

The environmental cost of land-use restrictions

MARK COLAS

Department of Economics, University of Oregon

JOHN M. MOREHOUSE

Department of Economics, University of Oregon

Cities with cleaner power plants and lower energy demand tend also to have tighter land-use restrictions; these restrictions increase housing prices and reduce the incentive for households to live in these lower greenhouse gas-emitting cities. We use a spatial equilibrium model to quantify the overall effects of land-use restrictions on the levels and spatial distribution of household carbon emissions. Our model features heterogeneous households, cities that vary in both their power plant technologies and their utility benefits of energy usage, as well as endogenous wages and rents. Relaxation of the current land-use restrictions in California to the level faced by the median urban household in the US leads to a 0.6% drop in national household carbon emissions and a decrease in the social cost of carbon of \$310 million annually.

KEYWORDS. Greenhouse gases, local labor markets, spatial equilibrium.

JEL CLASSIFICATION. Q4, R13, R31.

1. INTRODUCTION

Higher levels of carbon dioxide (CO₂, hereafter “carbon”) are associated with a multitude of global environmental issues. The amount of energy a household uses and, therefore, the amount of carbon emissions a household is responsible for, depends partially on where the household lives.¹ For example, Oklahoma City has high summer temperatures and relies heavily on coal-fired power plants while San Francisco has a moderate climate and uses electricity produced largely by hydroelectric plants. Households in Oklahoma City therefore consume large quantities of air conditioning using electricity generated by carbon-intensive power plants. Conversely, households in San Francisco use

Mark Colas: mcolas@uoregon.edu

John M. Morehouse: jmorehou@uoregon.edu

This paper was previously circulated as “Household Carbon Emissions and Land Use Regulations in a Spatial Equilibrium.” We would like to thank Trudy Cameron, Sebastian Findeisen, Andreas Gulyas, Joel Han, Kevin Hutchinson, Shoya Ishimaru, Jesús Fernández-Huertas Moraga, Francesc Ortega, Ed Rubin, Kegen Tan, Martin Watzinger, Woan Foong Wong, Nate Yoder, and Eric Zou for their input. Additionally, we would like to thank Josh Drukenbrod at the US Environmental Protection Agency for help with the National Emissions Inventory. We thank Amy Tran for excellent proofreading. This work benefited from access to the University of Oregon high performance computer, Talapas.

¹For example, Glaeser and Kahn (2010) showed that the median household in Memphis is responsible for nearly twice the amount of carbon emissions as the median household in San Francisco.

less electricity generated from more carbon-efficient power plants. As a result of these differences in electrical power plant technologies and climate, government policies that shape the distribution of households across cities may have important implications for national carbon emissions.

Local land-use restrictions that limit urban population density are often employed to improve the “greenness” of a city. However, these restrictions potentially limit population growth in many of the most desirable cities (Glaeser, Gyourko, and Saks (2005)). Additionally, Glaeser and Kahn (2010) document that cities with lower average household carbon emissions also have tighter land-use restrictions, suggesting that these restrictions may discourage people from living in lower carbon-emitting cities. The goal of this paper is to quantify the effect of local land-use restrictions on national household carbon emissions.²

To this end, we specify and estimate a spatial equilibrium model wherein power plant technologies and energy demand vary across cities. Heterogeneous and imperfectly mobile households choose which city to live in and how much housing, electricity, natural gas, and fuel oil to consume. Rents and wages are determined in equilibrium by the location and consumption choices of these households. Furthermore, cities vary in the tightness of local land-use restrictions. All else equal, tighter land-use restrictions imply higher costs of living, which will reduce the incentives for households to live in these cities.

Our model allows a household’s carbon emissions to vary across cities for two main reasons. First, to reflect climate differences, we allow the utility benefits derived from the use of electricity, natural gas, and fuel oil to vary by city. Second, due to spatial variation in the technologies employed in power plants, the carbon intensity of electricity production varies across cities. Land-use restrictions are tighter in cities with more carbon-efficient power plants and lower energy demand. Thus, households will be incentivized to live in “dirtier” cities.

For our analysis, we combine data from three main sources. We use the US Census and American Community Survey (ACS) for data on households’ city of residence, income, and rents, as well as expenditures on electricity, natural gas, and fuel oil for a large sample of households.³ We combine these household expenditure data with state-level energy prices from the Energy Information Association (EIA) to calculate implied household usage of each of these three energy types. Next, we use data from the Emissions & Generation Resource Integrated Database (eGrid) for power plant locations, output, and CO₂ emissions.

Following Glaeser and Kahn (2010), we use these data to document how the amount of carbon a household emits depends on the city in which they reside. One issue is that differences in emissions across cities may, to some extent, reflect sorting of households

²We take the relationship between land-use restrictions and carbon emissions as given and quantify the effect of land-use regulations on national carbon emissions. We leave it to future work to determine *why* carbon-efficient cities generally have tighter land-use restrictions.

³Throughout the paper, we use the 70 largest core based statistical areas as our definition of cities. The rest of the United States is aggregated up to the census division level.

with different propensities to use energy. Therefore, the observed differences in average emissions across cities may not reflect the direct effect of location on household emissions. In response to this issue, we employ the semi-parametric selection correction approach introduced by [Dahl \(2002\)](#) to estimate selection-corrected predicted energy usage and carbon emissions associated with living in various US cities.

We document substantial variation in carbon emissions across cities. For example, we estimate that if a household resides in Memphis, they would produce three times the annual carbon emissions had they resided in Honolulu. Additionally, we show that household carbon emissions are negatively correlated with a standard measure of land-use restrictions, the Wharton Land-Use Regulation Index—greener cities tend to have tighter land-use restrictions.⁴

Next, we combine these estimates of predicted energy usage with our other data sources to estimate our spatial equilibrium model. We estimate the parameters of household utility using the two-step estimator introduced by [Berry, Levinsohn, and Pakes \(2004\)](#) with data on household locations from the Census and ACS and our estimates of predicted energy usage.⁵ We use data from eGrid to estimate the carbon emissions associated with electricity production across regions and estimates from [Saiz \(2010\)](#) to calibrate the parameters of housing supply curves as a function of land-use restrictions.

California legislators recently voted down SB-50, a bill that would have relaxed local land-use restrictions in California cities.⁶ We use our estimated model to simulate the effects of such a policy. Specifically, we set land-use restrictions in California cities to the level faced by the median urban household in the United States. Due to the moderate climate and carbon-efficient power plants, California cities are associated with remarkably low carbon emissions. However, land-use restrictions are tight—the San Francisco CBSA is in the 86th percentile of the Wharton Index while Los Angeles CBSA is the 78th.⁷

As a result of relaxing these restrictions, we find that the long-run population in California cities increases. As demands for natural gas and electricity are lower in California, national usage of natural gas and electricity drop by 0.3% and 0.5%, respectively. Overall, this leads to a 0.6% decrease in national household carbon emissions, associated with a decrease in the social cost of carbon of \$310 million annually.⁸ This change is driven by a decrease in energy usage and an increase in the proportion of total electricity consumption coming from cleaner power plants in California. Furthermore, given that California

⁴The Wharton Land-Use Regulation Index was created from a survey sent to 6896 municipalities across the US, with questions that range from how many regulatory boards one must clear before construction to city-specific density and open space requirements.

⁵This approach has been utilized in a spatial setting by, for example, [Diamond \(2016\)](#), [Piyapromdee \(2021\)](#), and [Colas and Hutchinson \(2021\)](#).

⁶SB-50 was referred to as the “More HOMES Act” (Housing, Opportunity, Mobility, Equity, and Stability). The bill focused on the relaxation of density restrictions and reducing the number of areas zoned for single-family homes, particularly in areas near public transit and in commercial areas.

⁷[Quigley and Raphael \(2005\)](#) wrote, “California represents the most extreme example of autarky in land-use regulations of any U.S. state.”

⁸For the social cost of carbon calculations, we use the estimate of the social cost of carbon in the year 2020 from [Nordhaus \(2017\)](#) of \$44.4 per metric ton of CO₂ in 2020 dollars.

cities have high productivity levels, this leads to increases in average income for both unskilled and skilled workers.

Next, we entirely remove the existing negative correlation between land-use restrictions and carbon emissions by setting the Wharton Index in all cities to the level faced by the median urban household in the US. Households respond by leaving the Midwest and South and moving to the West Coast and Northeast. Demand for natural gas is higher and electricity demand is lower in the cold Northeast. These population shifts therefore increase national gas usage and decrease total electricity usage. Overall, these changes in the spatial distribution of households and their energy consumption lead to a nearly 3.5% drop in national carbon emissions, implying a drop in the social cost of carbon of \$1.7 billion annually.

The main goal of this paper is to better understand the relationship between differences in land-use restrictions across cities and national household carbon emissions. However, household energy consumption also contributes to local pollution. Therefore, we also use the model to analyze the effects of relaxing land-use regulations on exposure to local pollutants. We focus on particulate matter of 2.5 micrometers or smaller ($PM_{2.5}$), a common measure of local air quality and that has been at the helm of US air quality regulation for the last two decades.⁹ Similar to household carbon emissions, emissions of $PM_{2.5}$ vary across space due to differences in energy consumption and the spatial distribution of power plants. However, unlike carbon emissions, emissions of local pollutants differentially affect air quality across cities. We use the Intervention Model for Air Pollution (InMAP) source-receptor matrix (Tessum, Hill, and Marshall (2017), Goodkind, Tessum, Coggins, Hill, and Marshall (2019)) to map electricity production to ambient air concentration of $PM_{2.5}$ across locations in the United States.

We find that the air quality in most US cities improves as a result of relaxing land-use restrictions in California. This result reflects that households use less electricity in California and that power plants used in California, in addition to being carbon-efficient, emit low levels of local pollutants.¹⁰ However, $PM_{2.5}$ exposure for the average household increases slightly, as there is an increase in the number of households in the relatively polluted Southern California.

Related literature Our paper is related to two recent papers, Hsieh and Moretti (2019) and Herkenhoff, Ohanian, and Prescott (2018), who find that relaxation of land-use restrictions in high-productivity cities would lead to large increases in GDP.¹¹ Our model focuses on an entirely different set of outcomes and incorporates energy demand, energy production, and emissions. Households in our model are differentiated in terms of

⁹Exposure to fine-grained particulate matter is associated with various detrimental health and economic outcomes such as (but not limited to): higher infant mortality (Chay and Greenstone (2005)), increased cognitive decline in seniors (Ailshire and Crimmins (2014)), and reduced property values (Chay and Greenstone (2003)).

¹⁰These reductions in $PM_{2.5}$ concentration in most cities can be thought of as a “cobenefit” of the reduction in carbon emissions. See Aldy, Kotchen, Evans, Fowlie, Levinson, and Palmer (2020) for a discussion of cobenefits in air quality regulation.

¹¹Albouy and Stuart (2020) also found that relaxation land-use restrictions would lead to a large redistribution of households across cities, but are less concerned with the effect of these changes on national productivity.

education, family composition, age, race, and the state in which they were born. As a consequence, our model allows for rich substitution patterns across cities and allows us to analyze how changes in land-use restrictions affect both the population and demographic composition across cities.

Our work builds on the descriptive findings in [Glaeser and Kahn \(2010\)](#) (GK). GK measured predicted carbon emissions for households in different cities across the country and documented a negative correlation between carbon emissions and land-use restrictions. Relative to GK, the primary contribution of this paper is to build and utilize a structural model to quantify the effects of land-use restrictions on national carbon emissions. National carbon emissions are determined in equilibrium; household locational sorting, energy demand, and housing supply/demand all determine the extent to which land-use restrictions affect national carbon emissions. Estimating the effects of a counterfactual change in land-use restrictions necessitates a structural equilibrium model.

This paper is also related to a large literature on how exposure to environmental externalities varies by location (See [Chay and Greenstone \(2005\)](#), [Currie, Davis, Greenstone, and Walker \(2015\)](#), [Muehlenbachs, Spiller, and Timmins \(2015\)](#), [Holland, Mansur, Muller, and Yates \(2019\)](#), or [Fowlie, Rubin, and Walker \(2019\)](#), for example). In our setting, exposure to carbon emissions is independent of the household's location—the effects of carbon emissions, and thus climate change are felt globally, not just locally. However, the amount of carbon dioxide emitted by a household depends on where the household is located.

Another related literature analyzes the effects of population density on the local carbon emissions (see [Fragkias, Lobo, Strumsky, and Seto \(2013\)](#) or [Jones and Kammen \(2014\)](#), for example). Other recent research uses simulation methods to analyze the effects of various land-use restrictions and transportation policies on within-city locational sorting and local carbon emissions (See [Larson and Yezer \(2015\)](#) or [Borck \(2016\)](#), for example). Our work differs from these two themes in that we focus on the sorting of households *across* cities, rather than the determinants of emissions within a city.¹² This paper is also related to [Fan, Fisher-Vanden, and Klaiber \(2018\)](#), who use a spatial equilibrium model to analyze the effects of climate change on household location choices and welfare. Finally, in other complementary work, [Mangum \(2017\)](#) analyzed the effects of housing and land stock allocations on carbon emissions.¹³

2. DATA

This paper utilizes individual data on household sorting and energy expenditures from the Census and ACS, detailed data on power plants from eGrid, and state level energy pricing data from the EIA. In what follows, we briefly describe each of the main data

¹²[Gaigné, Riou, and Thisse \(2012\)](#) argued that analysis of the effects of density-increasing policies on carbon emissions must account for relocation of households and firms across cities.

¹³Compared to [Mangum \(2017\)](#), our paper focuses more on the households sorting across cities and energy usage. Mangum's focus on the housing construction process allows for a more nuanced understanding of how different land-use restrictions affect the housing stock.

sources and how they are used in our analysis. Further details on the data can be found in Appendix A in the Online Supplementary Material (Colas and Morehouse (2022)).

CBSA level data We utilize Core Based Statistical Areas (CBSAs) as our definition of a geographic area. CBSAs correspond to distinct labor markets and are the Office of Management and Budget's official definition of a metropolitan area. To measure land-use restrictions in each CBSAs, we utilize a standard metric developed by Gyourko, Saiz, and Summers (2008), the Wharton Land-Use Regulation Index (WLURI). This index was created from a survey sent to 6896 municipalities across the US with questions that range from how many regulatory boards one must clear before construction to city-specific density and open space requirements. A higher value of the Wharton-Index implies more stringent restrictions and higher costs of developing land and is associated with more inelastic housing supply curves (e.g., Saiz (2010), Albouy and Ehrlich (2018), or Diamond (2016)).

Household data We use household-level data from the US Integrated Public Use Microdata Series (IPUMS); we utilize the 1990 Census, the 2000 Census, 5% 5-year American Community Survey (ACS) from 2006–2010, and 5% 5-year American Community Survey (ACS) from 2013–2017 (Ruggles, Flood, Foster, Goeken, Pacas, Shouweiler, and Sobek (2021)). Since our model contains a rich-level of household heterogeneity, a large dataset is imperative for our analysis. IPUMS provides information on yearly, household-level expenditures data on natural gas, electricity, and fuel oil in addition to information on demographics, location, and housing expenditures.¹⁴

Our model is concerned with emissions generated at home, therefore, we focus on three primary *energy types*: natural gas, electricity, and fuel oil.¹⁵ We combine data on expenditures on these three energy types with state level price data from the US Energy Information Administration (EIA) to impute household consumption of natural gas, electricity, and fuel oil.¹⁶

Power plants and emissions For each of the three energy types we consider, we use linear conversion factors to map usage of each energy type to carbon emissions. We assume 117 lbs of CO₂ are emitted per thousand cubic feet of natural gas consumed and 17 lbs of CO₂ are emitted per gallon of fuel oil consumed.¹⁷ Carbon emissions associated with electricity usage depend on where the electricity is consumed—electricity used in areas that generate electricity via coal plants will lead to more carbon emissions than in areas that rely more heavily on renewable sources.

¹⁴Renters and households living in multifamily homes may not pay for their own energy. We describe how we correct for this in Appendix A.4 in the Online Supplementary Material.

¹⁵GK also imputed emissions produced by cars. They find that differential emissions from cars are less important than differences in emissions from electricity and natural gas in explaining total differences in emissions across cities. Furthermore, emissions from driving calculated in GK are strongly correlated with total household emissions from other sources ($\rho = 0.56$) and a decrease in population density. If removing land-use restrictions increases city density, this will also lead to decreases in emissions from driving as well. Therefore, including emissions from driving in our analysis would likely strengthen our main conclusions.

¹⁶<https://www.eia.gov/state/seds/>

¹⁷<https://www.eia.gov/tools/faqs/faq.php?id=73&t=11>

We therefore utilize power plant-level data from the Emissions and Generation Resource Integrated Database (eGRID). These data provide information on the location, primary fuel input, emissions rate, and total megawatt hours of electricity generated for every power plant in the US. To assign households to power plants, we use the nine North American Electric Reliability Council (NERC) regions.¹⁸ These regions can be thought of as closed electricity markets as transmissions of electricity within a region is common but electricity is rarely transferred across regions (Holland and Mansur (2008), Glaeser and Kahn (2010)).¹⁹

We calculate the emissions factor associated with each NERC region as the weighted average CO₂ emissions per megawatt hour of electricity of all plants in the NERC region. The emissions factors range from roughly 800 to 1550 lbs of CO₂ emitted per megawatt hour of electricity consumed.²⁰ All CBSAs within a NERC region have the same CO₂ conversion factor.

For information on local pollutants, we use data from the National Emissions Inventory (NEI) and the EPA's Air Quality System (AQS). The NEI contains information about power plants such as stack height and emissions velocity, and emissions of various local pollutants such as PM_{2.5}. We use these data to construct a "pollution-transfer" matrix, which maps electricity usage in any given city to changes in ambient air quality in all other cities.²¹ The AQS data provide hourly levels of total particulate matter by city. We average across all hours in 2017 within each CBSA to obtain our measure of local average PM_{2.5} concentration.²²

3. DESCRIPTIVE STATISTICS

In this section, we calculate selection-corrected predicted household usage of electricity, natural gas, and fuel oil usage across cities and the associated carbon emissions. We focus on results calculated using data from 2017; earlier years are primarily used for estimation of parameters in the structural model (described in detail in Section 5).

Our goal is to isolate the role of a household's location on their energy usage and, therefore, carbon emissions. This will allow us to understand how the distribution of households—and, therefore, policies that affect the distribution of households—interact with national carbon emissions. We therefore construct a measure of predicted energy usage per household in each CBSA, controlling for differences in household composition, demographics, and unobserved differences in propensity to consume energy. First, we calculate each household's imputed energy usage in natural gas, electricity, and

¹⁸We omit the region for Alaska since no city in our model is in Alaska.

¹⁹While trading can also occur within a larger interconnection region, Holland and Mansur (2008) argued that NERC regions are the appropriate level of aggregation for electricity markets due to transmission and other technical constraints.

²⁰For the full distribution of emissions factors, see Appendix A.3 in the Online Supplementary Material.

²¹The pollution-transfer matrix and its construction are described in detail in Section 4.2 and Appendix A.10 in the Online Supplementary Material, respectively. The pollution-transfer matrix maps total energy production in each NERC region to ambient air concentration in all cities in the model.

²²To obtain census region average PM_{2.5} concentrations, we average over readings for all counties that are not part of the CBSAs in the given region.

fuel oil as their reported expenditure on each of these energy types divided by the state-level price of each energy type. We then employ the selection-correction method developed by Dahl (2002) to compute the selection-corrected predicted usage of each energy type in each city.²³

With the predicted per-household energy use in hand, we can calculate predicted carbon emissions for each CBSA.²⁴ We multiply the selection-corrected predicted usage for each fuel type with the respective emissions factor. As discussed in Section 2, we assume a constant emissions factor for fuel oil and natural gas. The emissions factor for electricity use varies across NERC regions.²⁵

3.1 Predicting energy usage across cities

Consider the following equation for household i 's usage of energy type m conditional on living in location j :

$$E_{ij}^m = \alpha_j^m + \beta_j^m X_i + u_i^m, \quad (1)$$

where E_{ij}^m is household i 's usage of energy type m , conditional on living in location j , α_j^m is a CBSA-specific intercept term, X_i is a vector of household i 's observable characteristics, β_j^m is a vector of parameters which varies by location j , and u_i^m represents household i 's idiosyncratic propensity to use energy type m , representing, for example, household i 's unobservable preferences for using air conditioning.²⁶

One difficulty with estimating (1) is that households may sort across locations based on their idiosyncratic propensity to use energy such that u_i^m is not mean-zero conditional on households' chosen locations. For example, households with a low tolerance for cold temperatures might avoid living in cities with cold weather, and may also have a greater propensity to use heating and, therefore, natural gas. This would induce a correlation between the unobserved propensity to use natural gas and the probability of living in cold cities. Concretely, let $V_{ij} = \bar{V}(X_i, B_i) + \varepsilon_{ij}$ represent household's i 's return for living in city j , where $\bar{V}(\cdot)$ is a component of utility which depends on household i 's observable characteristics X_i and the birth state of the household head, B_i , and ε_{ij} represents household i 's idiosyncratic preference for living in location j .²⁷ A household chooses to live in location j if it provides the highest return, that is, if $j = \operatorname{argmax}_{j'}(V_{ij'})$. If $\mathbb{E}[u_i^m \varepsilon_{ij}] \neq 0$, then this will induce selection on unobservables: $\mathbb{E}[u_j^m | j = \operatorname{argmax}_{j'}(V_{ij'})] \neq 0$. This would occur, for example, if people who prefer to live

²³The results with no selection correction are included in Online Appendix B.2.

²⁴We repeat this analysis with methane emissions in Online Appendix B.9.

²⁵The emissions factor for each NERC region is displayed in Online Appendix A.3.

²⁶This term could also vary by location j and be written as u_{ij}^m . We have written it as a household level term rather than a household by location level term for expositional purposes.

²⁷Importantly, the household head's birth state is assumed to not differently affect the households energy usage. As such, birth state serves as an exclusion restriction which helps to identify selection on unobservables in the outcome equation. One concern with using birth state as an exclusion restriction is that the climate in which an individual grew up might influence their preferences for energy usage as an adult. In Online Appendix B.1, we consider alternative specifications with controls for average temperature in the state where the household head was born. The results are qualitatively similar.

in Houston also have stronger preferences for using air conditioning. We will refer to $\mathbb{E}[u_j^m | j = \text{argmax}(V_{ij'})]$ as the “selection bias” term.

To deal with this selection issue, we employ a semiparametric selection correction based on the method proposed by Dahl (2002) (henceforth “Dahl”). Dahl showed that there exists a function that maps the household’s choice probabilities to the selection bias term. Concretely, let \mathbf{P}_{iJ} give the vector of the household’s choice probabilities for all cities in the set J . Then there exists a function $M_j^m(\cdot)$, such that $M_j^m(\mathbf{P}_{iJ}) = \mathbb{E}[u_j^m | j = \text{argmax}_{j'}(V_{ij'})]$.²⁸ Therefore, the selection bias can be controlled for if the econometrician controls for the function $M_j^m(\cdot)$ such that the estimating equation becomes

$$E_{ij}^m = \alpha_j^m + \beta_j^m X_i + M_j^m(\mathbf{P}_{iJ}) + \hat{u}_i^m. \quad (2)$$

Dahl noted that full estimation of (2) is generally infeasible as $M_j^m(\cdot)$ is an unknown function of the choice probabilities for all J cities. Therefore, we introduce two additional assumptions. First, following Dahl, we make an “index sufficiency assumption”: that the function $M_j^m(\cdot)$ can be replaced with an alternative function which only takes a subset of the choice probabilities as arguments. Second, we assume that the amount of selection on unobservables is constant for all cities within the same state such that $\mathbb{E}[u_i^m | j = \text{argmax}_{j'}(V_{ij'})] = \mathbb{E}[u_i^m | \hat{j} = \text{argmax}_{j'}(V_{ij'})]$ for any j, \hat{j} in the same state. For example, this implies that, conditional on the vector of observables X_i , the expectation of the idiosyncratic term u_i^m is the same in Dallas as it is in Houston.²⁹

Taken together, these two assumptions imply that the control function can be written as a function of a subset of the *state* choice probabilities. Let $\mathbf{P}_{i\hat{S}}^{\text{st}}$ give the vector of the household’s choice probabilities for all states in the set \hat{S} , where \hat{S} is a subset of the full set of states. We can then estimate the parameters of (1) using the equation

$$E_{ij}^m = \alpha_j^m + \beta_j^m X_i + M_j^m(\mathbf{P}_{i\hat{S}}^{\text{st}}) + \hat{u}_i^m, \quad (3)$$

where $M_j^m(\mathbf{P}_{i\hat{S}}^{\text{st}})$ is a correction function, which depends on $\mathbf{P}_{i\hat{S}}^{\text{st}}$.

In practice, we specify the function $M_j^m(\cdot)$ as a function of the probabilities of choosing the three largest states by population, the probability of choosing the state containing city j , and the interactions between the probability of choosing the state containing city j and the probabilities of choosing the three largest states.³⁰ For estimates of the state choice probabilities, we use the same approach as Dahl. Specifically, we divide households into cells, which vary in their demographic characteristics and their state of birth and calculate state choice probabilities as the proportion of households within each cell that chooses a given state.

²⁸This is subject to an invertibility condition: that these choice probabilities contain the same information as differences in subutility terms across cities.

²⁹This second assumption is useful in generating predicted values of energy usage. As we explain below, to separately identify α_0^m from the intercept of the selection correction function, we need to extrapolate the control function to households for which the probability of choosing a given location is equal to one. As state choice probabilities are closer to one than city choice probabilities, using state choice probabilities reduces the range over which we extrapolate the control function.

³⁰We show the sensitivity of our estimates to different choices of the correction function in Online Appendix B.1.

Finally, one complication arises because the intercept term α_j^m is not separately identified from the intercept of the control function. We overcome this identification issue by using the intuition of “identification at infinity” (Chamberlain (1986), Heckman (1990)): suppose the econometrician observes households with demographics \hat{X} and birth state \hat{B} such that $\text{Prob}(j = \text{argmax}_{j'}(V_{ij'})|\hat{X}, \hat{B}) = 1$. Since households with these characteristics choose location j with certainty, there is no selection on unobservables for these households: $\mathbb{E}[u_i^m | j = \text{argmax}_{j'}(V_{ij'}), \hat{X}, \hat{B}] = 0$. In terms of the selection correction function, this implies that $M_j^m(\hat{\mathbf{P}}_{iS}^{\text{st}}) = 0$, where $\hat{\mathbf{P}}_{iS}^{\text{st}}$ is the vector of choice probabilities where the probability of choosing the state containing city j is equal to one and the probability of choosing all other states is equal to zero. We first estimate equation (3). Then we use the restriction that $M_j^m(\hat{\mathbf{P}}_{iS}^{\text{st}}) = 0$ to back out the intercept of the selection correction equation, and thus α_j^m .³¹

Finally, after estimating the parameters of the energy usage functions, we calculate predicted usage of energy type m as

$$\hat{E}_j^m = \hat{\alpha}_j^m + \hat{\beta}_j^m \bar{X},$$

where $\hat{\alpha}_j^m$ and $\hat{\beta}_j^m$ denote parameter estimates and \bar{X} gives a vector of the mean values of each demographic characteristic.

3.2 Selection-corrected predicted usage and emissions

The predicted yearly household usage and emissions from the 2017 aggregated ACS for selected cities are shown in Table 1. To calculate these predicted emissions, we multiply selection-correction predicted usage by the appropriate emissions factors. We show results for the six lowest emissions cities, the six highest emissions cities, and the six median cities. The third column (“Emissions”) shows the sum of predicted CO₂ emissions from natural gas, fuel oil, and electricity for the CBSA. Predicted household emissions vary considerably across cities. In Honolulu, predicted emissions are slightly over 14 tons per year, whereas in Milwaukee they are over 46 tons per year.

The next two columns show emissions from gas and fuel oil, respectively, which are equal to predicted usage multiplied by the appropriate emissions factor. Natural gas emissions are generally the largest in colder regions.³² Emissions from fuel oil are generally quite small in magnitude compared to emissions from the other two energy types. The last three columns show predicted electricity usage, the electricity emissions factor,

³¹As a simple example, consider the case where M_j^m is specified as a first-order polynomial of choosing the state in question: $M_j^m(\mathbf{P}_{iS}^{\text{st}}) = M_0 + M_1 P_{\text{is}(j)}^{\text{st}}$, where $P_{\text{is}(j)}^{\text{st}}$ is the probability of choosing the state containing city j . Then the intercept of this correction function M_0 can simply be calculated as $M_0 = -M_1$. Note that the identification at infinity argument relies on the econometrician observing households for whom the probability of choosing the state in question is close to 1. If we allow selection to occur at the city level within each state, we would need to observe households for whom the probability of choosing a given city is close to 1. This is a very strong assumption in the case of many cities, given that we only observe data on an individual's state of birth, not their city of birth.

³²In Section 3.3, as in GK, we show that colder winter temperatures are highly predictive of natural gas usage.

TABLE 1. Predicted CBSA level CO₂ emissions by fuel type for the six lowest emissions cities, the six median cities, and the six highest emissions cities in 2017. The third column (“Emissions”) shows the sum of predicted CO₂ emissions from natural gas, fuel oil, and electricity for the CBSA. The next two columns show emissions from gas and fuel oil, respectively, which are equal to predicted usage multiplied by the appropriate emissions factor. The last three columns show predicted electricity usage, the electricity emissions factor, and predicted electricity emissions, equal to predicted electricity usage multiplied by the emissions factor.

CBSA	Rank	Emissions (1000 lbs)	Gas Emissions (1000 lbs)	Fuel Emissions (1000 lbs)	Electricity Use (MWh)	Electricity Conversion (1000 lbs/MWh)	Electricity Emissions (1000 lbs)
<i>Lowest</i>							
Honolulu, HI	1	14.24	0.00	0.00	9.36	1.52	14.24
Oxnard, CA	2	14.67	6.19	0.27	10.26	0.80	8.21
Riverside, CA	3	16.37	6.33	0.27	12.21	0.80	9.76
San Diego, CA	4	16.71	6.75	0.33	12.04	0.80	9.63
Los Angeles, CA	5	17.14	6.73	0.20	12.76	0.80	10.21
Sacramento, CA	6	17.96	7.72	0.47	12.23	0.80	9.78
<i>Middle</i>							
Baton Rouge, LA	33	26.56	4.41	0.43	20.98	1.04	21.72
Birmingham, AL	34	26.86	5.79	0.21	20.15	1.04	20.86
Jacksonville, FL	35	26.90	0.62	0.06	25.92	1.01	26.22
New Orleans, LA	36	27.15	4.61	0.41	21.38	1.04	22.13
Pittsburgh, PA	37	27.41	12.02	2.43	11.73	1.11	12.97
Houston, TX	38	27.51	4.12	0.13	22.92	1.01	23.25
<i>Highest</i>							
Tulsa, OK	65	40.21	12.47	0.28	21.60	1.27	27.46
Oklahoma City, OK	66	41.59	11.81	0.27	23.21	1.27	29.50
Indianapolis, IN	67	43.67	23.20	0.30	18.26	1.11	20.18
Memphis, TN	68	43.81	10.56	0.23	31.89	1.04	33.02
Omaha, NE	69	45.49	17.31	0.31	22.84	1.22	27.87
Milwaukee, WI	70	46.19	21.84	0.34	21.72	1.11	24.01

and predicted electricity emissions, which is equal to predicted electricity usage multiplied by the emissions factor.

Spatial variation in household carbon emissions comes from multiple sources. For example, power plants utilized in Memphis emit less CO₂ than Oklahoma City (1.04 lbs per MWh in Memphis compared to 1.27 in Oklahoma City). However, electricity usage in Memphis is so much higher than in Oklahoma City that overall household emissions are higher in Memphis, despite greater consumption of fuel oil and natural gas in Oklahoma City.³³ Conversely, consider emissions resulting from electricity in Houston compared to Tulsa. Households in Houston use more electricity than those in Tulsa. However, power plants near Tulsa are less carbon-efficient than those near Houston. Therefore, emissions from electricity use are higher in Tulsa. This underscores an important feature of

³³Since consumption of natural gas and fuel oil have the same conversion factor nationally, a higher level of emissions in one city necessarily means a higher level of consumption.




FIGURE 1. The left panel displays predicted natural gas usage and average January temperature. The right panel displays predicted electricity usage and average August temperature. Temperature data were obtained from [weather.com](https://www.weather.com). Each point is a CBSA. Temperature corresponds to the midpoint of the average minimum and maximum daily temperature recorded in the month of interest. The size of each point reflects the population of the CBSA. Electricity usage is measured in MWh and natural gas usage is measured in 1000 ft³.

the data: spatial variation in household electricity emissions is driven by both differences in energy usage and heterogeneity in power plants across regions.

3.3 Energy usage and climate

To further understand the differences in energy usage across cities, we now examine the relationship between energy usage and climate. Figure 1 shows the CBSA level relationship between average August temperature and predicted electricity usage and between average January temperature and predicted natural gas usage. Similar to [Glaeser and Kahn \(2010\)](#), we find strong relationships between temperature and consumption of different fuel sources. Electricity usage has a strong positive relationship with August temperature. Similarly, as January temperature increases, natural gas use decreases. Taken together, these results suggest that differences in energy usage across cities are largely driven by differences in climate.

3.4 Policy and emissions

Spatial variation in household carbon emissions implies that any policy that affects the spatial distribution of households will also impact national carbon emissions. The primary policy we are interested in is land-use restrictions. Figure 2 shows a scatter-plot between CBSA-level predicted emissions and the Wharton Land-Use Regulation Index. The Wharton Index is displayed on the horizontal axis; higher values of this in-




FIGURE 2. Predicted per-household CO₂ emission measured in pounds and Wharton Lane Use Index. Each circle is a CBSA; California cities are highlighted. The size of each circle reflects CBSA population.

dex correspond to tighter land-use restrictions. The vertical axis displays predicted per-household CO₂ emissions, measured in pounds.³⁴

Generally, tighter land-use restrictions are associated with lower predicted carbon emissions. In particular, California cities have very low predicted carbon emissions due to a combination of temperate climate and clean power plants. Cities in California also have very tight land-use restrictions, as measured by the Wharton Index.

To be clear, the goal of this paper is not to explain what generates this relationship. Instead, our goal is to study the implication of this correlation, in the sense that tight land-use restrictions inflate housing prices and incentivize individuals to live away from California and other cities with low carbon emissions. We proceed by building a spatial equilibrium model to quantify the effects of land-use restrictions on national carbon emissions.

4. MODEL

We employ a static spatial equilibrium featuring heterogeneous households, with endogenous wages and rents, similar to those used in [Diamond \(2016\)](#), [Piyapromdee \(2021\)](#), and [Colas and Hutchinson \(2021\)](#). We extend this class of model by allowing locations to vary by carbon and local pollutant output of regional power plants and the marginal benefits of energy usage. Therefore, our model is able to map changes in the distribution of households across locations to carbon and local pollutant emissions.

Household sorting is a static, discrete choice—households choose the city which provides the highest utility in terms of wages, rents, amenities, and energy prices. Households purchase three *energy types*—electricity, natural gas, and fuel oil—which they use to produce *energy services* (e.g., air conditioning, home heating). The utility

³⁴This statistic is the same as that displayed in the “Emissions” column of Table 1.

benefit of energy services varies by city. For example, the benefit of air conditioning (and thus electricity) is high in Memphis while the benefit of home heating (and thus natural gas) is high in Minneapolis. Each location has an upward sloping housing supply curve whose elasticity depends on local land-use restrictions and the amount of land available for development. Cities with tighter land-use restrictions will have more inelastic housing supply curves and higher rents, all else equal. Competitive firms combine skilled and unskilled workers using a CES production function to produce a numéraire consumption good. Thus, local wages and rents are endogenous to the distribution of households across cities. Changes in land-use restrictions across cities will change housing supply curves across cities and impact the equilibrium distribution of households.³⁵

As we show in Online Appendix A.1, emissions vary across household types. We allow households in the model to vary in their education level, race, age group, marital status, and whether or not the household has children. These household types vary in their preferences over locations, energy services, and housing. Within this demographic group, households also receive a premium for living in cities close to their birth state. This allows for rich substitution patterns in response to policy changes—a decrease in rents in San Francisco, for example, will lead to larger inflows of households who are born in California. Finally, households receive an idiosyncratic preference draw over each location, where the variance of the draw depends on household demographics. Therefore, households are imperfectly mobile across cities.

The amount of carbon emissions generated by a household varies for two reasons in the model. First, the marginal benefit of each energy type varies by city. Cities with a higher marginal benefit of energy usage will have higher levels of energy usage, all else equal. Second, the production technology and carbon efficiency of power plants vary across NERC regions. Electricity used at a given city must be produced by a power plant in the associated NERC region. Therefore, electricity usage in cities located in NERC regions with more carbon-efficient power plants will lead to lower emissions.

In addition to carbon emissions, household electricity usage also leads to emissions of local pollutants and, therefore, deterioration in air quality—where we explicitly focus on $PM_{2.5}$ concentration as our measure of local air quality. We construct a pollution-transfer matrix using a state-of-the-art source-receptor matrix, the InMAP source-receptor matrix (ISRM) (Tessum, Hill, and Marshall (2017), Goodkind et al. (2019)).³⁶ We use this pollution-transfer matrix to map household electricity usage and, therefore, $PM_{2.5}$ emissions, in any given city to ambient air quality in all cities in the model. In our baseline model, we do not directly account for $PM_{2.5}$ concentration in

³⁵We define an equilibrium in this setting in Online Appendix A.6.

³⁶InMAP is a “reduced-complexity” air transport model. Other popular reduced complexity models include COBRA (EPA (2020)) and APEEP (Muller and Mendelsohn (2007)). All of these models make simplifying assumptions around atmospheric chemistry equations to ease the computational burden of estimating pollution transfer. More complex atmospheric dispersion models such as HYSPLIT (Stein, Draxler, Rolph, Stunder, Cohen, and Ngan (2015)) and WRF-Chem generally perform better for predicting pollution transfer over long distances. However, the added complexity makes them far more computationally expensive. InMAP performs similar to many other reduced-complexity models (Tessum, Hill, and Marshall (2017)) and is easy to implement, making it our preferred method for modeling pollution transfer. See Hernandez-Cortes and Meng (2020) for a discussion of the limitations of InMAP relative to HYSPLIT.

household utility. In Section 8.3, we show that our results are very similar when we include PM_{2.5} concentration in the utility function.³⁷

4.1 Households

Let j index cities and i index households.³⁸ Households are endowed with a demographic type d , which includes the household head's education, marital status, race, whether or not they have children, and age group.³⁹ Locations vary by amenities, which we denote as λ_{ij} . To solve their decision problem, the household chooses a location j that yields maximal utility from consumption of the numeraire c , housing H , and energy services \hat{E}^m (such as heating or air-conditioning), and from amenities.

We parameterize the household's utility function as

$$u_{ij} = \alpha_d^c \log c + \alpha_d^H \log H + \sum_m \alpha_{jd}^m \log \hat{E}^m + \lambda_{ij}, \quad (4)$$

where α_d^c , α_d^H , and α_{jd}^m are parameters, which scale the marginal benefit of consumption, housing, and energy services. We set $\alpha_d^c = 1$ for each demographic group d to normalize for scale. α_{jd}^m varies across cities to reflect differences in the marginal benefit of energy services across cities (perhaps due to differences in weather).⁴⁰

We assume energy services are produced by the household using a fixed proportions energy production function which maps energy types into energy services. Let E^m denote usage of energy type m , where $m \in \{\text{Elec, Gas, Fuel}\}$. The energy service production function takes the form:

$$\hat{E}^m = f(x^m) = a^m E^m, \quad (5)$$

where a^m is a parameter that maps units of energy into energy services.

Substitution of (5) into (4) yields

$$u_{ijt} = \log c + \alpha_d^H \log H + \sum_m \alpha_{jd}^m \log E^m + \lambda_{ij} + C,$$

where $C = \sum_m \alpha_{jd}^m \log a^m$ is an additive constant that does not depend on the household's choices.

³⁷Our model only endogenizes PM_{2.5} concentration arising from household electricity usage. The changes in PM_{2.5} arising from household electricity usage are small relative to differences in total PM_{2.5} concentrations across cities.

³⁸In Section 5, we introduce t subscripts to indicate which variables and parameters vary over time. We omit these t subscripts for clarity as we present the model.

³⁹We divide households into two age groups based on the age of the household head. Households with heads over 35 years old are defined as "older" households.

⁴⁰In Section 3.3, we show that local temperature is highly predictive of energy usage. We assume that these parameters are not a function of the local population or population density. In Online Appendix A.5, we provide evidence that the population is not a significant predictor of energy demand. Fragkias et al. (2013) came to a similar conclusion. Further, we assume that land-use restrictions do not directly affect energy use. Some targeted land-use restrictions, such as urban growth boundaries or endangered species habitats, may have positive environmental effects. See Lawler et al. (2014), for example, for a discussion.

The households' budget constraint is given by

$$I_{jd} = c + R_j H + \sum_m P_j^m E^m,$$

where I_{jd} is the income level of a household of demographic d living in city j and E^m is usage of energy type m . R_j and P_j^m represent the prices of housing and the price of energy of type m in city j . We normalize the price of consumption of c to one.

We decompose the amenity term, λ_{ij} , into five distinct components. In particular, we let

$$\lambda_{ij} = \gamma_d^{hp} \mathbb{I}(j \in B_i) + \gamma_d^{\text{dist}} \phi(j, B_i) + \gamma_d^{\text{dist}^2} \phi^2(j, B_i) + \xi_{jd} + \sigma_d \epsilon_{ij},$$

where $\mathbb{I}(j \in B_i)$ is an indicator for location j being in the head of household i 's birth state, $\phi(j, B_i)$ and $\phi^2(j, B_i)$ are the distance and squared distance, respectively, between the household head's birth state and location j . ξ_{jd} is a shared unobservable component of amenities and ϵ_{ij} is an idiosyncratic preference shock with dispersion parameter σ . Differences in ϵ_{ij} across individuals and cities reflect unobservable variation in attachment to a location that an individual might have. We assume that ϵ_{ij} follows a Type 1 Extreme Value distribution.⁴¹ We make an important assumption that unobserved amenities, ξ_{jd} , are taken as exogenous and are not a function of land-use restrictions, as is relatively standard.⁴²

Solving the household's maximization problem yields constant income shares on housing and energy of all types. We write a household of demographic group d 's optimal choice of housing, conditional on living in a city j as

$$H_{jd}^* = \frac{\alpha_d^H I_{jd}}{\alpha_{jd} R_j},$$

where, to simplify notation, we define $\alpha_{jd} \equiv 1 + \alpha_d^H + \sum_m \alpha_{jd}^m$. Optimal usage of energy type m is also a constant fraction of income:

$$E_{jd}^{m*} = \frac{\alpha_{jd}^m I_{jd}}{\alpha_{jd} P_j^m}.$$

We then solve for the indirect utility function associated with location j :

$$V_{ij} = (\alpha_{jd}) \log I_{jd} - \alpha_d^H \log R_j - \sum_m \alpha_{jd}^m \log P_j^m + \lambda_{ij}, \quad (6)$$

⁴¹We do not allow for the possibility of endogenous amenities, as in [Diamond \(2016\)](#). We also refrain from directly modeling the effects of carbon emissions on household's utility. Our model only includes US households and, therefore, cannot reliably capture the full social cost of carbon. In Section 8.3, we consider an extension in which local pollutants enter the household utility function.

⁴²See [Hsieh and Moretti \(2019\)](#), [Piyapromdee \(2021\)](#), or [Colas and Hutchinson \(2021\)](#), for example. [Diamond \(2016\)](#) and [Herkenhoff, Ohanian, and Prescott \(2018\)](#) allowed amenities to be endogenous to household composition, but not land-use restrictions directly. This assumption is supported by the findings in [Albouy and Ehrlich \(2018\)](#), who find that land-use restrictions increase the cost of housing production without improving local quality of life. Clearly, some land-use restrictions, such as park systems, have a direct and positive effect on amenities. We can think of our counterfactuals here as the relaxation restrictions that do not have a large direct effect on local amenities.

where we drop additive constants which have no affect on household decisions. The household's problem can be thought of as a discrete choice over all the locations, conditional on optimal housing and energy consumption. Given that the idiosyncratic preference draws are distributed as Extreme-Value Type I, we can write the probability that a household i chooses a location j as

$$P_{ij} = \frac{\exp(\bar{V}_{ij}/\sigma_d)}{\sum_{j' \in J} \exp(\bar{V}_{ij'}/\sigma_d)},$$

where $\bar{V}_{ij} = V_{ij} - \sigma_d \epsilon_{ij}$ is the household's indirect utility of choosing location j minus the idiosyncratic preference draw. We write the total number of households of demographic d who choose to live in location j as N_{jd} .

4.2 Energy production and emissions

Carbon emissions We allow for three types of energy in our analysis: natural gas, fuel oil, and electricity. We assume fuel oil and natural gas are purchased on an international market and treat the supply for these types of energy as perfectly elastic. We assume that the carbon byproduct of fuel oil and natural gas are constant regardless of where energy is consumed. Total household emissions of CO₂ from natural gas and fuel oil in city j is the sum of usage of the energy type multiplied by the appropriate conversion factor:

$$\text{CO}_2_j^m = \hat{\delta}^m \sum_d N_{jd} E_{jd}^m, \quad m \in \{\text{Gas, Fuel}\},$$

where $\sum_d N_{jd} E_{jd}^m$ is the total amount of fuel of type m consumed by people living in city j and $\hat{\delta}^m$ is the amount of CO₂ emissions per unit of fuel of type m .

We assume electricity is generated across NERC regions in the United States and then is transmitted to local labor markets within those regions.⁴³ Within each NERC region, perfectly competitive power plants produce electricity.⁴⁴ In our baseline specification, we assume that the marginal cost of energy production is constant.⁴⁵ In Section 8.2, we consider a model extension with increasing marginal cost. The results are qualitatively similar in both cases. We allow the conversion factor for electricity to vary

⁴³As noted by Holland, Mansur, Muller, and Yates (2016) and Zivin, Kotchen, and Mansur (2014), trading across NERC regions can bias estimates of marginal damages from electricity production. As the WECC Nerc region happens to overlap with the Western-interconnection, and our main counterfactual induces migration into California (part of the WECC region), our results are likely robust to aggregating electricity markets to the NERC or Interconnection level.

⁴⁴Electricity is a homogeneous good with a large number of many producers. However, when transmission constraints bind, generation companies may have local market power. See Joskow and Tirole (2000) for a discussion.

⁴⁵The short run supply of electricity is often modeled as a dispatch curve with constant marginal or linear marginal cost curves. However, as we are considering a long-run equilibrium, the supply curve is given by the long run marginal cost curve, allowing for the construction of new reactors or the entry of new plants.

by NERC region to reflect geographic variation in the carbon intensity of power plants.⁴⁶ For example, a larger percentage of power in the Western NERC region (WECC) comes from hydroelectric dams, whereas the Southern NERC region (SERC) relies more heavily on coal power.

Let $\hat{\delta}_{\mathcal{R}}^m$ represent the conversion factor of electricity to CO₂ emissions in NERC region \mathcal{R} and let $\mathcal{R}(j)$ map cities to their corresponding NERC regions. We write CO₂ emission resulting from electricity usage in CBSA j as

$$\text{CO}_2^m_j = \hat{\delta}_{\mathcal{R}(j)}^m \sum_d N_{jd} E_{jd}^m, \quad m \in \{\text{Elec}\}.$$

For simplicity, we use the following notation for emissions factors:

$$\delta_j^m = \begin{cases} \hat{\delta}^m, & m \in \{\text{Gas, Fuel}\}, \\ \hat{\delta}_{\mathcal{R}(j)}^m, & m \in \{\text{Elec}\}. \end{cases}$$

Concretely, the emissions factors for natural gas and fuel-oil are location-independent, while the emissions factors for electricity are constant within (but not across) NERC regions. Local CO₂ emission of each energy type m can then be written as

$$\text{CO}_2^m_j = \delta_j^m \sum_d N_{jd} E_{jd}^m.$$

We can then write national emissions from each energy type m as $\text{CO}_2^m = \sum_j \text{CO}_2^m_j$, and national emissions across all energy types is simply the sum of the energy specific emissions levels, $\text{CO}_2 = \sum_m \text{CO}_2^m$.

Next, we can examine how the distribution of households across cities affects the level of national carbon emissions.⁴⁷ We rewrite national emissions in terms of the covariance between the distribution of households and the efficiency of local electricity usage multiplied by the local energy usage:

$$\text{CO}_2 = \sum_m \sum_d (J \cdot \text{Cov}(N_{jd}, E_{jd}^m \delta_j^m) + N_d \mathbb{E}[E_{jd}^m \delta_j^m]),$$

where N_d , the total number of households of group d , and J , the total number of cities, are both model primitives. The expectation $\mathbb{E}[E_{jd}^m \delta_j^m]$ is taken over cities j . National emissions are increasing in the covariance of population and the product of energy usage and energy conversion factors. Therefore, policies that lead households to live in cities that are associated with higher energy usage and less carbon-efficient power plants will lead to increases in national carbon emissions.

As demonstrated in Section 3.4, the tightness of land-use restrictions is negatively correlated with local predicted CO₂ emissions levels. Furthermore, tighter land-use restrictions increase local rents and, in equilibrium, lead to lower population levels in

⁴⁶Note that $\hat{\delta}^m$ for fuel oil and natural do not vary by location as conversion factors for these types of fuel are independent of location.

⁴⁷The main cost of carbon emissions are felt globally and not modeled directly here.

these cities. In Section 7, we examine the quantitative implications of this relationship between land-use restrictions, energy demand, and power plant technology on national carbon output.

Local pollutants In addition to carbon emissions, the model features local pollutants. We focus on particulate matter of 2.5 micrometers or smaller (PM_{2.5}) as our primary measure of local pollution. We focus on particulate matter emissions from electricity only, as natural gas emits a negligible amount of PM_{2.5} (EPA (2019)).⁴⁸ In the model, we distinguish between PM_{2.5} *emissions* (measured in tons/year) and PM_{2.5} *concentration* (measured in micrograms per cubic meter (measured in μ/m^3)).

While PM_{2.5} is considered a local pollutant because it has negative consequences for those directly exposed, emissions of PM_{2.5} from a given location can impact air quality locally, regionally, and nationally. PM_{2.5} emissions are highly transportable; Morehouse and Rubin (2021) estimated that roughly 90% of particulate matter emissions from coal-power plants leave the state in which they were emitted within 48 hours. Electricity demand in California, for example, will lead to an increase in PM_{2.5} emitted from power plants in the associated WECC NERC region. These additional emissions affect air quality not only in California but potentially all Western states and—to a lesser extent—the rest of the United States.

To map emissions of PM_{2.5} from a given NERC region to concentration of PM_{2.5} for each city in the model, we employ a state-of-the-art “source-receptor” (SR) matrix derived from a recent integrated assessment model, the Intervention Model for Air Pollution (InMAP) (Tessum, Hill, and Marshall (2017), Goodkind et al. (2019)). The entries of the SR matrix provided by InMAP (henceforth, ISRM) are “transfer coefficients,” which give the marginal impact of particulate matter emissions (measured in tons/year) in any given location on the ambient concentration (measured in μ/m^3) in any other location. ISRM accounts for power-plant stack height, the velocity at which the particles were emitted, and local atmospheric conditions. We use ISRM to construct a “pollution-transfer” matrix, where each entry gives the conversion factor between electricity produced in NERC region R (measured in MWh) to pollution in CBSA j (again, measured in μ/m^3). With this matrix in hand, we can estimate the extent to which changes in household energy demand lead to changes in ambient air quality.⁴⁹ Let \mathbf{P} denote this pollution-transfer matrix, and let \mathbf{E}^{elec} give the vector of household electricity produced in each NERC region. The contribution of household electricity usage to PM_{2.5} concentration in each CBSA is given by

$$\mathbf{PM}^{\text{endog}} = \mathbf{E}^{\text{elec}} \times \mathbf{P}.$$

A given element of this vector, $\text{PM}_j^{\text{endog}}$, gives the concentration of PM_{2.5} in a city that arises from (national) household energy usage.

⁴⁸Furthermore, fuel-oil is used by households in relatively few states and will not have first-order consequences for overall PM_{2.5} concentration.

⁴⁹For details on this procedure, see Online Appendix A.10.

In addition to emissions from household electricity usage, $PM_{2.5}$ can originate from many sources (EPA (2019)).⁵⁰ To account for this, we assume that $PM_{2.5}$ concentration in a given city is the sum of $PM_{2.5}$ concentration resulting from household electricity usage and $PM_{2.5}$ produced by other sources, which we assume to be invariant to household location choices. Concretely, letting PM_j denote the overall level of $PM_{2.5}$ in city j , we assume

$$PM_j = \overline{PM}_j + PM_j^{\text{endog}}, \quad (7)$$

where \overline{PM}_j is the fixed, city-level pollution.

4.3 Housing supply

Each city has an upward sloping housing supply curve. The elasticity of the housing supply curve is allowed to vary by city as a function of the amount of available land and the strictness of land-use restrictions. Specifically, we follow Kline and Moretti (2014) and parameterize the inverse housing supply curve in city j as

$$R_j = z_j H_j^{k_j}, \quad (8)$$

where H_j is quantity of housing supplied, z_j is a scale parameter, and k_j is a parameter equal to the inverse elasticity of the housing supply curve (i.e., $\frac{\partial \log R_j}{\partial \log H_j} = k_j$). Taking logs of (8), we obtain

$$\log(R_j) = k_j \log(H_j) + \log(z_j).$$

The term k_j plays a crucial role in our analysis. Higher values of k_j imply more inelastic housing supply curves and higher rent levels. Therefore, cities with higher values of k_j will have lower equilibrium population levels, all else equal.

As shown by Saiz (2010), local land-use restrictions, as measured by the Wharton Land Use Index, and the fraction of land that is unavailable for development due to geographic constraints are strong determinants of more inelastic housing supply curves. We follow Saiz (2010) and parameterize k_j as a function of land-use restrictions and geographic constraints:

$$k_j = \nu_1 + \nu_2 \psi_j^{\text{WRI}} + \nu_3 \psi_j^{\text{GEO}},$$

where ψ_j^{WRI} is the Wharton Land Use Index and ψ_j^{GEO} measures the amount of land that is unavailable for development due to geographic restrictions.⁵¹ A higher value of ν_2 implies that cities with tighter land-use restrictions will have a more inelastic housing supply. As shown in Section 3.4, cities with higher values of ψ_j^{WRI} generally have lower carbon emissions per household. In the model, this disincentivizes households from living in cities with low carbon emissions.

⁵⁰Furthermore, household energy consumption form a relatively small proportion of total $PM_{2.5}$ emissions; for details, see Online Appendix B.3.

⁵¹Another option would be to use a more disaggregated measure of land-use restrictions. This would allow us to decompose the effects of various types of land-use restriction on carbon emissions.

Specifically, given that the idiosyncratic preferences draws are distributed as Extreme-Value Type I, the partial equilibrium elasticity of location choice with respect to rents is approximately equal to⁵²

$$\frac{\partial \log P_{ij}}{\partial \log R_j} \approx -\frac{\alpha_d^H}{\sigma_d}.$$

We can solve for the partial equilibrium effect of a household's choice probability with respect to land-use restrictions as

$$\frac{\partial \log P_{ij}}{\partial \psi_j^{\text{WRI}}} \approx -\nu_2 \frac{\alpha_d^H}{\sigma_d} \log(H_{jt}).$$

The partial equilibrium effect of land-use restrictions is proportional to the expenditure share on housing and the importance of land-use restrictions in dictating the housing supply elasticity ν_2 , and inversely proportional to σ_d , the dispersion in the idiosyncratic preference draw. Higher values of σ_d imply household location choices are less responsive to changes in rents; thus, variation in land-use restrictions will have smaller effects on household sorting.

4.4 Wages

Perfectly competitive firms in each city combine skilled and unskilled labor in a CES production function to produce the numéraire consumption good, where we define household heads with a college degree as skilled and household heads with less than a college degree as unskilled.⁵³ Therefore, wages for skilled and unskilled workers in each city are determined endogenously by the ratio of skilled to unskilled workers. Specifically, firms use a combination of skilled (S) and unskilled labor (U), as inputs in the following production function:

$$Y_j = A_j [(1 - \theta_j) U_j^{\frac{s-1}{s}} + \theta_j S_j^{\frac{s-1}{s}}]^{\frac{s}{s-1}},$$

where U_j and S_j are defined as the total efficiency units of labor supplied by unskilled and skilled workers in city j , respectively. A_j is the total factor productivity in city j and θ_j is the relative factor intensity of skilled workers. The elasticity of substitution between skilled and unskilled workers is given by s .⁵⁴

Firms take wages as given and choose skilled and unskilled labor quantities to maximize profits. We derive labor demand curves as a result of the firms skilled and unskilled

⁵²Differentiating P_{ij} with respect to rents yields $\frac{\partial \log P_{ij}}{\partial \log R_j} = -\frac{\alpha_d^H}{\sigma_d} (1 - P_{ij}) \approx -\frac{\alpha_d^H}{\sigma_d}$ for small values of P_{ij} .

⁵³Data on energy usage by firms are generally less readily available than data on household energy usage. As such, we choose to focus on household emissions. Glaeser and Kahn (2010) argued that commercial energy use and household energy use are likely to be highly correlated.

⁵⁴One straightforward way to introduce capital into the model is to assume that production is Cobb–Douglas in capital and a CES labor supply such that $Y_{jt} = A_{jt} K_{jt}^\eta ((1 - \theta_j) U_j^{\frac{s-1}{s}} + \theta_j S_j^{\frac{s-1}{s}})^{\frac{s}{s-1}})^{1-\eta}$ where η is a parameter. If capital supply is perfectly elastic, this production function implies wage equations that are equivalent to those here.

labor first-order conditions for profit maximization:

$$W_{js} = A_j \left(\frac{Y_j}{A_j} \right)^{\frac{1}{\zeta}} \theta_j S_j^{-\frac{1}{\zeta}},$$

$$W_{ju} = A_j \left(\frac{Y_j}{A_j} \right)^{\frac{1}{\zeta}} (1 - \theta_j) U_j^{-\frac{1}{\zeta}},$$

where W_{js} and W_{ju} are the wage rates for skilled and unskilled labor, respectively.

Within education groups, demographic groups are perfectly substitutable in production but vary in their productivity and, therefore, supply different amounts of efficiency units of labor. Income levels for an individual household are given by the number of efficiency units of labor supplied by the household multiplied by the appropriate wage rate. Income for a household of demographic group d living in city j is given by $I_{jd} = W_{ju} \ell_d$ for unskilled workers and $I_{jd} = W_{js} \ell_d$ for skilled workers, where ℓ_d represents the number of efficiency units supplied by agents of demographic group d .

5. ESTIMATION

In this section, we describe our estimation procedure. We focus most of our exposition on the estimation of household location choice and energy use parameters. Estimation of the housing supply and production are relatively standard and details are therefore relegated to Online Appendices A.8 and A.9. The carbon emissions factors are calculated as in Section 2.⁵⁵ We choose to use the 70 largest CBSAs, as defined by population in 1980. These 70 locations make up approximately 55% of the entire US population in 2017. We map individuals that do not live in one of these 70 areas into their corresponding census division, creating nine additional choices.

Note that by defining our locations as CBSAs, we are abstracting from household location choices across municipalities or neighborhoods within CBSAs. While neighborhoods or municipalities within a CBSA may differ in many dimensions, they are unlikely to differ substantially in their associated carbon emissions because climate and the set of power plants where electricity is produced are relatively constant within a given CBSA.

5.1 Households

We estimate household preferences using the two-step “BLP” procedure using repeat cross-sectional data from the 1990 Census, 2000 Census, 2010 aggregated ACS, and 2017 aggregated ACS (Berry, Levinsohn, and Pakes (2004)).⁵⁶ As we estimate household preferences using multiple cross sections of data, we introduce t subscripts to indicate

⁵⁵That is, we assume 117 lbs of CO₂ emitted per thousand cubic feet of natural gas consumed and 17 lbs of CO₂ emitted per gallon of fuel oil consumed. We calculate the weighted average CO₂ emissions of all plants in a NERC region. We then assign each of the CBSAs to a NERC region, thus assigning all individuals in our sample a carbon emissions factor for electricity.

⁵⁶This estimation procedure has been utilized extensively in the urban economics literature, for example, by Diamond (2016), Piyapromdee (2021), and Colas and Hutchinson (2021).

which variables and parameters vary over time and which parameters are assumed to be constant over time. We rewrite the household's indirect utility function as

$$V_{ijt} = (\alpha_{jd}) \log I_{jdt} - \alpha_d^H \log R_{jt} - \sum_m \alpha_{jd}^m \log P_{jt}^m + \gamma_{dt}^{hp} \mathbb{I}(j \in B_i) \\ + \gamma_{dt}^{\text{dist}} \phi(j, B_i) + \gamma_{dt}^{\text{dist}2} \phi^2(j, B_i) + \xi_{jdt} + \sigma_d \epsilon_{ijt}. \quad (9)$$

Therefore, the set of parameters to be estimated are α_d^H and α_{jd}^m , the parameters governing the budget shares of consumption, housing and energy spending, respectively; γ_{dt}^{hp} , $\gamma_{dt}^{\text{dist}}$, and $\gamma_{dt}^{\text{dist}2}$, the parameters governing the strength of home premium and the disutility of living further away from one's birth state; ξ_{jdt} , the unobserved city-level amenities; and σ_d , the parameters that govern the variance of idiosyncratic preference draws.

It will be useful to express indirect utility as the sum of the component of utility that varies by household and the "mean utility" that is constant for households of a given demographic group. Dividing (9) by σ_d , we can write indirect utility as

$$\hat{V}_{ijt} = \mu_{jdt} + \hat{\gamma}_{dt}^{hp} \mathbb{I}(j \in B_i) + \hat{\gamma}_{dt}^{\text{dist}} \phi(j, B_i) + \hat{\gamma}_{dt}^{\text{dist}2} \phi^2(j, B_i) + \epsilon_{ijt},$$

where

$$\mu_{jdt} = \frac{(\alpha_{jd})}{\sigma_d} \log I_{jdt} - \frac{\alpha_d^H}{\sigma_d} \log R_{jt} - \sum_m \frac{\alpha_{jd}^m}{\sigma_d} \log P_{jt}^m + \hat{\xi}_{jdt} \quad (10)$$

and where "hatted" parameters represent a given parameter divided by σ_d . Using (10), we write the probability that an household i chooses a location j as

$$P_{ijt} = \frac{\exp(\mu_{jdt} + \hat{\gamma}_{dt}^{hp} \mathbb{I}(j \in B_i) + \hat{\gamma}_{dt}^{\text{dist}} \phi(j, B_i) + \hat{\gamma}_{dt}^{\text{dist}2} \phi^2(j, B_i))}{\sum_{j' \in J} \exp(\mu_{j'dt} + \hat{\gamma}_{dt}^{hp} \mathbb{I}(j' \in B_i) + \hat{\gamma}_{dt}^{\text{dist}} \phi(j', B_i) + \hat{\gamma}_{dt}^{\text{dist}2} \phi^2(j', B_i))}.$$

In the first step of estimation, we estimate the birth state premium parameters and the mean utility via maximum likelihood. The log-likelihood function is given by

$$\mathcal{L}(\hat{\gamma}_{dt}^{hp}, \hat{\gamma}_{dt}^{\text{dist}}, \hat{\gamma}_{dt}^{\text{dist}2}, \mu_{jdt}) = \sum_{i \in \text{Set}_d} \sum_j \mathbb{I}_{ij} \log(P_{ijt}),$$

where \mathbb{I}_{ij} is an indicator equal to one if individual i lives in location j and zero otherwise, and Set_d is the set of households in demographic group d .⁵⁷

In the second step of estimation, we decompose the mean utility terms, μ_{jdt} . First, we define $\tilde{\alpha}_{jd}^m = \frac{\alpha_{jd}^m}{\alpha_{jd}}$. Given the Cobb–Douglas utility function, the expenditure share on fuel type m of demographic group d in city j is given by $\frac{E_{jd}^m P_j^m}{I_{jd}} = \tilde{\alpha}_{jd}^m$. We first choose the $\tilde{\alpha}_{jd}^m$ parameters to match the expenditure share on each fuel type by each demographic group in each city. Specifically, we calculate the expenditure share on each type of fuel

⁵⁷Computationally, we invert the choice probabilities using the contraction mapping in Berry (1994) to obtain the unique mean utility associated with every guess of the parameter vector $[\hat{\gamma}_d^{hp} \hat{\gamma}_d^{\text{dist}} \hat{\gamma}_d^{\text{dist}2}]$.

by city and demographic group using our selection-corrected energy usage estimates from Section 3.1.⁵⁸

As we show in Online Appendix A.11, we can rewrite mean utility as

$$\mu_{jdt} = \beta_d^w \tilde{I}_{jdt} + \beta_d^r \log R_{jt} + \hat{\xi}_{jdt}, \quad (11)$$

where $\tilde{I}_{jdt} = \frac{\log I_{jdt} - \sum_m (\hat{\alpha}_{jd}^m \log P_{jt}^m)}{1 - \sum_m \hat{\alpha}_{jd}^m}$, $\beta_d^w = \frac{(1 + \alpha_d^H)}{\sigma_d}$, and $\beta_d^r = -\frac{(\alpha_d^H)}{\sigma_d}$. We refer to \tilde{I}_{jdt} as “energy-budget adjusted income.” This is a household’s log income after adjusting for the fact that (1) the fraction of income that is spent on energy depends on local energy prices and (2) income is more valuable in locations with high marginal utility of energy.

To limit the number of parameters to be estimated, we place additional restrictions on α_d^H and σ_d . As we show in Online Appendix A.1, conditional on location, a household’s marital status, presence of children, and age of the household head are the most important determinants of emissions. Conditional on these characteristics, the education level and race of the household head play only a minor role in determining emissions. As such, we allow the σ_d and α_d^H parameters to vary by household marital status and the presence of children.⁵⁹ Therefore, we can write these parameters as $\alpha_{\text{Marr,Child}}^H$

and $\sigma_{\text{Marr,Child}}$ and let $\beta_{\text{Marr,Child}}^w = \frac{(1 + \alpha_{\text{Marr,Child}}^H)}{\sigma_{\text{Marr,Child}}}$ and $\beta_{\text{Marr,Child}}^r = -\frac{(\alpha_{\text{Marr,Child}}^H)}{\sigma_{\text{Marr,Child}}}$.

Plugging in these parameter restrictions and taking first differences of (11) over our four datasets yield our estimation equation:

$$\Delta \mu_{jdt} = \beta_{\text{Marr,Child}}^w \Delta \tilde{I}_{jdt} + \beta_{\text{Marr,Child}}^r \Delta \log R_{jt} + \Delta \xi_{jdt}. \quad (12)$$

In general, we expect that changes in unobserved amenities, $\Delta \xi_{jdt}$, will be correlated with changes in rents, $\Delta \log R_{jt}$, and adjusted incomes $\Delta \tilde{I}_{jdt}$. For example, consider an increase in unobservable amenities in city j : this will increase utility directly and induce households into city j . Mechanically, this leads to an increase in housing demand, and as a result equilibrium rents rise, thus causing a change in $\Delta \log R_{jt}$. A similar argument can be made for adjusted income. As such, we estimate (12) via two-step GMM, using instrumental variables to deal with these endogeneity issues.⁶⁰

First, we use the measure of labor-demand shifts introduced by Katz and Murphy (1992) to generate variation in income across cities.⁶¹ The instrument interacts histor-

⁵⁸That is, we use the estimates of selection-corrected emissions from Section 3.2 and calculate demographic specific energy usage as $E_j^m = \hat{\alpha}_j^m + \hat{\beta}_j^m X_d$, where X_d gives the vector of demographic characteristics of households in group d . We discuss the implications of using the selection-corrected estimates for the counterfactuals in Section 8.5.

⁵⁹Many papers in the literature focus on differences in mobility by education group, rather than by marital status and the presence of children (e.g., Bound and Holzer (2000) or Diamond (2016)). However, since education is not a strong predictor of household emissions, we found it much more important to focus on marital status and the presence of children, which play a large role in determining household emissions. We also considered specifications in which we allowed these to vary by marital status, the presence of children, and age of the household head. These results are included in Online Appendix B.4.

⁶⁰Estimates via two-stage least squares, continuously updating GMM and limited information maximum likelihood are very similar.

⁶¹The instrument has been used as an instrument for cross city wage changes in Piyapromdee (2021) and Notowidigdo (2020).

ical industry concentration patterns at the city level with national changes in hours worked across industry. Formally, letting ι index industries, and letting $e(d)$ denote the education group associated with demographic group d , the Katz–Murphy index for city j from the previous period t' to the period in question t can be written as

$$\Delta Z_{jdt} = \sum_{\iota} \omega_{\iota j e(d)}^{1980} (\text{Hours}_{\iota, e(d), -j, t} - \text{Hours}_{\iota, e(d), -j, t'}),$$

where $\omega_{\iota j e(d)}^{1980}$ is the share of total hours worked in industry ι in city j by education group $e(d)$ in 1980 as a share of total hours worked in city j by education group $e(d)$ in 1980. $\text{Hours}_{\iota, e(d), -j, t}$ are the national hours worked in industry ι , education $e(d)$, for all cities besides city j . Therefore, the term in parentheses gives the change in national hours worked in industry ι between the current time period and the previous time period. Cities with historical concentrations of growing industries will generally experience increases in income while cities with declining industries will experience decreases in income. The income changes generated by this instrument are assumed to be uncorrelated with $\Delta \xi_{jdt}$, changes in city-level unobservable amenities.⁶²

To generate variation in rents, we also include the ψ_j^{WRI} , our measure of land-use restrictions, and the interaction between the Katz–Murphy index and ψ_j^{WRI} as instruments. In essence, cities with tighter housing supply restrictions and, therefore, more inelastic housing supply curves will experience larger changes in rents, especially in response to changes in population. As an example, if two cities experience a positive labor-demand shock, captured by a positive value of the Katz–Murphy index, the city with the more inelastic housing supply curve will experience a larger rent increase. This variation in rents is assumed to be uncorrelated with changes in unobservable amenities.⁶³

5.2 Particulate matter concentration

Next, we estimate the fixed level of $\text{PM}_{2.5}$, $\overline{\text{PM}}_j$, in each city. Recall from (7) that total particulate matter concentration in a city, PM_j , is equal to the sum of particulate matter that is endogenous to household location choices and $\overline{\text{PM}}_j$, the level of particulate matter arising from other sources. We measure PM_j using data on the overall level of ambient particulate matter concentration for each CBSA from the EPA's Air Quality Systems data. With the total level and endogenous component in hand, the exogenous component of particulate matter concentration can be calculated as

$$\overline{\text{PM}}_j = \text{PM}_j - \text{PM}_j^{\text{endog}}.$$

6. PARAMETER ESTIMATES AND MODEL VALIDATION

Location choice parameters Table 2 shows our estimates of β_d^w and β_d^r , the parameters which determine the location choice elasticities with respect to adjusted income and

⁶²See Goldsmith-Pinkham, Sorkin, and Swift (2020) for a discussion of identification with “Bartik”-style instruments.

⁶³We consider alternative instruments and parameterizations in Online Appendix B.4.

TABLE 2. Parameter estimates. Standard errors in parentheses.

	Single	Married	
		No Children	With Children
β^w : Adjusted income	15.09 (2.80)	11.72 (2.19)	7.33 (1.47)
β^r : Rent	-9.03 (2.40)	-6.90 (1.89)	-4.82 (1.29)
σ : Idiosyncratic component	0.17 (0.03)	0.21 (0.04)	0.40 (0.11)
α^H : Housing parameter	1.49 (0.48)	1.44 (0.48)	1.92 (0.73)
Cragg–Donald Wald F Statistic	8.09	8.28	8.57

rents. The first column provides estimates for single households, the second for married households without children, and the third column provides estimates for married households with children. We estimate that β_d^w and β_d^r are largest in magnitude for single households, and lowest for married households with children, implying that single households will be the most responsive in their location decisions to changes in policy. This heterogeneity in location choice elasticities has important implications for the effects of relaxing land-use restrictions in California. In particular, this implies that single households will be the most likely to move to California in response to this policy change. However, as we demonstrate in Online Appendix A.1, single households have the lowest average carbon emissions of these three groups while married households with children have the largest. Households who are the most mobile also have the smallest impact on carbon output.

To the best of our knowledge, we are the first paper to estimate these parameters by marital status and by the presence of children. Therefore, it is difficult to directly compare our estimates to those in the literature. However, it is reassuring that the magnitude of our location choice elasticities are similar to those in Colas and Hutchinson (2021) and only slightly larger than those in Diamond (2016), who estimate parameters, which vary by education but not by marital status or presence of children.⁶⁴

The next rows of Table 2 translate these estimates to estimates of α_d^H and σ_d .⁶⁵ We can use these parameters to calculate the budget share of housing for each demographic group in each city as $\frac{\alpha_d^H}{\alpha_{jd}^H}$. We find an average budget share of housing across household groups of 0.49, which is consistent with what has been previously estimated in the literature.⁶⁶

⁶⁴One concern is that we might suffer from a slight weak instruments problem as is relatively common in this literature. We assess the robustness of our key results to these parameter estimates in Section 8.1.

⁶⁵This parameters cannot be directly compared with similar parameters in Colas and Hutchinson (2021), Piyapromdee (2021), or Diamond (2016) as α^H here does not equal the budget share of housing. Here, the budget share of housing is given by $\frac{\alpha_d^H}{\alpha_{jd}^H}$.

⁶⁶For example, Serrato, Carlos, and Zidar (2016) estimated a budget share of housing of 0.3, using data from the Consumer Expenditure Survey, Moretti (2011) estimated a budget share of housing of 0.41 using Census data, and Diamond (2016) calibrated an expenditure share of local goods of 0.62.

Birth state premium Online Appendix B.7 gives the estimates of γ_{dt}^{hp} , γ_{dt}^{dist} , and γ_{dt}^{dist2} , the parameters governing the strength of home premium and the disutility of living further away from one's birth state, for each year. For all years and demographic groups, households receive a large utility premium for choosing a location in their birth state. The utility value of a location is decreasing and convex in distance from birth state for all demographic groups.

Amenities Table 3 provides selected estimates of ξ_{jdt} , the shared unobservable component of amenities, for the year 2017. Recall that this parameter is allowed to vary by demographic group d and location j , meaning we estimate a separate value of ξ_{jdt} for each of our 24 demographic groups in each of our 79 locations for each year of our data. Table 3 displays the five cities with the highest and lowest values of ξ_{jdt} for households in the younger age group that vary in their race, education level, marital status, and the presence of children. The estimates for households with older household heads are similar and are included in Online Appendix B.8.

Across demographic groups, Miami, Los Angeles, and Seattle consistently rank among the highest amenity cities while upstate New York cities generally have low

TABLE 3. Demographic group city ranks according to the shared, unobservable component of amenities for households with younger household heads.

Rank	College or more		Less than College	
	Single (no kids)	Married (with kids)	Single (no kids)	Married (with kids)
<i>Panel (a): White</i>				
1	Portland, OR	Portland, OR	San Diego, CA	Seattle, WA
2	Miami, FL	Miami, FL	Miami, FL	Portland, OR
3	Los Angeles, CA	Seattle, WA	Portland, OR	Los Angeles, CA
4	San Diego, CA	Los Angeles, CA	Seattle, WA	Honolulu, HI
5	Seattle, WA	San Diego, CA	Oxnard, CA	San Diego, CA
66	Memphis, TN	Memphis, TN	Springfield, MA	Memphis, TN
67	Youngstown, OH	Worcester, MA	Worcester, MA	Springfield, MA
68	Syracuse, NY	Springfield, MA	Albany, NY	Worcester, MA
69	Springfield, MA	Syracuse, NY	Rochester, NY	Albany, NY
70	Worcester, MA	Youngstown, OH	Syracuse, NY	Syracuse, NY
<i>Panel (b): Nonwhite</i>				
1	Honolulu, HI	Los Angeles, CA	Los Angeles, CA	Los Angeles, CA
2	Los Angeles, CA	Honolulu, HI	Miami, FL	Honolulu, HI
3	Miami, FL	Seattle, WA	San Francisco, CA	Seattle, WA
4	Portland, OR	Miami, FL	San Diego, CA	San Francisco, CA
5	San Diego, CA	San Francisco, CA	Seattle, WA	Portland, OR
66	Rochester, NY	Knoxville, TN	Springfield, MA	Springfield, MA
67	Scranton, PA	Milwaukee, WI	Syracuse, NY	Albany, NY
68	Milwaukee, WI	Syracuse, NY	Albany, NY	Syracuse, NY
69	Youngstown, OH	Springfield, MA	Milwaukee, WI	Rochester, NY
70	Springfield, MA	Youngstown, OH	Rochester, NY	Milwaukee, WI

amenities. There is also interesting heterogeneity across demographic groups—Portland is especially popular among educated white households, while Honolulu is more popular among minorities. Our estimates of ξ_{jdt} can be compared to estimates of “quality of life” from the urban economics literature.⁶⁷ Compared to Albouy (2012), our estimates assign slightly higher amenities to higher population cities relative to lower population cities. Consistent with Kahn (1995), we find that Los Angeles and San Francisco have higher amenities than Chicago and Houston in each year.

Model fit Next, we assess how well our model fits the data. The results from 2017 are summarized in Figure 3. Panel (a) shows the log number of households in each city in the data and the baseline simulation. Each circle represents a CBSA. Given that we estimate a separate unobserved amenity value for each demographic group and each city (ξ_{jdt}), we can match these moments exactly. Next, we plot the simulated and observed log average distance between an agent’s birth state and chosen city for each CBSA. The results are displayed in panel (b) of Figure 3. Each circle represents a CBSA, and the size of the circle is proportional to its population. The model fits this aspect of the data fairly well.

Panels (c) and (d) of Figure 3 show the predicted and actual average usage of natural gas and electricity in each city. As we allow the benefit of energy usage (α_{jdt}^m) to vary by city and demographic group, we can match these moments exactly.

7. COUNTERFACTUALS

In this section, we use the estimated model to simulate changes in land-use restrictions. The results of the counterfactuals are summarized in Table 4. The first column shows the population distribution, fuel usage, emissions, and income in the baseline specification, with all parameters set at their baseline levels. The following columns present these statistics in each counterfactual.

7.1 Relaxation of land-use restrictions in California

California Senate Bill 50—which recently failed in the California legislature—would have overridden tight land-use restrictions in California cities. In this section, we examine the effects of California adopting such a policy and relaxing local land-use restrictions. As shown in Section 2, California cities are among the most carbon efficient in the country. However, they also have very tight land-use restrictions—San Francisco and Los Angeles are in the 86th and 78th percentiles in the strictness of land-use restrictions, respectively. Intuitively, the relaxation of land-use restrictions in California will lead to increases in California’s population and decreases in overall carbon emissions. However, the magnitude of the decrease is an empirical and quantitative question.

Specifically, we simulate setting land-use restrictions, ψ^{WRI} , in California cities to the level faced by the median urban household.⁶⁸ We display the main results in the second

⁶⁷For example, Blomquist, Berger, and Hoehn (1988), Kahn (1995), or Albouy (2012). See Lambiri, Biagi, and Royuela (2007) for a review.

⁶⁸In a previous version of the paper, we simulated relaxing restrictions to the level of the median city.




FIGURE 3. Model fit results. Each circle represents a CBSA. Panel (a) shows the log number of households in each city in the data and the baseline simulation. Panel (b) plots the simulated and observed log average distance between an agent's birth state and chosen city for each city. The size of the circle is proportional to a city's population. Panels (c) and (d) show the predicted and actual average usage of natural gas and electricity in each city.

column of Tables 4. Setting land-use restrictions in California to the level of the median urban household leads to a 20.5% increase in the total population in California cities, a 3.1% drop in the population of other locations in the West, and 1% to 2% drops in the Midwest, South, and Northeast.

Panels II and III of Table 4 show how these changes in the distribution of households translate to average usage and emissions. The relaxation of land-use restrictions leads to decreases in usage of all three types of fuel, as households move to the temperate California climate. Specifically, natural gas usage drops by 0.3%, electricity by 0.5%, and fuel oil usage drops by 1.5%. Panel III of Table 4 displays average emissions resulting from each type of fuel. Electricity emissions drop by over 0.7% despite only a 0.5% decrease in usage. As power plants utilized in California are relatively carbon-efficient, the drop

TABLE 4. Counterfactual results. Each panel shows the simulated percent of the total population living in various geographic areas, mean energy usage, mean emissions, and mean income in each specification. See text for details on each simulation.

	(1) Baseline	(2) Relax Cali	(3) Relax All
I. Percent total population			
California cities	9.1	11.0	7.2
Other West	13.6	13.1	17.8
Midwest	22.2	21.7	9.3
South	37.3	36.6	23.1
Northeast	17.9	17.6	42.6
II. Mean usage			
Gas (1000 cubic feet)	74.4	74.2	74.9
Electricity (MW h)	17.1	17.0	15.8
Fuel oil (gallons)	60.4	59.5	138.6
III. Mean emissions (lbs of CO₂)			
Gas	8711	8686	8772
Electricity	16331	16211	13242
Fuel oil	1622	1598	3723
Total	26664	26495	25737
(%)	100	99.4	96.5
IV. Income (relative to baseline)			
Skilled	100.0	100.5	113.0
Unskilled	100.0	100.0	100.4
All	100.0	100.2	104.8

in emissions from electricity is larger than the drop in electricity usage. All together, this implies a drop in national household carbon emissions of 0.6% or a \$310 million dollar drop in the social cost of carbon annually.⁶⁹

In addition to low emissions, cities in California are very productive. Panel IV of Table 4 shows the effects on average income, relative to the average income in the baseline. The average income of skilled workers increases by roughly 0.5% while the average income of unskilled workers increases slightly. This leads to an increase in income of 0.2% across all workers. Overall, the shift toward more productive and lower-emitting cities increases the output to emissions ratio by 0.7%.

Regional effects To better understand the regional impacts of the policy change, Table 5 gives the change in population distribution and prices across regions. Panel II shows the change in regional demographic composition. The change in land-use restrictions leads to increases in the share of unmarried households and households without chil-

⁶⁹As mentioned in the [Introduction](#), we use the estimate of the social cost of carbon in 2020 from [Nordhaus \(2017\)](#).

TABLE 5. Changes in the composition of population in response to reduction in California land-use restrictions.

	California Cities	Other West	Midwest	South	Northeast
I. Household distribution					
% Change population	20.5	-3.1	-1.9	-2	-1.7
II. Composition					
Change in single share	2.4	-0.5	-0.2	-0.2	-0.2
Change share without children	2.3	-0.4	-0.2	-0.1	-0.1
Change in college share	-0.4	0.0	-0.1	-0.1	-0.1
Change minority share	0.1	-0.4	-0.3	-0.4	-0.3
III. Prices					
% Change skilled income	-0.1	0.2	0.2	0.2	0.2
% Change unskilled income	-1.0	0.2	0.0	0.0	0.0
% Change average rents	-4.8	-1.2	-0.7	-0.7	-0.8

dren in California. As these groups are relatively lower usage groups, this composition effect leads to slightly smaller decreases in carbon emissions than the population change alone.

Panel III of Table 5 displays the change in average incomes and rents across regions. Interestingly, within California cities, average income decreases slightly, as the proportion of households in low-income Fresno increases while the proportion in high-income San Jose decreases. Equilibrium rents in Californian cities drop by roughly 5%, as a result of both the change the land-use regulations and the resulting increase in population in California. Average rents decrease by roughly 1% in the other regions, reflecting drops in regional housing demand as households move to Californian cities.

Distributional effects In Table 6, we explore the distributional implications of the relaxation of land-use regulations in California. The four columns give the percentage change in average income, rents, utility, and exposure to PM_{2.5} for different demographic groups. Utility is measured in log dollar equivalents.⁷⁰ Highly educated households benefit more than less educated households because the income premium in cities is generally larger for educated workers (Baum-Snow and Pavan (2013)). Further, single households experience larger utility gains than married households, as single households are more mobile and, therefore, better able to benefit from the drop in rents in Californian cities. Finally, all demographic groups see a small increase in their average

⁷⁰Given that the idiosyncratic preference draws are distributed as Extreme-Value Type 1, household i 's expected utility is given by $\log(\sum_{j \in J} \exp(\bar{V}_{ij}/\sigma_d))$ plus a constant. To translate this into log income equivalent, we first divide expected utility by α_{jd} in each city j . Note that $\frac{\log(\sum_{j \in J} \exp(\bar{V}_{ij}/\sigma_d))}{\alpha_{jd}}$ gives expected utility measured in log income equivalent for a household who lives in city j —an increase in 0.01 in this object, for example, provides a change in expected utility equivalent to a 1% increase in income for household who lives in city j . We then take the average across cities weighted by the household's choice probabilities in the baseline counterfactual.

TABLE 6. Changes in average income, rents, utility, and pollution exposure by demographic group.

	Percentage Change from Baseline			
	Income	Rents	Utility	PM _{2.5} Exposure
I. Education				
College education	0.5	-0.7	0.9	0.3
Less than college	0.0	-0.9	0.6	0.3
II. Family size				
Single	0.2	-0.8	0.7	0.4
Married w/o children	0.2	-0.7	0.6	0.2
Married w/ children	0.1	-1.0	0.3	0.1
III. Race				
White	0.2	-0.8	0.6	0.2
Nonwhite	0.2	-0.9	0.8	0.4

PM_{2.5} exposure as they increase their concentration in large cities and in particular, the relatively polluted cities in Southern California.

Local pollutants Figure 4 plots the changes in particulate matter concentrations by CBSA when land-use restrictions in California are relaxed. Specifically, the x -axis gives the change in PM_{2.5} concentration when we relax land-use regulations in California




FIGURE 4. Histogram of CBSA level differences in particulate matter concentration from electricity relative when land-use restrictions in California are relaxed relative to the baseline.

compared to the baseline. The y -axis gives the number of cities that fall into a given range of changes. The colors indicate the four Census regions.

Similar to carbon emissions, there is a national reduction in $PM_{2.5}$ concentration from relaxing land-use regulations in California. The mechanism is quite similar to carbon emissions. Households are induced to live in California where they use less electricity—due to California’s temperate climate—and the power plants they use are less $PM_{2.5}$ -intensive. This leads to a reduction in $PM_{2.5}$ in most CBSAs.

However, unlike CO_2 , the spatial distribution of $PM_{2.5}$ emissions is important as $PM_{2.5}$ is a local and not global pollutant. When households move to California, this increases electricity demand in the WECC NERC region—which overlaps closely with the western Census region. This leads to an increase in the level of $PM_{2.5}$ emissions in WECC and, therefore, a slight increase in $PM_{2.5}$ concentrations in cities in the Western region. In all other regions, CBSAs experience decreases in $PM_{2.5}$. This is because energy demand falls in these regions as households move away, and the offset in $PM_{2.5}$ from local energy demand is greater than the increase in $PM_{2.5}$ from far away sources (namely, California). This decreases both emissions of $PM_{2.5}$ and concentration of $PM_{2.5}$ in these regions.

7.2 *Removing the correlation between land-use restrictions and emissions*

The negative correlation between land-use restrictions and city level emissions has important implications for national carbon emissions. To further explore the implications of land-use restrictions on carbon output, we simulate setting land-use restrictions to the level faced by the median urban household in all cities.

The results are displayed in the third column of Table 4.⁷¹ The results in panel I indicate that changing land-use restriction in all cities leads to a dramatic relocation from the South and Midwest to the West and Northeast. Specifically, the population in the Northeast region increases from 18% of the total population to 43% while the population in the Midwest and the South decrease by roughly one third to one half.

Panel II and III show usage and emissions from each energy type. Demand for natural gas and fuel oil are high while demand for electricity is low in the cold Northeast. As a result, natural gas usage increases by 0.7%, while electricity usage decreases by 7.8%. As a result of this decrease in electricity usage and relocation toward cities with more efficient power plants, emissions from electricity decrease by 18.9%. Overall, this leads to a 3.5% decrease in national carbon output and over an 8.5% increase in the national carbon efficiency of output. This implies a drop in the social cost of carbon of \$1.7 billion annually.

8. ROBUSTNESS AND EXTENSIONS

8.1 *Sensitivity to alternative parameters*

In this section, we examine the robustness of our main results to alternative values of key parameters. In particular, we recalculate the reduction in national carbon output

⁷¹In Online Appendix B.3, we show how the distribution of local pollutants change in this counterfactual.




FIGURE 5. Percentage change in national emissions from relaxing land-use restriction in California for various parameter values. In panel (a), we display σ_d for single on the y (vertical) axis, and σ_d for married households on the x (horizontal) axis. Panel (b) shows the percent reduction in carbon emissions as a function of α_d^H for both types of households.

resulting from the relaxation of land-use restrictions in California cities for a range of parameter values. First, we examine the model's sensitivity to the scale parameter of the idiosyncratic preference draw, σ_d . Lower values of σ_d imply that household location choice is more elastic with respect to wages and rents.⁷² Therefore, households will be more likely to change their location decisions in response to changes in land-use restrictions.

The percentage reduction in carbon output relative to the baseline for a range of values of σ_d for single and married households is shown in panel (a) of Figure 5. Recall in our baseline specification that we estimate $\sigma_d = 0.17$ for single households, $\sigma_d = 0.21$ for married households without children, and $\sigma_d = 0.40$ for married households with children and we found a decrease in carbon emissions of 0.6%. In the figure, σ_d for single households is displayed on the vertical axis and the average of σ_d for married households is displayed on the horizontal axis. We vary σ_d for married households such that the ratio of σ_d for married households with children compared to married households without children is held constant at the baseline level. Darker colors imply smaller changes in carbon emissions while lighter colors imply larger changes. Figure 5 illustrates that the change in carbon emissions is decreasing in σ_d for both single and married households. In the extreme case when $\sigma_d = 0.1$ for both types of households, households are very responsive to changes in rents. As a result, carbon emissions drop by 0.8% when we relax land-use restrictions in California. When $\sigma_d = 0.8$ for both types of households, carbon emissions drop by roughly 0.3%.

⁷²For each counterfactual in which we change σ_d , we recalculate the amenity values ξ_{jd} such as to keep the mean utility of each demographic group in each city equal to its baseline level. Therefore, the distribution of households of each demographic group given the baseline levels of land-use restrictions will be equal to the baseline distribution with the original values of σ_d .

Next, we examine the model's sensitivity to the budget share of housing parameter, α_d^H . Recall that we estimated $\alpha_d^H = 1.49$ for single households, $\alpha_d^H = 1.44$ for married households without children, and $\alpha_d^H = 1.92$ for married households with children. Higher values of α_d^H imply households spend a larger fraction of their income on housing and, therefore, will be more sensitive to housing prices in their choice of where to live. The results are displayed in panel (b) of Figure 5. The vertical axis shows values of α_d^H for single and the horizontal axis shows α_d^H for married households. We change α_d^H for married households such that the ratio of the parameter for married households with children to married households without children is held at the baseline level. Larger values of α_d^H of both types of households imply larger decreases in carbon emissions. When $\alpha_d^H = 0.4$ for both types of households, carbon emissions drop by 0.3% when we relax land-use restrictions in California. If we set $\alpha_d^H = 3.2$ for both types of households, carbon emissions drop by roughly 0.8%.

8.2 Endogenous electricity pricing

In our baseline specification, we assume that electricity is produced at a constant marginal cost; therefore, the supply curve of electricity is perfectly elastic. In this section, we consider an extension in which the price of electricity is determined endogenously.

Specifically, we assume that electricity producers in each NERC region form an upward sloping long-run electricity supply curve, reflecting differences in costs or productivity of potential electricity production opportunities within a region. For low quantities, electricity can be produced at a low cost. As electricity production increases, increasingly less productive resources must be utilized, which therefore implies higher costs of production.

A number of papers examine the short run supply curves of electricity. The short run electricity supply curve is often modeled as a “dispatch curve” with constant or linear marginal costs, to reflect the unique way in which electricity is allocated in the very short term.⁷³ Essentially, electricity generators are ranked in terms of their marginal cost of producing electricity. As demand increases, plants are dispatched to produce power in increasing order of marginal cost. However, this type of modeling approach is likely not a good representation of the long run energy supply curve which we consider here. In the long run, energy producers may respond to changes in energy demand by opening new reactors and new plants. Therefore, we posit a more parsimonious long run electricity supply curve as

$$C_{\mathcal{R}} = v_{\mathcal{R}} X_{\mathcal{R}}^{\kappa},$$

where $X_{\mathcal{R}}$ is the total quantity of energy produced in region \mathcal{R} , κ is a parameter equal to the inverse elasticity of the energy supply curve, and $v_{\mathcal{R}}$ is a region specific cost shifter.

Electricity is then transmitted to a specific local labor market at an additive transmission cost, ϕ_j .⁷⁴ Given the assumption of perfectly competitive generation companies,

⁷³For an example, see Ma, Sun, and Cheung (1999).

⁷⁴This cost may directly reflect costs of transmissions or network congestion costs.

we can write the inverse energy supply curve to city j as

$$P_j^{\text{elec}} = b_j + \kappa \log(X_{\mathcal{R}(j)}),$$

where $b_j = \phi_j + \log(v_{\mathcal{R}(j)})$.

To calibrate this model extension, we first calibrate the inverse elasticity of the electricity supply curve as $\kappa = \frac{1}{1.27}$, based on the estimates in [Dahl and Duggan \(1996\)](#). We then choose the parameters b_j to match state level electricity prices.

The main counterfactual results with endogenous electricity pricing are summarized in Online Appendix B.6. Overall, the population distributions across all counterfactuals are quite similar to the counterfactuals with perfectly elastic electricity supply. Households spend a relatively small fraction of their income on electricity and, therefore, changes in electricity prices have little impact on their location choices. Natural gas and fuel oil emissions are also nearly identical to the case with a perfectly elastic electricity supply. However, the reductions in electricity usage and, therefore, overall carbon emissions are smaller in the case of endogenous electricity prices. Overall this leads to a 0.4% reduction in carbon emissions from the relaxation of land-use restrictions in California.

8.3 Local pollutants in utility function

In this section, we consider an extension of our model in which local pollutants enter the utility function. It is worth emphasizing that only local pollutants produced by household electricity usage are endogenous in our model, while pollutants produced by all other sources are held constant. A richer model would also include how the distribution of households in space would lead to an increase in pollution from changes in the spatial distribution of manufacturing firms, for example.

With this caveat in mind, let the households' utility function be given by

$$u_{ij} = \alpha_d^c \log c + \alpha_d^H \log H + \sum_m \alpha_{jd}^m \log \hat{E}^m + \alpha^{\text{PM}} \log \text{PM}_j + \lambda_{ij},$$

where PM_j is the concentration of $\text{PM}_{2.5}$ in city j . We set $\frac{\alpha^{\text{PM}}}{\sigma_d} = -0.255$ for all demographics groups based on the estimates from [Bayer, Keohane, and Timmins \(2009\)](#).⁷⁵ We recalculate the unobserved amenity parameters, ξ_{jd} , such as to keep the mean utility of each demographic group in each city equal to its baseline level. All other parameters are kept at their baseline levels. Mechanically, the inclusion of $\text{PM}_{2.5}$ in the utility function will lead to higher estimated values of ξ_{jd} , so the population shares in the model match the data shares.

The results are displayed in Online Appendix B.6. The results are very similar to the baseline simulations. This is expected, given that $\text{PM}_{2.5}$ emissions from power plants only constitute a small fraction of total $\text{PM}_{2.5}$ emissions in a given city.⁷⁶

⁷⁵We use the instrumental variables results with no additional controls (Table 5, Column 3). The results with other IV estimates are similar.

⁷⁶We display the contribution of household electricity to total $\text{PM}_{2.5}$ concentration across various cities in Online Appendix B.3.

8.4 Power plant substitution

One potential issue with our counterfactuals is that new power plants built in order to accommodate increases in demand for electricity may be cleaner or dirtier than the current stock of power plants in that region. Therefore, the carbon emissions factors we use in our analysis will change in response to increases in electricity demand. For example, our main counterfactual of the relaxation of land-use restrictions in California led to a substantial increase in population and energy usage in California. As a result, new power plants may be constructed in the corresponding WECC NERC region, which may be cleaner or dirtier than the current power plants in the region. If these new power plants are cleaner than the current stock of power plants, we will underestimate the reduction in carbon emissions. On the other hand, we will overestimate the reduction in carbon emissions if these new power plants are dirtier.

To investigate how endogenous changes in the composition of power plants might affect our results, we compare power plants built before and after 2000. We find power plants built after 2000 emit considerably less CO₂ per MWh than plants built prior. Specifically, for the WECC NERC region, we find that power plants built prior to 2000 emit 858 lbs of CO₂ per MWh of electricity, whereas plants built after 2000 emit only 597 lbs of CO₂ per MWh.⁷⁷ These results suggest that if new power plants were built in response to increases in California's population, these new plants would be more carbon-efficient than the current stock of plants.

8.5 Selection across locations

Throughout our structural model, we have assumed that all households of a given demographic group who live in the same location have the same energy use. However, in reality households may have idiosyncratic propensities for energy use. In Section 3.1, we discussed how selection across locations based on these idiosyncratic propensities may have implications for average carbon emissions across locations and for estimates of predicted carbon usage.

To get a sense of how selection on the propensity to use energy would affect our counterfactual results, consider a version of the model in which household i in city j uses E_{ij}^m energy of type m , where E_{ij}^m can vary within demographic groups. Total carbon output is then given by the sum of all household carbon output:

$$\text{CO}_2 = \sum_m \delta_j^m \sum_d \sum_j \sum_{i \in \text{Set}_d} \mathbb{I}_{ij} E_{ij}^m,$$

where, as before, \mathbb{I}_{ij} is an indicator equal to one if household i lives in location j and Set_d is the set of households in demographic group d . We will consider how total carbon emissions will differ across two equilibria: a baseline equilibrium, where land-use restrictions are set to their current levels, and a relaxed equilibrium, where we relax land-use restrictions in California as in Section 7.1. For simplicity and to focus attention on

⁷⁷Online Appendix B.5 has further information on the full distribution of emissions from power plants built before and after 2000.

the role of selection, we assume a household's energy use is only a function of their location and, therefore, that E_{ij}^m is fixed across counterfactuals. This is equivalent to assuming that changes in land-use regulations will not affect a household's energy use, conditional on their location.⁷⁸

The change in total carbon emissions going from the baseline equilibrium to the relaxed equilibrium is given by

$$\Delta\text{CO}_2 = \sum_m \delta_j^m \sum_d \sum_j \sum_{i \in \text{Set}_d} (\mathbb{I}_{ij}^{\text{RELAX}} - \mathbb{I}_{ij}^{\text{BASE}}) E_{ij}^m, \quad (13)$$

where $\mathbb{I}_{ij}^{\text{RELAX}}$ is a function indicating household i chooses location j in the relaxed equilibrium, and $\mathbb{I}_{ij}^{\text{BASE}}$ is a function indicating household i chooses location j in the baseline equilibrium. We can rewrite (13) in terms of cross-city relocation flows as

$$\begin{aligned} \Delta\text{CO}_2 = & \sum_m \delta_j^m \sum_d \sum_j (N_{jd}^{\text{inflow}} \mathbb{E}[E_{ij}^m | \text{inflow}_{ij} = 1] \\ & - N_{jd}^{\text{outflow}} \mathbb{E}[E_{ij}^m | \text{outflow}_{ij} = 1]), \end{aligned} \quad (14)$$

where $\text{inflow}_{ij} = \mathbb{I}((\mathbb{I}_{ij}^{\text{RELAX}} - \mathbb{I}_{ij}^{\text{BASE}}) = 1)$ is a function indicating household i chooses j in the relaxed equilibrium but not the baseline equilibrium, $\text{outflow}_{ij} = \mathbb{I}((\mathbb{I}_{ij}^{\text{RELAX}} - \mathbb{I}_{ij}^{\text{BASE}}) = -1)$ is a function indicating household i chooses j in the baseline equilibrium but not the relaxed equilibrium, $N_{jd}^{\text{inflow}} = \sum_{i \in \text{Set}_d} \text{inflow}_{ij}$ is the total number of “inflows” in demographic d , and $N_{jd}^{\text{outflow}} = \sum_{i \in \text{Set}_d} \text{outflow}_{ij}$ is the total number of “outflows.” Therefore, the change in total emissions is given by the flows of households who relocate across equilibria multiplied by the average emissions conditional on relocating.

In our main specification above, we assumed that energy use conditional on choosing location j was common for all households of a given demographic group. In this case, the total change in emissions is given by

$$\Delta\text{CO}_2 = \sum_m \delta_j^m \sum_d \sum_j (N_{jd}^{\text{inflow}} E_{jd}^m - N_{jd}^{\text{outflow}} E_{jd}^m),$$

where we estimated E_{jd}^m as the selection-corrected energy usage for a given demographic group in location j . Therefore, selection on idiosyncratic propensity to use energy will lead to large misspecification errors if $\mathbb{E}[E_{ij}^m | \text{outflow}_{ij} = 1]$ and $\mathbb{E}[E_{ij}^m | \text{inflow}_{ij} = 1]$, the energy usages conditional on relocating, differ systematically from E_{jd}^m , the selection-corrected predicted energy usage.

To get a sense of how selection based on idiosyncratic propensity to use energy would change our main results, Table 7 evaluates (14) using several alternative assumptions for energy usage of households who relocate across equilibria. For these calculations, we use the values of N_{jd}^{outflow} and N_{jd}^{inflow} from the counterfactuals in Section 7.1.

⁷⁸In the main specification above, changes in the distribution of households will change income levels, which will therefore change energy use conditional on location. Quantitatively, we find the effects of endogenous income changes on energy usage to be small. Our assumption here is equivalent to assuming that (1) these income effects are not present or (2) that equilibrium wages are fixed.

TABLE 7. Change in national carbon emissions from relaxing land-use restrictions in California under alternative assumptions on selection on the propensity to use energy.

	Change in National Emissions
I. Main specification	−0.60%
II. Birthstate and demographic propensities	−0.59%
III. California propensities	−0.59%

First, in column 1, we set both $\mathbb{E}[E_{ij}^m | \text{outflow}_{ij} = 1]$ and $\mathbb{E}[E_{ij}^m | \text{inflow}_{ij} = 1]$ equal to E_{jd}^m as in our main specification. As in our main results above, we find that national carbon emissions decrease by 0.6%.

In our main results, we find that relaxing land-use restrictions in California leads to large influxes of households into California. One reasonable possibility is that households who relocate in response to relaxing land-use restrictions in California are similar in their idiosyncratic propensities to use energy as households of the same demographics and birth state who choose to live in California in the data. Recall that in Section 3.2, we used control functions to estimate the expected idiosyncratic propensity to use energy of households who live in location j . These control functions are functions of location-choice probabilities, where we estimated location-choice probabilities as the proportion of households of a given demographic group and birth place who choose each location. We can therefore use these estimates of the control function and location-choice probabilities to assign each household the same expected idiosyncratic propensity as households with the same demographics and birth place who choose to live in California in the data. Formally, let $\hat{M}_{CA}^m(\mathbf{P}_{iJ})$ denote the estimated selection control function for California as function of household i 's location-choice probabilities. We set

$$\mathbb{E}[E_{ij}^m | \text{outflow}_{ij} = 1] = E_{jd}^m + \mathbb{E}[\hat{M}_{CA}^m(\mathbf{P}_{iJ}) | \text{outflow}_{ij} = 1]$$

and

$$\mathbb{E}[E_{ij}^m | \text{inflow}_{ij} = 1] = E_{jd}^m + \mathbb{E}[\hat{M}_{CA}^m(\mathbf{P}_{iJ}) | \text{inflow}_{ij} = 1].$$

and evaluate (14). The results are shown in the second row of Table 7. We find that in this case, national carbon emissions drop by 0.59%, nearly identical to the quantification given the main specification.⁷⁹

A more extreme possibility is that all households who choose to relocate have the same propensity to use energy as the unconditional average household who chooses to locate in California in the baseline. Again we use our estimates of the selection control function and set

$$\mathbb{E}[E_{ij}^m | \text{outflow}_{ij} = 1] = \mathbb{E}[E_{ij}^m | \text{inflow}_{ij} = 1] = E_{jd}^m + \mathbb{E}[\hat{M}_j^m(\mathbf{P}_{iJ}) | \mathbb{I}_{iCA}^{\text{BASE}}].$$

⁷⁹The reason why this is slightly lower than our main results is that households who select to live in California have a slightly lower propensity to use electricity than the average household. Therefore, their move to California is slightly less consequential than the move of a higher electricity-use household.

The results are shown in the third row of Table 7. Again, we find that national carbon emissions drop by 0.59%. Overall, the results in this section suggest that selection on idiosyncratic propensity to use energy is unlikely to dramatically affect our main results.

9. CONCLUSION

Household carbon emissions vary considerably across cities. Land-use restrictions, which are set by local governments, tend to be tighter in cities with low carbon emissions and, therefore, encourage households to live in cities with less moderate climates and higher greenhouse gas-emitting power plants.

We began by following Glaeser and Kahn (2010) and documented large spatial variation in both the carbon efficiency of power plants and energy consumption. Cities with more temperate climates (such as San Francisco) tend to emit substantially less carbon than other cities. Furthermore, these cities also tend to have very tight land-use restrictions. To examine the effects of land-use restrictions on national carbon emissions, we then estimated a model of household sorting, energy demand, and locations that vary by power plant technology. We found that the relaxation of land-use restrictions in California leads to a decrease in national carbon output of 0.6% and a decrease in the social cost of carbon of \$310 million annually. Our main conclusion is that the positive correlation between tighter land-use restrictions and the greenness of cities has large implications for national carbon output.

Relaxing land-use restrictions would likely impact the spatial distribution of firms in addition to households. In our model, firms only use labor in production. In reality, many firms use the same energy inputs as households—such as electricity and natural gas. Additionally, firms also use land in production. Since local factor prices (such as land) are first order to firm location decisions (Serrato, Carlos, and Zidar (2016)), relaxing land-use restrictions in California would lead to firms sorting into California to take advantage of the lower land prices. Furthermore, when energy is an input to production and more production shifts toward California—where the electricity is carbon-efficient—carbon emissions would fall. Future research could incorporate firm sorting and energy demand into our framework to estimate the effects of land-use regulations on industrial carbon emissions.

Additional work could use our model to analyze the spatial implications of the Clean Air Act. The model could also be used to analyze changes in carbon emissions as a result of improved energy infrastructure and, therefore, easier electricity transmissions across regions. Future research could extend the model to analyze the effects of improving insulation or policies that change the composition of power plants—such as renewable energy subsidies.

REFERENCES

Ailshire, J. A. and E. M. Crimmins (2014), “Fine particulate matter air pollution and cognitive function among older US adults.” *American Journal of Epidemiology*, 180 (4), 359–366. [182]

Albouy, D. (2012), “Are big cities bad places to live? Estimating quality of life across metropolitan areas.” NBER Working Paper 14472, <https://www.nber.org/papers/w14472>. [206]

Albouy, D. and G. Ehrlich (2018), “Housing productivity and the social cost of land-use restrictions.” *Journal of Urban Economics*, 107, 101–120. [184, 194]

Albouy, D. and B. A. Stuart (2020), “Urban population and amenities: The neoclassical model of location.” *International Economic Review*, 61 (1), 127–158. [182]

Aldy, J., M. J. Kotchen, M. Evans, M. Fowlie, A. Levinson, and K. Palmer (2020), *Cobenefits and Regulatory Impact Analysis: Theory and Evidence From Federal Air Quality Regulations*. University of Chicago Press. [182]

Baum-Snow, N. and R. Pavan (2013), “Inequality and city size.” *Review of Economics and Statistics*, 95 (5), 1535–1548. [209]

Bayer, P., N. Keohane, and C. Timmins (2009), “Migration and hedonic valuation: The case of air quality.” *Journal of Environmental Economics and Management*, 58 (1), 1–14. [214]

Berry, S. (1994), “Estimating discrete-choice models of product differentiation.” *RAND Journal of Economics*, 25 (2), 242–262. [201]

Berry, S., J. Levinsohn, and A. Pakes (2004), “Differentiated product demand systems from a combination of micro and macro data: The new car market.” *Journal of Political Economy*, 112 (1), 68–105. [181, 200]

Blomquist, G. C., M. C. Berger, and J. P. Hoehn (1988), “New estimates of quality of life in urban areas.” *American Economic Review*, 89–107. [206]

Borck, R. (2016), “Will skyscrapers save the planet? Building height limits and urban greenhouse gas emissions.” *Regional Science and Urban Economics*, 58, 13–25. [183]

Bound, J. and H. J. Holzer (2000), “Demand shifts, population adjustments, and labor market outcomes during the 1980s.” *Journal of Labor Economics*, 18 (1), 20–54. [202]

Chamberlain, G. (1986), “Asymptotic efficiency in semi-parametric models with censoring.” *Journal of Econometrics*, 32 (2), 189–218. [188]

Chay, K. Y. and M. Greenstone (2003), “The impact of air pollution on infant mortality: Evidence from geographic variation in pollution shocks induced by a recession.” *The Quarterly Journal of Economics*, 118 (3), 1121–1167. [182]

Chay, K. Y. and M. Greenstone (2005), “Does air quality matter? Evidence from the housing market.” *Journal of Political Economy*, 113 (2), 376–424. [182, 183]

Colas, M. and K. Hutchinson (2021), “Heterogeneous workers and federal income taxes in a spatial equilibrium.” *American Economic Journal: Economic Policy*, 13 (2), 100–134. [181, 191, 194, 200, 204]

- Colas, M. and J. M. Morehouse (2022), "Supplement to 'The environmental cost of land-use restrictions.'" *Quantitative Economics Supplemental Material*, 13, <https://doi.org/10.3982/QE1544>. [184]
- Currie, J., L. Davis, M. Greenstone, and R. Walker (2015), "Environmental health risks and housing values: Evidence from 1,600 toxic plant openings and closings." *American Economic Review*, 105 (2), 678–709. [183]
- Dahl, C. and T. E. Duggan (1996), "US energy product supply elasticities: A survey and application to the US oil market." *Resource and Energy Economics*, 18 (3), 243–263. [214]
- Dahl, G. B. (2002), "Mobility and the return to education: Testing a Roy model with multiple markets." *Econometrica*, 70 (6), 2367–2420. [181, 186, 187]
- Diamond, R. (2016), "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000." *American Economic Review*, 106 (3), 479–524. [181, 184, 191, 194, 200, 202, 204]
- EPA (2019), "Our nation's air: Status and trends through 2018." <https://gispub.epa.gov/air/trendsreport/2018/>. [197, 198]
- EPA (2020), "CO-benefits risk assessment (COBRA) health impacts screening and mapping tool." Environmental Protection Agency. <https://www.epa.gov/statelocalenergy/co-benefits-risk-assessment-cobra-health-impacts-screening-and-mapping-tool>. [192]
- Fan, Q., K. Fisher-Vanden, and H. A. Klaiber (2018), "Climate change, migration, and regional economic impacts in the United States." *Journal of the Association of Environmental and Resource Economists*, 5 (3), 643–671. [183]
- Fowlie, M., E. A. Rubin, and R. Walker (2019), "Bringing satellite-based air quality estimates down to Earth." *AEA Papers and Proceedings*, 109, 283–288. [183]
- Fragkias, M., J. Lobo, D. Strumsky, and K. C. Seto (2013), "Does size matter? Scaling of CO2 emissions and US urban areas." *PLoS One*, 8 (6), e64727. [183, 193]
- Gaigné, C., S. Riou, and J.-F. Thisse (2012), "Are compact cities environmentally friendly?" *Journal of Urban Economics*, 72 (2–3), 123–136. [183]
- Glaeser, E. L., J. Gyourko, and R. Saks (2005), "Why is Manhattan so expensive? Regulation and the rise in housing prices." *The Journal of Law and Economics*, 48 (2), 331–369. [180]
- Glaeser, E. L. and M. E. Kahn (2010), "The greenness of cities: Carbon dioxide emissions and urban development." *Journal of Urban Economics*, 67 (3), 404–418. [179, 180, 183, 185, 190, 199, 218]
- Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020), "Bartik instruments: What, when, why, and how." *American Economic Review*, 110 (8), 2586–2624. [203]
- Goodkind, A. L., C. W. Tessum, J. S. Coggins, J. D. Hill, and J. D. Marshall (2019), "Fine-scale damage estimates of particulate matter air pollution reveal opportunities for

location-specific mitigation of emissions." *Proceedings of the National Academy of Sciences*, 116 (18), 8775–8780. [182, 192, 197]

Gyourko, J., A. Saiz, and A. A. Summers (2008), "A new measure of the local regulatory environment for housing markets: The Wharton residential land use regulatory index." *Urban Studies*, 45 (3), 693–729. [184]

Heckman, J. (1990), "Varieties of selection bias." *American Economic Review*, 80 (2), 313–318. [188]

Herkenhoff, K. F., L. E. Ohanian, and E. C. Prescott (2018), "Tarnishing the golden and empire states: Land-use restrictions and the US economic slowdown." *Journal of Monetary Economics*, 93, 89–109. [182, 194]

Hernandez-Cortes, D. and K. C. Meng (2020), "Do environmental markets cause environmental injustice? Evidence from California's carbon market." NBER working paper 27205. [192]

Holland, S. P. and E. T. Mansur (2008), "Is real-time pricing green? The environmental impacts of electricity demand variance." *The Review of Economics and Statistics*, 90 (3), 550–561. [185]

Holland, S. P., E. T. Mansur, N. Z. Muller, and A. J. Yates (2016), "Are there environmental benefits from driving electric vehicles? The importance of local factors." *American Economic Review*, 106 (12), 3700–3729. [195]

Holland, S. P., E. T. Mansur, N. Z. Muller, and A. J. Yates (2019), "Distributional effects of air pollution from electric vehicle adoption." *Journal of the Association of Environmental and Resource Economists*, 6 (S1), S65–S94. [183]

Hsieh, C.-T. and E. Moretti (2019), "Housing constraints and spatial misallocation." *American Economic Journal: Macroeconomics*, 11 (2), 1–39. [182, 194]

Jones, C. and D. M. Kammen (2014), "Spatial distribution of US household carbon footprints reveals suburbanization undermines greenhouse gas benefits of urban population density." *Environmental Science & Technology*, 48 (2), 895–902. [183]

Joskow, P. L. and J. Tirole (2000), "Transmission rights and market power on electric power networks." *The Rand Journal of Economics*, 450–487. [195]

Kahn, M. E. (1995), "A revealed preference approach to ranking city quality of life." *Journal of Urban Economics*, 38 (2), 221–235. [206]

Katz, L. F. and K. M. Murphy (1992), "Changes in relative wages, 1963–1987: Supply and demand factors." *Quarterly Journal of Economics*, 107 (1), 35–78. [202]

Kline, P. and E. Moretti (2014), "People, places, and public policy: Some simple welfare economics of local economic development programs." *Annual Review of Economics*, 6 (1), 629–662. [198]

Lambiri, D., B. Biagi, and V. Royuela (2007), "Quality of life in the economic and urban economic literature." *Social Indicators Research*, 84 (1), 1. [206]

- Larson, W. and A. Yezer (2015), “The energy implications of city size and density.” *Journal of Urban Economics*, 90, 35–49. [183]
- Lawler, J. J., D. J. Lewis, E. Nelson, A. J. Plantinga, S. Polasky, J. C. Withey, D. P. Helmers, S. Martinuzzi, D. Pennington, and V. C. Radeloff (2014), “Projected land-use change impacts on ecosystem services in the United States.” *Proceedings of the National Academy of Sciences*, 111 (20), 7492–7497. [193]
- Ma, X., D. Sun, and K. Cheung (1999), “Energy and reserve dispatch in a multi-zone electricity market.” *IEEE Transactions on Power Systems*, 14 (3), 913–919. [213]
- Mangum, K. (2017), “The role of housing in urban carbon emissions.” Andrew Young School of Policy Studies Research Paper Series, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2957749. [183]
- Morehouse, J. and E. Rubin (2021), “Downwind and out: The strategic dispersion of power plants and their pollution.” Center for Growth and Opportunity Working Paper Series, <https://www.thecgo.org/research/downwind-and-out-the-strategic-dispersion-of-power-plants-and-their-pollution/>. [197]
- Moretti, E. (2011), “Local labor markets.” In *Handbook of Labor Economics*, Vol. 4 (O. Ashenfelter and D. Card, eds.), 1237–1313, Elsevier, Amsterdam. [204]
- Muehlenbachs, L., E. Spiller, and C. Timmins (2015), “The housing market impacts of shale gas development.” *American Economic Review*, 105 (12), 3633–3659. [183]
- Muller, N. Z. and R. Mendelsohn (2007), “Measuring the damages of air pollution in the United States.” *Journal of Environmental Economics and Management*, 54 (1), 1–14. [192]
- Nordhaus, W. D. (2017), “Revisiting the social cost of carbon.” *Proceedings of the National Academy of Sciences*, 114 (7), 1518–1523. [181, 208]
- Notowidigdo, M. J. (2020), “The incidence of local labor demand shocks.” *Journal of Labor Economics*, 38 (3), 687–725. [202]
- Piyapromdee, S. (2021), “The impact of immigration on wages, internal migration, and welfare.” *The Review of Economic Studies*, 88 (1), 406–453. [181, 191, 194, 200, 202, 204]
- Quigley, J. M. and S. Raphael (2005), “Regulation and the high cost of housing in California.” *American Economic Review*, 95 (2), 323–328. [181]
- Ruggles, S., S. Flood, S. Foster, R. Goeken, J. Pacas, M. Shouweiler, and M. Sobek (2021), “Integrated public use microdata series: Version 11.0 [Machine-readable database].” University of Minnesota, Minneapolis. [184]
- Saiz, A. (2010), “The geographic determinants of housing supply.” *Quarterly Journal of Economics*, 125 (3), 1253–1296. [181, 184, 198]
- Serrato, S., J. Carlos, and O. Zidar (2016), “Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms.” *American Economic Review*, 106 (9), 2582–2624. [204, 218]

Stein, A. F., R. R. Draxler, G. D. Rolph, B. J. B. Stunder, M. D. Cohen, and F. Ngan (2015), “NOAA’s HYSPLIT atmospheric transport and dispersion modeling system.” *Bulletin of the American Meteorological Society*, 96 (12), 2059–2077. [192]

Tessum, C. W., J. D. Hill, and J. D. Marshall (2017), “InMAP: A model for air pollution interventions.” *PloS One*, 12 (4), e0176131. [182, 192, 197]

Zivin, J. S. G., M. J. Kotchen, and E. T. Mansur (2014), “Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies.” *Journal of Economic Behavior & Organization*, 107, 248–268. [195]

Co-editor Peter Arcidiacono handled this manuscript.

Manuscript received 3 February, 2020; final version accepted 3 August, 2021; available online 31 August, 2021.