within narrowly defined industries, heterogeneity in energy and CO_2 productivity across plants is enormous. Second, heterogeneity in energy and CO_2 productivity exceeds heterogeneity in most other productivity measures, like labor or total factor productivity. Third, heterogeneity in energy and CO_2 productivity has important implications for environmental policies targeting industries rather than plants, including technology standards and carbon border adjustments.

To snow.

Contents

List of Figures

List of Tables

Acknowledgments

First and foremost, I am grateful to my dissertation advisor Reed Walker, who provided advice, feedback, guidance, and support from the first year of my PhD through every stage of the doctoral program and research process. I am also grateful to my other committee members – Emmanuel Saez especially for always encouraging me to think bigger, which helped me stay inspired in my work; Joe Shapiro especially for his grounded logic, which always made hard tasks feel more manageable; and Pat Kline especially for taking me in as a student despite our different research fields and vastly deepening my technical understanding.

They say it takes a village, and I am grateful that the rest of my academic village consisted of the Department of Economics and the Energy Institute at Haas. I've learned from interactions with too many people to name, but I am especially thankful to Alan Auerbach, Severin Borenstein, David Card, Meredith Fowlie, Cecile Gaubert, Hilary Hoynes, Enrico Moretti, Jesse Rothstein, Jim Sallee, Danny Yagan, Katherine Wagner, and Chris Walters for their generosity with their time and for their many helpful conversations and suggestions. I also owe a debt of gratitude to Ben Handel and Jon Kolstad, whose championship, kindness, and encouragement from the very beginning has been instrumental in getting me here and building my confidence as a researcher.

Several other people were central in helping me along my journey to Berkeley. Arun Chandrasekhar, Greg Lewis, Markus Mobius, and Ben Olken were among the first to teach me to ropes of economics research while I was at Microsoft Research. Sri Nagavarapu – who patiently answered my questions in office hours every week and showed me the path to an Economics PhD when I first decided to switch professions as a senior in college – is an exemplar of the impactful mentorship I hope to have the opportunity to pay forward some day. Jim Valles was the first researcher to take a chance on me when he let me join his research lab as a sophomore, and he continued to support my professional aspirations when I decided that fluid mechanics was not my calling. Last but certainly not least, my parents, Lilya and Misha, taught me from a very young age that learning is a gift, and that I should always interrogate my assumptions about the world.

I am lucky to have also had the support of many classmates and friends through these years. First year would not have been the same without my study group with Emily, Helen, and Ingrid. Teaching was a highlight of my PhD experience, thanks in part to the warmth and kindness of my co-GSI Arlen. My friendships with other graduate student researchers (past and present) at the Energy Institute consisted of a lovely balance of chats about energy economics, California backpacking trips, and food. Regular check-ins with Maria and Steph kept me connected while uncertainty loomed in the depth of the pandemic and beyond, and Nina kept me focused while simultaneously never letting me neglect to have fun. My friends from other corners of life – especially Alejandra, Ali, Bev, Cody, Emily, John, Katie, Kristina, Margiana and Simon – lifted me through the lows, celebrated with me in the highs, offered new perspectives, and reminded me why the work was worth doing. My partner Josh was steadfast as I navigated stress and uncertainty, and supported me in every way; from patiently explaining how to use github, to baking cupcakes when I needed a snack.

This research would not have been possible without funding and support from several institutions. I am grateful to have received generous support from the California Policy Lab, the Haas Center for Equity, Gender, and Leadership, the National Science Foundation Graduate Research Fellowship Program, UC Berkeley Burch Center, UC Berkeley Opportunity Lab and Smith Richardson Foundation, and Washington Center for Equitable Growth. Portions of the research in this dissertation were performed at a Federal Statistical Research Data Center under FSRDC Project Numbers 1406 and 2410. All results have been reviewed by the US Census Bureau to ensure that no confidential information is disclosed. All views expressed are my own and not those of the U.S. Census Bureau. I thank Angela Andrus for her patience and help in navigating disclosure review.

Chapter 1

The Role of People vs. Places in Individual Carbon Emissions

1.1 Introduction

Climate change, caused by carbon emissions and other greenhouse gases, is an urgent threat. Global average temperatures have increased by over 1C (1.8F) relative to pre-industrial levels (NASA 2020), and the International Panel on Climate Change has cautioned that even warming of 1.5C could lead to catastrophic consequences, including increased frequency and severity of extreme weather events, degradation of ecosystems, and population displacement. There is substantial spatial heterogeneity in household carbon emissions across cities and neighborhoods in the US, and researchers and policy makers have highlighted this variation as an opportunity for decarbonization, pointing to features of low carbon places, such as density and high-quality public transportation infrastructure, that higher-emissions places could adopt in order to lower household carbon emissions.

However, differences in mean carbon emissions across cities and neighborhoods reflect a combination of local amenities, household characteristics, and taste-based sorting. The relative contributions of these pieces is a central determinant of whether place-based interventions that change urban form would lead to meaningful and rapid reductions in carbon emissions. For instance, if places with large single-family homes and car-oriented transportation infrastructure are high emissions because the people who live there dislike multi-family homes and public transit, then deregulating zoning or building new rail lines would have little impact on household emissions. Conversely, if the lack of denser housing and transit options is a constraint on household choices, rather than a reflection of their preferences, then interventions that change these local public amenities could have the potential to decrease carbon emissions for many households at once.

In this paper, I decompose variation in household carbon emissions into a component driven by household characteristics and a component driven by place effects – i.e., the amount by which the same household's carbon emissions would differ from place to place due to differences in the underlying features of those places. To do so, I construct a longitudinal panel of residential and transportation energy use for over one million individuals from 20 years of restricted-access Decennial Census and American Community Survey (ACS) micro-data. The longitudinal nature of my data makes it possible to link individual survey respondents over time and across places and use a mover design. I use changes to household carbon emissions for over 250,000 movers across roughly 1,000 cities and 60,500 neighborhoods to estimate place effects and their contribution to heterogeneity in carbon emissions.

I begin my analysis by documenting observational patterns of city and neighborhood level variation in household carbon emissions in my sample. Detailed geographic identifiers in the administrative Census data make it possible to directly estimate these values. I estimate that households in cities with high average emissions emit 50% more than households in low emissions cities, and households in neighborhoods with high average emissions emit over two times more than households in low emissions neighborhoods. Accounting for variation driven by observed household characteristics such as household size and income decreases the dispersion across place estimates, but by less than 10%.

The heterogeneity that remains after accounting for observable household characteristics reflects some combination of unobserved household characteristics and causal place effects. Place effects could result from a variety of local amenities and supply-side factors that determine patterns of household energy use. For example, place effects could reflect aspects of urban form such as public transportation, bike and pedestrian infrastructure, highway networks, density, and/or zoning regulations. They could also be driven by natural amenities such as climate. Lastly, they could arise from supply-side factors that determine fuel shares and electricity emissions factors, both of which shift the carbon emitted for a given level of energy use. In the conceptual framework of this paper, I show how place effects relate to the parameters of a consumer energy demand model in which average energy demand, energy demand elasticities, energy prices, and average emissions factors vary across places.

My empirical strategy uses movers in my data to estimate the contributions of place effects and household characteristics to heterogeneity in household carbon emissions. The basic idea behind the mover design is to account for unobserved differences between households by comparing carbon emissions for the same household living in different places. The central assumption in the mover design is that changes to unobserved determinants of household carbon emissions are uncorrelated with mover destinations. A crucial advantage of undertaking this analysis with Census data is that I observe and can control for many time-varying household characteristics that could correlate with both potential emissions and destination choices, including income and children. I proceed in two steps.

First, I use an event study to estimate how much carbon emissions change after households move, as a share of the mean difference between their origin and destination. On average, when households move to a new city, their carbon emissions change by about 90% of the mean difference between their origin and destination cities. There is more sorting of households across neighborhoods, but even at the neighborhood level more than half of mean differences between neighborhoods is explained by place effects: when households move to a new neighborhood, their carbon emissions change by about 60% of the mean differences between their origin and destination neighborhoods. These shares are symmetric for moves to lower or higher emissions places, and they are stable across moves between places with large mean differences or small mean differences in carbon emissions. Under the baseline assumptions of the mover design, the estimates from the event study can be used to calculate the emissions externality of policies that drive existing patterns of household migration, for example policies that restrict housing supply in on average low-emissions cities. To interpret the event study shares as unbiased causal estimates of the effect of moving between any pair of places, it is necessary to impose an additional assumption that there is no systematic selection of certain types of households to certain types of places.

To allow for unrestricted patterns of sorting, next I estimate a non-parametric distribution of household and place effects from a two-way fixed effects model. I use these estimates to do a variance decomposition of overall heterogeneity, including heterogeneity across households living in the same place as well as heterogeneity not explained by the data. The weaker identifying assumptions afforded by the two-way fixed effects decomposition come at the cost of limited mobility bias: naive estimates of variance components are upward biased because some place effects are estimated from only a few movers (Andrews et al. 2008). I account for this upward bias using the heteroskedasticity-robust "leave-out" estimator proposed by Kline, Saggio, and Sølvsten (2020). I find low correlations between unobserved household and place effects, even at the neighborhood level. This suggests that unobserved sorting contributes to differences between places through "segregation" of households, but not in a way that is systematically correlated with unobserved neighborhood attributes. City effects explain roughly 15% of overall heterogeneity, while neighborhood effects explain roughly 25% of overall heterogeneity. Controlling for variation driven by climate, electricity emissions factors, and energy prices in order to isolate the dimensions of place that more likely reflect urban form decreases the place shares to 10% at the city level and 15% at the neighborhood level. While this leaves the majority over overall heterogeneity to other factors, my estimates nevertheless imply the potential for considerable reductions to household carbon emissions from interventions that decrease place effects: I estimate that if a neighborhood goes from having a place effect one standard deviation above the national mean to having a place effect one standard deviation below the national mean, household emissions for residents of that neighborhood would decrease by about 40%.

I characterize low and high emissions places by examining observable characteristics from within the Census data as well as from Walk Score[1], a private company that generates estimates of the walk-ability, transit-ability, and bike-ability of every address in the US.[2] Observable correlates of low emissions places for the most part mirror relationships in the observational data: low emissions places have higher density, more transportation alternatives to cars, and lower shares of single family homes.

Given that the amenities of low emissions places are those often found in urban areas, I consider the impact on carbon emissions of urbanization, generally defined. To minimize

[1] www.walkscore.com

[2] Data provided by Redfin Real Estate (www.redfin.com).

mean squared prediction error, I first adjust my estimates of place effects using a linear empirical Bayes (shrinkage) estimator to down-weight parameters that are noisily estimated. I examine three scenarios. The first two scenarios are variants on the question: what would happen to household carbon emissions if places were more like New York City? The third scenario evaluates how carbon emissions would change if every place across the US were more like the largest city in its metropolitan or micropolitan area. This exercise doesn't take into account general equilibrium effects, but it serves as a useful first order approximation of the effect on household carbon emissions of policies that "expand" cities – e.g. either through policies that make it possible for more people to live in the city (without changing its fundamentals), or through investment and regional development that limits suburban sprawl and increases the number of neighborhoods that have amenities similar to those of the largest nearby cities. I estimate that this type of "place-based climate policy" would result in 13% reductions in average household emissions from residential energy use and commuting.

Taken together, the results in this paper provide new evidence on the role of places in household carbon emissions. I provide direct estimates of nation-wide, neighborhood level variation in household carbon emissions, building on evidence that this variation is substantial (Jones and Kammen 2014; Ummel 2014; Green and Knittel 2020), but finding less heterogeneity than predicted from national data projected onto local place and household characteristics.[3] While related work has highlighted the consequences of spatial heterogeneity in household carbon emissions for allocative efficiency (Glaeser and Kahn 2010; Colas and Morehouse 2021) and distributional impacts and political economy of hypothetical climate policies (Cronin, Fullerton, and Sexton 2019; Sallee 2019; Green and Knittel 2020), the focus of my paper is on examining its causes. I show that place effects can be interpreted as summarizing the parameters of a model of heterogeneous energy demand, where average demand and price elasticities of demand vary across places as a result of different amenities and supply-side factors. Several papers have generated estimates of heterogeneous energy demand parameters, but have necessarily done so in spatially limited contexts (Auffhammer and Rubin 2018; Gillingham 2014; Nowak and Savage 2013; Spiller et al. 2014). Finally, my work builds a bridge to a set of papers that has used observational data paired with modeling techniques to estimate strong relationships between urban form and carbon emissions (e.g. Shammin et al. 2010; Timmons, Zirogiannis, and Lutz 2016; Ribeiro, Rybski, and Kropp 2019; Pomponi et al. 2021; Ko 2013).

My empirical approach builds on and contributes to a recent body of work that uses mover designs to estimate place effects on other individual outcomes, e.g. nutritional choices (Allcott et al. 2019), health outcomes and health care utilization (Eid et al. 2008; Finkelstein, Gentzkow, and Williams 2016, 2020), intergenerational mobility (Chetty and Hendren 2018), and wages (de la Roca and Puga 2018; Card, Rothstein, and Yi 2021). My paper is the first paper to use a mover design to estimate the role of places in determining household

[3]Differences in estimates could also be driven in part by the fact that these papers estimate household carbon footprints from all consumption, including indirect emissions from food and durable goods.

carbon emissions. These mover design papers fit into a broader literature in urban and spatial economics examining the role of places or specific place-based amenities in determining individual outcomes. For example, many papers in the urban and spatial literature have examined the role of density (e.g. see Duranton and Puga 2020, for a review) and transportation infrastructure (Tsivanidis 2019; Allen and Arkolakis 2021) – two amenities closely related to energy use and carbon emissions – in determining productivity, wages, and wage inequality, and several papers have studied the effect of transportation infrastructure on urban form and energy use specifically (e.g. Baum-Snow 2007; Duranton and Turner 2018). A related set of papers has studied the valuation of local public goods (Bayer, Ferreira, and McMillan 2007; Schönholzer 2021), endogenous formation of amenities (e.g. Diamond 2016), and the costs and benefits of using place-based policies to improve aggregate welfare (e.g. Kline and Moretti 2014; Gaubert, Kline, and Yagan 2019; Fagjelbaum and Gaubert 2020). The evidence on the relative importance of sorting vs. places is mixed, and in my paper a large share of variation in carbon emissions is driven by factors other than places; nevertheless, my results highlight considerable potential reductions in household carbon emissions from changes in the distribution of place effects, adding evidence on one channel through which places play a key role in individual outcomes.

The remainder of this paper proceeds as follows. In Section 1.2, I discuss my empirical setting and data, and show several stylized facts about carbon emissions in the US. In Section 1.3, I present my model, and discuss the interpretation of place effects. In Section 1.4, I describe my empirical strategy and identifying assumptions. I present my main findings on the role of unobserved place vs. person heterogeneity in carbon emissions in Section 1.5. I then describe correlates of unobserved heterogeneity in Section 1.6, and predict how aggregate carbon emissions would change under counterfactual distributions of place effects in Section 1.7. Section 1.8 concludes.

1.2 Data and Stylized Facts about Carbon Emissions in the US

I build a 20-year panel of individual and household-level data using the 2000 restricted access Decennial Census long form and the 2001-2019 American Community Survey (ACS). The Census long form is a stratified random sample covering one in six households in the US, and the ACS is a stratified random sample covering 1% of households in the US each year except for 2001-2005 when it covered roughly 0.4% of households (U.S. Department of Commerce 2014). I link individuals across surveys using Protected Identification Keys, which are unique person identifiers assigned by the Census Bureau based on names, addresses, dates of birth, other household members, and social security numbers (when available).[4]

[4]Neither the Decennial Census nor the ACS ask respondents for their social security number. Layne, Wagner, and Rothhaas (2014) use data with social security numbers to show that the error rate in assigning Protected Identification Keys without social security numbers is below 1%. See Wagner and Layne (2014) for

For every individual in my panel, I observe measures of residential and transportation energy use, and a rich set of demographic, household, workplace, and home characteristics, including detailed geographic identifiers. I supplement the Census and ACS with several external data sets in order to convert energy expenditures to energy services and emissions, and to characterize places. In the remainder of this section, I define my geographic units of analysis and outcome variables, provide a high-level overview of the key control and explanatory variables I use, and discuss the construction of my analysis sample. Additional details can be found in app:data.

Geographic Units of Analysis

Throughout the study, I analyze spatial heterogeneity at two levels of geographic granularity which are meant to represent roughly a city or labor market and a neighborhood.

My first geographic unit of analysis is a Core Based Statistical Area (CBSA). CBSAs are designated by the Office of Management and Budget and cover the population of metropolitan and "micropolitan" areas in the US. Each CBSA is a set of contiguous counties with strong commuting ties and at least one urban core area of at least 10,000 people. In addition to formally designated CBSAs, I define residual CBSAs by state from unassigned rural areas. My second geographic unit of analysis is a census tract. Census tracts are county subdivisions that typically cover contiguous areas, have populations of 1,200-8,000 people (4,000 on average), and are delineated with boundaries that follow identifiable physical features. They are designed to be relatively stable, but are split or merged every ten years if populations exceed or fall below the 1,200-8,000 window.[5]

Carbon Emissions

My primary outcome is metric tons of carbon emissions from residential energy and passenger vehicle use, which account for roughly one third of US greenhouse gas emissions.[6]

I estimate carbon emissions from residential energy use from household-reported expenditures on electricity, natural gas, and other home heating fuels in the last year, combined

detailed discussion of the assignment algorithm used by the Census. There is some variation in assignment succes rate across demographic groups – in particular white and higher income individuals are more likely to be successfully assigned a Protected Identification Key – but for all demographic subgroups the success rate is greater than 85%. See Bond et al. (2014) for additional discussion of the variation in assignment rates across population subgroups.

[5] Census geographic definitions vary over time to account for changes in administrative boundaries and populations. To ensure that I don't erroneously identify people who live in places where the designation changed as movers, I use the 2000-2010 census block concordance to assign 2010 geographic definitions to all years in the data.

[6] 75% of US greenhouse gas emissions are from burning fossil fuels. Of these, 20% are from residential energy use (including electricity), and another 20% are from light duty (i.e. passenger) vehicles (U.S. Energy Information Administration 2020).

with external data on local annual retail prices and fuel emissions factors. For electricity, I calculate county-level average prices using data from the U.S. Energy Information Administration (2020) Annual Electric Power Industry Report. This report contains sales, revenues, and total customers for every major utility in the US, by sector and state. It also delineates counties contained in each utility's service territory. I calculate county-level retail electricity prices using customer-weighted average prices (revenue divided by sales) across all utilities with service territories containing the county, and I compute household electricity consumption by dividing reported expenditures by my price estimates. I then assign households to one of 12 National Electric Reliability Council sub-regions using a tract-level crosswalk from the Homeland Security (2021) Infrastructure Foundation-Level Database, and compute emissions using the average annual emissions rates assigned to each sub-region by the U.S. Environmental Protection Agency (2021a) Emissions & Generation Resource Integrated Database. For natural gas and other home heating fuels, I obtain average retail prices at the state level from the Energy Information Administration (2020) State Energy Data System. If a household reports non-zero expenditures on "other home heating fuels", I impute the fuel used from their answer to the question "What was the primary fuel used for home heating?" Finally, I obtain fuel emissions factors from the U.S. Environmental Protection Agency (2018) Emission Factors for Greenhouse Gas Inventories.

I estimate carbon emissions from transportation energy from individually-reported commuting behavior. My outcome captures variation in carbon emissions driven by commute lengths, number of commutes, and mode of transit.[7] I estimate commute distance using the geodesic distance between home and place of work census blocks, and I estimated commute speed from estimated mileage and reported time-length of commute. I estimate gasoline usage using annual national average fuel economy from the U.S. Environmental Protection Agency and Energy (2020), accounting for the fact that in general fuel economy is roughly 30% higher on highways than in cities. Finally, I estimate the number of annual commutes using reported weeks worked last year and hours worked last week, and convert annual gallons of gasoline to carbon emissions using the motor gasoline emissions factor from Energy Information Administration (2020) State Energy Data System. Individuals who commute by rail, subway, streetcar, bus, bike, or walk, and individuals who work from home are assigned zero carbon emissions from commuting.[8] I examine the sensitivity of my results to using the Federal Highway Administration (2019) National Household Travel Survey (NHTS) to predict fuel economy and non-commute miles from household and geographic characteristics available in both the Census and NHTS. This is not my baseline approach, as it infers how much of variation in vehicle fleets and fuel economy observed in the NHTS is driven by individual preferences vs. place-based factors from cross-sectional variation.[9]

[7]Commuting accounts for about 28% of all vehicle-miles travelled, and 39% of person-miles travelled on transit systems (US Department of Transportation 2015), which means I underestimate CO_2 emissions from personal vehicle use for most people in my sample.

[8]This is a generous assumption that favors public transit. It is roughly correct on the intensive margin, but not on the extensive margin unless new investment is required to be zero-emissions.

[9]Place-based factors that contribute to variation in vehicle fleets could include social norms, perceptions

Individual and Household Characteristics

Throughout the analysis, I use demographic and household characteristics to control for variation driven by observable characteristics. My primary demographic and household controls are age, education (completion of a bachelor's degree), sex, race and ethnicity, household income (from salaries and wages, interest, social security, supplemental security, public assistance, retirement, and self employment), household size, and number of kids. I control for age using bins: 18-24, 25-30, 30-34, 35-39, 40-49, 50-64, and 65+. I control flexibly for number of kids using categorical variables for 0, 1, 2, or 3+ kids.

As highlighted in Card, Cardoso, and Kline (2016), the normalization choice for categorical variables does not affect the estimated size of the place variance component or the variance component of the sum of fixed and observable household effects, but it does affect the relative sizes of the place and unobserved household effects, as well as the estimated covariances. Throughout my analysis, I choose the age bin 40-49, no college degree, male, white & non-Hispanic as the omitted categories. Other than "white", these are the categories with the highest within-group variance in outcomes. Thus this normalization should err towards finding a larger unobservable person component relative to place component.

I also observe home-owner status, whether a household lives in a detached single family home, building age, and the number of vehicles in a household. Because these characteristics are intermediate outcomes, which affect CO_2 and likely reflect some combination of household preferences and place characteristics, I do not use these variables as controls throughout my analysis. I do, however, use them in the second half of the paper to explore correlates of unobserved place and household heterogeneity.

Place Characteristics and Amenities

In addition to individual-level data on home characteristics from the full sample in my microdata, I use several external sources of data to characterize amenities at the block, tract, city and regional level. My focus is on amenities that are directly relevant to energy consumption and carbon emissions in the residential and transportation sectors.

To capture variation in climate, I use data on annual heating degree days (HDDs) and cooling degree days (CDDs) at the state-climate division level from National Oceanic and Atmospheric Administration (2020). Degree days are computed as the annual sum of the daily difference between that day's temperature and 65F, and are meant to be a measure of the heating and cooling requirements of a place.

To account for neighborhood-level variation in transportation and leisure amenities, I use data from Walk Score, a private company that generates estimates of the walk-ability, transit-ability, and bike-ability of every address in the US. Walk Score® rankings capture proximity to different commercial amenities such as grocery stores, as well as street characteristics such as block lengths and intersection widths. Bike Score™ indices capture characteristics that

of safety (e.g. if everyone around you is driving a big car it is safer for you to drive a big car; certain types of cars may be able to handle adverse weather better), road widths, ease of parking, etc.

make biking more or less accessible, such as the existence of bike lanes, road connectivity, and hilliness. Transit Score® ratings capture proximity to different types of transit, and the frequency and connectivity of nearby options. For transit, I also observe the number of bus routes and rail routes within a half mile. Finally, I observe a set of amenity scores that measure proximity to parks and leisure and commercial amenities (e.g. grocery stores, restaurants, retail). Other than route counts, each score is an index from 0-100. I assign over 6 million unique Walk Score points reflecting data from early 2020, one to every populated census block in the US, by matching census block centroids to the nearest Walk Score latitude-longitude coordinate.

Analysis Samples

I restrict my analysis to individuals who are at least 18 years of age, who are not identified as the householder's child or grandchild, and who are not missing any of the outcome variables or key explanatory or control variables described above. I also impose several additional restrictions related to energy variables. I exclude from the sample individuals belonging to households whose residential energy costs are included in rent, or whose gas costs are included in their electricity bill, because I don't observe expenditures in those cases. I also exclude individuals in households where residential energy use is top coded or whose commute time is top coded, as the top-coding will obfuscate changes in individual consumption for the highest demand individuals. Lastly, I exclude individuals if the sum of their household residential energy expenditures is zero, if they are in the bottom 1% of non-zero residential energy cost observations, or if they are in the top 1% of commute distance observations as these outliers more likely reflect survey misreporting. My full sample consists of all individuals who meet these restrictions across the 48 continental states and the District of Columbia. This is over 16 million people across 12 million households (Table 1.1, column (1)). I use the full sample to estimate observational geographic and household heterogeneity.

I construct a panel sample by restricting the full sample to individuals for whom I have at least two observations, and who did not indicate in the ACS that they had moved within the last year.[10] This restriction ensures that I am assigning residential energy expenditures to the correct location. The panel sample consists of 1,062,000 people across 889,000 households (Table 1.1, column (2)).

Finally, I impose two additional sample restrictions which are necessary for the implementation of my empirical strategy. First, because residential energy is determined at the household level, and place effects are identified from the variation in outcomes of movers between places, I restrict the sample to only individuals who live with the same set of other full sample individuals across observations.[11] Second, I restrict CBSAs and tracts to the

[10] In the 2000 Decennial Census, the question asked whether respondents had moved within the last five years. Since this is significantly more restrictive, I don't drop these individuals.

[11] This restriction is weaker than requiring individuals live in a consistent household across observations. In particular, if someone lives with different roommates across observations, but their roommates aren't in the full sample because of e.g. missing variables, I do not drop them from the data. Moreover, because

"leave-out connected set" – the network of CBSAs or tracts that remain connected to each other by at least one mover when I drop all the observations in any given household (see app:l1o for an illustration). I do this after dropping tracts with fewer than 10 full sample household observations. The networks are constructed separately at the CBSA and tract level. This means it is possible for a household to be in the CBSA panel but not the tract panel if the tracts they live in are not in the leave-out connected set of tracts. The leave-out restriction drops a negligible share of (residual) CBSAs and roughly 13% of (disproportionately rural) tracts, yielding approximately a 5% sample size reduction (Table 1.1, Columns (4) and (6)).

CBSA movers are households in the CBSA panel that live in different CBSAs across observations (99,500 people in 87,500 households, Column (5)), and similarly, tract movers are households in the tract panel that live in different tracts (within or across CBSAs) across observations (275,000 people in 236,000 households, Column(6)). The CBSA panel, tract panel, CBSA movers, and tract movers make up my four primary analysis samples. The main analysis is implemented at the household level: carbon emissions are given by household residential emissions plus the sum of individual commuting emissions over all individuals in the household, and other household level characteristics are taken as averages over person characteristics. All estimates are weighted using Census sample weights.

Sample Statistics

Table 1.1 shows sample statistics for the full sample, unrestricted panel sample, the two geographically restricted panel samples, and the two mover samples.

A comparison across the samples yields three main take-aways. First, individuals in the panel are on average more likely to be white and have higher income than the full sample (columns (1) and (2)). This reflects known heterogeneity in Protected Identification Key assignment rates within the Census Bureau (Bond et al. 2014). The panel sample is also 6 percentage points less likely to live in a tract designated as urban by the Census, 8 percentage points more likely to live in a detached home, and 1 percentage point more likely to commute by car. Second, further restricting the baseline panel to the CBSA and tract panels (columns (3) and (4)) does not meaningfully change the distribution of demographics, (intermediate) outcomes, or place characteristics. Finally, movers (columns (5) and (6)) tend to be younger, more credentialed, and have higher income (conditional on age) than both stayers and the full sample. Movers also are more likely than stayers to live in urban tracts, less likely than stayers to live in detached homes, and they have higher rates of using electric heating and have lower emissions from residential energy, making them more comparable to the full sample on all of these dimensions.

Overall, about 80% of household carbon emissions in my sample are from residential energy, and about 20% are from commuting. Close to three quarters of the sample live in

people under the age of 18 are dropped from the full sample, this does not drop households that have new children or households in which children move out as they become adults.

a detached, single family home, a vast majority of the sample commutes by car, and on average households live within half a mile of only one bus route and only 0.1 rail routes.

Table 1.2 shows additional statistics for the panel samples. I observe the vast majority of households in my panel sample exactly twice, with on average 8-10 years in between observations. Movers tend to be younger than stayers the first time I observe them, and are much more likely to have had a child or more than 50% change in household income. Households tend to move to places with higher shares of detached single family homes and worse non-car transportation amenities. The majority of moves in my household are from urban to urban tracts, urban to suburban tracts, or suburban to suburban tracts. Finally, consistent with secular trends of mobility in the US, households are generally moving to places that are warmer (16-21% reductions in cooling degree days, and 6-11% reductions in heating degree days). For additional comparisons of movers vs. stayers, estimates of the likelihood of moving given shocks to household income or number of children, and the full set of transition probabilities across urban, suburban, and rural places, see Appendix Tables A.1, A.2, and A.3 respectively.

Table 1.1: Sample Statistics

		Panel Sample			Mover Sample	
	(1) Full	(2) All	(3) CBSA	(4) Tract	(5) CBSA	(6) Tract
Panel A: Demographics						
College	0.25	0.25	0.25	0.25	0.35	0.31
Age	44	46	46	46	43	43
White	0.82	0.89	0.89	0.90	0.89	0.88
Female	0.48	0.47	0.47	0.47	0.45	0.46
Household income	103,700	114,700	114,800	115,500	116,700	116,200
Household kids	1.0	1.0	1.0	1.0	1.0	1.0
Household size	2.8	2.9	2.9	2.9	2.8	2.9
Panel B: Outcomes						
Tons CO_2	18.7	19.9	19.8	19.9	18.8	18.7
Tons CO_2 − *Residential*	15.2	16.3	16.3	16.4	15.2	15.4
Tons CO_2 − *Commute*	3.5	3.5	3.5	3.5	3.5	3.4
Panel C: Intermediate Outcomes						
% Detached home	72.4	80.9	80.9	81.5	73.3	72.9
% Use electricity only	28.8	23.5	23.6	23.5	29.8	27.5
% Commute by car	94.9	96.3	96.3	96.6	95.8	96.2
Commute minutes	25.3	24.9	24.9	24.8	26.2	25.7
Panel D: Place Characteristics						
% Urban	32.2	26.3	26.4	25.6	30.6	32.0
% Suburban	46.3	43.9	44.0	44.4	44.1	47.7
% Rural	21.5	29.8	29.6	30.0	25.3	20.4
Walk Score	26.25	22.64	22.69	21.98	22.08	24.39
Bike Score	35.38	33.10	33.13	32.77	33.58	34.92
Transit Score	9.07	6.92	6.95	6.51	6.83	7.96
N Bus routes	1.58	1.16	1.16	1.07	1.24	1.35
N Rail routes	0.16	0.09	0.09	0.08	0.10	0.10
Annual CDD	1,364	1,224	1,226	1,214	1,361	1,339
Annual HDD	4,369	4,796	4,786	4,828	4,483	4,510
N People	16,200,000	1,062,000	1,040,000	1,006,000	99,500	275,000
N Households	12,190,000	889,000	836,000	807,000	87,500	236,000
CBSAs	1,000	1,000	1,000	1,000	1,000	1,000
Tracts	71,500	69,500	69,500	60,500	53,500	60,500

Note: Column (1) shows statistics for the full sample. Column (2) shows statistics for the panel sample, with no restrictions that individuals be in the same household or live in a connected geography. Columns (3) and (4) show the panel samples restricted to individuals in a consistent household overtime and the CBSA and tract leave-one-out connected sets, respectively. Columns (5) and (6) show statistics for the CBSA and tract mover samples. All means are weighted using census sample weights. Counts and shares are unweighted and rounded according to Census Bureau disclosure rules.

Table 1.2: Panel Statistics

	Panel Sample		Mover Sample	
	CBSA	Tract	CBSA	Tract
Panel A: Sample Characteristics				
% first observed in 2000	9.8	9.9	15.3	13.7
Years between obs	7.8	7.8	10.2	9.8
Panel B: Demographic Characteristics				
Age first observed	42.1	42.0	37.2	37.2
% $\|\Delta$ HH income $\| > 50\%$	28.0	27.9	44.6	40.3
Δ num. kids	-0.12	-0.12	0.07	0.08
% Δ num. kids $\neq 0$	18.6	18.7	29.8	29.8
Panel C: Mover Place Changes				
Δ Walk Score			-6.4	-6.5
Δ Bike Score			-3.9	-3.9
Δ Transit Score			-2.3	-2.7
Δ N Bus Routes			-0.5	-0.5
Δ N Rail Routes			-0.04	-0.04
Δ Tract % detached home			0.05	0.05
% Moves Urban-to-Urban			12.4	17.9
% Moves Urban-to-Suburban			15.3	13.5
% Moves Suburban-to-Suburban			20.6	28.4
%Δ CDD			21.4	16.4
%Δ HDD			-10.7	-6.1
N People	1,040,000	1,006,000	99,500	275,000
N Households	836,000	807,000	87,500	236,000
CBSAs	1,000	1,000	1,000	1,000
Tracts	69,500	60,500	53,500	60,500

Note: Columns (1) and (2) shows panel statistics for the CBSA and tract panel samples. Columns (3) and (4) show statistics panel statistics as well as summary measures of mobility patterns for the CBSA and tract mover samples. All means are weighted using census sample weights. Counts and shares are unweighted and rounded according to Census Bureau disclosure rules.

Observational Heterogeneity

Carbon emissions from residential energy and passenger vehicle use vary immensely across individuals in the full sample. Individuals one standard deviation above the national mean emit 3.4 times as much as individuals one standard deviation below the national mean. Patterns of energy use are strongly correlated with observable characteristics such as income, household size, race & ethnicity, and education. Appendix Figure A.1 shows relationships between carbon emissions and these characteristics. Accounting for observable characteristics decreases heterogeneity across individuals, but significant variation remains: carbon emissions of individuals one standard deviation above the mean are still three times higher than those of individuals one standard deviation below the mean, holding differences in individual observables fixed. A nonparametric regression of household carbon emissions on a set of fixed effects for age, college education, race and ethnicity, household income, household size, and number of children indicates that these characteristics can explain 15% of overall variation in carbon emissions.

There is also substantial spatial variation in carbon emissions across the United States. I estimate unconditional and conditional place means, μ_j, using an ordinary least squares regression of log of individual CO_2 onto place fixed effects, year fixed effects τ_t, and in the conditional regression, individual and household observable characteristics X_{it}:

$$\ln CO_{2it} = \mu_{j(i,t)} + X_{it} + \tau_t + \varepsilon_{it} \tag{1.1}$$

Per capita carbon emissions in CBSAs one standard deviation above the mean are about 54% higher than per capita carbon emissions in CBSAs one standard deviation below the mean, with that difference decreasing only slightly to 50% when accounting for differences in population compositions across areas. At the neighborhood level, individuals in high emissions neighborhoods emit on average 2.2 times what individuals in low emissions neighborhoods do, or 2.1 times more after accounting for differences in observables. Figure 1.1 shows normal distributions reflecting the mean and standard deviation of per capita carbon emissions over individuals, CBSAs, and tracts. The dotted gray line shows the raw distribution, and the solid lines show conditional means. Even after accounting for observational characteristics, significant spatial heterogeneity remains, particularly at the neighborhood level. For the remainder of this analysis, I refer to CBSA and neighborhood means conditional on observable characteristics as "obervational means", following the terminology used by Abaluck et al. (2021).

Figure 1.2 shows how carbon emissions differ across urban, suburban, and rural areas. Suburban and especially rural places have higher emissions than urban places. Controlling for heterogeneity driven by individual observable characteristics decreases the gap between urban and suburban households by almost half, from 2.5 tons to 1.5 tons, and also decreases the gap between urban and rural household by 1 ton, from 6.5 to 5.5.[12]

[12] Relatively small differences between urban and suburban means potentially reflect a pretty broad definition of urban by the Census.

Figure 1.1: Heterogeneity in Individual Carbon Emissions

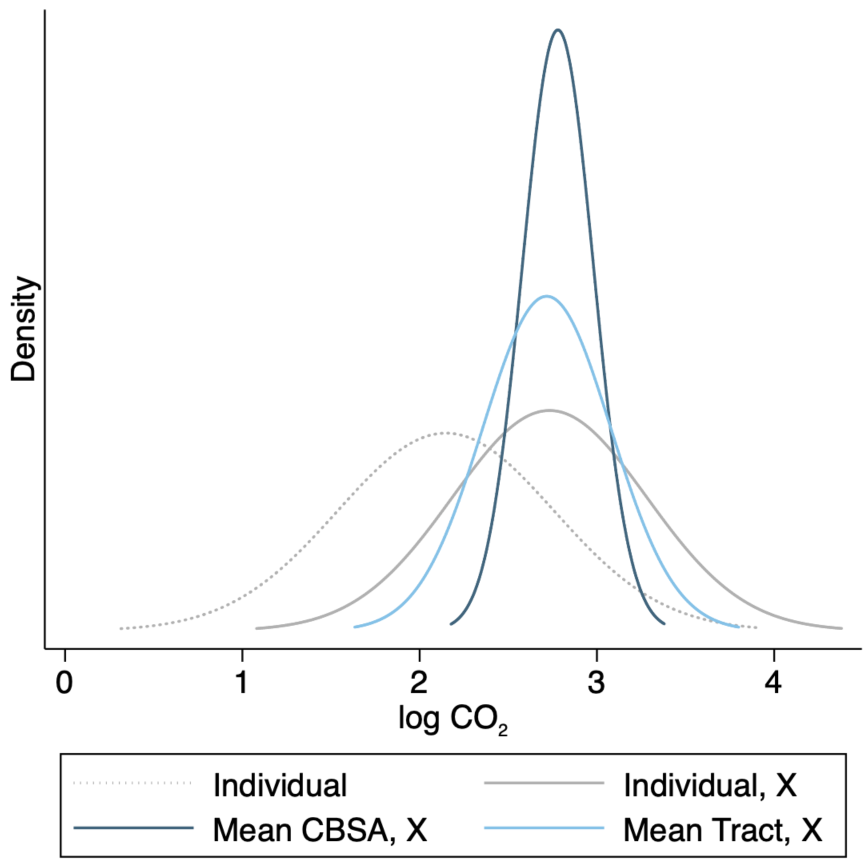

Note: This figure shows gaussian curves with means and normals reflecting the true distributions of per capita emissions, across individuals, CBSAs, and tracts. Raw distributions are not shown in order to facilitate Census disclosure review processes, but have higher kurtosis and are negatively skewed. The dotted gray line (labeled "Individual") corresponds to the distribution of individual CO_2, conditioning only on year FEs. The solid gray line (labelled "Individual, X") corresponds to the distribution of CO_2 over individuals conditional on year FEs and observable characteristics. The dark blue solid line ("Mean CBSA, X") and light blue solid line ("Mean Tract, X") correspond to the distributions of CBSA and tract (respectively) mean per capita CO_2 conditional on observable year FEs and characteristics. Observable characteristics include age, gender, race, education, household size, and number of children.

Figure 1.2: Household Carbon Emissions in Urban, Suburban, and Rural Places

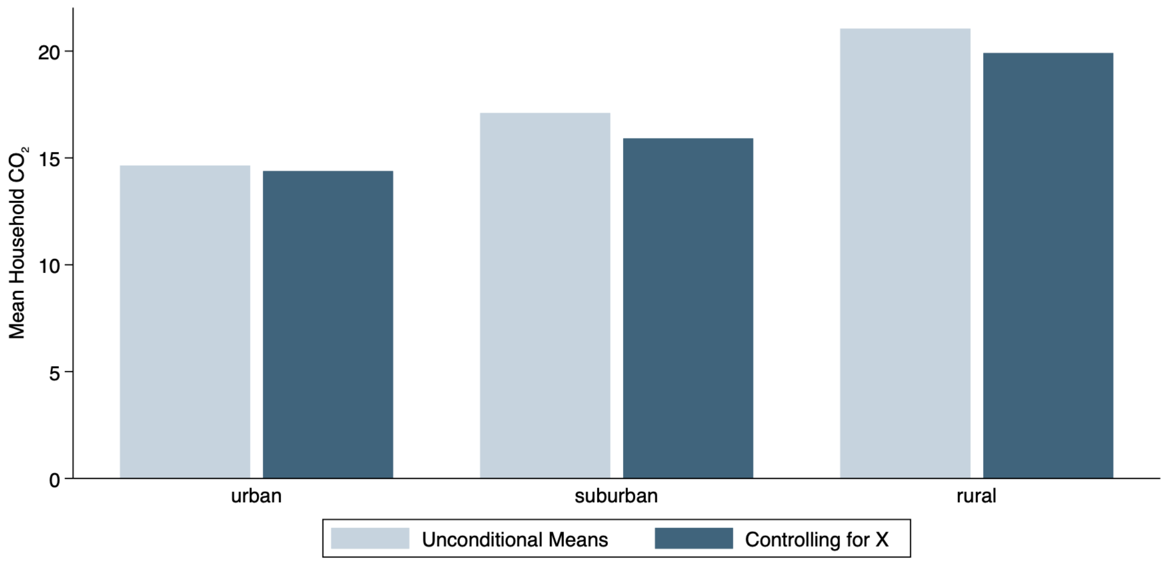

Note: This figure shows mean household CO_2 for urban, suburban, and rural areas. Observable characteristics include fixed effects for age, gender, race, education, household size, and number of children. Places are defined as urban if they are designated as an urban tract by the census. Places are defined as suburban if they are not designated as an urban tract by the census, but are contained within a CBSA. Rural areas are tracts outside of CBSAs. The unconditional regression has an R^2 of 0.08,and the conditional regression has an R^2 of 0.21.

Figures 1.1 and 1.2 highlight substantial spatial heterogeneity in carbon emissions. They also show that while observational heterogeneity in household carbon emissions is partially driven by sorting of households with different characteristics to different types of places, the majority remains unexplained. The fundamental goal of this paper is to understand how much of is remaining unexplained heterogeneity is driven by unobservable individual preferences, and how much is driven by causal place effects, i.e. the amount by which the same household's carbon emissions would differ from place to place, due to the underlying features of that place, holding household characteristics (including unobserved preferences or endowments) fixed.

1.3 Model

Individual i living in place j consumes energy E in the form of four categories of fuels (f). In the residential sector, they can consume electricity (e), natural gas (n), and other heating fuels (o). In the transportation sector they can consume motor gasoline (m).[13] Average demand a_j, price elasticities of demand ρ_j^f, and prices P_j^f are allowed to vary by place. Demand also depends on observable fixed and time varying characteristics (such as age, household size, and income) X_{it}, individual fixed unobserved determinants of demand, α_i, individual time-varying unobserved determinants of demand ε_{it}, and national annual trends τ_t. Thus, individual demand for residential and transportation energy is given by:

$$\ln E_{it} = a_j + \sum_{f \in \mathcal{F}} \rho_j^f \cdot \ln P_j^f + X_{it}\beta + \tau_t + \alpha_i + \varepsilon_{it} \tag{1.2}$$

Place-based differences in average energy demand and in price elasticities of demand could arise from a range of fixed and malleable characteristics of places. These include climate, built environment or urban form – e.g. public transit, pedestrian and bike infrastructure, proximity to highways and parking, density, and proximity to leisure and commercial amenities – and regulatory characteristics – e.g. zoning restrictions, building codes, or clean electricity standards. For example, average demand for heating fuels is higher in cold places, and average demand for motor gasoline is higher in places where households live farther from employers or commercial amenities. Price elasticity of demand for gasoline may be higher in places with better alternative transportation options, and price elasticity of demand for electricity may be higher in places with a larger variety of home sizes and styles.

For each energy type E^f, carbon emissions, CO_2, are a product of the quantity of fuel consumed and the fuel's carbon emissions factor ϕ^f. Emissions factors reflect the physical carbon content of fuel. They are constant over time and place for natural gas, oil, and motor gasoline, but vary for electricity as a result of differences in fuels used for electricity generation. Thus, household consumption of residential and transportation energy results in carbon emissions:

$$CO_{2it} = \phi_{jt}^e \cdot E_{it}^e + \phi^{ng} \cdot E_{it}^{ng} + \phi^o \cdot E_{it}^o + \phi^{mg} \cdot E_{it}^{mg} \tag{1.3}$$

For each fuel f, define s_i^f as household i's fuel share of fuel f. Also define s_j^f as the average fuel share of fuel f in place j, and $\tilde{s}_i^f$ as household i's relative fuel share, $\frac{s_i^f}{s_j^f}$. This allows me to rewrite household carbon emissions as a function of their total energy consumption, fuel emissions factors, and (relative) fuel shares

$$CO_{2it} = \sum_{f \in \mathcal{F}} \left(\tilde{s}_i^f s_j^f \cdot \phi_{jt}^f \right) \cdot E_{it} \tag{1.4}$$

[13]Electric vehicles are a negligible share of driving in my sample time frame. If someone has an electric vehicle, I over-estimate their emissions, because the electricity they use to charge their vehicle is included in residential energy (if they charge at home) but I also assign them gasoline emissions.

Combining this expression with Equation 1.2 yields

$$\ln CO_{2it} = \underbrace{\ln\left(\sum_{f\in\mathcal{F}} s_j^f \cdot \phi_{jt}^f\right) + a_j + \sum_{f\in\mathcal{F}} \rho_j^f \cdot \ln P_j^f}_{\psi_j} + X_{it}\beta + \tau_t + \alpha_i + \ln\left(\sum_{f\in\mathcal{F}} \tilde{s}_i^f \cdot \phi_{jt}^f\right) + \varepsilon_{it}$$
$$= \psi_j + X_{it}\beta + \tau_t + \tilde{\alpha}_i + \tilde{\varepsilon_{it}} \tag{1.5}$$

Equation 1.5 is the baseline estimating equation that I take to the data. The above derivation shows that place effects ψ_j capture place-based variation in average energy demand, variation in price elasticities of energy demand, variation in prices, average electricity emissions factors, and average fuel shares. Household effects α_i capture relative energy demand. In a simplified world in which household's relative fuel shares and electricity emissions factors are constant across places, household effects also capture household relative fuel shares (scaled by fuel emissions factors). However, in reality the model contains an interaction between household and place specific factors – households with the same fuel shares across places will have lower emissions in places with cleaner electricity, and similarly, moving to a place with cleaner electricity will lead to larger reductions among households with higher electricity shares. This means that there is some mis-specification built into the two-way fixed effects model.

It is worth noting here two additional simplifications that the model makes. First, I've allowed price elasticities of demand to vary across places (as a result of amenities that serve as complements or substitutes to energy consumption) but not to vary across households. Allowing for heterogeneity in demand elasticities across households introduces a second interaction term in the error, as household elasticities are interacted with place-specific prices. Both this interaction and the interaction between household fuel shares and place-specific electricity emissions factors motivate treating the errors as heteroskedastic. Second, in the baseline model place effects are fixed, meaning that any time variation (including e.g. in prices) is absorbed in the place effects, which reflect average differences between places over my sample time frame.

Variance Decomposition

Using the two-way fixed effects model derived in Equation 1.5, heterogeneity in household carbon emissions can be decomposed as below (lumping τ_t with X_{it} for brevity):

$$\begin{aligned} Var(y_{ij}) &= Var(\psi_j) + 2\cdot Cov(\alpha_i, \psi_j) + Var(\alpha_i) \\ &\quad + Var(X_{it}\beta) + 2\cdot Cov(\alpha_i, X_{it}\beta) + 2\cdot Cov(\psi_i, X_{it}\beta) + Var(\varepsilon_{it}) \end{aligned}$$

The focus of my analysis is on the first three terms: the variance component of place effects, the variance component of unobserved person effects, and their covariance, which captures the spatial heterogeneity that results from systematic sorting on unobserved preferences. Abusing notation, I re-define y_{it} as the residualized outcome, after having regressed household carbon emissions on time effects and observed household characteristics. For the remainder of this section, I discuss the variance decomposition for this residualized outcome.

$$Var(y_{it}) = Var(\psi_j) + 2 \cdot Cov(\alpha_i, \psi_j) + Var(\alpha_i) + Var(\varepsilon_{it}) \tag{1.6}$$

I follow Song et al. (2019) and further decompose unobserved heterogeneity into a between-place component $Var_j(\bar{y}_j)$, which captures the variation in mean household carbon emissions across places, and a within-place component $Var_i(y_{it} - \bar{y}_j | i \in j)$, which captures the heterogeneity in carbon emissions of households living in the same place:

$$\begin{aligned} Var(y_{it}) &= Var_j(\bar{y}_j) + Var_i(y_{it} - \bar{y}_j | i \in j) \\ &= \underbrace{Var(\psi_j) + 2 \cdot Cov(\bar{\alpha}_j, \psi_j) + Var(\bar{\alpha}_j)}_{\text{Between}} + \underbrace{Var(\alpha_i - \bar{\alpha}_j) + Var(\varepsilon_{ij})}_{\text{Within}} \end{aligned} \tag{1.7}$$

Equation 1.7 highlights that heterogeneity between places reflects variation in place effects, sorting of certain types of households to certain types of places (the covariance term), and what Song et al. (2019) refer to as segregation of households, i.e. the extent to which households of different types segregate across places, whether or not this pattern reflects systematic sorting on place types.[14] In addition to the between-place heterogeneity, overall heterogeneity reflects heterogeneity in household carbon emissions within places, as well as heterogeneity that cannot be explained by the two-way fixed effects model.

1.4 Empirical Strategy

My empirical strategy uses moves across places to estimate place effects and their contribution to spatial heterogeneity in carbon emissions. The intuition behind the mover design is the following: Suppose high-emissions places are high emissions because of a causal place effect, for example because there are no alternatives to commuting other than by car, or because zoning regulations constrain the types of homes households can live in. Then when a household moves from an on average high emissions place to an on average low emissions place, their carbon emissions should decrease because of lower-emissions alternatives now available to them. Conversely, if spatial heterogeneity is driven by strong preferences, then, households that currently live in detached single family homes and commute by car would continue to do so even given alternate options, and moving from on average high to low emissions places should have little effect on household carbon emissions.

I decompose carbon emissions heterogeneity using two versions of the mover design. The first is an event study that characterizes movers' changes in emissions as a share of origin-destination differences in mean carbon emissions. This approach gives a decomposition of heterogeneity between places, although it is not a decomposition of variance terms (and place shares are not constrained to fall between zero and one). It is also unbiased only if there is

[14] $\bar{\alpha}_j \equiv E[\alpha_i | i \in j]$

no systematic sorting of household types to place types. While this assumption is somewhat restrictive, the event study approach is much more efficient than estimating the full two-way fixed effects model. It yields causal estimates under these stronger assumptions on sorting, but is useful for prediction and as descriptive evidence under weaker assumptions. The second approach estimates a non-parametric distribution of household and place effects using the two-way fixed effect model derived in Section 1.3. This approach gives a decomposition of overall heterogeneity, and yields unbiased estimates under weaker assumptions on selection. In the remainder of this section I discuss the mover design identifying assumptions and then each of these variance decompositions in turn.

Identifying Assumptions

Estimates from both versions of the mover design are unbiased under three assumptions: (1) additive separability of place effects, or constant effects (2) non-persistence of outcomes, and (3) exogenous mobility, or conditional orthogonality.[15]

Assumption 1: Additive Separability of place effects, or constant effects.

A core modeling assumption of the two-way fixed effect design is that the outcome – log carbon emissions – is additively separable in person and place effects.

This specification implies that place effects increase and decrease CO_2 proportionally by the same amount for everyone. This is realistic for several potential mechanisms through which place effects could arise. For example, it is natural to model climate as scaling residential heating or cooling needs up or down by the same factor for everyone. To take a few examples centered on urban form: if place effects are driven by density, it may be reasonable to expect places with higher density to decrease the size of homes (and therefore residential energy requirements) or the length of commutes (and therefore transportation energy requirements) by the same factor for low and high baseline users. Similarly, an increase in transportation alternatives to cars might decrease the share of trips taken by car for all households proportionally.

Nevertheless, the two-way fixed effects model imposes a substantial restriction: it does not allow for heterogeneous treatment effects or match effects. I already showed in Section 1.3 that the model is mis-specified, as there is an interaction between individual fuel shares and place-specific emissions factors – a person who prefers electricity to natural gas for cooling and heating will experience a larger decrease in emissions when moving to a place with a clean electricity grid than a person who prefers natural gas to electricity, all else equal. In addition to this, heterogeneous place effects could arise if, for example, place effects are due to a public transit option that only low-income households use but doesn't change high-income household behavior, or if all households use the public transit option but low-income households get rid of their car and eliminate all car trips, while high-income

[15] These are discussed in much more depth in Hull (2018)

households eliminate only a share. Alternatively, heterogeneous place effects might arise if there is not a lot of variation across places in e.g. number of car trips taken or home sizes for low baseline users, but high users respond strongly to places with particularly good or bad amenities.

Given this, the two-way fixed effects model should be treated as an approximation, and the question becomes whether there is selection of certain types of households to certain types of places. If there is an interaction term in the error, as long as there is no selection on this interaction, the mover design will yield unbiased estimates of the average place effects. To rule out selection on heterogeneous effects, I follow Card, Heining, and Kline (2013) and test whether moving from a low emissions place to a high emissions place and moving from high emissions place to a low emissions place are associated with equal and opposite changes in household carbon emissions. Unlike in their setting, in which higher wages are unambiguously good, it is not ex-ante obvious whether we would expect selection to be assortative or disassortative. Nevertheless, testing for symmetry of moves provides evidence on the existence of either type of selection.

To see this, consider differences in potential outcomes across an origin o and destination d, allowing now for there to be an interaction $\eta(\alpha_i \cdot \psi_j)$ between person and place types:

$$E[CO_{2it}(d)] - E[CO_{2it}(o)] = (\psi_d - \psi_o) + \eta(\alpha_i \cdot \psi_d) - \eta(\alpha_i \cdot \psi_o)$$

Because of the multiplicative nature of the interaction term, for a high-type household h and a low type household l moving between the same origin and destination:

$$|\eta(\alpha_h \cdot \psi_d) - \eta(\alpha_h \cdot \psi_o)| > |\eta(\alpha_l \cdot \psi_d) - \eta(\alpha_l \cdot \psi_o)|$$

Thus, regardless of whether the interactive term is positive or negative, and regardless of whether sorting is assortative or disassortative, this type of interaction, paired with selection, would lead to asymmetries between changes in household carbon emissions from moves to higher on-average places vs. lower on-average places.

I group places into four quartiles based on observational averages of carbon emissions, and I estimate household carbon emissions for each origin-destination quartile pair, adjusting for annual trends and controlling for demographic and household characteristics. Results are shown in Figure 1.3. For parismony, the figure shows only moves from the lowest quartile emissions places to all 4 quartiles and vice versa, as well as moves within 1st quartile places and moves within fourth quartile places as bounds in gray. Moves across quartiles lead to equal and opposite changes in household carbon emissions, suggesting that the log-linear model of place effects is a good approximation. The figures also provide evidence of selection, with households that move from the lowest quartile to a different place in the lowest quartile having lower emissions on average than households that move from the lowest quartile to higher quartiles (and vice versa).

Figure 1.3: Changes in household CO_2 when moving across quartiles of Mean CO_2

(a) CBSA (b) Tract

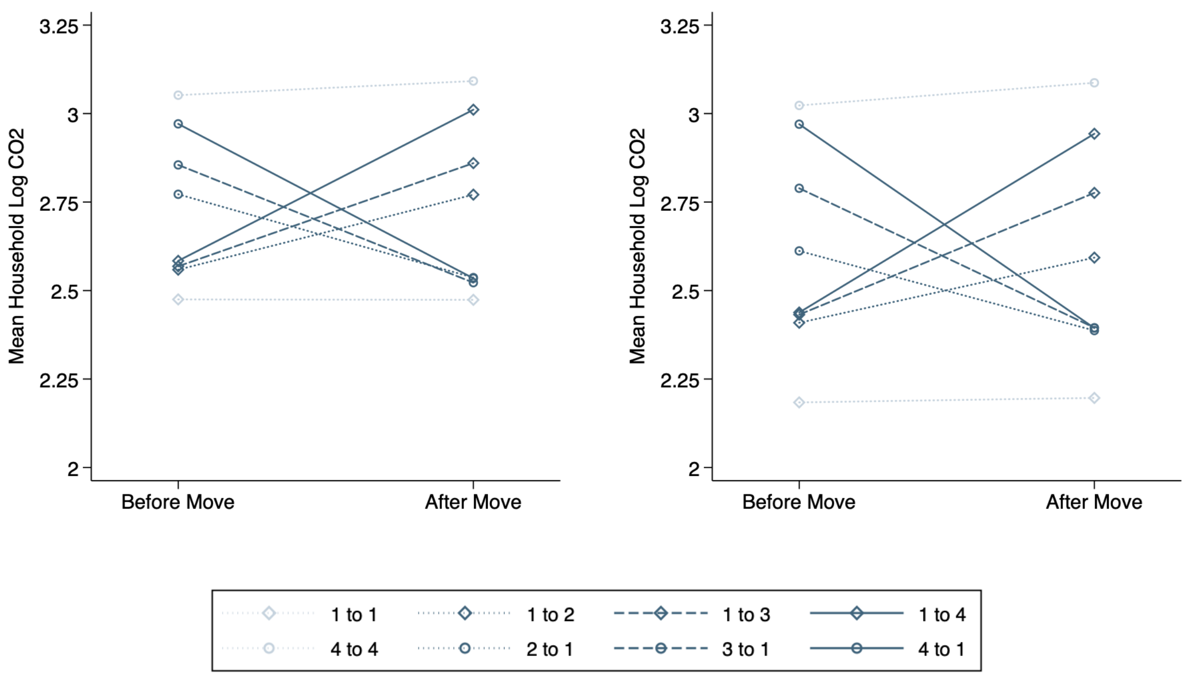

Note: This figure shows average household carbon emissions for movers across places classified into quartiles based on their mean carbon emissions in the full sample. Only the subset of moves to and from the lowest emissions places (quartile 1), as well as moves within the highest emissions places (quartile 4) are shown. Estimates are conditional on year fixed effects and the standard set of household characteristics used throughout this analysis.

Assumption 2: Non-persistent Outcomes.

As highlighted above, relative place effects are identified from pairwise comparisons of household carbon emissions between their origin and destination,

$$E[CO_{2it}(d)|\alpha_i, X_{it}, \tau_t] - E[CO_{2it}(o)|\alpha_i, X_{it}, \tau_t] = \psi_d - \psi_o$$

This expression holds for any two households moving between o and d, regardless of the history of places $\{j\}$ they lived in prior. Note, however, that this doesn't rule out that the place somebody was born may have a persistent effect on their preferences and carbon emissions. Because I include household effects in the model, and only include individuals over the age of 18 in the sample, any persistent effect of place of birth and upbringing on carbon emissions will be captured by the household fixed effect.

Assumption 3: Exogenous Mobility, or conditional orthogonality.

Finally, estimates of place effects from the mover design are unbiased only if moves are strictly exogenous; in other words if shocks to unobserved determinants of CO_2 are conditionally uncorrelated with destination choices.

$$E[\epsilon_{it}|\alpha_i, \psi_{j(i,t)}, X_{it}, \tau_t] = 0 \tag{1.8}$$

Note that the two-way fixed effects model allows for unrestricted sorting of households on fixed or time-varying observable characteristics and on fixed unobservable characteristics. For instance, entering middle age and having children is associated with an increase in energy consumption generally (Appendix Figure A.1), and also significantly increases the probability of moving to a suburb (Appendix Table A.2); however this endogeneity does not bias my estimates, because I observe age, household size, and number of children. Similarly, estimates of place effects are robust to household changes in energy demand and simultaneous moves to a new neighborhood that might arise from an increase (or decrease) in income, because I observe household income. Finally, if people have heterogeneous, but fixed, preferences for neighborhood amenities – e.g. if people have a particular distaste for public transit, a strong preference for large homes, or a particular love for walking or biking – and their choice of what neighborhood to live in reflects those preferences, estimates of place effects are unbiased because these unobserved but fixed determinants of CO_2 are captured by individual fixed effects. The ability to account for these time invariant unobserved preferences is a crucial benefit of the mover estimation strategy.

Thus, the main threat to identification is the possibility that moves correspond to *changes* in unobserved preferences – either single idiosyncratic shocks or evolving. A standard approach for ruling out this source of endogeneity is to test for parallel trends between movers and stayers prior to the move. A limitation of my data is that I observe the majority of my sample only twice, which makes it impossible to test for parallel trends. In Section 1.5, I show that the effect of moving appears to be stable across duration between observations, meaning estimates from households observed less than 5 years apart are similar to estimates from households observed more than 15 years apart. If moves were endogenous to preferences evolving, or "drifting" over time, you would expect that my heterogeneous parameter estimates would evolve in a parallel way. While somewhat comforting, this does not rule out the possibility of moves corresponding to a single idiosyncratic shock to preferences. To rule this out, I use data from the Panel Study of Income Dynamics (PSID), over the same sample period, and assess whether movers in the PSID exhibit any changes to energy expenditures prior to their move. While I do not know where households move from or to, I find that energy expenditures are flat leading up to a move and increase afterward, consistent with life-cycle trends presented in Table 1.2 of people moving to places with larger homes and fewer non-car transportation amenities, and with the secular trend over my sample frame of people moving to places with higher cooling needs. This result is shown Appendix Figure A.4.

Event Study

The first decomposition I estimate is based on an event study, as in e.g. Finkelstein, Gentzkow, and Williams (2016). Consider a household i that moves from origin o to destination d. Household i's expected change in carbon emissions is given by:

$$E[\ln CO_{2it}(d) - \ln CO_{2it}(o)|\alpha_i, X_{it}, \tau_t] = \psi_d - \psi_o$$

I re-express the change in place effects in terms of the share of differences between observational means, $\bar{y}_d - \bar{y}_o$, attributable to differences between place effects:

$$\begin{aligned} \psi_d - \psi_o &= \frac{\psi_d - \psi_o}{\bar{y}_d - \bar{y}_o} \cdot (\bar{y}_d - \bar{y}_o) \\ &\equiv \theta_{o,d} \cdot (\bar{y}_d - \bar{y}_o) \end{aligned}$$

Plugging this expression into the two-way fixed effect model yields an event study, which I use to estimate the share of differences *between* places attributable to place effects, θ:

$$\begin{aligned} \ln CO_{2it} &= \alpha_i + \psi_{j(i,t)} + \tau_t + X_{it}\beta + \varepsilon_{it} \\ &= \alpha_i + \psi_o + 1[moved] \cdot (\psi_d - \psi_o) + \tau_t + X_{it}\beta + \varepsilon_{it} \\ &= \tilde{\alpha}_i + 1[moved] \cdot \theta \cdot (\bar{y}_d - \bar{y}_o) + \tau_t + X_{it}\beta + \varepsilon_{it} \end{aligned} \tag{1.9}$$

Relative to the unrestricted two-way fixed effects model, the event study approach vastly reduces the dimensionality of the estimation problem, as now the place share of heterogeneity is characterized by a single parameter θ as opposed to the full distribution of J place effects. However, this efficiency comes at the cost of an additional assumption, that heterogeneity in θ cannot be correlated with other parameters in the model. In other words, because place types are inferred from observational means, the event study limits selection of households to places so that there is no systematic sorting of e.g. high type households to high type places. In Equation 1.7, this amounts to requiring the covariance term to be equal to zero.[16]

Variance Decomposition

The second decomposition I estimate is the one described in Equation 1.6 (and shown again below), which is based on estimation of the full two-way fixed effects model, allowing for unrestricted correlations between place effects and household characteristics.

$$Var(y_{it}) = Var(\psi_j) + 2 \cdot Cov(\alpha_i, \psi_j) + Var(\alpha_i) + Var(\varepsilon_{it})$$

In contrast to the event study decomposition, the two-way fixed effects decomposition allows unrestricted sorting of households across places, and the variance share attributable

[16] The household share term, which captures segregation of household types that is uncorrelated with unobserved place-based heterogeneity, is given by $\frac{E_d[\alpha_i + X_{it}\beta|\tau] - E_o[\alpha_i + X_{it}\beta|\tau]}{\bar{y}_d - \bar{y}_o}$.

to household heterogeneity reflects not only the between component (i.e. how households differ across places on average) but also the within component (i.e. how much variation in carbon emissions there is across observably similar households within the same place).

A well-documented challenge to estimating variance components in two-way fixed effect models is limited mobility bias (Andrews et al. 2008): estimates of place effects are noisy because they are estimated from a small sample of movers to and from each place. This creates an upward bias in the plug-in variance estimate relative to the true variance of place effects, even if estimates of place effects themselves are unbiased. To address this, I estimate variance components using the heteroskedasticity-unbiased leave-out estimator proposed by Kline, Saggio, and Sølvsten (2020), henceforth KSS. The KSS estimator uses a leave-out estimate of standard errors to correct estimates of the variance components for sampling variability.

I implement the leave-out estimator at the household level, leaving out all observations corresponding to a household match, not just an individual match. In the mover sample, the KSS estimator is robust to unrestricted heteroskedasticity and serial correlation within each match. Because it is not possible to leave out matches for stayers without dropping all their observations, if there is serial correlation in the error term, KSS estimates of the person variance component in the panel sample are an upper bound on the true value. To reduce the computational burden of the estimator, I use the Johnson-Lindenstrauss approximation (JLA) algorithm introduced by KSS to estimate the statistical leverages of each match, i.e. the amount by which estimates change when leaving out the match. KSS show that using JLA introduces an approximation error of roughly 10^{-4} relative to estimating statistical leverages directly. See Appendix A.3 for some additional detail on the implementation of the empirical approach, and KSS for a complete discussion of the leave-out estimator and JLA algorithm.

1.5 Results

This section presents the core results of my paper: estimates of the share of spatial heterogeneity attributable to place effects. I begin this section by showing results from the event study specification, which – even if the stronger assumptions on selection are violated – serve as additional descriptive evidence and can be used to predict how household carbon emissions will change for movers under existing patterns of mobility. I then present results from the unrestricted two-way fixed effect model. I conclude the section with a discussion on interpreting the two versions of the analysis, as well as several sensitivity analyses.

Event Study

This section presents estimates from the event study derived in Section 1.4

$$\ln CO_{2it} = \tilde{\alpha}_i + 1[moved] \cdot \theta \cdot (\bar{y}_{d_{-i}} - \bar{y}_{o_{-i}}) + \tau_t + X_{it}\beta + \varepsilon_{it} \tag{1.10}$$

Table 1.3: Share of Spatial Variation in Mean CO_2 Attributable to Place Effects

	Panel			Movers	
	(1)	(2)	(3)	(4)	(5)
CBSA					
$\bar{y}_d - \bar{y}_o$	0.90***	0.90***	0.90***	0.90***	0.90***
	(0.007)	(0.007)	(0.013)	(0.008)	(0.017)
N	1,715,000	1,715,000	664,000	179,000	44,000
R^2 (adj.)	0.72	0.75	0.77	0.69	0.68
Tract					
$\bar{y}_d - \bar{y}_o$	0.77***	0.65***	0.61***	0.62***	0.60***
	(0.003)	(0.003)	(0.006)	(0.004)	(0.008)
N	1,656,000	1,656,000	640,000	483,000	127,000
R^2 (adj.)	0.73	0.75	0.77	0.72	0.72
Controls		X	X	X	X
No big life events			X		X

Note: This table reports event study estimates of place shares of spatial heterogeneity in household CO_2. Columns (1) and (6) report estimates from the panel sample with no controls apart from year fixed effects. Columns (2) and (7) add controls for the standard set of household characteristics. Columns (3) and (8) restrict the estimation sample to movers only, to allow movers to differ systematically from stayers. Columns (4)-(5) and (9)-(10) use the subset of the panel and mover samples that did not have a change to the number of kids in their household or a larger than 50% increase or decrease to income. All estimates are weighted using Census sample weights.

$\bar{y}_{j-i}$ are sample means calculated from the full sample, leaving out the household observation.[17]

Table 1.3 presents estimates of the place share, $\hat{\theta}$, from the event study. Column (1) shows estimates with no controls other than year fixed effects. Adding controls (column (2)) does not change the CBSA estimate, but decreases the share of heterogeneity attributable

[17] To the extent that there is sampling variability in the distribution of observational means, my estimate of the relationship between origin-destination mean changes and individual changes in log CO_2 may be biased. In practice, using a linear empirical Bayes estimator to adjust observational means for sampling variability as in, e.g. Abaluck et al. (2021) or Finkelstein, Gentzkow, and Williams (2020) does not materially change the results.

by tract effects by 12 percentage points. This is consistent with evidence that household sorting across neighborhoods plays an important role in neighborhood-level variation in carbon emissions, while moves across CBSAs are more likely to be driven by other factors such as new job opportunities or proximity to friends or family.

In column (2), the effect of changes to observable characteristics – e.g. having kids – on carbon emissions is estimated from both stayers and movers. However, it may be the case that households who move after having children do so in part because having children changed their preferences more than having children changed the preferences of households who ended up staying where they were. If the decision of whether to move or not is driven (at least in part) by such heterogeneous preference shocks, then any differential effect of the preference shock to movers would be incorrectly attributed to place effects, biasing my estimates. To address this, I re-estimate the event study with movers only (column (4)), which allows movers to differ systematically from stayers. Once again, this does not change the estimates in the CBSA specification, but further decreases the the share of spatial heterogeneity attributable to tract effects by 3 percentage points.

Even after accounting for differences between movers and stayers, mover destinations might still reflect heterogeneous preference shocks. For example, the decision of one household to move from a city to a suburb after having children could reflect a different shock to preferences than that of a household that moves from one neighborhood within a city to another after having children. While I cannot rule this out entirely, I explore the extent to which such selection patterns might bias my results by re-estimating the event study on a sample restricted to only households who did not experience a large shock to observable characteristics. Namely, I restrict the sample to only households who never had a change in the number of kids living in their home, and never had more than a 50% increase or decrease in household income between observations. If heterogeneity in unobserved time-varying preferences leads households to choose different types of neighborhoods, then estimating the event study using different sets of households with *observably* different preference shocks should lead to different results. Reassuringly, estimates from this approach (in columns (3) and (5)) are similar to estimates using the mover-only sample.

I explore two additional dimensions of heterogeneity: the magnitude and sign of changes in mean carbon emissions between movers' origin and destination, and the duration between mover observations. Figure 1.4 shows changes in mover households' carbon emissions by decile of origin-destination differences in observational means, with all demographic and household controls. The gray line is 45-degrees, and would correspond to place effects accounting for 100% of variation in mean differences across places. The slope of the solid line corresponds to the estimate of θ from the full mover samples (column (4) of Table 1.3). I find that both for moves across CBSAs and moves across neighborhoods, estimates of the share of heterogeneity attributable to place effects are symmetric and linear across origin-destination mean changes. The fact that the place share estimate is stable across deciles of move types is suggestive evidence that my results aren't being driven by only a subset of movers or mover destinations. It also provides additional validation for the log-linear model specification, serving as kind of an extension of the symmetry check presented in Figure 1.3.

Figure 1.4: Place Share of Spatial Variation in Mean CO_2, by Move Type

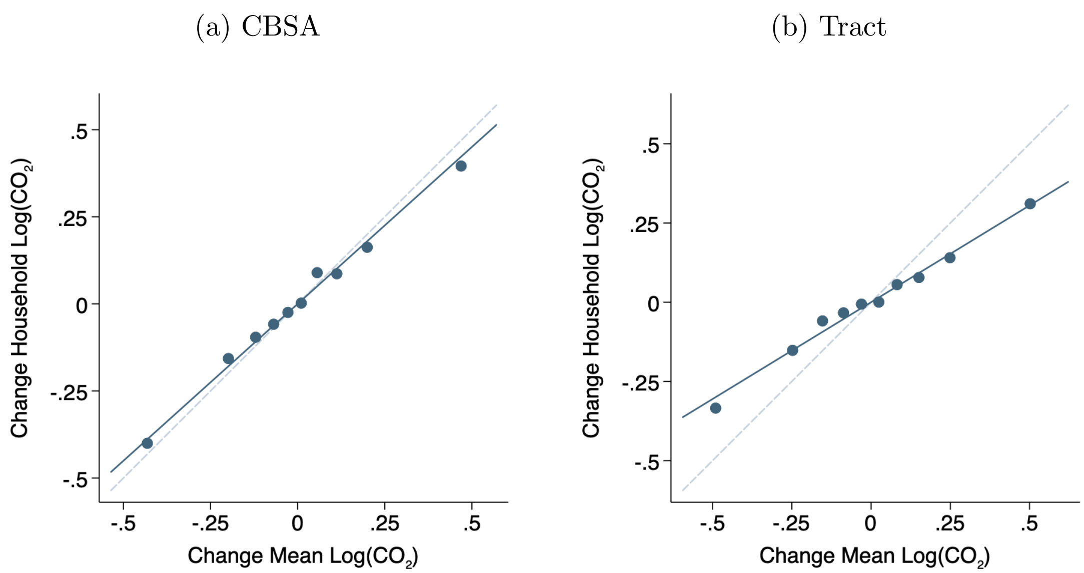

Note: This figure shows event study estimates of the share of spatial variation in mean carbon emissions that can be explained by place effects, by size of origin-destination differences in mean household carbon emissions for movers. Movers are split into ten deciles, according to the size of the gap in mean carbon emissions across their origin and destination. All estimates are from models that control for observable household characteristics and year fixed effects. The solid lines show the regression estimates from the pooled model,and the dotted gray line denotes 45, i.e. the scenario in which moving to on average higher or lower emissions places leads to a 1-for-1 increase in own carbon emissions. All estimates are weighted using Census sample weights.

Figure 1.5 shows tract-level estimates by duration between observations, with all demographic and household controls. This exercise allows me to evaluate two possible sources of bias in the model. First, it provides evidence on the extent to which my place effect estimates may be biased by life-cycle patterns of energy and CO_2 demand. If my estimates are unintentionally capturing changes to preferences over different stages of life (ages) of household members, I would expect estimates to be larger for households I observe 15 years apart than those I observe 5 years apart. Second, if households select where to move based on preferences that drift over time in a way that *isn't* captured by age or other life-cycle effects, my estimates of place effects would capture a combination of true causal effects and selection, and the longer the gap between observations, the larger I would expect the selection component to be. This would result in estimates of place effects that are increasing or decreasing with the duration between moves, depending on the direction of selection.

I show estimates for the full panel sample (light gray), and the restricted panel of only households with no significant changes to income or household composition (dark blue). The pooled estimate is contained in the 95% confidence interval of all but two-duration specific estimates, and coefficients appear to be mostly stable – the estimate from households

Figure 1.5: Place Share of Spatial Variation in Mean CO_2, over Time

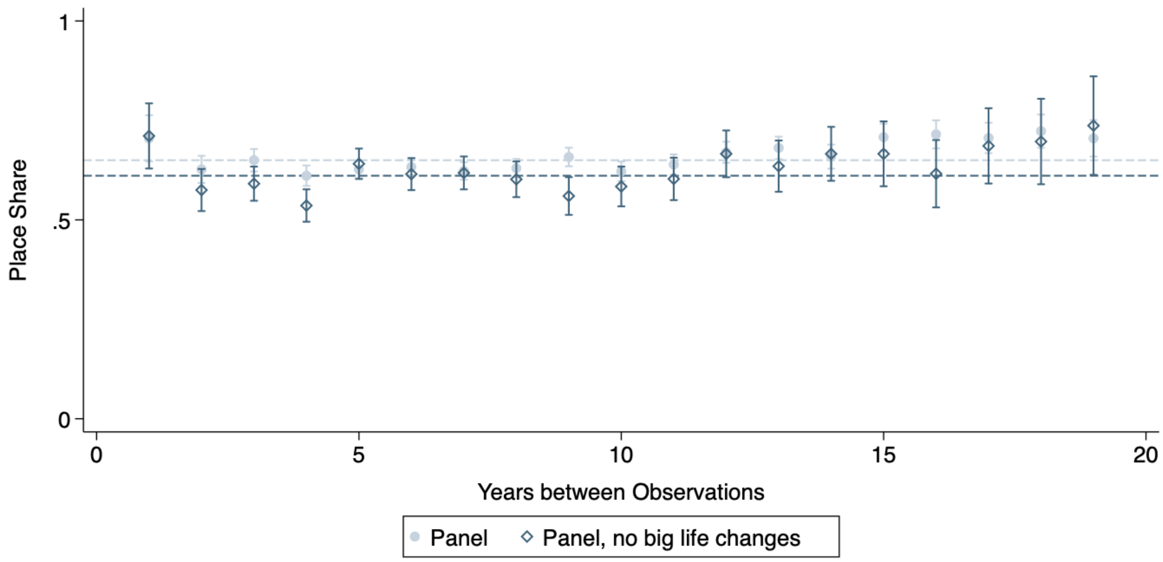

Note: This figure shows event study estimates of the share of spatial variation in mean carbon emissions that can be explained by place effects, by duration between mover observations. In other words, each coefficient is the estimate for place effects generated from the sub-sample of households that I observe X years apart. Coefficients plotted in light gray are estimated from the model using the full panel of stayers and movers. Coefficients plotted in the dark blue are estimated from the model using the sub-sample of stayers and movers with no changes in the number of children and less than 50% change in household income between observations. All estimates are weighted using Census sample weights.

observed one year apart is higher than the pooled estimate, and there is a slight but not statistically significant upward trend for estimates from households observed 16 years apart or longer. Given that these are also the duration bins with the fewest observations, I do not interpret this as strong evidence of place effect estimates being biased by drifting preferences. Analogous CBSA estimates are shown in Appendix Figure A.5, and exhibit a similar pattern.

One additional result that comes out of the analysis of the duration-specific event study is that household carbon emissions appear to change instantaneously. This suggests that place effects are driven by attributes that directly impact carbon emissions demand, rather than characteristics such as peer effects or habit formation, which I would expect to lead to gradual changes in behavior over time.[18]

[18] I do not observe how long ago households moved, but the expected value of how long ago someone moved is increasing in the duration between observations.

Variance Decomposition

Estimates from the KSS bias-corrected variance decomposition are shown in tab:kss. The table presents overall variance for the sample, the share of variance attributable to each of the unobserved heterogeneity components, and the KSS-adjusted coefficient of determination (R^2). Panel A presents estimates from the entire panel of movers and stayers, while Panel B presents estimates from the mover only sample. Estimates from a variance decomposition with no bias correction can be found in app:akm.

In the baseline analysis with year fixed effects and the standard vector of household controls, I estimate that CBSA effects account for 16-19% of overall heterogeneity, and tract effects account for 24-26% of overall heterogeneity. To estimate the share of CO_2 heterogeneity more likely attributable to just local built environment and public amenities, I re-estimate variance components, partialing out measures of climate and electric grid intensity, and then additionally, prices. Specifically, in columns (2) and (6) I control for heating degree days, cooling degree days, and electricity emissions factors (all in logs), and in columns (3) and (7) I also add controls for lagged fuel shares interacted with national retail prices.

I find that controlling for climate and electric grid intensity decreases the place share of spatial heterogeneity by roughly 10 percentage points, to 10-16% of overall heterogeneity. This decrease is consistent with a well-understood, robust relationship between climate and energy use (e.g. Goldstein, Gounaridis, and Newell 2020) and the mechanical relationship between electricity emissions factors and CO_2. However, remaining neighborhood attributes explain a larger share of variation than climate and grid intensity, underscoring the importance of residual place characteristics such as urban form. Accounting for cross-sectional fuel price variation does not further change the results.

Finally, it is possible that place effects evolve over time in ways that differ from national average trends in carbon emissions. For instance, the governments in certain states or cities may be particularly concerned about climate change and enact regulations or make place-based investments aimed at reducing emissions for their residents. In addition to transit and zoning examples I've highlighted throughout the paper, such policies could include regulatory efforts more directly targeting energy sources, such as renewable portfolio standards, state or regional cap and trade programs, or laws banning gas stoves in new homes. More generally, changes to place effects could arise from local or regional planning initiatives motivated by factors completely unrelated to decision-makers' climate objectives. For instance, the Phoenix metropolitan area – one of the fastest growing metropolitan areas in the US – has grown by nearly 1.6 million residents since 2000. This period of growth has been accompanied by a mix of suburban expansion, urban development, the opening of a new light rail system, and several high way expansions.[19]

To account for time variation in place effects that differs from national trends, I also estimate time-varying place effects ψ_{jt} at the CBSA level.[20] I follow Lachowska et al. (2020)

[19]See e.g. *The Phoenix Metro Area* (2020).

[20]For parsimony in the census disclosure review process, I am not disclosing time-varying tract effects as these involve a different leave-out sample, whereas the CBSA leave-out-connected set is the same in the

Table 1.4: Unobserved Heterogeneity in CO2 – Variance Decomposition

	CBSA				Tract		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A: Panel Sample							
$Var(\log CO_{2ij})$	0.29	0.29	0.29	0.29	0.28	0.28	0.28
Share $Var(\psi_j)$	0.188	0.095	0.099	0.205	0.257	0.152	0.154
Share $Var(\alpha_i)$	0.505	0.502	0.503	0.306	0.377	0.374	0.373
Share $2 \cdot Cov(\alpha_i, \psi_j)$	-0.001	0.004	0.004	-0.001	-0.006	0.009	0.010
R^2	0.69	0.61	0.61	0.51	0.62	0.54	0.55
B: Mover Sample							
$Var(\log CO_{2ij})$	0.32	0.32	0.32		0.31	0.31	0.31
Share $Var(\psi_j)$	0.163	0.098	0.102		0.239	0.156	0.163
Share $Var(\alpha_i)$	0.112	0.091	0.0966		0.102	0.103	0.104
Share $2 \cdot Cov(\alpha_i, \psi_j)$	0.010	0.013	0.013		0.010	0.016	0.016
R^2	0.30	0.22	0.22		0.36	0.29	0.30
Amenities		X	X			X	X
Prices			X				X
TV-FE				X			

Note: This table reports results from the heteroskedasticity-robust KSS estimation of variance components. All specifications include demographic and household controls as well as time fixed effects. Columns (1) and (5) report the baseline variance decompositions at the CBSA and tract levels. Columns (2) and (5) add controls for local mean heating degree days, cooling degree days, and electricity emissions factors (all in logs). Columns (3) and (6) additional control for a price index, constructed from lagged fuel shares interacted with national retail prices. Finally, column (4) computes time-varying CBSA place effects using 5-year windows (2000-2004, 2005-2009, 2010-2014, and 2015-2019), using stayer observations across time windows to identify time variation in place effects, while movers, as before, identify cross-sectional variation.

and estimate time-varying fixed effects using stayers to identify variation across time within place. To maintain connectivity in my set of places, and because for the most part places evolve slowly, I allow these to vary at 5-year intervals. Thus, there's a different time-varying place effect for each period 2000-2004, 2005-2009, 2010-2014, and 2015-2019. Results are shown in column (9) – allowing CBSA effects to evolve slightly increases their variance share to 21%.

The interpretation of my results is somewhat complicated by the fact that the contribution of unobserved household characteristics to overall heterogeneity is highly sensitive to whether the model is estimated on movers and stayers or movers only. In the panel sample, unobserved household heterogeneity accounts for 51% of overall heterogeneity when measuring place at the CBSA level, and 38% when measuring place at the neighborhood level. This share is stable to partialing out exogenous amenities and prices, but the CBSA estimate decreases to 31% when allowing CBSA effects to change over time. Using the mover-only sample substantially decreases the unobserved household contribution across specifications, to 11% in the CBSA specification and 10% in the mover specification. Similarly, estimates of the covariance between unobserved place and household characteristics are also sensitive to the sample choice. In both panel specifications, the covariance is slightly negative but effectively zero – the larger of the two correlation coefficients is -0.02. In contrast, in the mover sample I find a positive (though still small) correlation coefficient of .07 at the CBSA level and .06 at the tract level.

There are two reasons we might expect estimates from the panel and mover samples to differ. The first is from fundamental differences across stayers and movers, and the second is that KSS cannot correct bias induced by serial correlation in the error term among stayers. To try to shed light on the relative importance of these pieces, it is useful to compare estimates from the KSS decomposition to estimates from the naive, uncorrected (AKM) decomposition in app:akm. If results are driven by differences between the panel and mover sample, such differences should also be evident in the AKM estimates, even though we expect estimates of both variance components in AKM to be higher than in KSS because of limited mobility biased. In contrast, if results are driven by serial correlation in stayers' error term, then we would expect the relative contributions of the unobserved heterogeneity components in the AKM estimation to be fairly similar, with differences being introduced only in the KSS correction. The AKM estimates suggest that the relative place and person shares are almost identical across the panel and mover samples in the CBSA analysis. In the tract analysis, the relative size of the household variance component does drop several percentage points, but not nearly as dramatically as it does in the KSS analysis, suggesting that the estimated household variance components in the panel sample of KSS are an upper bound on the true value, with the upward bias driven by serial correlation in the stayer error term.

Serial correlation in the error term could arise as a result of several sources of measurement error in my outcome variable. While an advantage of using the Census for this analysis is that it allows me to observe many household characteristics that are unobservable in standard

time-varying case as it was in the baseline.

administrative datasets on energy use, thereby making it possible to control for changes to household characteristics that are correlated with both changes to energy demand and move propensity and destinations and decrease potential bias from unobserved preference shocks, a disadvantage is that the survey nature of the data means that my outcomes are constructed from a combination of survey responses and local external data. In particular, I use local average prices and local average emissions factors to convert reported energy expenditures and commute times into carbon emissions. Both of these could introduce serial correlation into my estimates of stayer outcomes. Additional detailed discussion of measurement error within the residential and transportation sectors, as well as implications for interpreting results, can be found in app:data$_res$.

How do these results inform the interpretation of the event study decomposition? Here, there are also two things to note. First, recall that the event study estimates are unbiased only if heterogeneity in the share parameter is uncorrelated with observed and unobserved household characteristics. I've shown evidence that my event study estimates are stable several observable dimensions of heterogeneity in the data. I also showed in the KSS decomposition that the covariance between unobserved components of heterogeneity are very close to zero – the largest correlation coefficient across the four baseline estimates (CBSA vs. tract & panel vs. mover) is 0.06. Together, this suggests that bias from this assumption on selection should be minimal.

Even with no bias from selection, the event study yields estimates of shares of mean differences between places attributable to place effects, while the KSS estimates yield a variance decomposition of *overall* variation, and this can lead to meaningful discrepancies in magnitudes. To see this, imagine two places, one ψ_{low} and one with ψ_{high}, and identical populations across the two places. If there is high variation in carbon emissions across populations and a small difference between ψ_{low} and ψ_{high}, the event study would yield a share coefficient of 1 (since populations are identical across places, all between differences are driven by place effects), but the KSS decomposition would yield a place variance component of close to zero (because of a very large within component to the variance). In practice, this is very close to what happens at the CBSA level – the vast majority (90%) of differences between CBSAs can be attributable to variation in place effects and not household attributes, but there is much more variation in household carbon emissions within CBSAs than there is across, leading to a variance component of 16-19% in the KSS estimation, about half of which is attributable to climate and electric grid intensity. At the neighborhood level, household sorting contributes more to variation between places, dropping event study estimates of the place share to 62%; accounting for variation within places (using the panel sample with KSS) further decreases the place variance share to 26% of overall heterogeneity (or 42% of heterogeneity explained by the model, calculated from re-scaling by the R^2).

Specification Tests and Robustness

As an additional specification test, in Appendix A.6, I show binned scatter plots similar to the one presented for event study results (Figure 1.4), but now with deciles of changes in estimated place effects, rather than observational means, on the x axis. I plot these against two sets of changes in household mean outcomes: changes for the full mover sample, and changes in the sample restricted to only households with no big life changes. In a correctly specified model, changes in place effects should lead 1-to-1 to changes in household carbon emissions, though attenuation bias from noisily estimated place effects should decrease the slope. Crucially, I find no difference across the two samples, suggesting (as in the event study analysis) that selection on heterogeneous preference shocks isn't a first order threat to identification in my analysis.

To evaluate the sensitivity of my results to my outcome definitions, Appendix A.4 shows estimates from a KSS decomposition using residential energy only, and using total energy but imputing carbon emissions from transportation energy using the National Highway Transportation Survey (NHTS). One version of the NHTS imputation uses a LASSO regression to predict heterogeneous fuel economy from household, geographic, and commute characteristics that are common to both the Census and NHTS surveys, and then uses predicted relationships to estimate carbon emissions from commuting accounting for variation in fuel economy. The second version additionally predicts total miles travelled, and uses both heterogeneous fuel economy and heterogeneous relationships between commuting and total miles travelled to estimate carbon emissions from car travel generally. I find that restricting the analysis to residential energy increases the overall variance by 6 points ($\sim$ 25%), and increases the share of heterogeneity attributable to place effects by 6 percentage points at the CBSA level, and by one percentage point at the tract level. Estimated place shares do not appear to be highly sensitive to re-defining the transportation outcome variable. Additional discussion of these results can be found accompanying Appendix Table A.4.

1.6 The Characteristics of Low and High Emissions Places

With estimates of place effects in hand, I move on to characterizing the local amenities that are associated with high and low emissions places. As highlighted in the conceptual model, place effects reflect a mix of differences in demand for energy, energy prices, energy demand elasticities, fuel mixes, and emissions factors.

The urban and planning literature has identified many place-based characteristics that could contribute to differences in energy demand and energy demand elasticities. For example, on the residential energy side, larger homes tend to use more energy, as do single-family homes; there's a strong relationship between carbon emissions and density, though it is potentially not monotonic because of the effect of density on micro-climates (e.g. through heat island effects); and parks, plants, and green surface coverage are all negatively correlated with energy use (see e.g. Ko 2013, for a review). In transportation, car use is lower in places with more alternative transportation options, fewer parking minimums, and more directly connected roads (e.g. Transportation Research Board 2009; Barrington-Leigh and Millard-

Figure 1.6: Correlates of Unobserved Heterogeneity

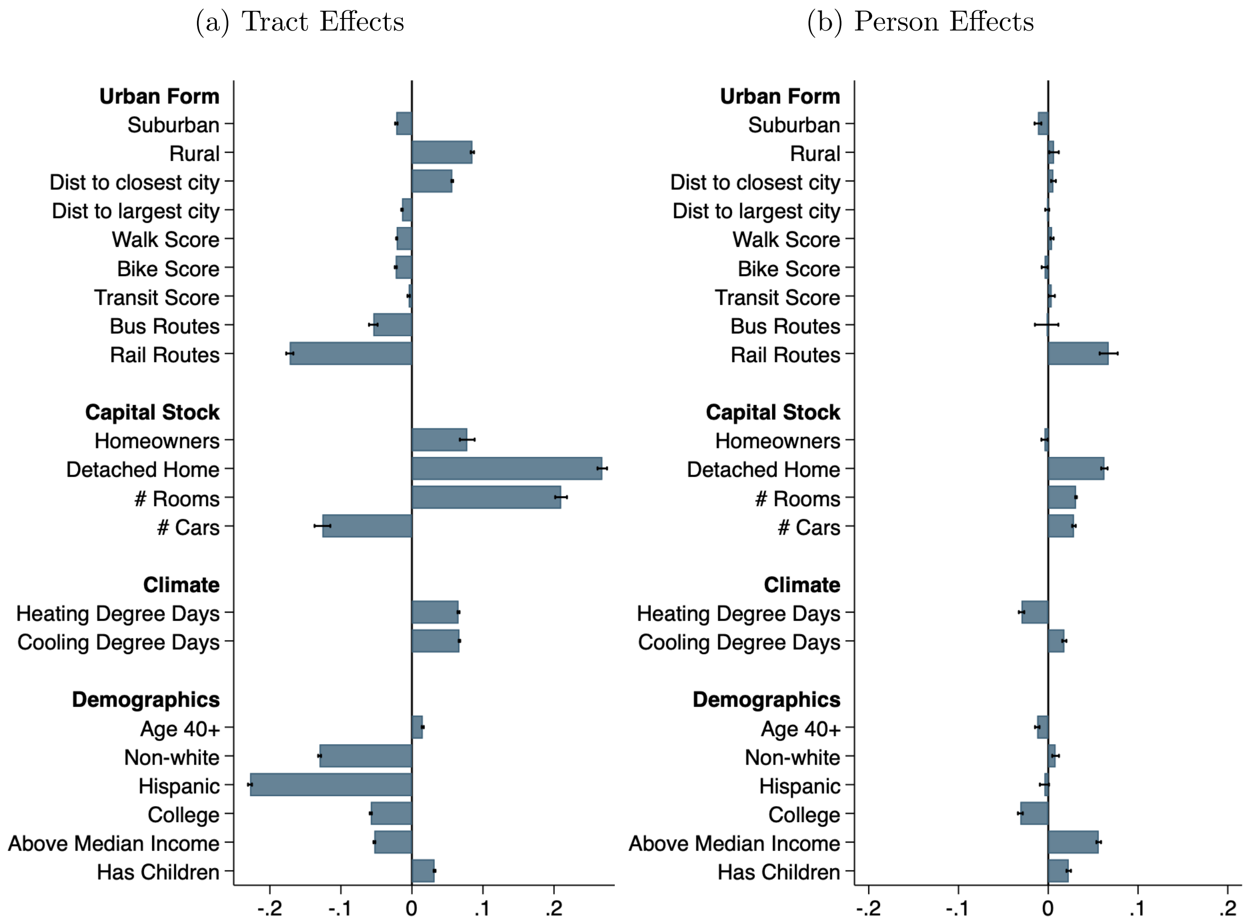

Note: This figure presents estimates from the multivariate ordinary least squares regressions of place effects onto a set of local amenities (urban form, capital stock, climate, electric grid intensity, and density), place effects onto a set of local average demographics, person effects onto a set of local amenities, and local effects onto a set of local average demographics. All amenity variables are measured in logs (+1 for variables where zero is possible), except rural and suburban indicators, and – for the person effect regression only – the homeowner and detached single family home indicators.

Ball 2017). It is worth noting that residential and transportation energy use are related. For instance, denser neighborhoods tend to reduce the distance someone needs to travel to get between destinations, and conversely, public transit and pedestrian infrastructure are more convenient transportation options in dense neighborhoods than they are in areas where destinations are spread out.

In Figure 1.6, I show the result from projecting place and person effects estimated in the KSS mover sample onto a set of some of energy-relevant amenities, as well as onto a set of demographic characteristics. I categorize place amenities into three groups. Amenities in the urban form category are ones that households effectively take as given. I include in these indicators for whether a tract is classified as urban, suburban or rural by the Census, geodesic distances between tract centroids and the centroid of the closest city and the largest city, walk scores, bike scores, transit scores, and the average number of bus routes and rail routes within half a mile of the census block centroids contained within the tract. Amenities in the capital stock category are ones that reflect a combination of the options available in a neighborhood and household preferences – these include the share of homeowners, the share of detached single family homes, the average number of rooms per house, and average number of cars per household. Climate is (for the purposes of this paper) exogenous, and captured using annual heating degree days and cooling degree days in NOAA subdivisions. Bars show coefficient estimates from a multivariate regression of fixed effects onto all three groups of amenities, as well as electric grid intensity and density which I don't report in the figure because I estimate effects larger (in absolute value) than one, so their inclusion makes the other values more difficult to see.[21] For the tract effect estimation, capital stock variables reflect tract-level means; for the household effect they reflect the household's own choice. Finally, all amenities except for suburban and rural indicators (and homeowner and detached home indicators for the household regression only) are measured in logs, so correlates should be interpreted as the percent increase or decrease in place effects associated with a one percent increase or decrease in the amenity.

The results show that tracts with a large share of detached single family homes, bigger homes, and a larger share of homeowners tend to have significantly higher carbon emissions place effects. Un-intuitively, tracts with more cars per household have lower place effects. Tracts with more bus and rail routes within a half mile, and those with better walk and bike scores have lower place effects, as do tracts that are closer to the closest city within the CBSA. Tracts closest to the largest city are higher emissions, and conditional on all these regressors, suburbs are no longer higher emissions than urban areas. One possible explanation for these last two correlations is longer commutes due to congestion. The regression of person effects onto place amenities shows much weaker (or zero) correlations for many of the amenities, consistent with minimal sorting on unobserved characteristics that I estimate in KSS, but intuitively high type households are more likely to live in larger, detached single family homes, and have more cars. They also live in places near more rail routes, potentially reflecting some combination of suburban commuter rail networks, and the high cost of living in central cities.

In the regression of unobserved characteristics on demographics, I find that non-white and Hispanic households are more likely to live in low carbon emissions places, as are college educated and above median households. Households with children live in higher carbon emissions places, as do older households (though the effect here is small). On the household side, college educated households tend to have lower unobserved preferences for carbon

[21] I estimate a coefficient of 4 on log electric grid intensity. This suggests that places that have clean energy grids are also making other investments or decisions that make them lower carbon emissions. I estimate a coefficient of close to -10 on density, suggesting a very strong association between density and place effects.

emissions, as do younger households (but again with a small effect), and above median income households and households with kids have higher unobserved preferences. Additional results on the correlates between observable person and observable place characteristics are presented in Appendix A.6.

1.7 Implications for Aggregate Carbon Emissions

I use estimates of place effects to consider some back of the envelope calculations of how aggregate carbon emissions would change under counterfactual scenarios where the distribution of place effects differs from its current realization. Because place effects are noisily estimated, I use linear Empirical Bayes, i.e. a shrinkage estimator, to forecast place effects that reflect the best (minimum mean squared error) linear prediction of the true values, given my estimates from the KSS analysis. Many papers in the public and labor literatures have used this approach to predict e.g. teacher value add or neighborhood effects in other contexts (Chetty, Friedman, and Rockoff 2014a, 2014b; Angrist et al. 2017; Chetty and Hendren 2018; Finkelstein, Gentzkow, and Williams 2020; Abaluck et al. 2021). Although the linear approximation only corresponds to the true Empirical Bayes posterior when errors are normal and homoskedastic, Kline, Rose, and Walters (2021) show that even when errors are heteroskedastic, the linear shrinkage estimator doesn't do much worse than non-parametric Empirical Bayes. The shrinkage estimates are given by:

$$\hat{\psi}_j^{EB} = \lambda_j \hat{\psi}_j + (1 - \lambda_j) \frac{1}{J} \sum_j \hat{\psi}_j \tag{1.11}$$

where the weights $\lambda_j = \frac{\hat{\sigma}_j^2}{s_j^2 + \hat{\sigma}_j^2}$ capture the signal-to-noise ratio of each estimate and downweight noisy estimates to the grand mean.

I use this approach to estimate counterfactual carbon emissions under three different scenarios: What if the top 10 most populous CBSAs in the US all had the place effects of the New York City CBSA? What if the principal cities of the top 10 most populous CBSas all had the place effects of Manhattan? And what if cities and towns all had the place effects of the principal cities in their (nearest) CBSA? The goal of this exercise is to get a sense for how carbon emissions would evolve under interventions to the built environment of places, without attributing a causal effect to any single amenity, since my correlational analysis doesn't parse those causal effects out.

A naive comparison of household carbon emissions in the New York metropolitan area and the nine other largest metropolitan areas in the US (see Appendix A.8 for the full list) suggests that household emissions from residential energy and commuting are about 8% lower in the New York metropolitan (17.9 tons, annually, per household as compared to 19.2). However, I find that assigning the average New York Metropolitan area place effect to each of these other metropolitan areas actually increases average household emissions slightly, to 19.3 tons per household, highlighting the high place effects of the suburbs around

NY. In contrast, if each principal city of the other top nine metropolitan areas had the place effect of Manhattan, household emissions from residential energy and commuting for current residents of those cities would decrease by over 50%, from 15.6 tons per household to 7.0. This is not as large as the naive decrease to 3.9 tons per household – some of Manhattan's low emissions can be explained by household sorting – but still the Manhattan place effect is significantly lower than the effect of the other 9 largest cities in the US, on average.

Manhattan is unique in its density and transit infrastructure within the US, so the last scenario I consider is intended to capture more closely the spirit of what might happen under some of the regional zoning and transit-oriented development proposals that are emerging across the US.[22] If each place had the place effect of the principal city in its CBSA, annual household carbon emissions would go down from residential energy use and commuting would go down by on average 13%, from 20.6 to 17.9 tons. Again, a naive comparison (20.6 vs. 15.04) overstates the difference between central city and surrounding areas, but my estimates suggest that changing places could yield meaningful reductions in household carbon emissions. For comparison, the Waxman-Markey bill, which failed to pass in 2009 but was, until 2021, the largest federal legislative effort to decrease carbon emissions in the US, was projected to decrease economy-wide emissions 17% in 2020 relative to 2005 (Center for Climate and Energy Solutions 2009).

This exercise lends insight into how development that shifts population shares across place types by "expanding" places with lower place effects (either by making their neighbors look more like them, or by allowing more people to live in the place without changing its fundamentals), could affect emissions in the future. My estimates yield only a first-order, partial equilibrium approximation to the effect of such interventions, as in practice there would be some re-sorting of populations, changing the distribution of household types living in each place and therefore changing aggregate carbon emissions.

1.8 Discussion

Overall, my results suggest that roughly 15-25% of heterogeneity in household carbon emissions from residential and transportation energy use across the US can be explained by place effects, or about 10-20% can be explained by place effects after partialing out variation driven by climate and electric grid intensity. While this leaves the majority of variation either to unobserved household characteristics or unexplained factors in my model, I find that over half of mean differences between places can be explained by place effects, and my estimates

[22]For example, in 2018, Minneapolis was the first city in the US to ban exclusionary zoning (which restricts land to be used for single-family homes only) city-wide. In 2021, California passed State Assembly Bills 9 and 10, which reduce administrative hurdles to "up-zoning" residential land zoned for single family homes only to allow up to four units, as well as land near transit corridors. There have also been attempts to create incentives for up-zoning at the federal level. For example, President Biden's original infrastructure bill proposal in March 2021 included grants to cities who got rid of exclusionary zoning.

suggest the potential for meaningful reductions in carbon emissions from "place-based" interventions that make the distribution of place effects across the US more urban.

Whether such place-based interventions would be welfare maximizing would depend on the costs of implementing them relative to the cost of business as usual or other climate mitigating policies (e.g. a carbon tax).[23] Infrastructure in the US is notoriously expensive to build, making it unlikely that big expansions of new rail (e.g. building a NYC style subway system in Houston) would pass the cost-benefit test in current circumstances. However, correlates of place effects include many amenities that could re-purpose existing built environment without expensive new additional investments – bus lines, bike lanes, pedestrian infrastructure, and dense housing (which, through compactness could decrease related infrastructure and service costs) are all more likely potential contenders. Incorporating cost estimates for a marginal value of public funds analysis (Hendren and Sprung-Keyser 2020) is an important avenue for future research.

The welfare benefits of such interventions would also of course depend on the causal relationships between local amenities and place effects, and household preferences for local amenities. The correlations I presented between amenities and place effects don't identify causal relationships, but they highlight a strong association between many local public goods and carbon emissions, suggesting an important potential role played by local public goods. While Tiebout (1956) posits that residential sorting allows for efficient provision of local public goods, his framework only applies to amenities without scale economies. Moreover, there is reason to believe that residential sorting is not efficient due to frictions or exclusionary policies (e.g. Rothstein 2017; Hausman and Stolper 2020; Christensen and Timmins 2021; Avenancio-León and Howard 2020). Estimating causal relationships between local public amenities and household carbon emissions and quantifying whether emissions-relevant local public amenities are at an efficient level are additional important directions for future work.

Finally, there are several limitations of my empirical analysis that should be taken into consideration while interpreting my results. The first is that due to the survey nature of my data, carbon emissions are noisily measured. This leads to lower explanatory power of the model than is standard in papers in the labor literature using these methods to estimate firm wage premia. The relatively low explanatory power of the model could also reflect model mis-specification, but with only two observations per household for the majority of estimates, the number of specification tests I can do is limited. Second, there is relatively little variation in urban form across the US – 95% of commuters in my sample commute by car, and 75% of residential land in the US is zoned for single family homes only (Badger and Bui 2019). Moreover, place effects are identified from movers, who differ from the general US population in meaningful ways. The external validity of my results is contingent on estimates being stable to widening the distributions of place and person types that they are estimated on.

[23]The welfare impacts would also depend on other externalities or agglomeration benefits of such interventions, which have been studied extensively in the environmental and urban economics literatures. For example, the types of interventions considered in my paper could also impact local air pollution, congestion, traffic fatalities, and labor market productivity.

Chapter 2

The Race Gap in Residential Energy Expenditures

2.1 Introduction

This paper provides estimates of the Black-white residential energy expenditure gap in the US. I use publicly available data from the American Community Survey (ACS) from 2010 to 2017 to show that annual residential energy expenditures – defined as the sum of expenditures on electricity, natural gas, and other home heating fuels – are both statistically and economically significantly higher for Black households than for white households.[1]

Unconditional differences in residential energy expenditures could be driven by many factors including regional variation in climates, prices, and public support for energy efficient investments; household variation in income, wealth, credit access, and home ownership; and local variation in housing stock. After controlling for year, income, household size, and city of residence, Black renters pay \$273 more a year than white renters (16% of the sample average of \$1,705), and Black homeowners pay \$408 more a year than white homeowners (15% of the sample average of \$2,649). Energy expenditures for both groups are decreasing between 2010-2017, and the conditional gap in annual expenditures decreases by about \$150 for the average household, but continues to be economically significant at about \$200 for renters and \$310 for homeowners in 2017. The gap is fairly stable in levels across most income deciles, except it closes at the very top of the income distribution. Therefore, as a percent of income (and baseline residential energy expenditures), the gap is largest for low income households.

Given the long history of discriminatory housing policy, lending practices, and racial

[1]Residential energy expenditures are distinct from transportation energy expenditures such as gasoline purchases. In 2019, residential energy use made up about 20% of energy consumption in the US, and transportation energy use made up about 30%. These two sectors are the largest sources of emissions from individual energy consumption, and understanding the energy expenditure gap in both sectors is crucial for assessing the impacts of possible climate policies. I leave the analysis of transportation energy expenditure gaps to future research.

segregation in the United States, differences in housing stock and accumulated wealth are possible explanations for the remaining residential energy expenditure gap. Controlling for home type or vintage does not eliminate, or even significantly reduce, the gap. This may be because neither variable is a complete measure of housing quality. Evidence from the 2015 Residential Energy Consumption Survey (RECS) is consistent with this interpretation: Conditional on income, Black households are more likely to report that their home is drafty. They also report fewer Energy Star qualified appliances and home features, and are less likely to have received a rebate or tax credit for having upgraded an appliance. These differences exist despite the fact that Black households in the RECS sample are just as (if not slightly more) likely to have gotten an energy audit.

This paper contributes to a growing body of work on energy burden. For example, Reames (2016), Bednar, Reames, and Keoleian (2017) and Kontokosta, Reina, and Bonczak (2020) find that energy burden is higher in high minority share neighborhoods than low minority share neighborhoods in a few cities across the US, and Hernandez, Aratani, and Jiang (2014) study differences in energy insecurity by family characteristics, including race, in the 2011 ACS. Carley and Konisky (2020) review implications of these differences for a clean energy transition. The energy economics literature has to date focused on energy expenditure differences along other dimensions, especially income (e.g. Kolstad and Grainger 2010), and increasingly, geography (e.g. Cronin, Fullerton, and Sexton 2019). In terms of racial differences in burden from the current energy system, the focus has mostly been on differential exposure to resulting pollution. Research shows that Black people are much more likely to live near pollution point sources and be exposed to neighborhoods with higher particulate matter (e.g. Peach 1983, Tessum et al. 2019). (Rothstein 2017) argues that disproportionate exposure to pollution is due to discriminatory siting of sources, and Christensen, Sarmiento-Barbieri, and Timmins (2020) show evidence from an experiment that discrimination, which restricts housing choice sets, causes disproportionate sorting of Black families into neighborhoods near polluting point sources. Hausman and Stolper (2020) argue that hidden information about pollution, even when constant across all households, also leads to disproportionate sorting across neighborhoods because pollution is correlated with other disamenities.

More broadly, this paper builds on insights from a large body of work on the persistent effects of systemic racism on other outcomes. Black people have less wealth and are less likely to own homes (e.g. Rothstein 2017), they are more likely to face high cost loans, even when controlling for credit score and other risk factors (e.g. Bayer, Ferreira, and Ross 2018), and they pay higher property taxes for the same home values (Avenancio-León and Howard 2020). Aaronson, Hartley, and Mazumder (2019) provide evidence that many of the above outcomes were meaningfully affected by the Home Owners' Loan Corporation (HOLC) redlining maps in the 1930s, which restricted credit access in Black neighborhoods. Beyond facing discrepancies in home prices, Hardy et al. (2018) show that Black Americans face higher year-to-year income volatility, and Ganong et al. (2020), show that as a result of wealth differentials, Black consumption is more sensitive to income shocks. These differences in wealth, home ownership, income volatility, and credit access all serve as potential barriers to living in higher quality, more energy efficient homes or to making necessary energy efficiency upgrades.

Tying these literatures together, this paper contributes to a broad set of evidence that Black Americans bear a disproportionate burden of the current energy system, both through disproportionate pollution exposure, and as I highlight, through disproportionate costs, likely at least in part as a result of persistent disparities in wealth and housing. In the remainder of this paper, I outline the data and methodology, present descriptive results, and discuss conclusions and next steps.

2.2 Data and Methodology

I use the American Community Survey (ACS) Public Use Microdata Sample from 2010 to 2017 (Ruggles et al. 2020). The ACS is a nationally representative survey of about 1% of the US population every year. I restrict the sample to households that are either entirely Black or entirely white. I drop households with missing or negative income. I calculate residential energy expenditures as the sum of self-reported electricity expenditures, natural gas expenditures, and other home heating fuel expenditures, and I drop households whose residential energy bills are included in their rent payments or condo fees.[2] I deflate all dollar amounts to 2012 dollars using the consumer price index (CPI) from the Bureau of Economic Analysis (BEA), and express income in thousands of dollars. ACS household incomes and energy expenditures are censored at the 99.5th percentile by state-year, and I additionally censor household size at 10 people. After all restrictions, the pooled sample consists of 7,906,852 people. All estimates are weighted by the ACS's household weight. Black households make up 12.9% of the weighted sample.[3]

I compute unconditional and conditional annual residential energy expenditure gaps by regressing household residential energy expenditures on an indicator for household race, year fixed effects, and an increasing set of household controls:

$$y_{it} = \delta \cdot 1[Black_i] + \tau_t + X_{it}\beta + \epsilon_{it}$$

where y_{it} is annual household (i) energy expenditures, τ_t are year fixed effects, and X_{it} includes characteristics such as household income, size, and geographic characteristics. Residential energy expenditures do not actually increase linearly in either household income or household size, so I have also run these specifications controlling for household income deciles and household size dummies on the right hand side, as well as in log-log form. In both versions, level estimates and implied percentage gaps are very similar to those in my main specification, so I report the linear specifications for simplicity. I estimate specifications separately for renters and homeowners, as renters may face principal-agent problems that

[2]This is 6% of white households and 9% of Black households after all other sample restrictions.

[3]This percentage is slightly increasing over the course of the sample, from 12.7% in 2010 to 13.2% in 2017.

prevent them from making optimal energy efficiency investments. For estimates that include Metropolitan Statistical Area (MSA) fixed effects, I use a state fixed effect for observations where MSA is not identified. My preferred specification includes city fixed effects, which are the most granular geographic control I can add using the publicly available microdata sample. Since these are meant to be narrow geographies, I only include observations with an identified city. This decrease the sample size and changes the sample composition significantly: The city sample is 901,580 households (about 13% of the weighted full sample), and the weighted share of Black households in this sample is 27%. Errors are clustered at the state level in all specifications.

To understand residential energy expenditure patterns in more depth, I expand on my preferred specification to look at how annual expenditures have changed over time. I also compute income deciles for the full sample population each year, and look at how the gap differs across income deciles, and how that distribution has evolved between 2010 and 2017.

Lastly, I explore possible mechanisms. Continuing to use the ACS sample, I add flexible controls for home type (single-family detached, single-family attached, van or mobile home, 2-4 plex, 5+ unit apartment building) and home vintage (decade fixed effects) to test whether either of these variables reduces the gap. I also supplement my analysis with the 2015 Residential Energy Consumption Survey (RECS). RECS is administered by the Energy Information Administration every 4-6 years to a small, nationally representative set of housing units. RECS asks a detailed set of questions about energy use and investments through a combination of surveys and in-person interviews. I restrict the RECS sample to mirror sample restrictions in the ACS.[4] The final sample consists of 4,805 respondents, with Black respondents making up 12% of the weighted sample. I use RECS to test differences by race, conditional on income, in receipt of energy assistance and audits, self-assessed home quality, and availability of Energy Star appliances and other energy-efficient home features. I also test differences in energy burden, as measured by whether a household reported reducing or forgoing on basic necessities to pay an energy bill, whether a household reported keeping the home at an unhealthy temperature in order to pay an energy bill, or whether a household received a disconnect notice due to inability to pay a bill. All estimates are weighted by RECS sample weights, and errors are clustered at the census division level.

[4] In RECS, I only know the race of the respondent, not everyone in the household. I keep only Black or white respondents. RECS reports incomes in 8 categories (in thousands: < 20, [20,40) [40,60), [60-80), [80-100), [100-120), [120-140), ≥ 140) and only reports Census Divisions for geography. I drop households whose residential energy bills are included in their rent (9% of white respondents and 14% of Black respondents in the weighted sample). Lastly, many of the questions asked by RECS allow respondents to answer "I don't know" or refuse to answer. I treat these answers as missing. I treat "N/A"s as "No"s, except for when estimating the share of Energy Star -rated appliances/features, in which case I treat "N/A"s as missing.

2.3 Results

Evidence on Expenditures

Table 2.1 shows evidence from the ACS that there is a statistically and economically significant residential energy expenditure gap across Black and white households in the years 2010-2017.[5]

Column (1) shows the unconditional mean difference: on average, Black households in my sample pay about $54 more a year in energy bills than white households do, although this unconditional difference is not significantly different from 0. The gap becomes statistically significant and economically meaningful after controlling for income (column 2): Black households pay about $193 more a year than white households do. This is 8% of the sample average annual expenditures. The gap persists with controls for household size (column 3), and is driven by both renters and homeowners (columns 4 and 5). The gap for homeowners is bigger in levels ($381 relative to $258), but as a percent of sample averages the gaps are comparable (14% for renters and 15% for homeowners). Accounting for sorting across climates by controlling for MSA (columns 6 and 7) decreases the gap somewhat for both renters and homeowners but it is still economically and statistically significant at 10% and 11% of average expenditures, respectively.[6]. Columns 8 and 9 add city fixed effects. This is my preferred specification because it most precisely controls for location-specific characteristics. Within the same cities, Black renters spend $273 more a year than white renters (16% relative to average), and Black homeowners spend $408 a year more than white homeowners (15% relative to average).[7] In Table 2.2, I report estimates for all specifications using just the city sample; they are significantly bigger than those in the full sample. This suggests that the large gap when I include city fixed effects is driven by the restriction of the sample to people living in cities, likely as a result of the fact that wealthy suburbs that use a lot of energy tend to be white.

Figure 2.1 shows that average energy expenditures conditional on income, household size, and city have decreased between 2010 and 2017 for both Black and white households. The conditional energy expenditure gap has also decreased in this period, by about $150 for the average household, although I cannot reject a constant gap over time.[8] In 2017 the gap

[5]I have also analyzed the Black-white energy expenditure gap with the Consumer Expenditure Survey (CEX) and Residential Energy Consumption Survey (RECS). The patterns and orders of magnitude are broadly consistent. Both those surveys do not have as much geographic detail and so I exclude those results and focus on the ACS for brevity.

[6]Controlling more flexibly for weather by interacting MSA FEs with year FEs does not change these estimates

[7]Individual regressions of electricity costs, natural gas costs, and other home heating fuel costs suggest that the gap is driven by electricity and natural gas. If anything, Black households spend less on home heating fuel than white households do, but this difference goes to 0 within cities.

[8]A useful avenue for future work is to explore what has driven this decrease. Of particular interest could be the role of the American Recovery and Reinvestment Act of 2009, which directed significant funds into energy efficiency investments.

Table 2.1: Gap in Annual Residential Energy Costs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Black	53.5	192.7***	163.8***	258.1***	380.5***	179.1***	289.9***	272.6***	408.0***
	(37.00)	(33.64)	(35.47)	(46.94)	(40.60)	(25.84)	(29.87)	(53.95)	(46.34)
HH income		4.783***	3.855***	1.413***	3.263***	2.019***	3.008***	1.387***	2.131***
		(0.219)	(0.176)	(0.169)	(0.192)	(0.273)	(0.128)	(0.293)	(0.435)
HH size			240.7***	271.2***	224.9***	277.4***	224.0***	268.6***	248.6***
			(13.22)	(18.52)	(12.67)	(15.90)	(10.13)	(17.57)	(20.07)
Constant	2592.9***	2226.1***	1702.6***	1203.0***	1976.5***	1176.9***	2004.9***	1055.2***	1989.4***
	(83.11)	(69.99)	(54.87)	(50.99)	(60.70)	(23.06)	(27.55)	(44.95)	(32.13)
Sample Mean	2373.5	2373.5	2373.5	1811.2	2615.5	1811.2	2615.5	1705.1	2648.9
Year FE	X	X	X	X	X	X	X	X	X
Renters				X		X		X	
Homeowners					X		X		X
MSA FE						X	X		
city FE								X	X
R-squared	0.00871	0.0697	0.117	0.114	0.103	0.165	0.167	0.180	0.198
N	7,906,852	7,906,852	7,906,852	1,936,533	5,970,319	1,936,533	5,970,319	363,715	537,865

Note: This table reports annual residential energy expenditure gaps in the ACS, pooled across 2010-2017. All values are reported in 2012 dollars, and household income is reported in $1000s. Standard errors are clustered on state.

Table 2.2: Gap in Annual Residential Energy Costs, City Sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Black	312.8***	451.5***	414.7***	434.6***	646.0***	289.2***	432.9***	272.6***	408.0***
	(52.55)	(54.31)	(62.25)	(82.12)	(52.63)	(55.13)	(47.72)	(53.95)	(46.34)
HH income		3.851***	3.089***	0.826***	2.312***	1.340***	2.207***	1.387***	2.131***
		(0.167)	(0.211)	(0.183)	(0.265)	(0.294)	(0.463)	(0.293)	(0.435)
HH size			269.1***	266.7***	256.8***	267.4***	245.6***	268.6***	248.6***
			(23.81)	(20.63)	(30.99)	(19.17)	(21.44)	(17.57)	(20.07)
Constant	2323.1***	2015.8***	1427.1***	1037.9***	1866.6***	1045.6***	1963.3***	1055.2***	1989.4***
	(102.0)	(104.6)	(76.68)	(77.22)	(75.22)	(42.34)	(43.12)	(44.95)	(32.13)
Sample Mean	2209.8	2209.8	2209.8	1705.1	2648.9	1705.1	2648.9	1705.1	2648.9
Year FE	X	X	X	X	X	X	X	X	X
Renters only				X		X		X	
Home-owners only					X		X		X
MSA FE						X	X		
City FE								X	X
R-squared	0.0140	0.0551	0.116	0.122	0.107	0.174	0.185	0.180	0.198
N	901,580	901,580	901,580	363,715	537,865	363,715	537,865	363,715	537,865

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table reports annual energy-expenditure gaps in the ACS, pooled across 2010-2017, restricting the sample in each specification to only households in a city identified by the ACS. All values are reported in 2012 dollars, and household income is reported in $1000s. Standard errors are clustered on state.

remains significantly different from zero, at close to $200 a year for renters and $310 a year for homeowners.

Figure 2.2 shows that the gap is fairly stable across income deciles, for both renters and homeowners, except at the very top of the income distribution where it closes. Given that energy expenditures are a larger share of lower-income households' consumption, this means energy burden is especially heightened for low income Black households.

Possible Mechanisms

I first test whether the residential energy expenditure gap can be explained by differences in home type. Columns (1) and (2) of Table **??** show results. As expected, single-family detached homes have the highest energy expenditures, and apartments in large buildings have the lowest expenditures. However, controlling for home type does not decrease the gap for renters, and it only decreases the gap for home-owners by about $50 relative to the main specification. I next test whether the residential energy expenditure gap can be explained by home vintage by controlling for home vintage with decade fixed effects. Columns (3) and (4) of Table **??** show results. Despite the fact that newer homes are broadly speaking associated with lower residential energy expenditures, controlling for home vintage does not change the residential energy expenditure gap. Controlling for home type and vintage may not have an effect on the residential energy expenditure gap because both variables are imperfect proxies for energy efficiency, since they do not capture renovations or investments into energy efficient appliances.

To explore these mechanisms, I turn to the 2015 RECS. I compare survey responses about home quality, appliance quality, and energy burden across race, conditioning on income categories. A few key patterns emerge in Table **??**. First, Black respondents are about 13 percentage points more likely to report that their home was at least somewhat drafty. Out of a set of several appliances and home features[9], Black respondents have a 7 percentage point lower share that were Energy Star rated, and they are 3 percentage points less likely to report having received a rebate or tax credit for upgrading an appliance. If anything, Black respondents are slightly more likely to have gotten an energy audit, suggesting that this isn't a matter of differential information, though this result is not statistically different from 0. Moreover, Black respondents were about 50% more likely to report having reduced or forgone basic necessities at least one month in the last year in order to afford their energy bill, were about 40% more likely to report having kept the home at an unhealthy temperature at least one month in the last year in order to afford their energy bill, and were about twice as likely to have received a disconnect notice due to inability to pay a bill at least one month in the last year. These estimates suggest that energy costs are highly salient, and are evidence of a striking disparity in energy burden.

[9]RECS asks about 8 appliances/features: clothes washer, clothes dryer, dishwasher, fridge, freezer, water heater, light bulbs, and windows.

Figure 2.1: Energy Expenditures Over Time

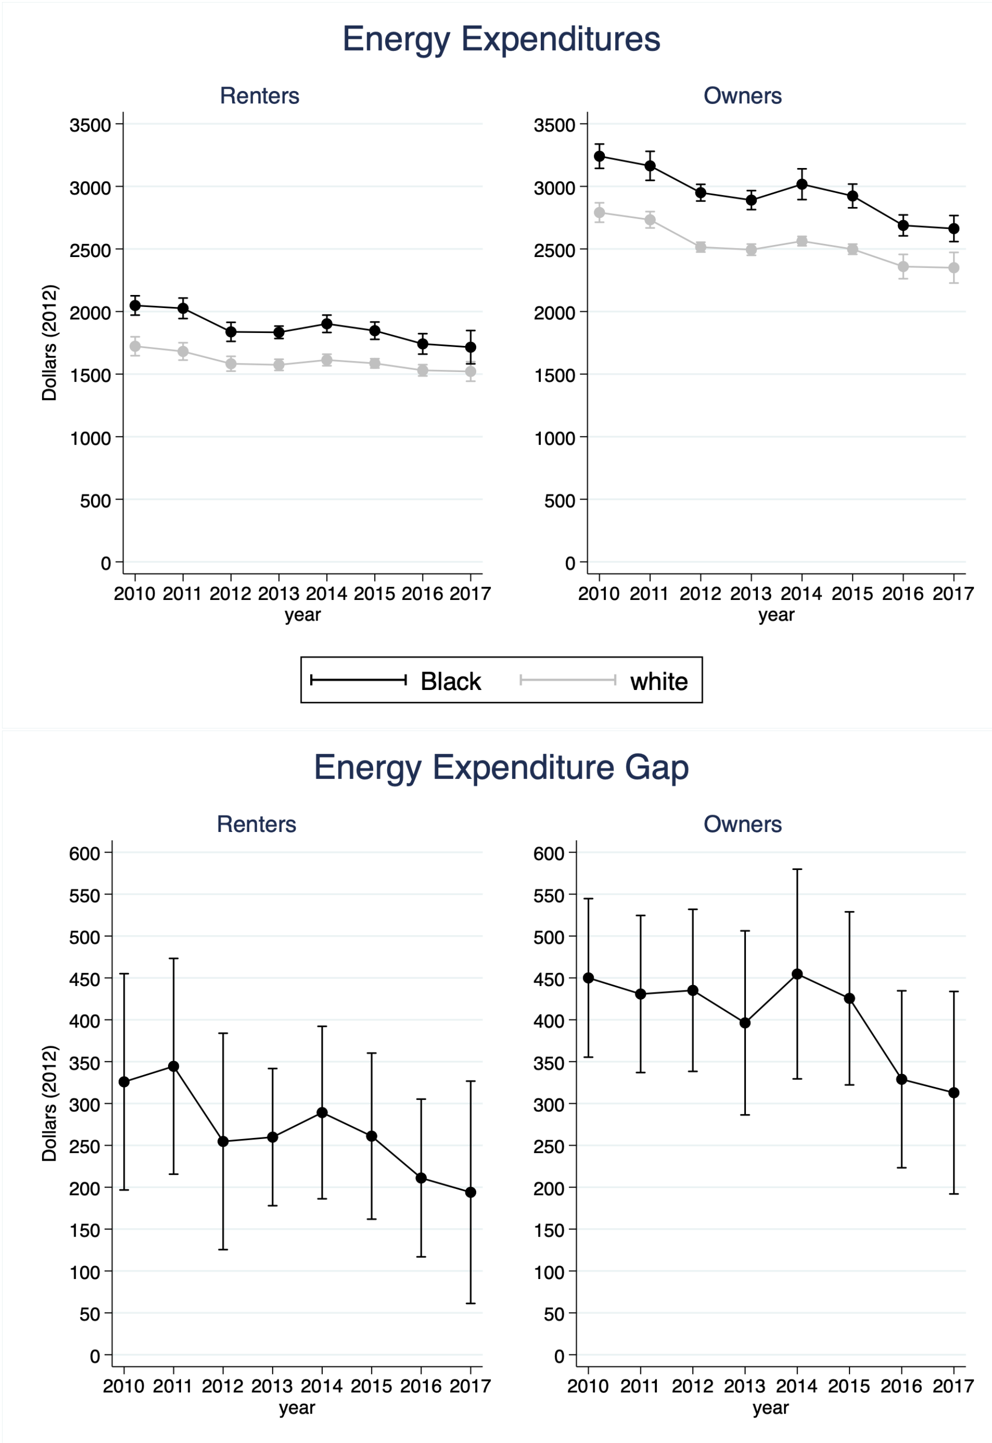

Note: The top panel of this figure shows the evolution over time of mean annual energy expenditures in the ACS conditional on income, household size, and city FE. The bottom panel of this figure shows the evolution over time of the conditional gap between Black and white expenditures. All values are reported in 2012 dollars. Standard errors are clustered on state. Bars are 95% confidence intervals.

Figure 2.2: Energy Expenditures by Income

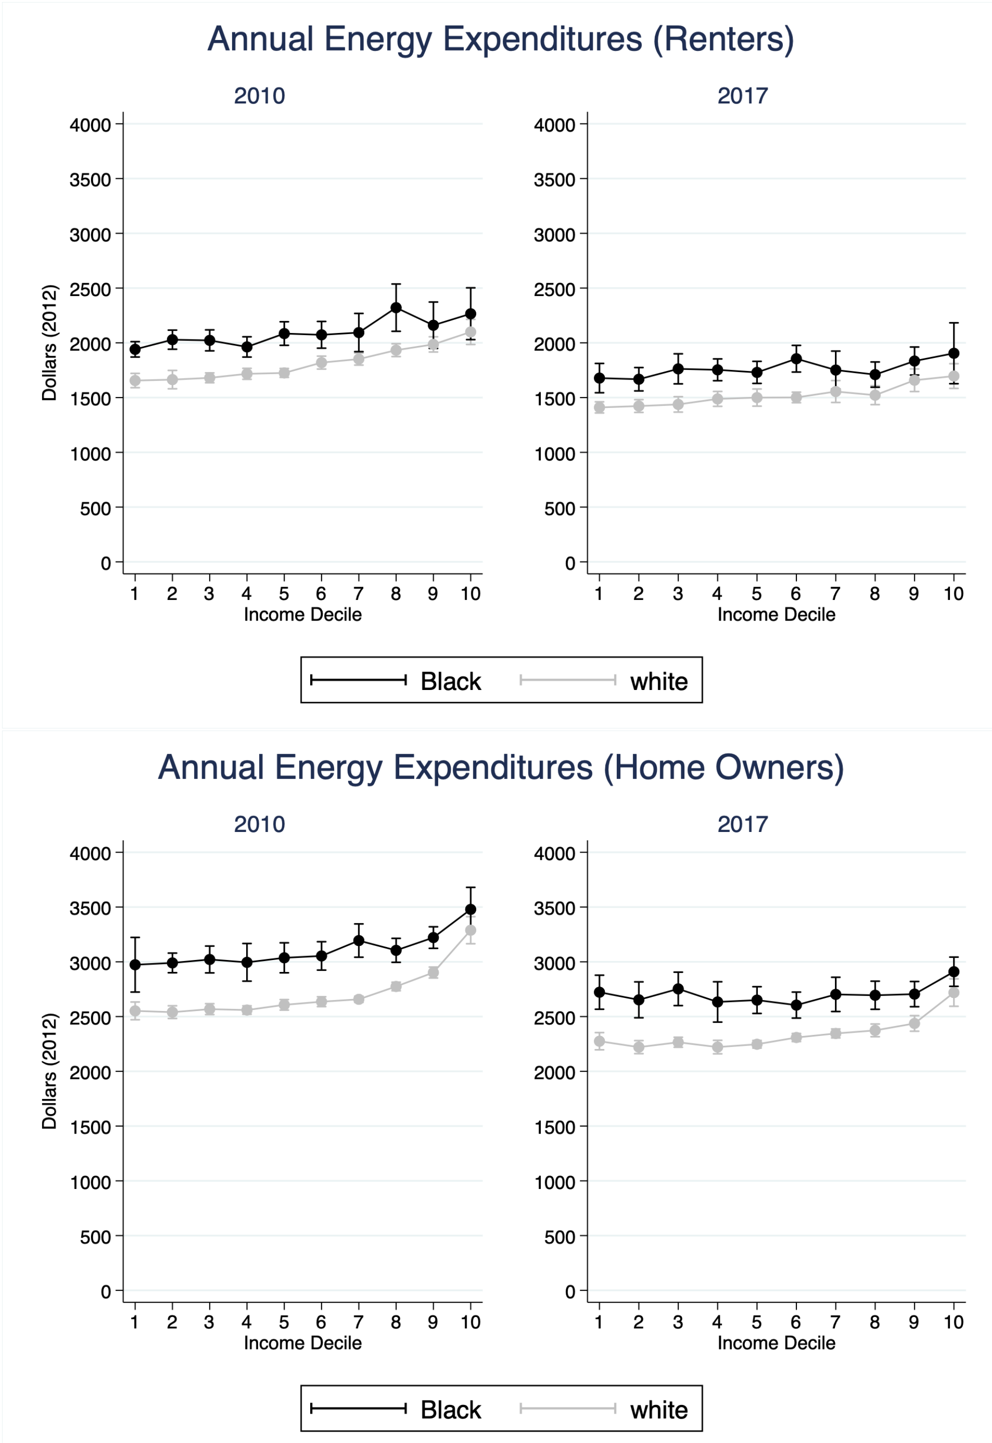

Note: This figure shows mean annual energy expenditures in the ACS, conditional on household size and city, by income decile for Black and white households. The top panel shows expenditures for renters, in 2010 on the left and 2017 on the right, and the panel row shows expenditures for home owners. All values are reported in 2012 values. Standard errors are clustered on state. Bars are 95% confidence intervals.

Table 2.3: Gap in Annual Residential Energy Costs, Controlling for Home Vintage

	(1)	(2)	(3)	(4)
Black	275.2***	353.0***	276.9***	409.7***
	(50.44)	(36.61)	(55.21)	(44.88)
HH income	1.127***	2.221***	1.490***	2.217***
	(0.155)	(0.279)	(0.315)	(0.439)
HH size	186.1***	208.8***	265.4***	248.6***
	(12.98)	(12.36)	(16.91)	(20.35)
Single-Family Attached Home	-373.1***	-462.3***		
	(55.10)	(32.35)		
Van or Mobile Home	-385.0***	-375.5***		
	(59.32)	(70.10)		
2 - 4 plex	-638.2***	-186.3*		
	(28.77)	(86.47)		
5+ Unit Apt. Building	-1151.1***	-1451.7***		
	(51.31)	(204.4)		
Vintage: 1970 - 1979			-204.8***	-152.1***
			(35.56)	(24.37)
Vintage: 1980 - 1989			-206.1***	-246.6***
			(49.28)	(28.94)
Vintage: 1990 - 1999			-188.3**	-98.06**
			(54.53)	(31.01)
Vintage: 2000 - 2009			-194.4**	-242.4***
			(68.78)	(40.73)
Vintage: 2010 - 2017			-370.5***	-448.9***
			(71.12)	(41.17)
Constant	1995.2***	2253.5***	1140.0***	2053.0***
	(42.75)	(38.46)	(27.94)	(29.92)
Sample Mean Energy Expenditures	1705.1	2648.9	1705.1	2648.9
Year FE	X	X	X	X
Renters only	X		X	
Home-owners only		X		X
City FE	X	X	X	X
R-squared	0.275	0.236	0.185	0.201
N	363,715	537,865	363,715	537,865

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table reports annual energy-expenditure gaps in the ACS, pooled across 2010-2017, controlling for home type and vintage. Columns 1 and 2 control for home types. The omitted category is single-family detached homes. Columns 3 and 4 control for home vintage. The omitted category is homes built before 1970. All specifications include city fixed effects (and are comparable to columns 8 and 9 in Table 2.1). Standard errors are clustered on state.

Table 2.4: Conditional Differences in Housing Stock and Energy Burden

	White	Black	Diff
Received energy assistance in 2015	0.031	0.042	0.011
			(0.015)
Got an energy audit	0.086	0.099	0.013
			(0.014)
Said home was well insulated	0.320	0.329	0.009
			(0.027)
Said home was drafty	0.506	0.640	0.134***
			(0.023)
Share of Energy Star appliances or features	0.443	0.370	-0.073***
			(0.008)
Received an appliance rebate or tax credit	0.105	0.070	-0.034*
			(0.015)
Has solar PV	0.013	0.013	-0.000
			(0.003)
Has smart meter	0.332	0.297	-0.035
			(0.037)
Has smart thermostat	0.033	0.047	0.014
			(0.011)
Had to reduce/forgo basic necessities	0.202	0.312	0.111*
			(0.037)
Kept home at unhealthy temperature	0.103	0.144	0.041*
			(0.015)
Received disconnect notice	0.137	0.277	0.141***
			(0.013)
N	4,282	523	

mean coefficients; standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table tests differences across race in self-reported responses about home quality, appliance quality, and energy burden, conditioning on income categories. Data is from the 2015 RECS. Standard errors are clustered on census division.

2.4 Discussion and Conclusions

This paper provides estimates of the Black-white residential energy expenditure gap in the US. These estimates suggest that Black households face a higher energy burden than white households at almost every position in the income distribution. Understanding the differential energy burden is critical when designing policies that will affect energy prices, such as much-needed policy to reduce greenhouse gas emissions. This is especially true given that this gap may be another of many outcomes that has been affected by the persistent effects of systemic racism in the United States, mediated in particular by differences in housing stock and wealth.

This paper has some important limitations. The results are suggestive but not causal, and energy expenditures are self reported on an annual basis. In future versions of this paper, I will use residential billing data in the state of California for this analysis. This will eliminate any recall error, and will also allow me to observe differences in prices, payment of late fees, and participation in low-income assistance programs. Billing data also make it possible to control for weather more directly, and provide more spatial granularity, which I will use to estimate the long-term impacts of residential segregation policies such as redlining.

Chapter 3

Regulating Mismeasured Pollution: Implications of Firm Heterogeneity for Environmental Policy

3.1 Introduction

This paper provides the first estimates of within-industry heterogeneity in energy and CO_2 productivity for industries spanning the entire US manufacturing sector. We define energy and CO_2 productivity as log dollars of output per dollar of energy input or per ton of CO_2 emitted.[1] Three key takeaways emerge. First, within narrowly defined industries, heterogeneity across plants in energy and CO_2 productivity is enormous. For example, given one dollar of energy inputs, a plant at the ninetieth percentile of a typical industry's energy productivity distribution produces 580 percent more output than a plant at the tenth percentile of the same industry. Second, these values significantly exceed heterogeneity in most other measures of productivity. For example, the corresponding 90–10 differences for labor and total factor productivity are 400 percent and 150 percent, respectively. Third, heterogeneity in energy and CO_2 productivity has important implications for industry-based environmental regulations. Many countries have considered pairing a carbon tax on domestic output with a tariff on imports that is proportional to the carbon content of the imports. We show that an industry-based carbon tariff, which abstracts from within-industry heterogeneity, will substantially differ from the correct plant-level Pigouvian tax for many plants. Many existing environmental regulations and standards apply uniformly across plants within an industry. For example, the US Clean Air Act requires plants in regulated industries and regions to meet an industry-level technology standard by installing "Best Available Control Technologies." Similarly, the Clean Water Act's Industrial Effluent Guidelines require plants to meet an industry-level technology standard. Several tradable permit markets use industry-level

[1] Energy and CO_2 productivity are the inverse of energy and CO_2 intensity. We use the former metric to facilitate comparisons to other single-factor and total-factor productivity measures.

rebates to compensate firms.[2] Due to substantial data requirements, researchers and policymakers have a limited understanding of the extent of producer heterogeneity in energy and CO_2 productivity. We show that plants within an industry have very different pollution emissions rates, and therefore, such industry-level regulations will be too stringent for some plants and too lenient for others.

We use confidential, plant-level data from the US Census of Manufactures CM and the Manufacturing Energy Consumption Survey MECS to explore this heterogeneity. We distinguish about 375 six-digit NAICS industries. One industry, for example, manufactures carbon black; another makes ethyl alcohol. Our main results calculate plant-level energy expenditures on raw fuels and electricity as reported in CM and MECS. We also calculate plant-level CO_2 emissions by converting each fuel consumption choice to CO_2 equivalents using emissions factors (e.g., tons CO_2 emitted per ton.

A few estimates near the paper's end analyze carbon tariffs. These estimates account for energy consumption and emissions required to produce intermediate inputs that are used for final good production, sometimes called "indirect emissions." For example, in most of the paper, emissions for the cookware industry include coal, gas, oil, and electricity used to shape a pan. Indirect emissions for the cookware industry also include fossil fuels used to make aluminum, which is then purchased as an intermediate input to make a pan. We calculate indirect emissions in two separate ways. The first is standard: we invert the US input-output table to compute the dollars of coal, oil, and natural gas inputs required to produce a dollar of output in each industry. This accounts for energy used to produce inputs, energy used to produce the inputs to inputs, etc. Our second measure of indirect energy is nonstandard: we use plant-level data on the dollar value of each individual material input the plant uses, along with associated industry codes for each material, which are all part of the CM Materials Trailer. We combine this information with the inverted input-output table to calculate indirect energy and emissions separately for each plant.

This paper builds on several literatures. One explores the implications of firm heterogeneity for environmental policy and either argues for market-based instruments like pollution taxes or cap-and-trade markets (Carlson et al. 2000; Goulder and Parry 2008) or analyzes industry-based regulation in Melitz-type settings when firms are heterogeneous (Shapiro and Walker 2018). Several papers within this literature specifically analyze border adjustments (Cosbey et al. 2019; Kortum and Weisbach 2017). This paper also relates to work analyzing the efficiency of imperfectly targeted environmental policies (Jacobsen et al. 2020). A related literature shows that total factor productivity is heterogeneous within narrowly-defined or homogeneous industries (Syverson 2011); other work interprets heterogeneity as factor misallocation (Hsieh and Klenow 2009). Existing analysis of heterogeneity in energy productivity is limited, though includes studies of a subset of energy-intensive, trade-exposed sectors (Gray and Metcalf 2017). The remainder of the paper discusses data, methodology,

[2] California's AB-32 cap-and-trade distributes additional permit allocations to energy intensive, trade-exposed industries using an industry-level assistance factor to help combat against regulatory leakage. These assistance factors are applied at the industry-level when determining permit allocations for a facility.

and results.

3.2 Data & Methodology

We measure plant-level energy inputs using data from the 2007 CM and the 2006 MECS. The CM includes about 350,000 US manufacturing plants operating in 2007, while MECS includes a probabilistic sample of around 15,000 plants. We join MECS and the CM at the plant level, using a unique plant identifier. Our MECS estimates use survey weights to make statistics represent the broader manufacturing sector. The CM reports each plant's value of shipments, capital stock, production hours, and expenditure on electricity, fuels, and materials. We exclude "administrative records" since many of their values are imputed. We also exclude records where output, fuel expenditures, or electricity expenditures are imputed.

The CM and MECS both report plant-level expenditure on fuels and on electricity, which we use to compute CO_2 emissions. MECS further reports physical quantities and expenditures for each fuel, which we convert to CO (See chapter appendix for details . Since the CM does not report expenditures by fuel type, we use MECS to calculate industry-level averages of CO_2 per dollar of fuel expenditure, and we multiply each CM establishment's fuel expenditure by these averages. For electricity inputs, we use the EPA's eGrid database, which assigns annual total output emissions rates CO_2 per KWh to 26 regions of the country, to calculate the mean marginal emissions based on plant location of electricity consumption.

We account for indirect emissions only in our estimates of carbon tariffs. We do this in two separate ways. First, we use the 2007 US benchmark input-output data of the Bureau of Economic Analysis. We invert the input-output table to compute the total dollars of coal, oil, and natural gas inputs required to produce a dollar of output in each industry. We apply emissions coefficients from the Energy Information Agency and Environmental Protection Agency EPA to calculate the total CO_2 emitted per dollar of output in an industry. Our second measure of indirect emissions comes from the CM Materials Trailer, which provides plant-level detail on the dollar value of each material input, along with associated input industry codes. We multiply these expenditures by the corresponding industry emissions rate from the inverted input-output table. Thus, while emissions rates are constant across intermediate input industries, plant-level variation in intermediate input intensity generates additional heterogeneity in energy and emissions productivity.

We use all these data to construct multiple measures of energy and emissions productivity. For comparability with common productivity measures, we construct productivity measures as the log of the value of shipments per dollar of direct energy input, or per metric ton of CO_2 emitted. We also discuss estimates that define productivity as log dollars of value added per unit of energy input or CO_2 emissions in the Appendix. We calculate value added by subtracting expenditures on capital, labor, materials, and energy from the plant's total value of shipments.

For each industry, we measure productivity heterogeneity by calculating the ninetieth and tenth percentile of energy and CO_2 productivity across plants within the same industry.[3] We also compute the within-industry standard deviation of all productivity measures for each of the 375 industries. Lastly, we summarize these industry-level dispersion measures by taking the unweighted mean across all industries. This latter statistic provides some insight on within-industry heterogeneity in productivity for the mean industry.

3.3 Results

Table 3.1 shows the mean and dispersion of seven different productivity measures. This table accounts for only "direct" CO_2 or other inputs at a plant. Columns 1–6 calculate productivity as the log of the plant's value of shipments divided by some measure of a plant's factor demand, CO_2 emissions, or intermediate inputs. Columns 2 and 3 report value of shipments per ton of CO_2 produced, where CO_2 is calculated using the CM and MECS samples, respectively. Columns 4–6 report other single factor productivity measures, as indicated in the column headings. Column 7 presents statistics from a total factor productivity index.[4]

Panel A of Table 3.1 shows mean productivity levels. For example, column 1 implies that energy costs are roughly 1.5 percent of output value 0.015 1 exp 4.16 for the mean plant in our sample, since the log of output per dollar energy input is 4.16.[5] Panel B of Table 3.1 summarizes the industry-level dispersion measures. The first row presents the mean of the within-industry 90–10 ratio, taken across all industries in our sample. The second row of panel B shows the standard deviation of the within-industry 90–10 ratio, taken across all industries. The third row shows the difference between the ninetieth percentile industry and the tenth percentile industry of this within-industry 90–10 dispersion measure. Panel C shows similar values, but using within-industry standard deviations.

Panels B and C of 3.1 show substantial heterogeneity in output per dollar of energy expenditure or per ton of CO_2 emitted, which is the paper's first main finding. The top-left entry in panel B, for example, shows that given a dollar of energy inputs in the industry with the mean energy productivity dispersion, a plant at the ninetieth percentile of the within-industry energy productivity distribution produces 580 percent more output than a plant at the tenth percentile of that within-industry distribution does. Dispersion in CO_2 productivity is even wider, at 2.27 log points 870 percent difference. The standard deviation of energy and of CO_2 productivity within the average industry is 0.75 to 0.89 log points, respectively.

Panels B and C also show the paper's second main finding: dispersion in CO_2 and energy productivity is larger than dispersion in most other productivity measures. Both panels

[3] To respect confidentiality requirements for 90–10 statistics, we use each industry's mean and standard deviation of the respective productivity measure to simulate the ninetieth and tenth percentile using a normal distribution. Estimates using the simulated data are nearly identical to those from the underlying microdata.

[4] This index uses a Cobb-Douglas production technology with three inputs: labor, capital, and materials. Output elasticities for each input are constructed from industry-level revenue shares (Syverson 2011).

[5] We report this calculation for energy productivity but not other columns because not all other inputs are measured in dollars (e.g., labor is in terms of worker hours).

Table 3.1: Single and Total Factor Productivity Statistics

	Direct energy (1)	CO_2 [CM] (2)	CO_2 [MECS] (3)	Labor (4)	Capital (5)	Materials (6)	TFP (7)
Panel A. Industry-wide statistics							
Mean	4.16	8.42	8.80	4.51	1.01	0.95	1.81
SD	0.94	1.16	1.15	0.83	0.95	0.66	0.56
Panel B. Within-industry 90-10 difference in productivity							
Mean	1.92	2.27	2.27	1.63	2.22	1.34	0.92
SD	0.47	0.57	1.17	0.45	0.50	0.61	0.39
p90-10	1.21	1.46	3.01	1.16	1.27	1.58	0.99
Panel B. Within-industry standard deviation of productivity							
Mean	0.75	0.89	0.89	0.64	0.87	0.52	0.36
SD	0.18	0.22	0.46	0.17	0.19	0.24	0.15
p90-10	0.47	0.49	1.14	0.44	0.43	0.55	0.33

Notes: Panel A means and SD are computed from plant-level CM and MECS observations. Panel B statistics are calculated using the 375 within-industry 90–10 dispersion measures. Panel C statistics are calculated using the 375 within-industry standard deviation measures. See text for details.

show that dispersion in energy and CO_2 productivity is more than twice as large as dispersion in total factor productivity TFP. Typically, single-factor productivity measures are more dispersed than TFP, but Table 3.1 shows that dispersion in energy and CO_2 productivity exceeds dispersion in other single-factor productivity measures like labor or material productivity.[6] Dispersion in energy and capital productivity is more similar, though worth interpreting cautiously since the durability of capital investments makes the value of the capital stock difficult to measure.

All pairwise t-tests not shown for space reject the hypothesis that dispersion in energy and CO_2 productivity equals dispersion in the other productivity levels. Appendix Table B.1 shows similar conclusions from value-added productivity measures. It may be unsurprising that CO_2 productivity varies so much, since differences in fuel inputs, variation across the grid in the CO_2 intensity of electricity generation, and related forces make CO_2 more variable than energy expenditure. It is more surprising that energy productivity varies more than other single-factor productivity measures, since even though some fuels are dirtier than others, one might expect plants to use similar amounts of energy to produce a single unit of

[6]The greater dispersion of single factor productivity compared to TFP stems from cross-plant differences in factor intensities. For example, if one plant has a greater labor share than another plant due to lower local wages, the two plants may have the same TFP but different labor productivity. Differences across plants in factor prices e.g., wages generally affect single-factor productivity but not TFP (Syverson 2011).

output. Panel A of Table 3.1 shows that mean productivity for energy and labor are similar, so the difference in dispersion is not driven by scale effects.

Figure 3.1: Dispersion of Within-Industry 90-10 Productivity Measures

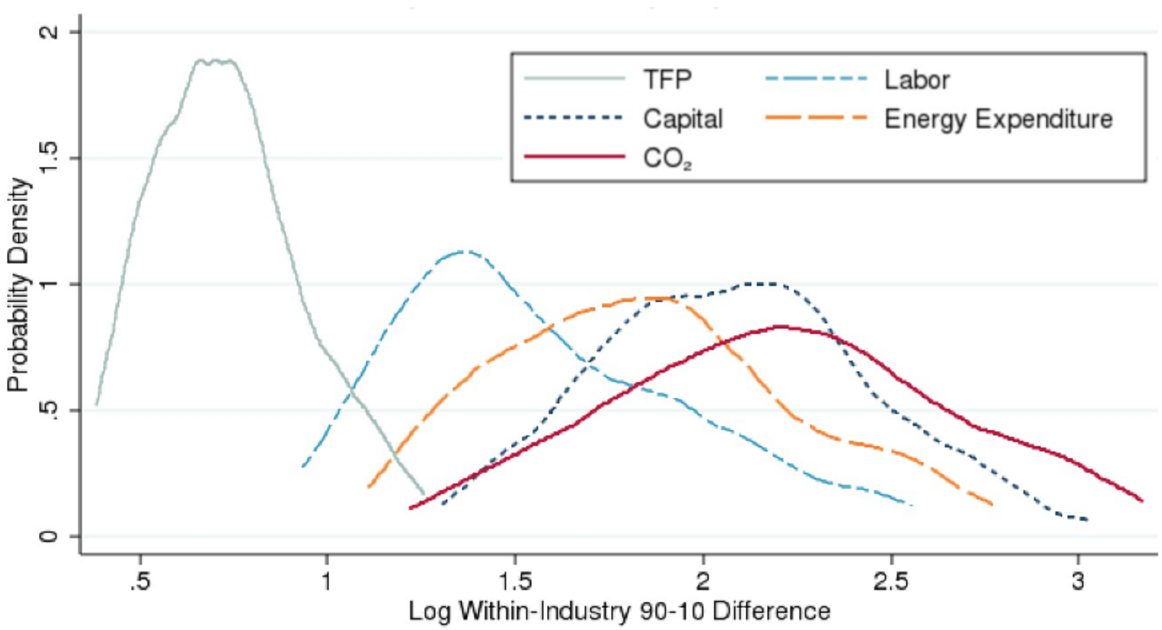

Note: Each kernel density plot was created using the approximately 375 6-digit NAICS dispersion measures for the corresponding productivity measure. Kernel densities have been censored at the 5th and 95th percentiles in accordance with U.S. Census Bureau disclosure avoidance.

Figure 3.1 demonstrates the paper's first two conclusions. This graph plots the distribution of industry-level 90–10 dispersion measures. Each of the roughly 375 observations underlying one of these lines is an industry; the value of each observation equals the within-industry 90–10 productivity ratio. The mean of the CO_2 distribution dark solid line lies above the mean of all other productivity dispersion measures, demonstrating that CO_2 dispersion for the average industry is greater than dispersion in the other productivity measures. The greater width of the CO_2 distribution relative to the TFP and labor distributions shows that within-industry dispersion in CO_2 productivity is more variable across industries than within-industry dispersion in TFP or labor productivity.

Implication for Carbon Tariffs

In many countries, policymakers have proposed import tariffs proportional to the carbon content of imported goods in order to guard against emissions leakage.[7] These are often referred to as carbon border adjustments or carbon tariffs.

[7]This type of policy was in the Waxman-Markey bill that passed the US House but not the Senate in 2009. In 2017, France, Mexico, and Canada discussed imposing one on the United States after the Trump

Table 3.2: Social Costs of Carbon Per Dollar of Output

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Industry-wide statistics						
Mean	0.019	0 .041	0.034	0.044	0.077	0.071
SD	0.048	0.050	0.032	0.062	0.110	0.075
Panel B. Within-industry 90-10 Differences in SCC/$						
Mean	0.060	0.060	0.051	0.089	0.142	0.120
SD	0.204	0.204	0.102	0.205	0.221	0.128
p90-10	0.523	0.523	0.261	0.526	0.567	0.329
Panel C. Within-industry standard deviation of SCC/$						
Mean	0.023	0.023	0.020	0.035	0.056	0.047
SD	0.080	0.080	0.040	0.080	0.086	0.050
p90-10	0.035	0.035	0.045	0.053	0.073	0.097
Direct source	CM	CM	MECS	CM	CM	MECS
Indirect source		BEA	BEA	CM	CM	CM
Leontief inverse		X	X		X	X

Notes: Panel A means and SD are computed from plant-level CM and MECS observations. Panel B statistics are calculated using the 375 within-industry 90–10 dispersion measures. Panel C statistics are calculated using the 375 within-industry standard deviation measures. Each column computes SCC per dollar of output using different inputs, as indicated in the column headings and table footers. A column represents either direct or total emissions, where direct emissions come from either the CM or MECS, and indirect emissions come from either the BEA I-O table or the CM Material trailer. See text for details.

Table 3.2 reports the level and distribution of the external cost of CO_2 emissions per dollar of output. If another country imposed a carbon tariff on imports from the United States, the social cost of carbon SCC per dollar output provides one measure of the relevant tariff. We assume a standard SCC of \$40 per metric ton of CO_2. Each column represents a different method of calculating CO_2 per dollar of output. Column 1 presents direct emissions from fuels plus electricity per dollar of output using CM data. Column 2 adds indirect emissions to the direct emissions estimates from column 1, where indirect emissions are calculated by inverting the industry-level input-output table.[8] Adding industry-level indirect emissions changes the mean externality panel A but not the within-industry dispersion panels B–C.

Administration announced it was withdrawing from the Paris Treaty on Climate Change. California has just implemented such a measure for government purchase decisions the Buy Clean California Act.

[8]Total emissions are the sum of direct and indirect emissions. Direct emissions come from plant-level data. Indirect emissions come from the input-output table. Note that the input-output table provides both direct and indirect emissions for an industry. We subtract the industry-level direct emissions from total industry-level emissions to get our measure of indirect emissions.

Column 3 shows the same direct plus indirect emissions estimates from column 2 but uses MECS rather than CM to measure direct plant-level emissions. While MECS is a smaller sample than CM, it contains plant-level information on the types of fossil fuels used. Column 4 replaces the industry-level indirect emissions estimates used in columns 2 and 3 with indirect emissions calculated using the CM Materials Trailer. Column 4 uses plant-level information on input purchases to calculate indirect emissions. For each input material, however, it only accounts for the industry average of direct emissions of that input material and not its indirect emissions. Column 5 is similar to column 4, but for each material input, it calculates total (not just direct) emissions of each input using the inverted input-output table. Column 6 is similar to column 5 but uses MECS to measure direct emissions.

Panel A of Table 3.2 shows that the external cost of CO_2 emissions for the mean plant is 2 to 8 percent of product value. Column 2 suggests that a uniform Pigouvian carbon tariff imposed on imports from US manufacturers should be around 4 percent. In the first row of panels B–C, column 2 shows that the mean industry has a 90–10 SCC difference of 0.06. This implies that even if a carbon tariff were imposed based on industry-specific means instead of the economy-wide 4 percent, many plants would have a carbon tariff which is well below the appropriate plant-level tax, whereas others would face a tax rate that is far too high. Comparing columns 3 and 5 of panel B shows that using plant-level records of intermediate good purchases from CM, rather than industry-level records from the input-output table, approximately doubles both the 90–10 and standard deviation measures of dispersion.

Figure 3.2 plots the distribution of industry-level 90–10 differences in SCC per dollar output. This shows the main conclusions from Table 3.2 visually. Many industries have high 90–10 differences, and this distribution of dispersions has a long right tail which is understated by our censoring at the fifth and ninety-fifth percentiles. Thus, a Pigouvian tax based on industry averages would still miss significant heterogeneity in true SCC per dollar of output. Detailed analyses of carbon tariffs have noted many challenges, ranging from legal ambiguity to information burdens. This paper uses plant-level data to highlight another trade-off—while a plant-specific tariff would impose a large information burden, an industry-level tariff would have substantial targeting errors stemming from firm heterogeneity.

3.4 Discussion and Conclusions

The records used for this paper are the most detailed data we are aware of that cover the entire US manufacturing sector. The plant-level granularity and detailed information on plant-level input purchases reveal significant heterogeneity in energy and CO_2 productivity, which exceeds heterogeneity for most other measures of single-factor and total factor productivity.

However, there are at least three reasons why our approach may understate the true extent of heterogeneity. First, we do not observe the full upstream set of plants that contribute to final output for a given plant in our data. Instead, we assign industry-level emission and

Figure 3.2: Dispersion in SCC Per Dollar Output

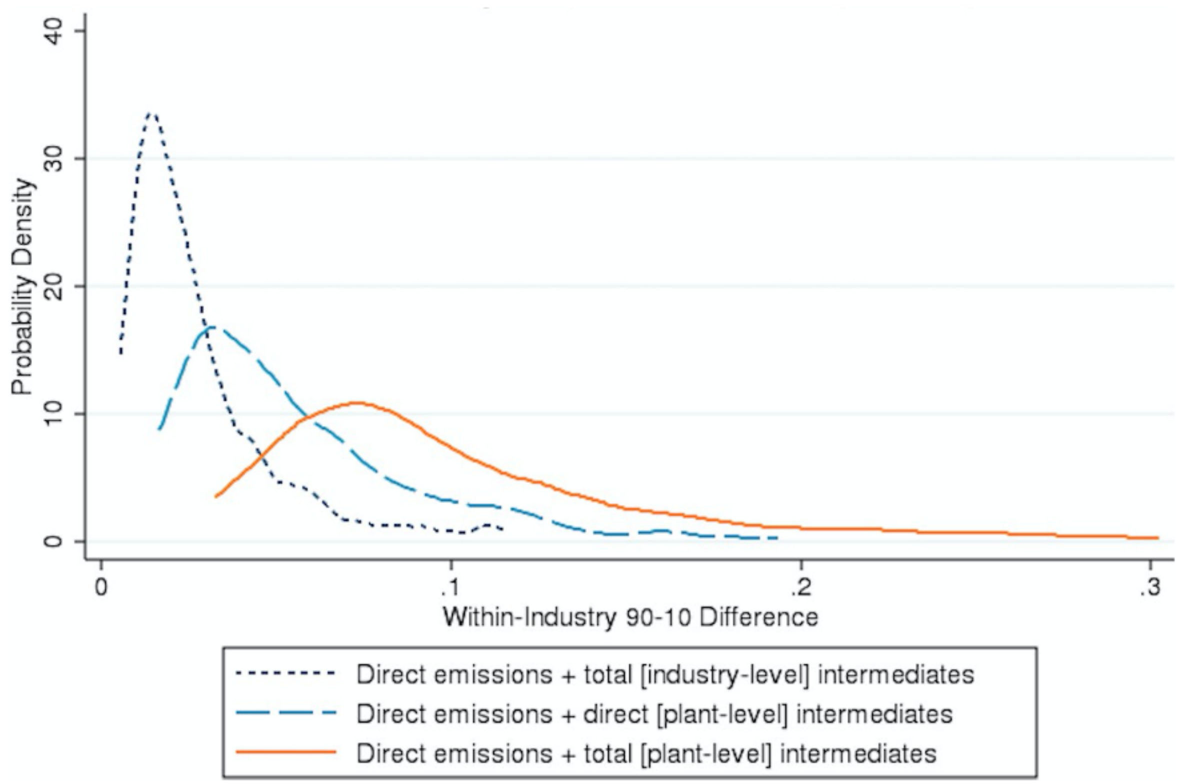

Note: Each kernel density plot was created using the approximately 375 6-digit NAICS dispersion measures for the corresponding emissions intensity measure. Densities are censored at the 5th and 95th percentiles.

energy intensities to construct our indirect emission and energy measures.[9] If supplying plants are significantly different in terms of emissions or CO_2 productivity, then we would understate heterogeneity. Second, our productivity estimates are based on revenues and not quantities. This should lead to underestimates of dispersion since more productive plants tend to have lower prices. Lastly, by excluding "administrative records" and other imputes from the CM, we are missing many of the smallest manufacturing establishments which might contribute to even more within-industry heterogeneity.[10]

How large are the welfare consequences of this heterogeneity for policies like technology standards or carbon tariffs that target industries and not plants? What are the economic reasons why energy productivity is more widely dispersed than labor or total factor productivity? How would decreasing factor misallocation across firms affect CO_2 emissions?

[9]Even when we observe plant-level input purchases, we only observe the industry of those inputs and not the specific plant.

[10]It is worth noting that while these reasons suggest we are understating true heterogeneity, any remaining measurement error after excluding imputed observations could lead to overstatement of true heterogeneity.

Finally, what does heterogeneity in CO_2 productivity imply about heterogeneity in marginal abatement costs? We leave these important questions for future work.

Bibliography

Aaronson, Daniel, Daniel Hartley, and Bhash Mazumder. 2019. "The Effects of the 1930s HOLC "Redlining" Maps." *Working Paper.*

Abaluck, Jason, Mauricio Caceres Bravo, Peter Hull, and Amanda Starc. 2021. "Mortality Effects and Choice Across Private Health Insurance Plans." *Quarterly Journal of Economics.*

Allcott, Hunt, Rebecca Diamond, Jean-Pierre Dubé, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell. 2019. "Food Deserts and the Causes of Nutritional Inequality." *Quarterly Journal of Economics* 134 (4): 1793–1844.

Allen, Treb, and Costas Arkolakis. 2021. *The Welfare Effects of Transportation Infrastructure Improvements.* Working Paper.

Andrews, M. J., L. Gill, T. Schank, and R. Upward. 2008. "High wage workers and low wage firms: negative assortative matching or limited mobility bias?" *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171 (3): 673–697.

Angrist, Joshua D., Peter D. Hull, Parag A. Pathak, and Christopher R. Walters. 2017. "Leveraging Lotteries for School Value-Added: Testing and Estimation." *The Quarterly Journal of Economics* 132, no. 2 (February): 871–919.

Auffhammer, Maximilian, and Edward Rubin. 2018. *Natural Gas Price Elasticities and Optimal Cost Recovery Under Consumer Heterogeneity: Evidence from 300 million natural gas bills.* Working Paper, Working Paper Series 24295. National Bureau of Economic Research.

Avenancio-León, Carlos, and Troup Howard. 2020. *The Assesment Gap: Racial Inequalities in Property Taxation.* Working Paper.

Badger, Emily, and Quoctrung Bui. 2019. "Cities Start to Question an American Ideal: A House With a Yard on Every Lot." *New York Times.*

Bailey, Martha, Hilary Hoynes, Maya Rossin-Slater, and Reed Walker. 2020. *Is the Social Safety Net a Long-Term Investment? Large-Scale Evidence from the Food Stamps Program.* Working Paper.

Barrington-Leigh, Chris, and Adam Millard-Ball. 2017. "More Connected Urban Roads Reduce US GHG Emissions." *Environmental Research Letters* 12 (4).

Baum-Snow, Nathaniel. 2007. "Did Highways Cause Suburbanization?" *Quarterly Journal of Economics* 122 (2): 775–805.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan. 2007. "A Unified Framework for Measuring Preferences for Schools and Neighborhoods." *Journal of Political Economy* 115 (4): 588–638.

Bayer, Patrick, Fernando Ferreira, and Stephen Ross. 2018. "What Drives Racial and Ethnic Differences in High-Cost Mortgages? The Role of High-Risk Lenders." *The Review of Financial Studies* 31 (1): 175–205.

Bednar, Dominic J., Tony Gerard Reames, and Gregory A. Keoleian. 2017. "The intersection of energy and justice: Modeling the spatial, racial/ethnic and socioeconomic patterns of urban residential heating consumption and efficiency in Detroit, Michigan." *Energy and Buildings* 143:25 –34.

Bond, Brittany, J. David Brown, Adela Luque, and Amy O'Hara. 2014. *The Nature of the Bias When Studying Only Linkable Person Records: Evidence from the American Community Survey.* CARRA Working Papers. Center for Economic Studies, U.S. Census Bureau.

Borenstein, Severin, and Jim Bushnell. 2019. "Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency." *Energy Institute Working Paper.*

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *Quarterly Journal of Economics* 131 (2): 633–686.

Card, David, Joerg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128:967–1015.

Card, David, Jesse Rothstein, and Moses Yi. 2021. *Location, Location, Location.* Working Paper.

Carley, Sanya, and David M. Konisky. 2020. "The Justice and Equity Implications of a Clean Energy Transition." *Nature Energy.*

Carlson, C., D. Burtraw, M. Cropper, and K.L. Palmer. 2000. "Sulfur Dioxide Control by Electric Utilities: What are the Gains from Trade?" *Journal of Political Economy* 108, number =.

Center for Climate and Energy Solutions. 2009. "American Clean Energy and Security Act of 2009."

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2014a. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates." *American Economic Review* 104 (9): 2593–2632.

———. 2014b. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review* 104 (9): 2633–79.

Chetty, Raj, and Nathaniel Hendren. 2018. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." *The Quarterly Journal of Economics* 133 (3): 1107–1162.

Christensen, Peter, Ignacio Sarmiento-Barbieri, and Christopher Timmins. 2020. "Discrimination and the Toxics Exposure Gap in the United States: Evidence from the Rental Market." *NBER,* no. 26805.

Christensen, Peter, and Christopher Timmins. 2021. *The Damages and Distortions from Discrimination in the Rental Housing Market.* NBER working paper 29049.

Colas, Mark, and John Morehouse. 2021. *The Environmental Cost of Land Use Restrictions.* Working Paper.

Cosbey, Aaron, Susanne Droege, Carolyn Fischer, and Clayton Munnings. 2019. "Developing Guidance for Implementing Border Carbon Adjustments: Lessons, Cautions, and Research Needs from the Literature." *Review of Environmental Economics and Policy* 13, no. 1 (January): 3–22. ISSN: 1750-6816.

Cronin, Julie Ann, Don Fullerton, and Steven E. Sexton. 2019. "Vertical and Horizontal Redistribution from a Carbon Tax and Rebate." *Journal of the Association of Environmental and Resource Economists* 6 (S1).

de la Roca, Jorge, and Diego Puga. 2018. "Learning by Working in Big Cities." *Review of Economic Studies* 84:106–142.

Diamond, Rebecca. 2016. *American Economic Review* 106 (3): 479–524.

Duranton, Gilles, and Diego Puga. 2020. "The Economics of Urban Density." *Journal of Economic Perspectives* 34 (3): 3–26.

Duranton, Gilles, and Matthew A. Turner. 2018. "Urban form and driving: Evidence from US cities." *Journal of Urban Economics* 108 (C): 170–191.

Eid, Jean, Henry G. Overman, Diego Puga, and Matthew Turner. 2008. "Fat City: Questioning the Relationship between Urban Sprawl and Obesity." *Journal of Urban Economics* 63 (2).

Energy Information Administration. 2020. *State Energy Data System.*

Fagjelbaum, Pablo D., and Cecile Gaubert. 2020. "Optimal Spatial Policies, Geography, and Sorting." *Quarterly Journal of Economics* 135 (2): 959–1036.

Federal Highway Administration. 2019. *National Household Travel Survey.*

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2020. "Place-Based Drivers of Mortality: Evidence from Migration." *American Economic Review* forthcoming.

———. 2016. "Sources of Geographic Variation in Health Care: Evidence from Patient Migration." *Quarterly Journal of Economics* 131 (4): 1681–1726.

Ganong, Peter, Damon Jones, Pascal Noel, Diana Farrell, Fiona Greig, and Chris Wheat. 2020. "Wealth,Race, and Consumption Smoothing of Typical Income Shocks." *NBER.*

Gaubert, Cecile, Patrick Kline, and Danny Yagan. 2019. "Place-Based Redistribution."

Gillingham, Kenneth. 2014. "Identifying the Elasticity of Driving: Evidence from a Gasoline Price Shock in California." *Regional Science and Urban Economics* 47 (4): 13–24.

Glaeser, Edward L., and Matthew E. Kahn. 2010. "The Greenness of Cities: Carbon Dioxide Emissions and Urban Development." *Journal of Urban Economics* 67 (3): 404–418.

Goldstein, Benjamin, Dimitrios Gounaridis, and Joshua P. Newell. 2020. "The Carbon Footprint of Household Energy Use in the United State." *Proceedings of the National Academy of Sciences* 117 (32): 19122–19130.

Goulder, Lawrence H., and Ian W. H. Parry. 2008. "Instrument Choice in Environmental Policy." *Review of Environmental Economics and Policy* 2, no. 2 (July): 152–174.

Gray, Wayne B., and Gilbert E. Metcalf. 2017. "Carbon tax competitiveness concerns: assessing a best practices carbon credit." *National Tax Journal* 70, no. 2 (June): 447–468.

Green, Tomas W., and Christopher R. Knittel. 2020. "Distributed Effects of Climate Policy: A Machine Learning Approach." *The Roosevelt Project Special Working Paper Series.*

Hardy, Bradley, Jonathan Morduch, William A. Darity Jr., and Darrik Hamilton. 2018. "Reinforcing Inequalities: Income Volatility and its Overlap with Wealth, Income, Race, and Ethnicity." *Working Paper.*

Hausman, Catherine, and Samuel Stolper. 2020. *Inequality, Information Failures, and Air Pollution.* NBER working paper 26682.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *Quarterly Journal of Economics* 135 (3): 1209–1318.

Hernandez, Diana, Yumiko Aratani, and Yang Jiang. 2014. *Energy Insecurity among Families with Children.* Technical report. National Center for Children in Poverty, Columbia University Mailman School of Public Health.

Homeland Security, U.S. Department of. 2021. *Homeland Infrastructure Foundation-Level Data.*

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India*." *The Quarterly Journal of Economics* 124, no. 4 (November): 1403–1448.

Hull, Peter. 2018. "Estimating Treatment Effects in Mover Designs."

Jacobsen, Mark R., Christopher R. Knittel, James M. Sallee, and Arthur A. van Benthem. 2020. "The Use of Regression Statistics to Analyze Imperfect Pricing Policies." *Journal of Political Economy* 128 (5): 1826–1876.

Jones, Christopher, and Daniel M. Kammen. 2014. "Spatial Distribution of U.S. Household Carbon Footprints Reveals Suburbanization Undermines Greenhouse Gas Benefits of Urban Population Density." *Environmental Science and Technology* 48:895–902.

Kline, Patrick, and Enrico Moretti. 2014. "People, Places and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs." *Annual Review of Economics* 6:629–662.

Kline, Patrick, Evan Rose, and Chris Walters. 2021. "Systemic Discrimination Among Large U.S. Employers." *Quarterly Journal of Economics.*

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. 2020. "Leave Out Estimation of Variance Components." *Econometrica* 88:1859–1898.

Ko, Yekang. 2013. "Urban Form and Residential Energy Use: A Review of Design Principles and Research Findings." *Journal of Planning Literature* 28 (4).

Kolstad, Charles, and Corbett Grainger. 2010. "Who Pays a Price on Carbon?" *Environmental and Resource Economics* 46 (3): 359–376.

Kontokosta, Constantine E., Vincent J. Reina, and Bartosz Bonczak. 2020. "Energy Cost Burdens for Low-Income and Minority Households." *Journal of the American Planning Association* 86 (1): 89–105.

Kortum, Samuel, and David Weisbach. 2017. "The design of border adjustments for carbon prices." *National Tax Journal* 70, no. 2 (June): 421–446.

Lachowska, Marta, Alexandre Mas, Raffaele D. Saggio, and Stephen A. Woodbury. 2020. *Do Firm Effects Drift? Evidence from Washington Administrative Data.* NBER Working Paper 26653.

Layne, Mary, Deborah Wagner, and Cynthia Rothhaas. 2014. *Estimating Record Linkage False Match Rate for the Person Identification Validation System.* CARRA Working Papers. Center for Economic Studies, U.S. Census Bureau.

NASA. 2020.

National Oceanic and Atmospheric Administration. 2020. *US Climate Division.*

Nowak, William, and Ian Savage. 2013. "The Cross Elasticity between Gasoline Prices and Transit Use: Evidence from Chicago." *Transport Policy* 55 (10): 885–910.

Peach, Dexter. 1983. "Siting Of Hazardous Waste Landfills and Their Correlation with Racial and Economic Status of Surrounding Communities." *US General Accounting Office.*

Pomponi, Franceso, Ruth Saint, Jay H. Arehart, Niaz Gharavi, and Bernardino D'Amico. 2021. "Decoupling Density from Tallness in Analysing the Life Cycle Greenhouse Gas Emissions of Cities." *npj Urban Sustainability* 1 (33).

Reames, Tony Gerard. 2016. "Targeting energy justice: Exploring spatial, racial/ethnic and socioeconomic disparities in urban residential heating energy efficiency." *Energy Policy* 97:549 –558.

Ribeiro, Haroldo V., Diego Rybski, and Jurgen P. Kropp. 2019. "Effects of changing population or density on urban carbon dioxide emissions." *Nature Communications* 10 (3204).

Rothstein, Richard. 2017. *The Color of Law: A Forgotten History of How Our Government Segregated America.* New York ; London: Liveright Publishing Corporation. ISBN: 978-1-63149-285-3. http://id.lib.harvard.edu/alma/990149136710203941/catalog.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. 2020. *IPUMS USA: Version 10.0 [dataset].* Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D010.V10.0.

Sallee, James M. 2019. *Pigou Creates Losers: On the Implausibility of Achieving Pareto Improvements from Efficiency-Enhancing Policies.* Working Paper, Working Paper Series 25831. National Bureau of Economic Research.

Schönholzer, David. 2021. *Measuring Preferences for Local Public Goods.* Working Paper.

Shammin, Md, Robert Herendeen, Michelle Hanson, and Eric Wilson. 2010. "A multivariate analysis of the energy intensity of sprawl versus compact living in the U.S. for 2003." *Ecological Economics* 69 (October): 2363–2373.

Shapiro, Joseph S., and Reed Walker. 2018. "Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade." *American Economic Review* 108, no. 12 (December): 3814–3854.

Song, Jae, David J. Price, Faith Guvenen, Nicholas Bloom, and Till Von Wachter. 2019. "Firming up Inequality." *The Quarterly Journal of Economics* 134 (1).

Spiller, Elisheba, Heather Stephens, Christopher Timmins, and Allison Smith. 2014. "The Effect of Gasoline Taxes and Public Transit Investments on Driving Patterns." *Environmental and Resource Economics* 59 (4): 633–657.

Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49, no. 2 (June): 326–365.

Tessum, Christopher W., Joshua S. Apte, Andrew L. Goodkind, Nicholas Z. Muller, Kimberley A. Mullins, David A. Paolella, Stephen Polasky, et al. 2019. "Inequity in consumption of goods and services adds to racial–ethnic disparities in air pollution exposure." *Proceedings of the National Academy of Sciences* 116 (13): 6001–6006.

The Phoenix Metro Area. 2020. Technical report. Maricopa Association of Governments.

Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditures." *Journal o Political Economy* 65 (5): 416–424.

Timmons, David, Nikolaos Zirogiannis, and Manuel Lutz. 2016. "Location matters: Population density and carbon emissions from residential building energy use in the United States." *Energy Research Social Science* 22 (December): 137–146.

Transportation Research Board. 2009. *Driving and the Built Environment: The Effects of Compat Development on Motorized Travel, Energy Use, and CO_2 Emissions.* Technical report. National Research Council of the National Academies.

Tsivanidis, Nick. 2019. *Evaluating the Impact of Urban Transit Infrastructure:Evidence from Bogota's TransMilenio.* Working Paper.

Ummel, Kevin. 2014. *Who Pollutes? A Household-Level Database of America's Greenhouse Gas Footprint.* Working Paper 381. Center for Global Development.

U.S. Department of Commerce, Bureau of the Census. 2014. *American Community Survey Design and Methodology.* Technical report.

US Department of Transportation, Federal Highway Administration. 2015. *2015 Status of the Nation's Highways, Bridges, and Transit: Conditions Performance.* Technical report.

U.S. Energy Information Administration. 2020. *Form EIA-861: Annual Electric Power Industry Report.*

U.S. Energy Information Administration. 2020. *U.S. energy-related carbon dioxide emissions fell in 2019, mainly in electric generation.* Technical report.

U.S. Environmental Protection Agency. 2018. *Emission Factors for Greenhouse Gas Inventories.* Technical report.

———. 2021a. *Emissions & Generation Resource Integrated Database.*

———. 2021b. *U.S. Transportation Sector Greenhouse Gas Emissions 1990-2019.* Technical report.

U.S. Environmental Protection Agency, Office of Energy Efficiency, and Renewable Energy. 2020.

Wagner, Deborah, and Mary Layne. 2014. *The Person Identification Validation System (PVS):Applying the Center for Administrative Records.* CARRA Working Papers. Center for Economic Studies, U.S. Census Bureau.

Appendix A

Appendix for Chapter 1

A.1 Data Appendix

Additional Details on Variable Construction

- **Missing and imputed variables:** I follow Chetty and Hendren (2018) and Bailey et al. (2020) and treat all imputed variables as missing, unless otherwise described. Dollar values are inflated to 2019 values using the CPI. Throughout the analysis I use demographic and household characteristics to control for selection on time-varying observables, I use work characteristics to construct commuting variables, and I use home characteristics in the second half of the paper to characterize places and study associations between built environment and place effects.

- **Flags**: In 2014 the ACS flags a lot of variables as "allocated" (to 0) if they checked a box indicating that they did not use natural gas or fuel use and then left the expenditure question blank. Because of this, I make an exception to the allocation flag and allow for residential energy to be allocated to 0 based on the checkbox question.

- **Work characteristics:** For each individual I have employment status, industry and occupation, place of work, weeks worked last year, and hours worked last week. I allow place of work tracts or more detailed geographies to be missing, but I drop observations if county of work is missing (unless the individual works from home, in which case I impute their place of work from their home, or if they are unemployed). I also allow current employment status to be missing if weeks worked last year and hours worked last week are not missing and not imputed. In 2008-2018, the weeks worked variable is binned; I follow Chetty and Hendren (2018) and assign the midpoint to all individuals in the bin. Since these variables are an input into my measure of commuting energy use, I use the midpoint from the bin for all years to keep the variable definition consistent.

- **Carpooling**: I divide CO_2 by the number of car-poolers for individuals who report carpooling.

- **Commuting Distance**: I estimate commute mileage using the GPS distance between reported home and place of work census blocks. To account for the fact that geodesic distances don't capture the indirect nature of roads, I rescale my mileage estimates to match the national average commuting distance reported in the NHTS (12 miles). For individuals who only report their county of work but not their census block of work, I impute miles travelled using reported commute time and average commute speeds for people with similar residence-job geographic pairs. I use a similar imputation for individuals for whom the travel speeds implied by dividing estimated miles by commute time are infeasible – over 80 mph on average in a car or motorcycle, and over 150 mph in a train.[1]

- **Number of annual commutes**: I estimate commuting days per week using reported hours worked last week and assuming people work 8 hours a day up to 5 days a week, assuming people worked 5 days if they worked 40-50 hours a week, 6 days if they worked 50-60 hours in a week, and 7 days if they worked more than that. I assume everyone commutes twice a day, and that commuting behavior is the same for all the weeks worked last year.

- **Vehicle fuel economy:** I assume individuals that commute by car or taxi do so in a vehicle with annual national average fuel economy, using data from the Federal Highway Administration (2019). For motorcycles, I scale mpg by 2 (US Department of Transportation 2015). This is a minor point as motorcycles account for only roughly 0.6% of vehicle miles driven (U.S. Environmental Protection Agency and Energy 2020). I also account for the fact that in general fuel economy is roughly 30% higher when driving on highways than in cities by adjusting mpg up by 19% relative to the national average for drivers whose average commuting speed is greater than 55 mph, and down by 9% relative to the national average for drivers whose average commuting speed is lower than 40 mph (U.S. Environmental Protection Agency 2021b).

- **Emissions from public transportation:** I assign zero emissions to commutes by public transportation, walking, or biking.

- **Identifying kids:** I designate a household member a child and drop them from the analysis sample if they are under the age of 18, or if they are identified as a child via the Census' relationship to householder code.

- **Building age**: I allow building age to be unknown in my analysis sample

Measurement Error in Household Carbon Emissions

There are several sources of measurement error in household carbon emissions from residential and transportation energy use. This could introduce bias in either estimates of household and place effects, estimates of the variance components, or both.

[1] This is the fastest speed a train ever goes in the US, along a small segment of the Northeast Corridor.

Note that if errors are random but serially correlated within a household, both a naive variance decomposition and a KSS variance decomposition on a sample consisting of stayers and movers will overstate the share of heterogeneity attributable to households; however, when I restrict to the mover only sample, the KSS correction accounts for serial correlation in the error term and gives unbiased estimates of variance components.

Below, I discuss sources of measurement error, as well as potential biases that arise in my estimates as a result.

Household reporting of residential energy expenditures:

Households may not accurately remember or report their energy expenditures. Inaccurate reporting could arise for example due to inattention to bills, or due to bias driven by the seasonality of energy expenditures – e.g. if household use their last monthly bill to proxy for annual expenditures.

If household inattention is fixed it will be absorbed by the household effect. If inattention leads high types to overstate their expenditures, and low types to understate their expenditures, this would lead to an upward bias in the household component of heterogeneity and vice versa. It is also reasonable to think that inattention may be random but serially correlated within household.

With fixed or random inattention, estimates of place effects themselves are unbiased. However, if moves are correlated with changes in attention, this could lead to bias in estimates of place effects. For example, if households move after positive income shocks, and higher income households pay less attention to their energy bills, *and* this inattention leads to systematic under- or over-estimation of expenditures, estimates of place effects with more inattentive residents would be biased.

Seasonality is unlikely to bias my estimates because surveys are sent out randomly, and therefore the season households were surveyed shouldn't be correlated with other components of the model.

Electricity prices:

I estimate electricity prices from total utility revenues divided by total utility customers, by county. This introduces three sources of measurement error in electricity prices.

First, in counties served by more than one utility, I cannot match customers to the actual utility they are served by. If customers in an area can select their residential energy provider, this could lead to bias in the household component of heterogeneity. For example, if higher type customers are selecting into lower average price utilities, I will underestimate the household component of heterogeneity. Similarly, if there are several utilities serving different neighborhoods within the same county, this could lead to bias in the place component of heterogeneity. In particular, I will over-estimate consumption in neighborhoods served by more expensive utilities, and under-estimate consumption in neighborhoods served by cheaper utilities. If more expensive utilities generally serve lower consumption neighborhoods, this will lead me to underestimate the place component of heterogeneity.

Second, residential customers generally face a two part tariff consisting of a fixed charge and a marginal volumetric charge, where the marginal price can either be increasing or decreasing in consumption depending on the utility. Because I am using average prices, calculated from utility residential revenues and quantities sold, I overestimate the average volumetric price and in turn underestimate consumption for everyone (more so for households in high fixed charge service territories). Moreover, for some utilities, marginal prices are either increasing or decreasing in consumption. When prices are increasing in consumption, I under-estimate prices faced by high-demand customers and over-estimate prices faced by low-use customers. This means I over-estimate quantities consumed by high-demand customers and under-estimate quantities consumed by low-demand customers, leading to an upward bias in my estimates of the household variance component. Conversely, if prices are decreasing in consumption, I underestimate the household variance component. Borenstein and Bushnell (2019) estimate that in the US, roughly 37% of customers face increasing block pricing, and roughly 21% face decreasing block pricing, though in all cases the rate structure is fairly narrow. They also estimate that across territories, utilities that utilize increasing-block pricing generally serve lower demand customers on average. Thus, my estimates likely somewhat over-estimate variation across households within utility territories, and underestimate variation across territories. Overall, unobserved rate structures should lead me to estimate a lower bound on place-based heterogeneity and estimate an upper bound on preference-based heterogeneity.

Finally, residential rates can vary within utilities, and I don't observe which rate a household has selected. This leads to the same biases as not being able to observe which utility a customer chooses, discussed above. Additionally, I do not observe if a household has solar, and in many states solar customers face different price schedules with significant subsidies for selling generated power back to the grid. This lowers their average price per kwh, causing me to underestimate quantity consumed and in turn CO_2 from electricity purchased from the grid by these customers.

Electricity carbon emissions factors:

I estimate carbon emissions intensity of electricity using average emissions factors at the NERC subregion. This does not capture the fact that electricity is generated from different fuels throughout the course of the day (e.g. solar peaks in the afternoon) and across seasons (e.g. there is less solar in the winter). The error in household carbon emissions that results from this is likely serially correlated within household, and can be accounted for in the mover-only KSS specification. However, if consumption profiles are also correlated with these patterns, my estimates of household carbon emissions will be biased. For example, if low type users consume more electricity when marginal emissions are higher then I would tend to under-estimate their carbon emissions and over-estimate the household component of heterogeneity.

Natural gas and other residential heating fuel prices and CO_2:

Many of the same price measurement errors arise with natural gas as with electricity, but generally individuals have less choice over their utility, fixed charges are larger, and there is less prevalence of block pricing. Unlike electricity, fuel emissions factors for other fuels are the same regardless of where a household lives. However, in the case of natural gas a significant source of emissions is upstream methane leaks, which I don't consider in this analysis.

Assignment of heating fuel:

I estimate carbon emissions from fuel use by assigning all expenditures on "other home heating fuels" to the fuel reported as the primary fuel. If a household has non-zero other fuel expenditures, but they don't list a primary fuel, I impute their primary fuel based on the most commonly used primary fuel among other survey respondents in their state and year (out of residual oil, propane, and wood). If households use more than one heating fuel, or use a heating fuel other than the one I imputed for them, there will be error in my measurement of carbon emissions, both as a result of dividing expenditures by the wrong fuel price, and as a result of assigning the wrong carbon emissions factor. I will overestimate household carbon emissions if reported or imputed fuel prices are lower than actual average fuel prices faced by the household, or if reported or imputed fuel types have higher emissions factors than the fuels actually used.

If I tend to overestimate carbon emissions from heating fuels for otherwise high-type households and underestimate carbon emissions from heating fuels for otherwise low-type households, then my household variance component will be biased upward, and vice versa. Moreover, if moves are correlated with shocks to unobserved fuel components, this could lead to bias in my estimates of place effects. For example, if a household uses the same heating fuel everywhere they live but doesn't report this fuel, if they move to a place where their neighbors use an on average higher emissions heating fuel, I would overestimate the place effect. In practice, the share of households reporting non-zero energy expenditures on heating other than electricity or natural gas is small, and my estimates are not meaningfully affected when I exclude other heating from the calculation.

Commuting Distances:

Because I estimate commute miles from geodesic distances between coordinates, I will underestimate speed and miles travelled for individuals who have less direct commutes. If the directness of a commute is the result of place-based constraints (e.g. the result of living in a gated community or a neighborhood with many winding roads and cul-de-sacs), and if these types of neighborhoods tend to be farther from employment centers and have longer commutes to begin with, then I will underestimate the place component of spatial heterogeneity.

Additionally, I impute miles for the people for whom I don't observe census block of work using average mph for home and place of work county pairs. This will lead me to overstate commute distances for people with slower than average commutes, and understate commute distance for people with faster than average commutes. If faster than average commutes are also longer than average, then I will underestimate the person component of spatial heterogeneity. My estimates are not sensitive to using a simpler measure of commute distance, calculated from simply dividing reported commute time by the average national commute speed, 32 mph (Federal Highway Administration 2019), suggesting that errors in commute speeds are unlikely to bias my estimates.

Total Commuting Miles:

I use weeks worked last year to estimate total commuting from typical commuting behavior last week. This assumes that hours worked are stable, that people work at the same place all year, and that information about commutes reported for last week is representative of commutes generally. Any deviations along these dimensions introduces measurement error into my outcome. It is likely that such errors are more likely to arise for lower income households with less job stability, but it is unlikely that it results in a systematic over- or under-estimate of commute miles on average.

Commuting Emissions:

I assume everyone drives a vehicle with the annual national average fuel economy, using data from the NHTS. This is a significant oversimplification, as it ignores patterns of heterogeneity in fuel economy both across commute lengths and across regions. If people with longer commutes drive more fuel efficient vehicles, I will overstate heterogeneity. On the other hand, if people who want to conserve on gas both buy more fuel efficient vehicles and choose to have shorter commutes, I will understate heterogeneity. The bias in my estimates of relative shares is more ambiguous. If these patterns are driven solely by individual preferences, I will over/understate the relative importance of the person component in spatial variation. On the other hand, if they are driven by local norms or place characteristics such as e.g. the availability of parking, I will over/understate the relative importance of the place effect.

Another source of error arises in the assignment of emissions to other modes of transit. In practice, most public transit in the US is not zero-emissions right now. Assigning zero emissions to public transit over-states heterogeneity across transit vs. car commuters.

Finally, if households change their mode of transit over the year, or if they use multiple modes of transit in a single commute, I do not capture this variation. For example, if households report taking public transit as their primary mode, but in reality they drive part of the distance of their commute, I will under-estimate their carbon emissions and overstate overall heterogeneity.

Non-Commuting Transportation Emissions:

I don't observe transportation other than commuting. In particular, I don't observe local travel for errands or leisure, nor do I observe airplane travel. Thus, I (weakly) underestimate carbon emissions magnitudes. If commuting is a rank-preserving share of total transportation emissions, my results will be qualitatively correct but off in magnitudes. However, if for example places with long commutes have lower other transportation emissions (because everybody spends leisure time in their back yard) whereas places with short commutes have higher other transportation emissions (because people go away for the weekend), then my estimates cannot be used to infer anything about transportation emissions overall.

A.2 The Leave-One-Out Connected Set

Consider the following data: Household 1 moves from NY to Chicago, household 2 moves

Individual & Household Geographic Locations

Year	Household	Place
1	1	NY
2	1	CHI
1	2	CHI
2	2	NY
1	3	SF
2	3	CHI

from Chicago to NY, and household 3 moves from San Francisco to Chicago. This data can be visualized as a network, where each place is a node, each household is a node, and edges connect households to each place they've lived in.

Household + Place Network

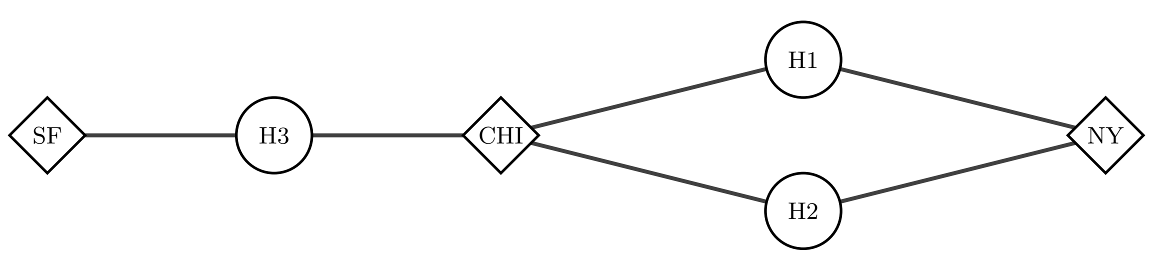

In this figure, San Francisco, Chicago, and New York are all connected by movers – this is a connected set. The leave-out connected set is the set of places that remains connected after dropping any household from the data. In this example, San Francisco is *not* in the leave-out connected set, because it is only connected to the rest of the network through H3.

A.3 Computational Appendix

For parsimony, I proceed in two steps, regressing $\log(CO_2)$ on observable characteristics and year fixed effects, and residualizing so that I am left with

$$\tilde{y}_{ij} = \alpha_i + \psi_j + \varepsilon_{it}$$

The share of overall variance attributable to place effects can then be captured by the variance component of place effects,

$$Var(\psi_j) \equiv \sigma^2_\psi = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}(\psi_{j(i,t)} - \bar{\psi})^2$$

and the covariance component between place effects and person effects

$$Cov(\alpha_i, \psi_j) \equiv \sigma^2_{\alpha,\psi} = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}(\psi_{j(i,t)} - \bar{\psi}) \cdot \alpha_i$$

KSS provides an estimate for the standard error $\psi_i^2 = Var(\varepsilon_i)$ based on a leave out estimate of σ_i^2:

$$\hat{\sigma}_i^2 = y_i(y_i - x_i'\hat{\beta}_{-i}) = y_i\frac{(y_i - x_i'\hat{\beta})}{1 - P_{ii}}$$

where $P_{ii} = x_i'(x_i x_i')^{-1}x_i$ is the observation leverage.

A.4 Additional Figures and Tables

Additional Figures

Figure A.1: CO_2 Profiles by Demographic Characteristics

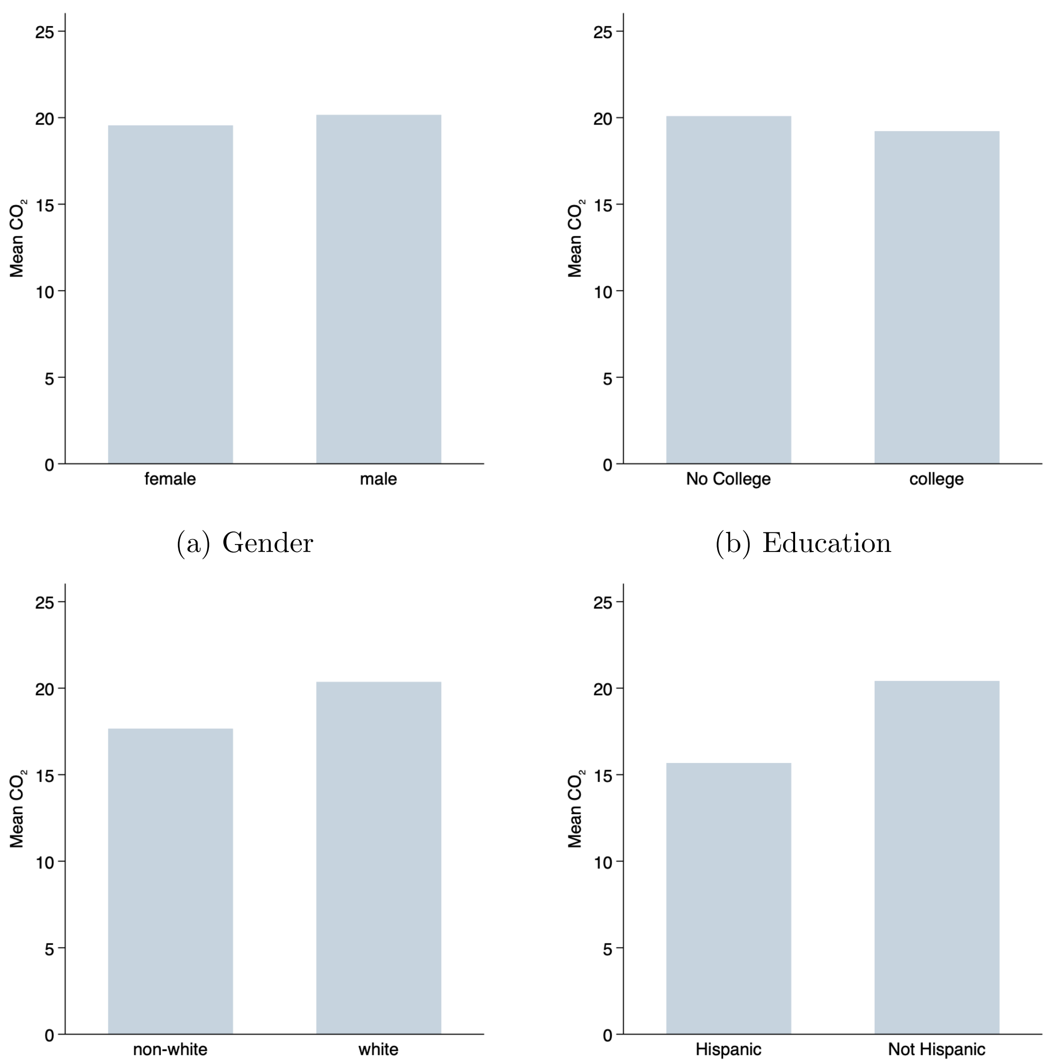

(a) Gender (b) Education

(c) Race (d) Ethnicity

Note: This figure shows variation in household carbon emissions by household member demographics. Panel (a) shows that households with more women (age 18+) have slightly lower emissions (consistent with women having fewer and shorter commutes). Panel (b) shows that college educated households have slightly lower emissions. Panel (c) and (d) show large differences by race and ethnicity – white households and non-Hispanic households have higher emissions on average than non-white and Hispanic households. All estimates reflect the full sample, pooled 2000-2019, weighted by Census sample weights.

Figure A.2: CO_2 Profiles by Age and Income

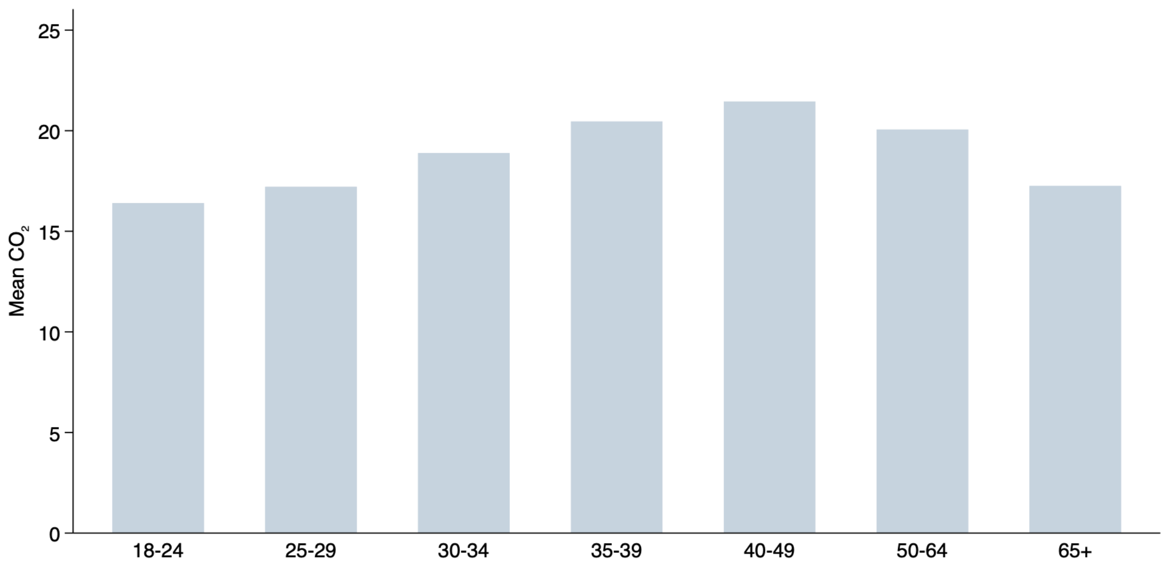

(a) Age

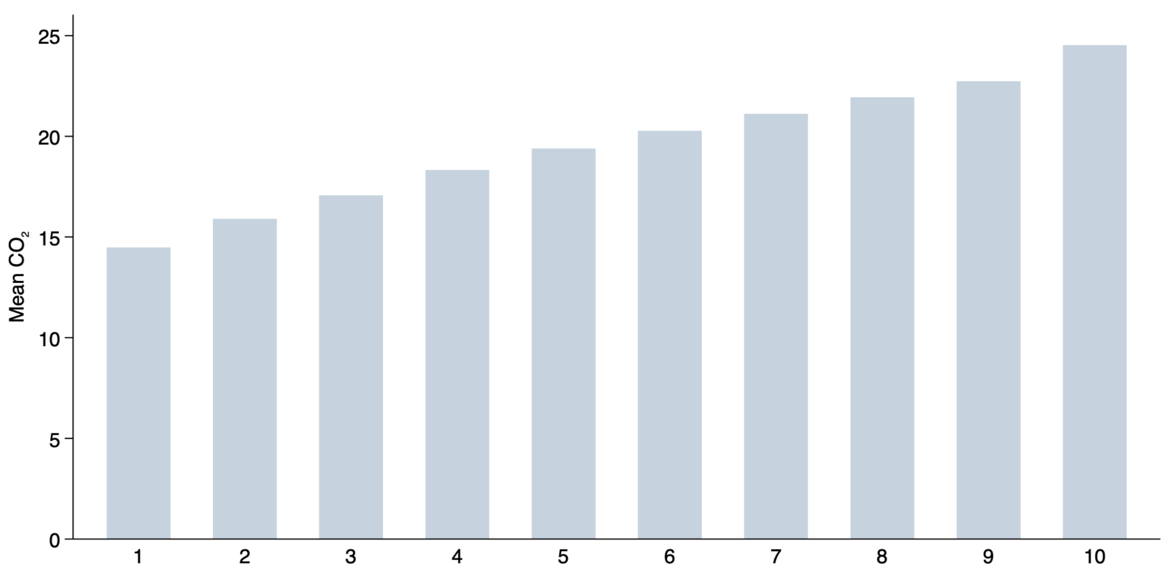

(b) Income Decile

Note: This figure shows variation in household carbon emissions by household member age and household income deciles. Panel (a) shows a non-linear relationship between the adult age of household members and mean carbon emissions which increases through people's 40s and then decreases again (likely reflecting a combination of higher incomes and children still being in the home). Panel (b) shows an increasing relationship between household income decile and carbon emissions. All estimates reflect the full sample, pooled 200-2019, weighted by Census sample weights. Household income is CPI-adjusted.

Figure A.3: CO_2 Profiles by Household Size

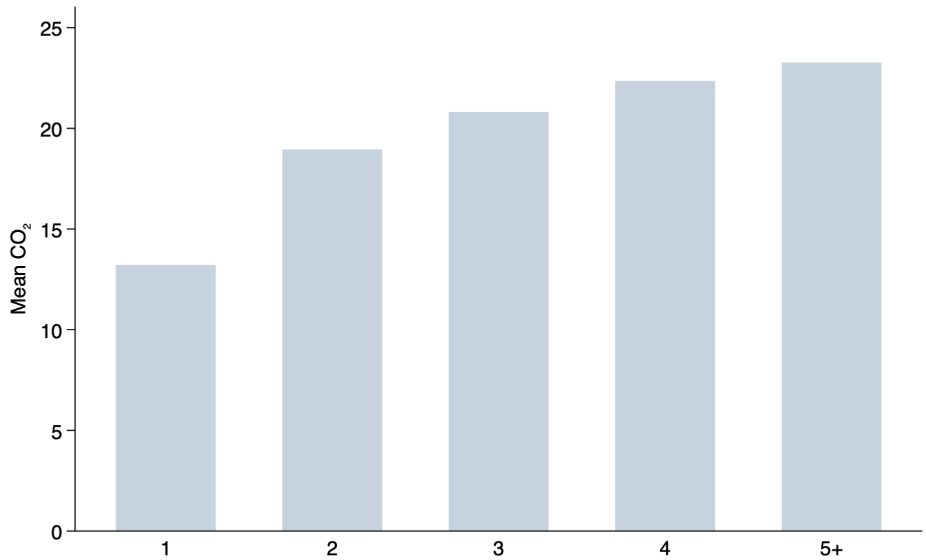

(a) Household Size

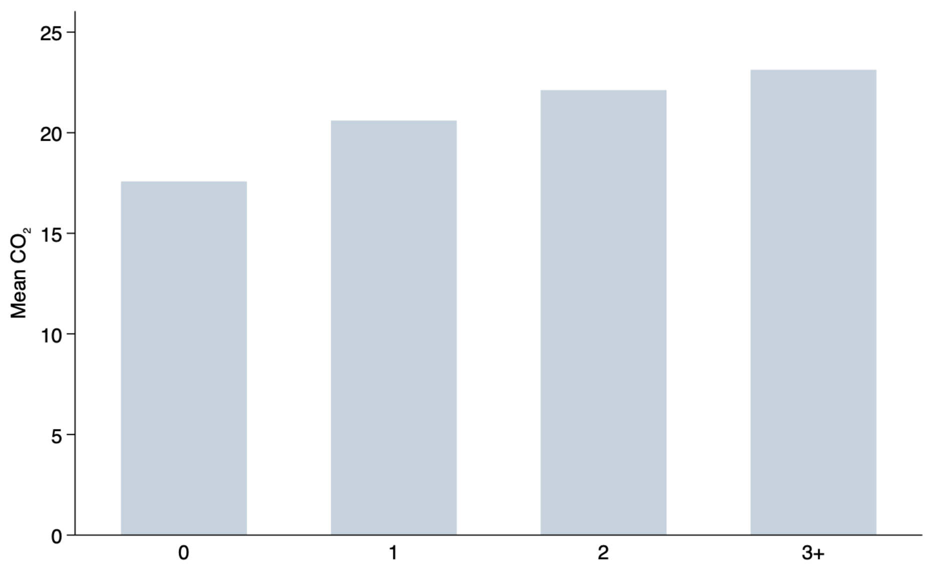

(b) Number of Kids

Note: This figure shows variation in household carbon emissions by household size (a) and number of children (b). Carbon emissions increase with household size and with the number of children, but less than proportionally, and the increase is fairly small going from 4 to 5+ people, or 2 to 3+ kids. All estimates reflect the full sample, pooled 200-2019, weighted by Census sample weights.

Figure A.4: Energy Expenditures in Mover Households in the PSID

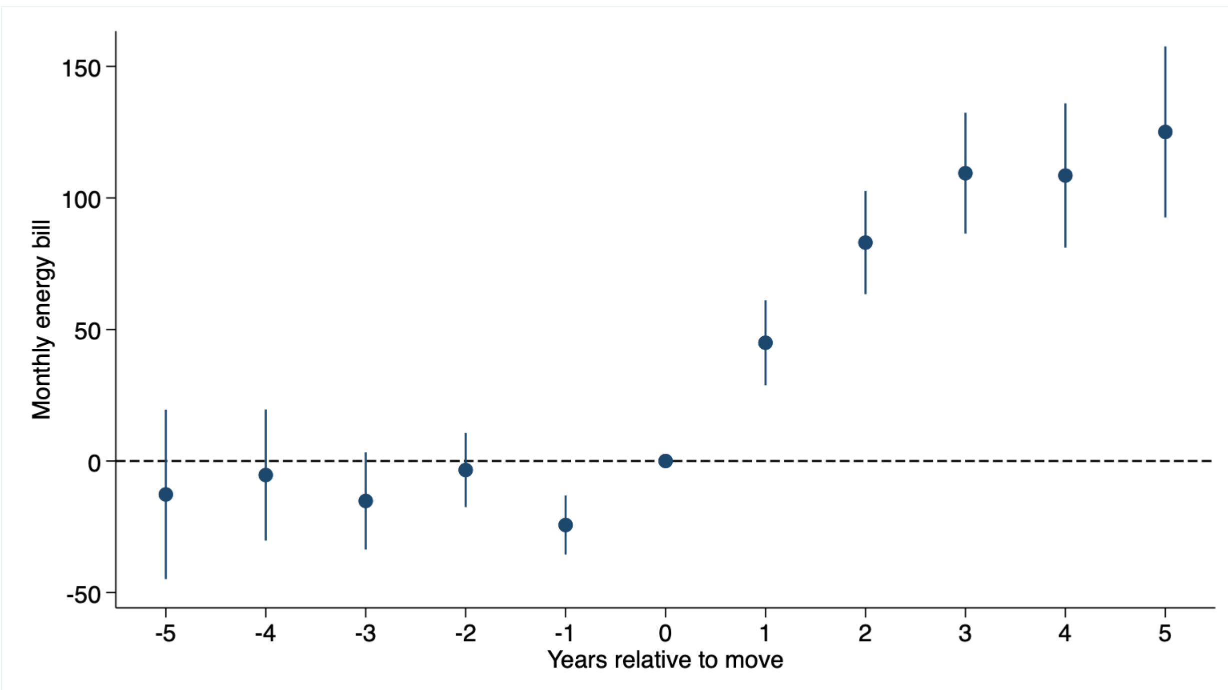

Note: I examine whether there are pre-trends in energy consumption for movers using data from the PSID, given data limitations in my baseline data. In particular, I test whether there are significant changes to monthly energy bills in the years prior to a move, after controlling for household characteristics such as income and household size. If anything, I find a slightly countervailing pre-trend for movers, with energy bills decreasing in the year before a move, and then increasing in the several years after (consistent with a secular trend of households moving to higher emissions places).

Figure A.5: Event study by duration – CBSA

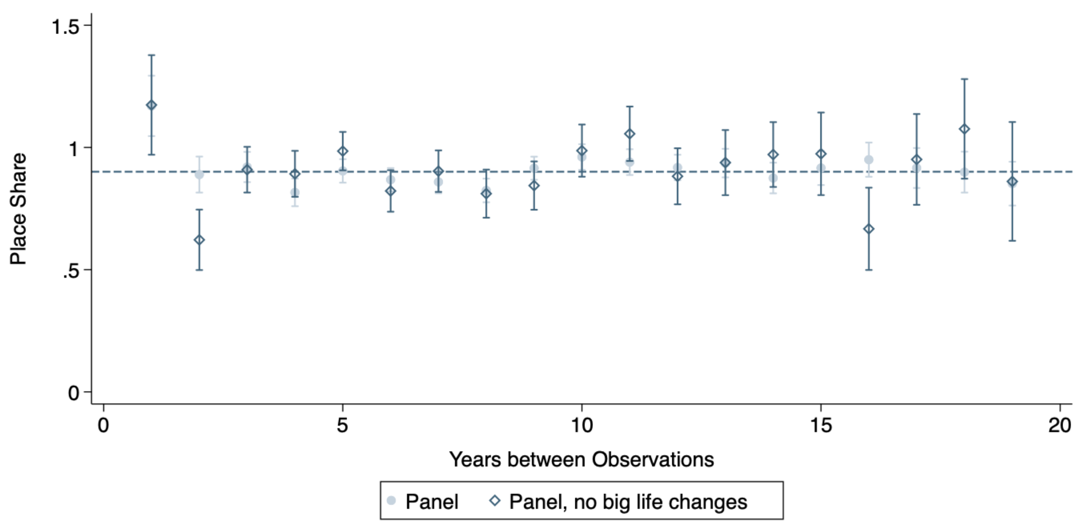

Note: This figure shows event study estimates of the share of spatial variation in mean carbon emissions that can be explained by place effects, by duration between mover observations. In other words, each coefficient is the estimate for place effects generated from the sub-sample of households that I observe X years apart. Coefficients plotted in light gray are estimated from the model using the full panel of stayers and movers. Coefficients plotted in the dark blue are estimated from the model using the sub-sample of stayers and movers with no changes in the number of children and less than 50% change in household income between observations. All estimates are weighted using Census sample weights.

Figure A.6: Place Effects vs. Household Carbon Emissions

(a) CBSA (b) Tract

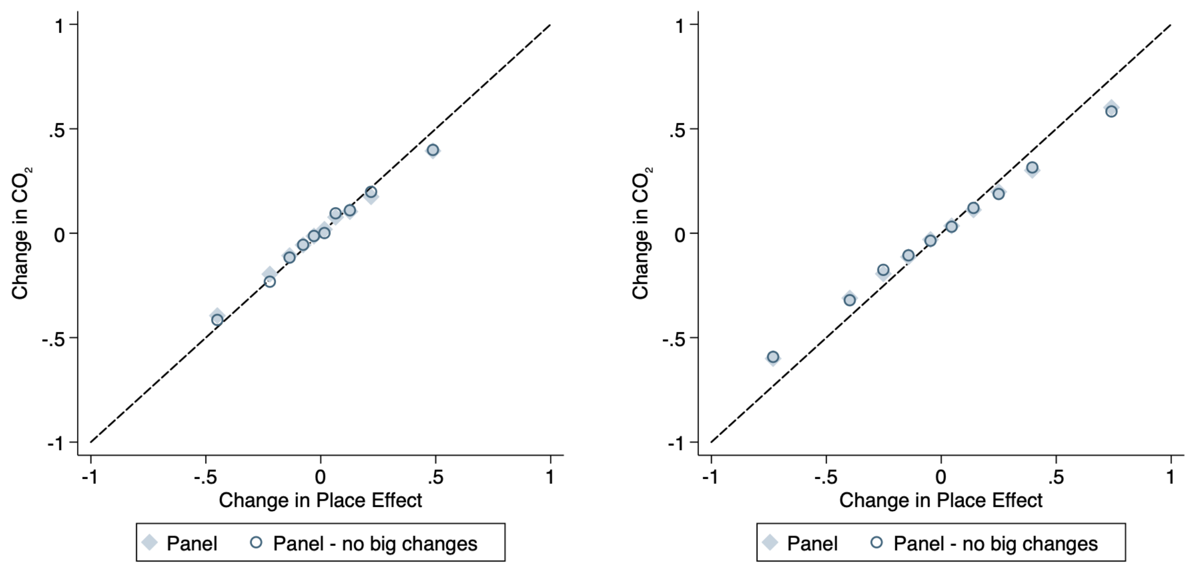

Note: This figure shows event study estimates of the share of spatial variation in mean carbon emissions that can be explained by place effects, by size of origin-destination differences in mean household carbon emissions for movers. To two sets of points compare the full sample of movers (solid diamond) to the sample of movers with no significant changes to income or number of children (empty circle). The dotted black line shows the 45line. All estimates are weighted using Census sample weights.

Additional Tables

Table A.1: Mean CO_2 – Movers vs. Stayers

	CBSA Panel				Tract Panel			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Moved	0.05***	0.00			0.08***	0.02***		
	(0.002)	(0.002)			(0.001)	(0.001)		
From			-0.11***	-0.04***			-0.07***	-0.03***
			(0.002)	(0.002)			(0.001)	(0.001)
To			-0.03***	-0.04***			-0.02***	-0.03***
			(0.001)	(0.001)			(0.001)	(0.001)
Cons.	2.85***	2.82***	2.86***	2.81***	2.85***	2.82***	2.88***	2.83***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
R^2 (adj.)	0.719	0.741	0.191	0.345	0.717	0.738	0.342	0.449
Controls	No	Yes	No	Yes	No	Yes	No	Yes

Note: This table compares household carbon emissions for movers and stayers. Columns (1)-(2) and (5)-(6) compare movers overall to stayers overall, with and without controls. Movers have higher carbon emissions than stayers, but the differences is smaller after controlling for differences in income and other demographic characteristics. Columns (3)-(4) and (7)-(8) present within-comparisons of stayers in a given place to movers from that place and movers to that place. The results here highlight that generally, movers have lower emissions than stayers, both at their origin and destination locations.

Table A.2: Probability of Moving

	CBSA	Tract
− N kids	0.007***	0.017***
	(0.0004)	(0.0007)
+ N kids	0.050***	0.150***
	(0.0005)	(0.0008)
Δ HH inc < -50%	0.035***	0.078***
	(0.0006)	(0.0010)
Δ HH inc > 50%	0.070***	0.154***
	(0.0005)	(0.0008)
Constant	0.044***	0.142***
	(0.0003)	(0.0004)
R^2 (adj.)	0.018	0.046
N	1,715,000	1,656,000

Note: This table shows that households with a change in the number of children or a larger than 50% (in absolute value) change in income are much more likely to move than stay. This is especially true of positive increases in both of these outcomes, and particularly for moves across neighborhoods.

Table A.3: Mover Origin and Destination Types

(a) CBSA Movers

	To Rural	To Suburban	To Urban	Total Share
From Rural	0.11	0.09	0.05	0.25
From Suburban	0.10	0.21	0.11	0.42
From Urban	0.06	0.15	0.12	0.33
Total Share	0.27	0.45	0.28	1.00

(b) Tract Movers

	To Rural	To Suburban	To Urban	Total Share
From Rural	0.09	0.07	0.03	0.19
From Suburban	0.08	0.28	0.08	0.44
From Urban	0.04	0.14	0.18	0.36
Total Share	0.21	0.49	0.29	1.00

Note: This table shows shares of origin-destination tract types for CBSA movers (panel (a)) and tract movers (panel (b)). Close to half of households move to suburban tracts. The most common type of move (among both CBSA and tract movers) is from a suburban tract to a suburban tract. Tract movers are less likely to move either from or to a rural neighborhood, in part because rural tracts are less likely to be in the leave-out connected tract set.

Table A.4: Place-Based Heterogeneity in CO_2 – Sensitivity to Outcome Definition

	CBSA				Tract			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Share $Var(\psi_j)$	0.188	0.242	0.195	0.197	0.257	0.270	0.262	0.238
Share $Var(\alpha_i)$	0.505	0.462	0.498	0.473	0.377	0.355	0.371	0.378
Share $2 \cdot Cov(\alpha_i, \psi_j)$	-0.001	-0.001	-0.001	0.006	-0.006	-0.000	-0.005	-0.002
R^2	0.69	0.70	0.69	0.68	0.62	0.62	0.62	0.61
$Var(\log CO_{2ij})$	0.29	0.35	0.29	0.20	0.28	0.35	0.29	0.19
Baseline	X				X			
Residential Only		X				X		
NHTS Commute			X				X	
NHTS Total mi.				X				X

Note: This table presents KSS decomposition estimates testing the sensitivity of my results to different outcome definitions. Columns (1) and (5) present baseline estimates again, to ease comparisons. Columns (2) and (6) present estimates using residential energy use only as the outcome. Results highlight that there is more heterogeneity overall in residential energy use than in commuting, and a larger share is attributable to place effects – 24% at the CBSA level and 27% at the tract level. Evidently, residential energy use drives more of the spatial heterogeneity across CBSAs than commuting, while the two sectors contribute in approximately equal parts at the tract level. In columns (3)-(4) and (7)-(8) I test the sensitivity of my results to changing my estimate of emissions from the transportation sector, and using the combined residential+transportation energy outcome. The "NHTS commute" approach uses a penalized Lasso regression to predict vehicle fuel economy from individual and household demographic characteristics (age, race, household size, household income, gender, number of vehicles, commute mode of transit, commute length) and geographic characteristics (CBSA, state, urbanity) and adjust carbon emissions from commuting for estimated fuel economy. The "NHTS Total miles" approach uses the same variables to predict total annual vehicle miles travelled. Taking these approaches decreases the overall variance in my outcome, perhaps evidence that households with longer commutes drive more fuel efficient vehicles and/or drive less for other purposes, but doesn't substantially change the place share of heterogeneity – the largest change is from using the total miles measure, which decreases the tract share of variance from 26% to 24%.

Table A.5: Place-Based Heterogeneity in CO_2 – No Bias Correction

	CBSA				Tract		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A: Panel Sample							
$Var(\log CO_{2ij})$	0.29	0.29	0.29	0.29	0.28	0.28	0.28
Share $Var(\psi_j)$	0.201	0.107	0.111	0.220	0.558	0.452	0.454
Share $Var(\alpha_i)$	0.567	0.564	0.564	0.562	0.765	0.759	0.759
Share $2 \cdot Cov(\alpha_i, \psi_j)$	-0.014	-0.008	-0.008	-0.015	-0.277	-0.262	-0.261
R^2	0.74	0.66	0.66	0.75	0.77	0.69	0.69
B: Mover Sample							
$Var(\log CO_{2ij})$	0.32	0.32	0.32		0.31	0.31	0.31
Share $Var(\psi_j)$	0.177	0.112	0.115		0.491	0.406	0.413
Share $Var(\alpha_i)$	0.505	0.503	0.503		0.588	0.584	0.584
Share $2 \cdot Cov(\alpha_i, \psi_j)$	0.001	0.004	0.004		-0.159	-0.152	-0.152
R^2	0.69	0.62	0.63		0.76	0.69	0.69
Amenities		X	X			X	X
Prices			X				X
TV-FE				X			

Note: This table reports results from the biased AKM estimation of variance components. All specifications include demographic and household controls as well as time fixed effects. Columns (1) and (5) report the baseline variance decompositions at the CBSA and tract levels. Columns (2) and (5) add controls for local mean heating degree days, cooling degree days, and electricity emissions factors (all in logs). Columns (3) and (6) additional control for a price index, constructed from lagged fuel shares interacted with national retail prices. Finally, column (4) computes time-varying CBSA place effects using 5-year windows (2000-2004, 2005-2009, 2010-2014, and 2015-2019), using stayer observations across time windows to identify time variation in place effects, while movers, as before, identify cross-sectional variation.

Table A.6: Place Correlates w/ Observable Characteristics of Urban Form

	Above Median Inc. (1)	White (2)	College (3)	Has Kids (4)
Suburban	0.02***	0.01***	-0.01***	0.02***
	(0.001)	(0.001)	(0.001)	(0.001)
Rural	-0.06***	0.10***	-0.06***	-0.02***
	(0.002)	(0.001)	(0.002)	(0.002)
Dist. to Closest City	-0.05***	0.02***	-0.02***	-0.01***
	(0.001)	(0.001)	(0.001)	(0.001)
Dist. to Largest City	0.07***	-0.04***	0.02***	0.02***
	(0.001)	(0.001)	(0.001)	(0.001)
Walk Score	-0.01***	0.01***	-0.01***	-0.00
	(0.001)	(0.000)	(0.000)	(0.001)
Bike Score	0.01***	-0.00***	0.00***	0.01***
	(0.001)	(0.001)	(0.001)	(0.001)
Transit Score	0.00***	-0.03***	0.00	0.01***
	(0.001)	(0.001)	(0.001)	(0.001)
Bus Routes	0.05***	0.03***	0.04***	-0.05***
	(0.005)	(0.003)	(0.004)	(0.005)
Rail Routes	0.20***	0.05***	0.10***	0.04***
	(0.004)	(0.003)	(0.003)	(0.004)
Tract Share Detached Homes	-0.38***	-0.17***	-0.26***	0.13***
	(0.005)	(0.004)	(0.005)	(0.005)
Tract Share Homeowners	0.16***	0.42***	-0.19***	-0.07***
	(0.008)	(0.006)	(0.007)	(0.008)
Tract Mean Cars/HH	0.38***	-0.13***	-0.34***	0.50***
	(0.008)	(0.006)	(0.008)	(0.009)
Tract Mean Rooms/House	1.15***	0.36***	1.00***	0.21***
	(0.006)	(0.004)	(0.005)	(0.006)
Block Density	6.69***	-6.47***	1.01***	4.17***
	(0.149)	(0.111)	(0.136)	(0.158)
Constant	-2.09***	0.18***	-1.09***	-0.51***
	(0.011)	(0.008)	(0.010)	(0.012)
R^2 (adj.)	0.09	0.06	0.04	0.02

Note: This table reports correlation coefficients between several demographic categories and a detailed vector of place characteristics.

Table A.7: Place Correlates w/ Observable Climate and Energy Supply Characteristics

	Above Median Inc. (1)	White (2)	College (3)	Has Kids (4)
Cooling Degree Days	-0.00***	-0.00***	0.04***	0.00
	(0.001)	(0.000)	(0.001)	(0.001)
Heating Degree Days	0.03***	0.10***	0.08***	-0.01***
	(0.001)	(0.001)	(0.001)	(0.001)
Electric Grid Intensity	-1.29***	0.82***	-2.55***	0.28***
	(0.011)	(0.008)	(0.010)	(0.012)
Constant	0.61***	-0.10***	-0.21***	0.53***
	(0.011)	(0.008)	(0.010)	(0.012)
R^2 (adj.)	0.01	0.06	0.04	0.001

Note: This table reports correlation coefficients between several demographic categories and a vector of exogenous place characteristics.

Table A.8: 10 most populous CBSAs (2020)

Rank	CBSA
1	New York-Newark, NY-NJ-CT-PA
2	Los Angeles-Long Beach, CA
3	Chicago-Naperville, IL-IN-WI
4	Dallas-Fort Worth, TX-OK
5	Houston-The Woodlands, TX
6	Washington-Baltimore-Arlington, DC-MD-VA-WV-PA
7	Philadelphia-Reading-Camden, PA-NJ-DE-MD
8	Miami-Port St. Lucie-Fort Lauderdale, FL
9	Atlanta-Athens Clarke County-Sandy Springs, GA-AL
10	Boston-Worcester-Providence, MA-RI-NH-CT

Appendix B

Appendix for Chapter 3

B.1 Data

Direct Energy and Emissions

The main text describes how we construct the analysis sample. Here we describe a few additional sample restrictions designed to limit measurement error. For each variable in the raw data, the final sample excludes observations that are more than 100 times larger than the 99th percentile value. We do not apply this rule to ratios, e.g., this restriction is applied to CO_2 and to output, where output is inventory adjusted, but not to CO_2 productivity. The final sample also excludes plants that report zero or have missing values for any of our variables,[1] and plants that do not report positive values for at least one material in the materials trailer. Finally, the sample excludes establishments that are unique in their industry after all the above restrictions, since we cannot compute 90-10 dispersions or standard deviations for these industries.

We calculate emissions from fuel use by multiplying each establishment's consumption by fuel-specific emissions factors. We assign these emissions factors using data from the EPA when possible and from the EIA otherwise. We treat acetylene, hydrogen, and waste and byproduct gases as zero emissions. For emissions from electricity, we assign CO_2 per MWh using the EPA's eGRID dataset. We match eGRID regions to counties and compute emissions from electricity at the establishment level by multiplying each establishment's electricity consumption with the corresponding emissions factor from eGRID. In cases where a county overlaps with several eGRID regions, we take an unweighted mean of emissions intensities across the relevant eGRID regions. For observations in the CM that are missing the county variable, we take the unweighted mean of emissions factors across counties within the state. We do not account for process emissions.

[1] In cases where electricity kWh variables are missing in MECS but not CM, we use CM values to calculate total emissions in MECS.

Indirect Emissions

We use the BEA's 2007 benchmark Make Table, Use Table, and Import Table to construct an industry-level input-output (I-O) table. The BEA tables distinguish between industries and commodities to reflect the fact that some industries produce commodities other than the primary product of that industry (known as secondary commodities). We use tables after redefinitions, which in certain cases reallocate secondary commodity outputs to the industry in which they are the primary product, because this makes industries more homogeneous.[2] In practice most I-O codes in the benchmark analysis represent both a commodity and an industry. Exceptions to this are four commodities which are not industries (scrap goods, non-comparable imports, used and second-hand goods, and rest-of-world adjustment), and nine industries corresponding to different types of government enterprises. In cases where a government industry has an analog in private industry – for example federal electric utilities – the BEA assigns both the public and private industries' commodity outputs to the private industry's commodity code. The make table is an industry-by-commodity table, with each element mij represent- ing industry i's output of commodity j, in nominal dollars. The use (and import) tables are commodity-by-industry tables, with each element uij representing the total (imported) amount of commodity i used in industry j's production, also in nominal dollars. In addition to the commodity-by-industry pairs, the use table contains three value added rows (com- pensation of employees; taxes on production and imports less subsidies; and gross operating surplus) and 20 final demand columns. These additional rows and columns play an impor- tant role in ensuring that total inputs equal total outputs, but they are not rows or columns of the final I-O table. The use and import tables are available from the BEA at producer values and purchaser values – we use producer values throughout to maintain consistency with the make table. We construct a domestic use table by subtracting import values from the use table.

The BEA combines crude oil and natural gas extraction into one industry (code 211000). We split this industry into two, in order to treat oil and natural gas extraction separately. We assign all of the petroleum refineries commodity produced by the original industry to the new crude oil industry, and we assign all of the industrial gas manufacturing commodity produced by the original industry to the new natural gas industry. The rest of commodity output is assigned such that total production of gas and crude oil are proportional to their overall production according to the EIA. We assume that the commodity input mix for each of the two new industries is the same, with levels proportional to industry output, and we maintain oil and gas extraction as one commodity.

We normalize elements of the make table by commodity totals to generate a "market shares" table, in which each element sij is the proportion of commodity j produced by industry i. Analogously, we normalize elements of the domestic use table by industry totals to generate a direct requirements table, in which each element dij is the proportion of industry j's production made up by commodity i. Because we are interested only in combustible fuel

[2]The BEA reallocates secondary output from an industry to the industry in which it is the primary product when the two industries' input structures differ significantly.

use, we adjust direct requirements values by proportions of fuel used for combustion using EIA values.[3]

We generate the industry level I-O matrix by multiplying the market share matrix by the direct requirements matrix. The elements of this matrix are how much of each input an industry uses to produce one dollar of output. Thus equilibrium is defined by:

$$X = AX + Y$$

where X is an industry-length vector of gross production, Y is an industry-length vector of final demand, and A is the I-O matrix. We can rearrange to get

$$X = (I - A)^{-1} Y$$

$(I - A)^{-1}$ is referred to as the Leontief inverse. Using the Leontief inverse, we can calculate how much output is necessary in total from every industry to meet a given vector of final demand.

Thus, we calculate total emissions embedded in the production necessary to meet a unit of demand for goods from a given industry by left multiplying the Leontief Inverse by a row vector of the raw emissions intensities for coal, crude oil, and natural gas, which we get from the EPA. Since we are using CM data to calculate a more granular measure of direct emissions from production, we calculate indirect embedded emissions by subtracting emissions from the direct requirements from the total emissions:

$$\text{IndirectEmissions}_j = \text{TotalEmissions}_j - \sum_i (\text{DirectEmissions}_j \times \text{InputOutput}_{ji})$$

where the direct emissions vector is calculated from the total emissions vector, resetting all values to 0 except those corresponding to utilities and fuel industries. After creating the BEA-level emissions intensities, we convert from BEA industry defini- tions to NAICS industry definitions using the concordance provided by the BEA. If multiple BEA industries correspond to a single NAICS industry, we take BEA output-weighted means to calculate a unique NAICS industry value. If a BEA industry gets split into multiple NAICS industries, all NAICS industries get the same value. There are several BEA industries that don't have corresponding NAICS codes—importantly, the BEA considers government util- ities and private utilities separately, and only the private utility gets mapped to a NAICS utility code.

We use the indirect emissions calculated from the BEA to account for the full embedded emissions of production in two ways. One is through addition of the intermediate emissions intensities, by industry, to direct emissions intensities from CM. The second uses the CM Materials Trailer, which identifies material inputs into production by establishment. We use the BEA emissions intensity values to calculate the direct and the total emissions embedded in material inputs. The direct emissions capture the industry averages for emissions from

[3]These are calculated as a proportion of first use energy consumption and not total energy consumption.

fuel and electricity use in the production of materials. The indirect emissions use the full Leontief inverse to capture all emissions generated throughout the economy in the production of the materials, on average by industry. We add these to CM emissions intensities to calculate two versions of total emissions productivity based on material inputs. In compiling data from the CM Materials Trailer, we assign zero emissions to unspecified materials inputs (the "other industry" category). The fact that these "other industry" inputs represent a reasonable share of all inputs provides another reason why our estimates understate true dispersion in energy and CO_2 productivity.

B.2 Additional Tables

Table B.1: Descriptive Statistics of Industry-Level Characteristics – Value Added

	Direct Energy (1)	CO_2 [CM] (2)	CO_2 [MECS] (3)	Labor (4)	Capital Stock (5)	Materials (6)
Panel A. Industry-wide statistics						
Mean	3.63	7.89	8.25	3.95	0.47	0.37
SD	1.07	1.28	1.30	0.77	0.98	1.00
Panel B. Within-industry 90-10 Differences in Productivity						
Mean	2.20	2.53	2.47	1.76	2.43	2.33
SD	0.54	0.59	1.28	0.56	0.53	0.76
p90-10	1.39	1.51	3.27	1.43	1.37	1.95
Panel C. Within-industry standard deviation of Productivity						
Mean	0.86	0.99	0.96	0.68	0.95	0.91
SD	0.21	0.23	0.50	0.22	0.21	0.30
p90-10	0.50	0.53	1.28	0.51	0.44	0.63

Notes: Panel A means and SD are computed from plant-level CM and MECS observations. Panel B statistics are calculated using the 375 within-industry 90-10 dispersion measures. Panel C statistics are calculated using the 375 within-industry standard deviation measures. See text for details.

Table B.2: Social Cost of Carbon per Dollar of Value Added

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Summary Stats, CM						
Mean	0.06	0.13	0.11	0.33	0.74	0.76
SD	2.72	2.72	0.31	4.48	8.78	6.51
Panel B. Within-industry 90-10 Differences in Productivity						
Mean	0.55	0.55	0.24	3.44	7.70	3.50
SD	4.22	4.22	1.49	28.37	65.82	16.33
p90-10	10.81	10.81	3.82	72.73	168.77	41.88
Panel C. Within-industry standard deviation of Productivity						
Mean	0.21	0.21	0.10	1.34	3.0	1.37
SD	1.64	1.64	0.58	11.06	25.67	6.37
p90-10	0.18	0.17	0.14	1.37	3.31	1.63
Direct Source	CM	CM	MECS	CM	CM	MECS
Indirect Source		BEA	BEA	CM	CM	CM
Leontief Inverse		X	X		X	X

Notes: Panel A means and SD are computed from plant-level CM and MECS observations. Panel B statistics are calculated using the 375 within-industry 90-10 dispersion measures. Panel C statistics are calculated using the 375 within-industry standard deviation measures. Each column computes SCC per dollar of output using different inputs, as indicated in the column headings and table footers. A column represents either direct or total emissions, where direct emissions come from either the CM or MECS, and indirect emissions come from either the BEA I-O table or the CM Material trailer. See text for details.