4-1-2017

The Role of Housing in Carbon Emissions

Kyle Mangum
Georgia State University, kmangum@gsu.edu

W. J. Usery Workplace Research Group Paper Series

Working Paper 2017-4-1
April 2017

The Role of Housing in Carbon Emissions

Kyle Mangum
Georgia State University

Andrew Young School
of Policy Studies

The Role of Housing in Carbon Emissions

Kyle Mangum*

April 2017

Abstract

This paper studies the impact of regional housing and land stock allocations on carbon emissions and the role for policy therein. It first measures the role of home sizes and population densities in determining differences between U.S. urban areas in carbon footprint of their residents, which the literature has shown to be large. Then it develops a dynamic spatial equilibrium model of housing stock evolution between connected, heterogenous markets to measure the impact of housing policies on aggregate emissions. The main finding is that policies incentivizing an intensive use of housing have direct effects on carbon emission by increasing energy usage and creating lower density cities, and indirectly, these tilt population allocation towards higher emissions cities, which offer more housing consumption on average. The paper derives emission-equivalent values for the user cost of housing, land use regulations, and carbon taxes.

Keywords. carbon emissions, housing supply, land use, energy use, dynamic spatial equilibrium

JEL codes: R11, R52, Q54, R31

*Department of Economics, Georgia State University. Correspondence to: kmangum@gsu.edu, or Department of Economics, Box 3992, Atlanta, GA 30302-3992. Previous versions have circulated under the title "The Global Effects of Housing Policy." I thank the Duke Urban and Environmental Economics Working Group for data access and hosting. I gratefully acknowledge the Georgia State Research Initiation Grant for support of in the early stages of this research. I also thank Joe Aldy, Jim Alm, Peter Brummund, Bill Gentry, John Gibson, Christian Hilber, Steve Malpezzi, Erin Mansur, Juan Moreno-Cruz, and Tony Yezer for helpful discussions, as well as numerous seminar and conference participants. Taha Kasim and Bret Hewett have provided capable research assistance. Errors are my own.

1 Introduction

Climate change from carbon emissions is one of modern society's greatest challenges. As a long term, intergenerational externality problem, it has understandably received ample attention from the economics literature on several fronts.[1] One emerging strand of literature concerns the role of urban form in the determining the severity of human environmental impacts; that is, the way in which structures and populations are organized over space can exacerbate or mitigate emissions. The general findings are that denser urban populations tend to be more efficient in the sense of generating less carbon footprint per user (Larson et al. (2012), Larson and Yezer (2015), and Borck (2016)), though this is not theoretically unambiguous (Gaigné et al. (2012), Borck and Pflüger (2015)). Moreover, there are marked differences between regions in per capita emissions rates, even between large urbanized areas in developed countries. Glaeser and Kahn (2010) show that the carbon footprint of a typical household varies dramatically between U.S. metro areas, with households in some cities generating twice the carbon of those in others. This suggests there could be important compositional considerations–if consumption of goods were to occur in certain areas rather than others, aggregate carbon emissions might be lower.

These findings suggest roles for both the intensive margin (how much structure or land is being used) and extensive margin (where the consumption is occurring) in the aggregate generation of carbon emissions.[2] However, there has not been a unified treatment of both margins together in an equilibrium framework. This paper fills that crucial gap in the literature.

To study the distribution of people across locations and on the structure of the urban form, this paper focuses on the allocation of housing. Housing is an important topic in part because of its ubiquity, but also because policies on many political levels affect its form and consumption. In the U.S., there are numerous federal policies affecting the affordability of housing; in particular, owner occupied housing is given preferential tax treatment.[3] On the local level, various forms of land use controls comprise important policy levers.[4] Both types of policies could affect how much housing and land are consumed and in what location, but there has yet to be a study of their environmental impacts to the economy at large.

[1] These include measuring social costs, intertemporal discounting and uncertainty (see, for example, Stern (2008), Nordhaus (2014), Aldy and Stavins (2012), Greenstone et al. (2013) Weitzman (2014), Golosov et al. (2014), Dietz et al. (2015)), the role of growth and technological change (see, for example, Metcalf (2009), Gans (2012), Acemoglu et al. (2016)), the impacts of globalization, trade, and leakage (see, for example, (Elliott et al. (2010), Eichner and Pethig (2011), Baylis et al. (2014)), and the political economy of multilateral cooperation (see, for example, Nordhaus and Yang (1996), Battaglini and Harstad (2016)).

[2] Glaeser and Kahn (2010) advocate for some of both, writing, "If the urban population lived at higher population density levels closer to city centers in regions of the country with warmer winters and cooler summers in areas whose electric utilities used less coal for producing power, then household greenhouse gas production would be lower" (p. 416).

[3] For studies of the possible intensive and extensive margin effects of user cost subsidies, see Hanson (2012), Hilber and Turner (2014), Albouy and Hanson (2014).

[4] Fischel (2001), in a conference volume on local governance, reflects on the literature on local land use regulation and states, "Zoning is the most important municipal function."

To approach these questions, this paper develops a model of housing supply across an economy of interrelated heterogenous markets. Locations are heterogenous in how much land they can supply for housing at a given point of time, though this can be mediated through structure density (i.e. building homes using less land input) as well as the withholding of land, to extents also heterogeneous across markets. These features rationalize the intensive margin differences between locations in both the amount of housing services and the amount of land consumed per capita. Consumption levels then in turn help determine the extensive margin (how many people are located where) in the economy's equilibrium. Thus, in a spatial equilibrium across separate markets, the intensive and extensive margins are jointly determined, because demand for a given location will depend in part on how much housing service it has to offer, and also on what is being offered in the outside option (the rest of the locations).

A land development problem in a large equilibrium model imposes a challenging dynamic problem, so this paper relies on techniques developed in Mangum (2017) to solve and estimate the model. Estimation is facilitated by a rich dataset describing the living area and land intensity of housing stocks and flows in all large U.S. metro areas over a three-decade period. Combined with standard data on incomes and housing prices, the model is richly parameterized to fit the multiple forms of heterogeneity between metro areas in land and housing stocks. The data allow for a rigorous, empirically relevant application to the actual U.S. urban economy.

The purpose of the model is to use counterfactual simulations to evaluate the carbon emission effects of several important housing policies. One is to alter the user cost of housing (e.g. simulating the removal of a federal mortgage interest deduction). Other simulations alter local housing supply functions in order to evaluate local land use policies such as density restrictions (e.g. minimum lot sizes) and general regulatory burdens inhibiting land development. Both demand and supply policy simulations alter the patterns of land and housing consumption within and between locations. Then, applying data on energy use and carbon emissions rates by location, the model predicts the change to aggregate carbon emission from the economy.

The results indicate that policies' effects on intensive margins are especially important in determining aggregate carbon emissions. Policies that have reduced the user cost of housing have increased consumer demand for housing and land. The result is more energy used by in-home sources (more heating, cooling, lighting, etc.) as well as less dense cities and thus more driving and gasoline consumption. Moreover, housing subsidies tilt population towards cities that offer housing as an outsized portion of their utility bundle relative to amenities and other consumption. This shifts more population into higher carbon emission cities, though quantitatively the extensive margin effect is second order.

Similarly, local land use policies affecting the intensive margin can increase carbon emission. Density restrictions place constraints on the amount of land input to housing construction, which on average lowers population densities and increases housing consumption. These effects increase carbon from in-home and transportation sources. Land supply regulations can also affect the intensive margin. Simulations show that relaxing regulation in lower carbon places–a local land use "subsidy," in a sense, to low carbon locales–does little to reduce the aggregate emissions, in large part because lower rents increase the consumption of land and housing within those cities and across the economy. Increased restriction of land use in high carbon emitting places–a land use "tax"–has the opposite effect and can decrease aggregate emissions.

The simulations indicate that both the intensive and extensive margins matter when accounting for aggregate emissions, but quantitatively the intensive margin is far more important. This is surprising given the apparently large gaps between cities in carbon emissions rates, and there are three reasons the extensive margin effects are modest. First, as a simple accounting matter, population reallocations have to be very large to "move the needle" on the aggregate. Second, the equilibrium prices and quantities determine where population goes, and this assignment may not perfectly correlate with local carbon footprints. Some population exiting the highest carbon areas goes to only modestly less carbon intense places. Finally, a substantial portion of the gap between cities in carbon footprint is actually due to how much housing and land are consumed therein. This demonstrates that the extensive and intensive margins are closely linked, and therefore need to be studied together in a unified model. With endogenous housing prices, changes to housing or land stocks affect how many people can live in a location and how much housing and land they will consume there. At the same time, changing housing and land consumption will affect how many people want to live in a location relative to their options elsewhere in the economy. The model in this paper is able to account for the multiple dimensions meted out jointly in spatial equilibrium.

Finally, the paper considers the imposition of a carbon tax on the economy. This simulation serves two purposes. First, it reverses the thinking of the previous experiments to study the effect that carbon policy might have on housing, land, and population allocations. Second, it helps to denominate the magnitude of the housing policy simulations. Results indicate that had the U.S. implemented a carbon tax comparable to Sweden's, housing consumption would be nearly 12 percent lower and population densities higher by about the same proportion. Population would shift out of housing (and carbon) intensive inland markets and move to the Northeast and West Coast. Then, aggregate carbon calculations measure that a penny per pound carbon tax is emissions-equivalent to a user cost subsidy of four percent, a quarter acre decrease to lot size restrictions, and a half standard deviation increase to regulations in high carbon areas.

The conclusion of the simulations is that housing policy is quantitatively relevant for carbon emissions. In particular, housing-friendly polices in the U.S. have increased carbon emissions to a nontrivial degree. While the analysis is short of a complete welfare comparison of all available policy levers, it is clear from the results that housing policies local and national should enter the conversation of environmental policy, especially in regard to climate change.

The paper proceeds as follows. Section 2 describes the data employed in the paper and draws some stylized facts from it, with emphasis on the differences between cities in housing and land intensity and carbon footprint and national trends in these variables. Section 3 introduces the model motivated by the stylized facts, and section 4 describes the procedure for recovering key parameters. Section 5 reports on the baseline simulation. Section 6 presents results of counterfactual simulations for housing and carbon policies. Section 7 concludes and suggests avenues for future work.

2 Data

I begin by discussing empirical patterns on the connection between housing and carbon emissions. This paper uses data collected from many sources. This section describes construction of the data and presents motivating facts therein. Some additional sources of data used in estimation will be introduced after discussing the model.

2.1 Geography

The "city" level of geography is the metropolitan area as given by the U.S. Census definition of Core-Based Statistical Area (CBSA).[5] I focus on the 49 largest CBSAs by 2000 Census population and aggregate all others into a single residual location for a total of 50.[6] Identified cities comprise approximately two-thirds the national urban population, and the residual the remaining third. Key pieces of data are unavailable for rural areas, so the paper focuses on metropolitan areas. I first discuss the housing data and then energy data.

2.2 Housing Stock

The focus on housing's intensive margin and density in this paper is facilitated by detailed data on home living area and lot sizes by location and year of construction. This information comes from microdata of U.S. county tax assessor records compiled by real estate data firm Dataquick. The data are effectively a census of housing stock in their year of collection, a cross section of urban counties from 2011 to 2012. Furthermore, the property records contain year of construction, so I can measure the housing and land intensity per unit by the place and time of entry into the housing stock. The measure of housing-to-land density is the actual flooring

[5] Throughout, I use the term "city" to refer to an entire metropolitan area, not merely the central city.

[6] Fifty is an arbitrary number of locations that seemed a reasonably large cross section for a first pass. The smallest specified cities in my data are Salt Lake City, UT and Rochester, NY; the largest to be aggregated to the residual location are Bridgeport, CT and Tulsa, OK. New Orleans, LA is excluded because of the disruption to housing stock and construction caused by Hurricane Katrina.

area ratio (FAR), living area to lot size, for a particular structure. I refer to FAR as "structure density" to distinguish from population density.

This level of detail on the housing stock at such a wide geographic coverage is unique to this paper. However, the tradeoff in using these data is that for multifamily buildings, the coverage of units is irregular and lot sizes are not generally available. Throughout the analysis, I focus on single family homes. This limitation is unfortunately common, as lot size and FAR for apartment buildings are not widely available (see also the discussion in Larson et al. (2012), Larson and Yezer (2015)).[7] The concern that this data limitation imposes is that, without multifamily buildings, this paper's results on population density will tend to be understated.[8] On the other hand, the results will find a role for density, and this suggests that considerations of population density are relevant to single family lot size and coverage, not merely "skyscrapers" (Borck (2016)).

For measures of city population size by structure type, I use annual county population estimates from the Census,[9] aggregated to CBSAs, and scaled by the fraction of the metro area living in single-family homes taken from a five percent subsample of the decennial census (Ruggles et al. (2015)).[10]

To measure the housing stock for each metro area over time, I use single family housing units by county from the decennial Censuses of 1980, 1990, 2000, and 2010, aggregated to the CBSA level. For intercensal years, I allocate the decadal change in the housing stock by the level of building permit activity in the CBSA, as collected by the Census and provided by Housing and Urban Development's State of the Cities Database (SOCD). As in Glaeser et al. (2014), I use permitting activity as an index because these are a noisy estimate of actual building activity, and do not necessarily sum to the change in housing stock. The index allocates permits by their expected arrival to the inventory of housing, which may vary spatially and temporally, using annual regional summaries of permit-to-completion rates and times from the Census. The allocated units comprise "construction" by city-year. Then, to get the total stock of living area, H, and land employed in housing, A, I multiply the number of units added each city and year by the living area and lot sizes measured from the tax assessor records for that city and year of

[7]It is standard for tax assessors to collect the lot size of single family properties, but not the parcel dimensions of condominium or apartment buildings. Construction activity data also tend to be more precise for single family than multifamily, as there is less measurement error in the counting of units.

[8]By comparing census data with the lot size information, I find that across cities, total population density, population density among single family homes, and the share of population in multifamily homes are all positively correlated. Hence, it does not appear that density via multifamily housing is used as a substitute for single family density, but rather these are the outcomes of local land prices and policies.

[9]Population density is calculated over residential land and does not include commercial, industrial, or infrastructure (e.g. roads) land; hence, the estimates reported will in general be higher than, for instance, Census estimates which include all land.

[10]The share of population in single family homes varies between cities, but negligibly within a city over time. I use the simple average proportion. Nationally, the fraction of urban residents in single family homes is 62 percent (s.d. 7 percent).

construction:

$$H_{jt} = (1-\delta)H_{j,t-1} + units_{j,t} \cdot liv\_area_{j,t} \;, \quad A_{jt} = (1-\delta)A_{j,t-1} + units_{j,t} \cdot lotsize_{j,t}$$

where j indexes locations and t time. δ is a depreciation rate calibrated as described below. I treat stock in 1980 as the initial condition because it is the first year of permitting data in the SOCD.

Table 1 provides summary of cross sectional differences between metro areas, organized into two parts: stocks, the characteristics of the housing stock as of the end of the data in 2011, and flows, changes in the housing stock over the available data period, 1980-2011. The weighted national average appears in the first row, and some unweighted cross sectional statistics are reported in the bottom rows. Heterogeneities in cities' housing stock run the gamut. First, the table shows differences in housing and land intensities. Some cities offer residents a relatively large amount of housing as part of their consumption bundle (column 1). For example, a typical Atlanta resident consumes about 200 square feet more living area (one-third of the mean) than a typical San Francisco resident. Across this sample of major major metro areas, a standard deviation is one-sixth and the range is 60 percent of the mean. In addition, this housing can come at quite different structural densities (column 2). While the average single family structure has living area to lot size at a rate of 5.6 percent, many cities are notably more dense, some at ratios in excess of 20 percent (Las Vegas, Los Angeles, San Francisco). Together, housing per person and housing per unit of land determine land per person, i.e. the inverse of population density (column 3). The two components are correlated, but only weakly, and the spread in population density across cities is substantial–a standard deviation is two-thirds the mean, and the most and least dense are an order of magnitude apart. I emphasize that these differences are among single family residents in large urban areas.

The other columns of the table report on changes to the stock in the three decades since 1980. Again, there is ample heterogeneity. While all cities exhibited stock expansion (column 4), and most more than doubled their available living areas, there were some rapidly growing locations (Las Vegas, Phoenix, Orlando, Atlanta, much of Texas), while some Rust Belt (Pittsburgh, Buffalo, Detroit) and constrained coastal cities (Los Angeles, New York, Boston) saw much less. Housing stock expansion is of course largely correlated with population growth (column 5), but not perfectly so. That is, housing per capita changed to varying extents across cities. While houses grew bigger in all cities by about 33 percent, some cities much added larger homes than others, with housing per capita growing above 40 percent in several markets. The trend toward larger homes came at slightly higher structural densities (column 7), but the cross sectional heterogeneity is apparent. Substantially larger but only slightly denser homes means that population densities declined in most cities, although again with considerable heterogeneity, with a standard deviation as large as the mean.

How do these heterogeneous flows add up nationally? The national trends are summarized in Figure 1. The left figure reports trends in housing and land consumption per capita, the while right figure plots the structure density of construction flow and the accumulated housing stock. All series are indexed to one in 1980. For each series, there is an actual weighted average, as well as an average using initial 1980 housing stock as weights in order to distinguish between common trends within cities versus compositional differences in the national average over time (i.e. which cities are growing above national rates). For the most part, the reweighting has small effects, meaning that despite the aforementioned heterogeneities between cities and imbalances in stock and population growth, the trends are present across locations.

Clearly, there is a strong upward trend in housing per capita nationally. Though structure densities of the 2000s were higher than those prior, this only marginally brought up the density of the accumulated housing stock. A small structure density increase combined with a large increase in consumption of structure per capita leads to population densities falling by about 20 percent. Again, the trends are scarcely affected by reweighting, primarily because the increase in per capita housing consumption is so pervasive.

Figure 1: Trends in National Housing Stock Characteristics

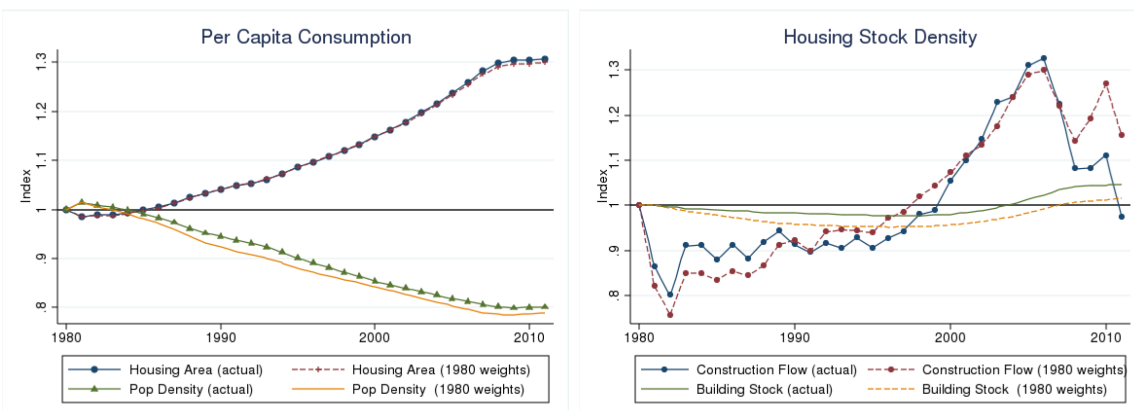

NOTES: The figure reports indices of the national aggregates of the series describing housing and land stocks and construction flows. Data sources and assembly are described in the main text. The series are indexed to one in 1980. The series label "1980 weights" refers to counterfactual calculation that fixes population weights to their 1980 shares rather.

The patterns in the housing data suggest that the model to be developed should feature a great degree of flexibility between cities, allowing stock to be added at different rates and densities, with consumers compensated for (lack of) housing in some locations with other consumption and/or local amenities. It should also allow for stock accumulation to account for trends in per capita consumption.

Table 1: Housing and Land: Summary Statistics by Metro Area

		1	2	3	4	5	6	7	8
			Stock, 2011				Flow, 1980-2011		
Size Rank	City	Housing Per Capita (sf per cap)	Structure Density (FAR)	Population Density (pop/sq mi)	Construction (%Δ stock)	Pop Growth (%Δ)	Change in Housing pp (%Δ)	Construction Density (FAR)	Change in Pop Dens. (%Δ)
	National Avg	596.3	0.056	2,612.9	76.2	45.6	130.7	0.057	−19.9
0	Residual	574.7	0.035	1,709.4	136.6	44.0	27.3	0.036	−16.6
1	New York, NY	606.8	0.117	5,355.6	88.9	16.2	22.2	0.099	−32.9
2	Los Angeles, CA	521.0	0.202	10,786.1	94.2	37.6	7.0	0.207	−2.9
3	Chicago, IL	592.7	0.112	5,275.9	127.6	18.0	48.8	0.104	−41.8
4	Philadelphia, PA	450.4	0.097	6,023.6	120.2	14.4	47.2	0.095	−33.7
5	Dallas, TX	694.3	0.120	4,805.3	237.7	116.3	23.7	0.122	−12.4
6	Miami, FL	539.8	0.189	9,774.7	157.1	76.1	13.5	0.207	4.9
7	Washington, DC	528.4	0.070	3,667.4	197.9	67.9	39.1	0.069	−27.3
8	Houston, TX	681.0	0.141	5,774.9	208.9	93.4	26.2	0.149	−1.7
9	Detroit, MI	623.6	0.074	3,324.9	95.2	−1.6	50.5	0.060	−52.3
10	Boston, MA	679.0	0.056	2,285.9	96.8	16.6	27.6	0.046	−40.0
11	Atlanta, GA	731.2	0.049	1,868.7	303.4	130.3	41.4	0.053	−1.2
12	San Francisco, CA	541.1	0.228	11,752.6	110.1	35.1	18.5	0.253	0.1
13	Riverside, CA	557.5	0.130	6,477.1	306.0	176.3	18.9	0.156	43.0
14	Phoenix, AZ	698.9	0.179	7,131.6	372.7	166.4	45.6	0.193	−4.6
15	Seattle, WA	611.9	0.086	3,904.2	155.4	67.2	18.8	0.088	−10.7
16	Minneapolis, MN	540.7	0.038	1,973.1	144.3	51.0	24.9	0.032	−38.1
17	San Diego, CA	550.0	0.061	3,083.8	147.2	68.7	13.3	0.049	−52.9
18	St Louis, MO	625.2	0.054	2,407.9	110.3	12.5	43.2	0.046	−41.3
19	Baltimore, MD	429.8	0.080	5,159.2	151.2	24.1	56.1	0.082	−31.2
20	Pittsburgh, PA	658.9	0.062	2,614.1	75.9	−10.9	47.1	0.051	−44.9
21	Tampa, FL	617.2	0.142	6,417.9	162.9	75.1	17.1	0.153	−1.0
22	Denver, CO	592.6	0.099	4,643.0	195.0	79.2	29.9	0.098	−21.0
23	Cleveland, OH	700.6	0.070	2,799.0	99.0	−4.9	58.7	0.064	−43.4
24	Cincinnati, OH	686.3	0.054	2,212.8	138.0	21.9	50.9	0.050	−40.8
25	Portland, OR	582.2	0.102	4,874.3	140.5	68.7	10.4	0.109	2.3
26	Kansas City, MO	724.2	0.094	3,634.5	144.1	36.5	40.0	0.096	−20.3
27	Sacramento, CA	627.8	0.105	4,666.0	229.1	97.9	32.1	0.116	3.8
28	San Jose, CA	471.1	0.193	11,411.3	93.1	41.3	2.4	0.229	16.6
29	San Antonio, TX	627.8	0.083	3,706.7	231.9	90.1	39.4	0.082	−29.4
30	Orlando, FL	638.3	0.116	5,076.5	306.7	169.8	21.3	0.129	17.6
31	Columbus, OH	607.8	0.059	2,728.1	149.8	46.3	32.4	0.056	−31.8
32	Providence, RI	746.8	0.051	1,911.1	100.4	12.5	34.6	0.042	−43.5
33	Norfolk, VA	661.2	0.091	3,849.1	157.7	39.2	43.1	0.092	−29.8
34	Indianapolis, IN	692.6	0.061	2,462.7	169.4	47.2	44.3	0.071	−9.1
35	Milwaukee, WI	570.2	0.117	5,697.3	98.2	11.8	34.9	0.115	−27.3
36	Las Vegas, NV	718.0	0.277	10,751.8	738.2	325.4	65.9	0.299	8.1
37	Charlotte, NC	687.7	0.061	2,464.2	248.5	109.9	33.6	0.076	26.2
38	Nashville, TN	666.9	0.032	1,325.3	188.5	77.2	28.3	0.036	−1.4
39	Austin, TX	676.9	0.101	4,139.1	362.6	204.8	23.0	0.104	−4.0
40	Memphis, TN	694.2	0.045	1,814.3	150.6	32.8	47.1	0.044	−36.2
41	Buffalo, NY	653.4	0.044	1,858.6	83.3	−8.8	50.8	0.036	−48.6
42	Louisville, KY	612.1	0.016	735.0	103.8	22.9	27.2	0.017	−11.9
43	Hartford, CT	782.6	0.045	1,605.5	117.1	15.4	43.7	0.040	−42.5
44	Jacksonville, FL	641.5	0.117	5,072.8	230.3	84.4	42.8	0.127	−11.5
45	Richmond, VA	690.8	0.035	1,412.4	184.8	51.2	46.9	0.033	−40.4
46	Oklahoma City, OK	638.4	0.044	1,923.4	139.4	46.6	26.3	0.044	−21.2
47	Birmingham, AL	645.6	0.025	1,099.9	102.3	21.7	27.6	0.024	−23.6
48	Rochester, NY	561.9	0.030	1,504.2	105.5	8.7	43.6	0.027	−38.5
49	Salt Lake City, UT	476.1	0.134	7,842.8	182.5	74.9	26.4	0.137	−12.8
	Mean	622.6	0.092	4,295.9	175.8	62.4	33.7	0.095	−19.1
	Std. Dev.	78.5	0.056	2,816.0	108.3	62.4	14.1	0.064	21.9
	Coef. Variation	0.12	0.60	0.65	0.61	1.00	0.41	0.66	-1.1
	Range (max-min)	352.8	0.261	11,017.6	662.3	336.3	63.5	0.281	95.9

NOTES: Data cover the period 1980-2011. Data sources and assembly are described in the main text. Cities are ranked by total census population, but the column "Pop." and "Pop. Density" use estimated population living in single family homes. Population density is calculated over *residential* land and does not include land in commercial, industrial, or infrastructure (e.g. roads) uses.

2.3 Energy Use

What is the impact of these housing characteristics on energy use and carbon emissions? Housing and land are consumed at different rates across locations and over time, but energy usage *per unit* of housing or land can also vary across space and over time. Ultimately, it is the goal of the model to predict energy usage and carbon emissions according to any allocation of housing stock, so I need information on energy usage rates per unit of stock, by location. This subsection briefly describes the data sources and discusses the relationship between housing consumption and carbon emissions.

There are several reasons why energy use per unit of housing or land could vary between locations, as discussed in detail by Glaeser and Kahn (2010). Two obvious potential reasons are differences in local price of energy and the area's climate–whether a household frequently runs a heater or air conditioner. Moreover, the carbon content of such energy use will depend on the fuel source of the region's electricity (e.g. coal versus solar) or households' available appliances and utilities (e.g. electric heat versus natural gas or an oil furnace). Similarly, rates of gasoline consumption will depend on road and public transit infrastructure, price of fuel, driving habits and car choice, and so on. I note from the outset that this paper is not a study of precisely why energy use rates may vary between locations, but rather a measurement of the connection between the amount and location of housing and the total emission of carbon. Thus, the goal of the energy measures is to remain inclusive of attributes outside the model (like weather) and predict usage based on the extensive and intensive margins on housing and land. For example, if residents of city A use more electricity than residents of city B because city A has more housing per person, the model captures this margin directly. However, if the residents of A use more electricity *per unit of housing* than B, because of, say, hotter summers or lower utility prices, the model simply summarizes it in the location-specific usage rate.

The data sources on energy use mostly follow Glaeser and Kahn (2010), expanded to longer time horizons. Gasoline consumption for households comes from the National Highway Transportation Survey (NHTS) in five surveys from 1983-2009. The outcome of interest is total gallons of fuel consumed per household by year. I limit analysis to unleaded gasoline vehicles. The NHTS includes specific metropolitan area location for the large cities used in my data, and for smaller cities (those in the residual locale), metro versus rural status is known, so rural drivers are excluded. To find average usage per household by city, I pool the data and run a regression of gasoline usage on location and time dummies. Usage rates for years in between surveys are linearly interpolated. The NHTS is representative of households with personal vehicles, but some urban residents do not own cars, and household sizes may vary by location. To find gasoline consumption per capita for each location, I scale the usage rate by persons per household and the percentage of households in the metro area who own personal vehicles, as taken from public use census files (Ruggles et al. (2015)).

In-home energy use is collected by the Residential Energy Consumption Survey (RECS). I focus on the common sources of energy–electricity, natural gas, and fuel oil. Not all homes use all energy sources. To account for differences in energy source usage, zeros are included in the averages.[11] The RECS data include the living area and age of the home, so I measure the rate of energy use per square foot of housing, by vintage, $e^{source}_{jt,vint}$, which allows, for instance, newer homes to be more energy efficient. The vintages are grouped into pre-1980, 1980s, 1990s, and 2000s. There are seven surveys from 1987-2009; usage rates for years between surveys are linearly interpolated. Geographic detail is more limited in RECS than NHTS, as the survey reports only metropolitan status (an indicator) and census region or, in the 2009 survey, sub-region, which can be as fine as state (e.g. California) or small groups of states (e.g. Indiana and Ohio). However, this will account for the first-order differences between locations in climate, energy source, and price. To measure metro-level energy usage rates, I limit the sample to homes in metro areas and assign these to their respective census sub-regions, using the sub-region means in 2009 to infer differences in within-region means in earlier survey years. The total energy consumed by metro area j at time t is then the sum of usage rate e times stock in each vintage,

$$energy^{source}_{jt} = \sum_{vint} e^{source}_{jt,vint} H_{jt,vint}$$

for each of the three sources of in-home energy.

Finally, the energy sources must be denominated in terms of carbon footprint. For gasoline, natural gas, and fuel oil, which are produced and traded in national markets, I borrow the carbon factors reported by Glaeser and Kahn (2010): 23.46 pounds per gallon for gasoline, 120.6 pounds per 1000 cubic feet for natural gas, and 26.86 pounds per gallon for fuel oil. The carbon content of electricity depends on the source of energy of power plants serving the region. The North American Electric Reliability Corporation (NERC), a nonprofit regulatory authority subject to U.S. and Canadian governmental oversight, publishes the carbon content of a kilowatt hour of electricity by state. I convert electricity usage to carbon using the "NERC factor" for each metro area's state, making housing unit weighted averages for cross-state metros.

First, I look at a snapshot of the cross sectional differences as of 2009.[12] Table 2 reports energy use and resultant carbon for an average resident of each metro area in the data; for comparison, a national urban weighted average is reported in the top row and the each column's coefficient of variation at the bottom. The table is sorted from lowest to highest carbon footprint. Columns 1 to 6 report carbon by source of consumption. Cities of the west, and in particular, California, dominate the top of the list, owing to their relatively low in-home usage rates (partly

[11]Notably, fuel oil appears predominantly in the Northeast and mostly older homes. I exclude propane and kerosene, which are rare and have no material impact on the results.

[12]2009 was the most recent year with both residential and gasoline energy surveys. Glaeser and Kahn (2010) focused on 2001.

due to mild weather) and lower carbon content of electricity. In these low carbon places, a larger share of carbon comes from gasoline, although many low carbon cities actually use gasoline at rates below the national average, despite several with reputations as "car cities." Moving down the list, one sees that additional carbon footprint comes from a variety of sources, but the higher carbon cities tend to be geographically central and southern, with hot, humid summers and, in the Midwest, cold winters. In-home usage rates increase notably compared to the west coast, and the carbon content of electricity (column 2 divided by column 1) is relatively higher. Across cities, there is more variation in in-home emissions than gasoline usage, especially since the carbon content of electricity varies over space. Together, these add up to some striking differences between cities in per capita usage rates (column 7), as first reported in Glaeser and Kahn (2010). The carbon footprint of residents in highest carbon cities is over double that of lowest carbon locations. Column 8 reports the percentage difference in emission rates relative to the national average.

How much of these differences are due to energy consumption per unit of land or housing (extensive margin, the place component) versus differences in the consumption of housing and land (intensive margin, the level component) that were shown in Table 1? In order to distinguish the two margins, I conduct an accounting counterfactual of the emission rates if all cities consumed housing and land at the national rate. For the three in-home energy sources, this is a straightforward projection of energy per square foot in each city times its relative housing consumption. For gasoline, I predict consumption by altering the city average by its relative population density times an estimate of gasoline consumption with respect to metro area density from Karathodorou et al. (2010).[13] Column 8 reports the predicted total carbon output if each city had per capita housing consumption and population density of the national average, and column 9 reports the percent differences. The results show there is some role for each margin. The low carbon cities at the top of the list still have lower than average carbon footprints, but the gaps are diminished, substantially in some cases. For example, San Francisco residents are responsible for 31 percent less carbon than an average U.S urban resident, but would emit only

[13]The projection is $\hat{gas} = gas_j + (dens_{natl} - dens_j) \times 0.33$. This leaves the local residual gas consumption intact, so that if a city had, say, an especially efficient subway system and therefore lower gas consumption than would be expected conditional on its density, this is preserved as extensive margin. There are several studies measuring the relationship between gasoline use and density, including Bento et al. (2005), Brownstone and Golob (2009), Su (2011), Larson et al. (2012), Larson and Yezer (2015). Karathodorou et al. (2010) was the most directly applicable to the current context because it measures the reduced form elasticity of gasoline use with respect to metro-level density. The NHTS includes the population density of the census tract of the respondent, but as Bento et al. (2005) point out, a particular household's location is likely a result of selection on driver type to neighborhoods within a metro area, and using neighborhood density would not produce a reliable estimate of the elasticity of gasoline use with respect to metro-level density. In general, concerns about differences in household energy use being due to household sorting (e.g. suburban households with preferences different from central city households) are less severe when comparing between metro areas than within them. All cities have some denser central neighborhoods and sparser suburbs. The point of emphasis here is the difference between local markets in how large are the homes and lots on average across the urban area.

five percent below average if they consumed housing and land at average rates. Likewise, higher carbon footprint cities at the bottom of the list still consume at above-average rates, but at smaller premia. For example, Atlanta residents emit carbon at a rate 29 percent higher than the national average, but if homes in Atlanta were smaller and denser at the national average, residents would emit at only eight percent more than average. Overall, the table shows that higher carbon cities have some attributes that cause higher emission rates per unit of housing or land, but they also tend to be less dense and offer more housing consumption. At uniform housing and land consumption, the variance between cities falls by half and the coefficient of variation by one-third.

What then are the implications of the national trends of rising housing consumption and falling population density? Figure 2 shows the national per capita rates of carbon emission from in-home sources (left plot) and gasoline (right plot). On each plot, there is a series of the actual emission rates as well as two counterfactuals: one using each city's actual emission rate but 1980 population weights to aggregate (to account for shifting population composition across cities), and another using each city's emission rates but limiting to 1980 levels of housing and land consumption. (The jaggedness of these series is due to the sporadic surveys and interpolation between. Solid dots on the actual series represent survey years.)

The left plot shows there is a general upward trend in carbon from in-home sources, despite newer vintages being more efficient per unit of living area. The 1980 population weighted series shows that there was a shift of population share to higher emitting locations, but this effect is small in aggregate. However, the rising home size has had a large effect on carbon from in-home sources. The 1980 home size counterfactual shows that usage rates would have been about 15 percent lower, and actually at a declining trend, had housing consumption stayed at constant levels. The right plot shows that carbon from gasoline has trended down with more fuel efficient vehicles, though this decline has been mitigated by falling population density. Had density stayed at 1980 levels, emission rates would be about 10 percent lower.

3 Model

The previous section showed that there are many margins in play: energy usage per unit of land or housing, housing per unit of land, housing per person, and land per person, and housing, land, and people per location. A serious treatment of the net effects of housing allocation on aggregate carbon emission should account for how all these various interactions and heterogeneities will add up. This calls for a unified model that incorporates housing consumption and land intensity in housing for heterogeneous locations, markets with inter-related populations, and locally specific rates of energy usage and carbon throughput. This section describes such a model.

Table 2: Average Per Capita Energy Use and Carbon Emission by Metro Area (2009)

		1	2	3	4	5	6	7	8	9	10
		Components						Actual		Counterfactual	
Size Rank	City	Electricity (kwh)	Carbon From Elec. (lbs)	Natural Gas (cu. ft.)	Fuel Oil (gal)	Carbon, All In-home (lbs)	Gasoline (gal)	Total Carbon (lbs)	Relative to Natl. (%Δ)	Total Carbon (lbs)	Relative to Natl. (%Δ)
	National Avg	3,660.1	4,655.9	13,183.3	14.4	6,633	293.9	13,531	0.00	13,531	0.00
28	San Jose, CA	2,035.9	1,217.4	12,021.0	0.0	2,667	265.2	8,891	−34.29	12,543	−7.30
13	Riverside, CA	2,271.8	1,358.6	13,249.5	0.0	2,956	253.4	8,904	−34.20	10,793	−20.24
2	Los Angeles, CA	2,178.8	1,302.9	12,992.6	0.0	2,870	264.3	9,073	−32.95	12,386	−8.46
12	San Francisco, CA	2,294.8	1,372.3	13,553.0	0.0	3,007	271.8	9,387	−30.63	12,808	−5.35
25	Portland, OR	4,773.6	1,686.1	14,817.6	0.0	3,473	273.3	9,887	−26.93	11,223	−17.06
27	Sacramento, CA	2,575.7	1,540.3	15,047.7	0.0	3,355	284.7	10,038	−25.82	11,103	−17.94
17	San Diego, CA	2,317.9	1,386.1	13,591.8	0.0	3,025	304.5	10,173	−24.82	10,846	−19.84
15	Seattle, WA	4,993.1	1,368.1	15,443.5	0.0	3,230	302.5	10,330	−23.66	11,137	−17.69
4	Philadelphia, PA	2,164.3	2,873.9	8,965.7	42.9	5,108	258.7	11,179	−17.38	14,528	7.36
1	New York, NY	2,200.1	1,689.3	19,241.3	59.2	5,599	255.7	11,599	−14.28	12,910	−4.59
48	Rochester, NY	1,841.4	1,465.8	14,954.4	69.6	5,139	279.9	11,709	−13.47	10,961	−18.99
35	Milwaukee, WI	1,869.5	3,144.5	15,071.6	0.0	4,962	313.8	12,328	−8.90	14,461	6.87
19	Baltimore, MD	2,781.1	3,932.5	4,375.6	23.7	5,096	311.0	12,395	−8.40	16,043	18.56
41	Buffalo, NY	2,161.0	1,720.2	17,458.7	78.3	5,929	278.3	12,461	−7.91	11,341	−16.19
3	Chicago, IL	2,859.0	3,345.0	25,533.7	0.0	6,424	267.6	12,704	−6.11	14,272	5.47
6	Miami, FL	4,987.7	6,628.7	993.7	0.0	6,749	255.9	12,754	−5.74	15,847	17.11
49	Salt Lake City, UT	2,489.4	5,093.2	16,663.3	0.0	7,103	254.5	13,076	−3.37	16,890	24.82
20	Pittsburgh, PA	3,175.5	4,055.1	9,948.9	69.7	7,127	263.7	13,317	−1.59	12,758	−5.71
16	Minneapolis, MN	2,434.3	4,005.1	15,266.0	6.2	6,013	315.5	13,418	−0.84	13,290	−1.78
36	Las Vegas, NV	4,248.4	5,217.0	18,904.1	0.0	7,497	261.8	13,640	0.81	15,287	12.97
0	Residual	3,711.3	4,817.2	12,683.8	12.9	6,694	297.9	13,687	1.15	12,987	−4.02
14	Phoenix, AZ	5,344.5	6,659.2	5,972.9	0.0	7,379	271.0	13,740	1.54	14,730	8.86
10	Boston, MA	2,136.9	2,071.2	11,550.2	120.7	6,705	307.1	13,912	2.82	12,751	−5.77
22	Denver, CO	2,098.1	4,007.3	22,839.8	0.0	6,762	307.0	13,967	3.22	15,195	12.29
21	Tampa, FL	5,569.8	7,402.3	1,108.2	0.0	7,536	293.0	14,412	6.51	16,125	19.17
32	Providence, RI	2,340.0	2,567.8	8,252.7	161.0	7,888	288.1	14,649	8.26	12,458	−7.93
43	Hartford, CT	2,415.5	1,763.3	5,730.3	188.5	7,518	307.5	14,734	8.89	11,884	−12.18
7	Washington, DC	3,547.4	5,773.5	6,385.5	20.4	7,092	326.7	14,760	9.08	16,353	20.85
45	Richmond, VA	4,652.8	5,597.3	9,525.5	14.3	7,131	328.4	14,838	9.65	12,366	−8.61
33	Norfolk, VA	4,549.1	5,714.7	9,266.3	13.3	7,190	328.8	14,907	10.17	15,282	12.94
30	Orlando, FL	5,734.7	7,621.4	792.9	0.0	7,717	309.1	14,973	10.65	15,943	17.83
31	Columbus, OH	3,187.5	6,091.4	17,608.2	0.0	8,215	288.8	14,994	10.81	14,907	10.16
9	Detroit, MI	2,894.5	4,336.0	27,809.3	0.0	7,690	317.8	15,149	11.96	15,560	14.99
18	St Louis, MO	3,777.2	5,767.2	16,422.2	0.0	7,748	317.0	15,187	12.24	14,713	8.73
44	Jacksonville, FL	5,709.6	7,588.1	916.3	0.0	7,699	319.8	15,204	12.36	16,324	20.64
29	San Antonio, TX	5,212.3	7,203.4	10,622.7	0.0	8,484	293.9	15,382	13.68	15,669	15.80
23	Cleveland, OH	3,656.3	6,987.2	20,633.2	0.0	9,475	293.3	16,359	20.90	15,302	13.09
8	Houston, TX	5,646.8	7,803.9	11,670.4	0.0	9,211	307.4	16,425	21.39	17,154	26.77
5	Dallas, TX	5,822.4	8,046.6	11,811.9	0.0	9,471	310.3	16,753	23.81	16,778	23.99
39	Austin, TX	5,701.1	7,879.0	11,050.0	0.0	9,211	322.3	16,776	23.98	16,651	23.05
47	Birmingham, AL	5,706.2	7,983.0	8,779.4	0.0	9,042	334.7	16,897	24.87	14,054	3.87
40	Memphis, TN	6,371.5	8,520.9	9,405.5	0.0	9,655	311.5	16,966	25.38	14,797	9.35
37	Charlotte, NC	6,313.9	8,239.6	9,197.0	0.0	9,349	343.4	17,409	28.66	15,912	17.59
11	Atlanta, GA	5,480.1	8,127.0	12,816.6	0.0	9,673	334.4	17,521	29.49	14,731	8.86
34	Indianapolis, IN	3,582.0	7,765.7	19,724.3	0.0	10,144	316.1	17,564	29.80	16,001	18.25
42	Louisville, KY	3,834.1	8,396.9	15,466.8	0.0	10,262	315.9	17,675	30.62	14,297	5.66
26	Kansas City, MO	4,199.6	7,912.1	20,609.2	0.0	10,397	312.4	17,729	31.02	16,764	23.89
38	Nashville, TN	6,114.1	8,767.6	8,632.2	0.0	9,809	341.5	17,824	31.72	14,971	10.64
24	Cincinnati, OH	3,828.1	8,018.9	18,844.3	0.0	10,291	334.8	18,148	34.12	16,465	21.68
46	Oklahoma City, OK	6,265.0	9,836.1	15,253.8	0.0	11,676	342.5	19,714	45.69	18,004	33.05
	Coef. Variation	0.38	0.53	0.47	2.34	0.33	0.08	0.19	4.76	0.13	2.46
	Variance ratio,(9)/(7)									0.50	

NOTES: The reports on carbon emissions by energy source and location for residents of the 50 selected locations (49 largest metro areas and a combined residual urban area). The rows are sorted by column 7, the calculated carbon footprint per resident. The "counterfactual" columns project emissions had the city been at the national average housing and land consumption rate.

Figure 2: Trends in National Carbon Emissions

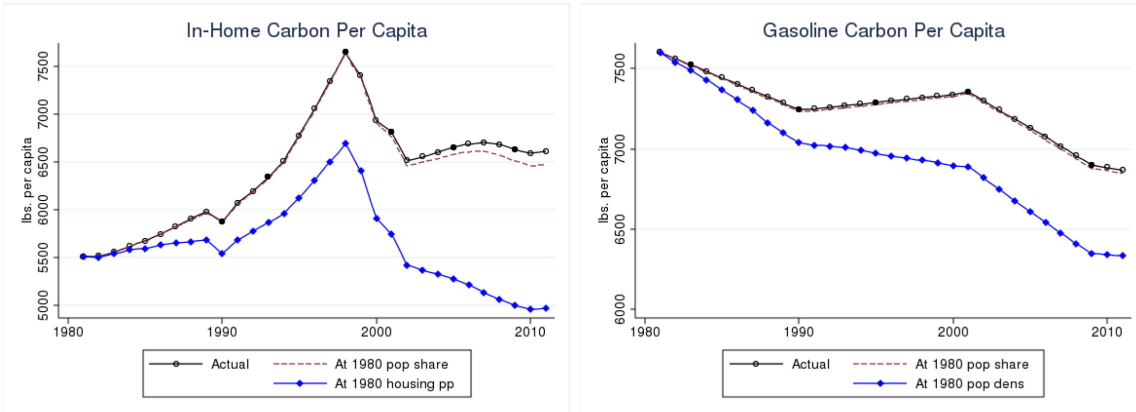

NOTES: The figure reports actual and counterfactual trends in national carbon emissions from the designated sources. Data sources and assembly are described in the main text. In the "Actual" series, shaded dots denote survey years and open dots denote interpolated years. "At 1980 pop share" fixes population weights to the 1980 distribution of population across cities. "At 1980 housing pp" or "At 1980 pop dens." uses the actual population weights, but projections emissions in housing and land consumption intensities remained at their 1980 levels.

3.1 Consumers

A closed economy populated by a measure P of identical, mobile individuals consists of J distinct markets, or "cities." An individual residing in city j draws the local income y without uncertainty, and obtains utility from numeraire consumption c, housing services h, and a local amenity μ, subject to a static budget constraint.[14] Housing services are nontradeable across locations and completely divisible within. The stock is rented to residents at the local price r. Utility is specified as

$$u(c,h,\mu) = ln(c) + \gamma ln(h) + \mu \;\; s.t. \; c + rh = y \tag{1}$$

where γ governs the elasticity of substitution of housing and consumption goods. (Section 4 discusses the choice of this particular utility function, which is convenient but not essential.) The important feature of the utility function is that the amount of housing services matters for utility and can therefore compensate for lack of numeraire or amenity goods. Therefore consumers in different locations may consume different amounts of housing (not a uniform "unit").

The consumer's first-order condition for optimization yields $\frac{h}{c} = \frac{\gamma}{r}$. Combining this with the budget constraint gives the individual's inverse demand equation for housing services:

$$r = \frac{\gamma}{1+\gamma}\frac{y}{h} \tag{2}$$

[14] Time and location subscripts are omitted until necessary.

3.2 Housing Market Equilibrium

The housing stock in a given period, H, is the result of past construction decisions by the builder. I describe the housing supply decision in more detail below, but it is predetermined for the purpose of determining equilibrium rents and population in a given period.

Within an arbitrary city, two conditions must hold. First, the residents must be at their optimal tradeoff of housing services and the consumption good, given by (2). Second, the housing market must clear: the housing services per resident sum to the total stock of (divisible) housing in the city:

$$p\,h = H \tag{3}$$

where p is local population. Equations (2) and (3) can be combined to obtain the city-wide demand curve

$$r = \frac{\gamma}{1+\gamma} y \frac{p}{H} \tag{4}$$

Ceteris paribus, rents rise in income and population, but fall with the supply of housing stock.

3.3 Spatial Equilibrium

As in the standard spatial equilibrium setting (Roback (1982)), full mobility implies zero arbitrage in utility so that residents must be indifferent between any two arbitrary cities j, k in terms of indirect utility U.

$$U_j(c_j^*, h_j^*, \mu_j) = U_k(c_k^*, h_k^*, \mu_k) \tag{5}$$

where * denotes quantities at the consumers' optimum. A spatial equilibrium obtains when (4) and (5) hold for all cities j, k. In equilibrium, indirect utility is equalized across locations, although the relative contributions of sources of utility (c, h, μ) may vary between locations j and k. Denoting the economy's indirect utility at a point in time as $\bar{U}$, and using the utility function (1) at a consumer's optimum in local market equilibrium (4), we have the condition

$$ln\frac{H}{p} = \bar{U} - ln(\frac{1}{1+\gamma}y) - \mu \tag{6}$$

Intuitively, (6) shows that consumers are willing to suffer less housing when local income and/or amenities are greater (when rents are higher). Moreover, combining (2) and (6) shows that rents will be lower, ceteris paribus, when the outside option is higher (when other places are relatively more attractive to consumers).

3.4 The Builder's Problem

The previous subsection described the equilibrium conditional on housing stock. This subsection describes how the builder supplies the housing stock, borrowing heavily from Mangum (2017).

Each location consists of one agent responsible for housing supply. The housing supplier is profit maximizing, lives infinitely over discrete periods, and discounts future profits at rate β. The agent is endowed with ownership of undeveloped land and decides whether and when to convert vacant land into housing, which can be sold (including its underlying land input) to consumers; hence, I will refer to this agent as "the builder." The stock of housing H is built atop a stock of land area A, and stocks depreciate at rate δ. The builder can add to the stock but cannot intentionally remove it. It is this irreversibility which makes the builder's problem dynamic: vacant land has continuation value, so its development exercises a real option. The builder's problem is written recursively as a Bellman equation,

$$V(X) = max_x[\pi(X, x) + \beta E(V(X'|X, x))] \tag{7}$$

where x are choice variables and X are state variables, evolving under the influence of x. I next define the choices and states.

3.4.1 Production and Input Factors

The builder makes housing services by combining land and capital. The production function is given by

$$i = F(a, k, \phi, \alpha) \tag{8}$$

where i is housing services added to stock, a is land input, k is capital, and ϕ and α are parameters governing housing productivity and the elasticity of substitution between inputs, respectively. In the empirics below, I use the standard Cobb-Douglas production function, $i = \phi a^{\alpha} k^{1-\alpha}$.[15]

The builder decides how much of each input to use each period. He adds to the housing stock according to (8), and the relative quantities of k and a determine density of building. A unit of capital comes at cost κ. This represents materials (wood, brick, glass, etc.) and the labor and equipment costs of their installation. A unit of land comes at cost ρ, and additionally, the builder faces a cost $c(\cdot)$ convex in the amount of housing added in one period. These costs represent the cost of converting virgin land to suitability for building.[16] Conceptually, the convex cost captures anything inelastically supplied in the city in a short time horizon (frictions in procuring new

[15]Recent work on the housing production function has found a constant returns to scale Cobb-Douglas function to be a reasonably good approximation. See Epple et al. (2010), Combes et al. (2012), Ahfeldt and McMillen (2014), Albouy and Ehrlich (2016).

[16]Recall that the builder already owns the land and resells it in post-production form. Any land costs are therefore in the acquisition, holding, and conversion of land, not its market value.

land, delays at the permit office, opposition to "excessive" new building, etc.). For exposition, I refer to it as the land assembly cost (the difficulty of getting all the housing together at the same point in time). Convexity is mathematically important as well. It ensures a finite solution, and effectively allows for within-city heterogeneity in costs and returns obscured by a single builder, meaning that under reasonable parameterizations, there is nonzero construction flow in each period, which is always a feature of the data for geographies as large as metro areas.

The housing services output is sold at a price R to be explained below. The builder's flow profits are

$$\pi = Ri - \kappa k - \rho a - \frac{c(\cdot)}{1+\nu} i^{1+\nu}$$

where $\nu \geq 0$ governs the extent of the convexity. Within a small area like a city, land is in limited supply. To reflect the exhaustibility of land, I specify the land assembly cost function to be increasing with the amount of land already employed, $c = c(A),\ c'(A) \geq 0$.[17] This means that the cost of adding more housing within a period is increasing in the amount of land already used up; or in other words, the short run elasticity is decreasing as land disappears. Intuitively, this could reflect increasing difficulty in land assembly, or greater likelihood of facing opposition from existing residents. Empirically, many large U.S. cities have seen housing prices increase faster than construction flows, indicative of decreasing price elasticity.

Assuming differentiability of the value function allows for the derivation of analytical first order conditions for optimization.[18] Taking first order conditions yields, after some manipulation,

$$\frac{F_a}{F_k} = \frac{\rho - \beta V_A}{\kappa}$$

For example, when F is CRS Cobb-Douglas,

$$\frac{F_a}{F_k} = \frac{\phi \alpha a^{\alpha-1} k^{1-\alpha}}{\phi(1-\alpha) a^{\alpha} k^{-\alpha}} \implies \frac{k}{a} = \frac{1-\alpha}{\alpha} \frac{\rho - \beta V_A}{\kappa} \tag{9}$$

which provides a condition for the derived demand for capital as a function of land (or vice versa). Using $\Delta = \frac{\rho - \beta V_A}{\kappa}$, the production function becomes

$$i = \phi(\Delta a)^{1-\alpha} a^{\alpha} = \phi(\frac{1-\alpha}{\alpha}\Delta)^{1-\alpha} a \tag{10}$$

which shows that the density of construction depends on the production function parameters

[17] This is not the only way to capture exhaustibility, but one specification which I found to be consistent with the construction flows data. I use the term "exhaustible" for exposition, though technically I am not imposing a maximum amount of land available.

[18] The flow profit function is smooth and continuously differentiable, so the value function should inherit these properties (Stokey et al. (1989), Thm. 9.10). The policy function may be discontinuous at a threshold; at low enough output prices, the builder can decide to construct zero.

and the relative costs of land and capital, Δ. Differences between locations in costs and/or function parameters will yield differences in structure density, motivated by Table 1. Differences in the shadow cost of land can be one component of this, as Mangum (2017) describes.

The factor demand condition (9) allows for the re-expression of (7) as a continuous choice of a single input. Using $\Phi = \phi(\frac{1-\alpha}{\alpha}\Delta)^{1-\alpha}$, (7) has become

$$V(X) = max_a[R\Phi a - (\Phi\kappa + \rho)a - \frac{c(A)}{1+\nu}(\Phi a)^{1+\nu} + \beta E(V(X'|X, a))]$$

Maintaining the assumption of the differentiability of V, the above yields an analytical expression for the policy functions.

$$a =(\frac{R\Phi - (\Phi\kappa + \rho) + \beta(V_H\Phi + V_A)}{c(A)\Phi^{1+\nu}})^{\frac{1}{\nu}} \tag{11a}$$

$$i =\Phi a \tag{11b}$$

The supply of new housing and land is jointly governed by the amount of land exercised, (11a), and the structure density, (11b), which is derived from primitives shown in (9). Each of these are functions of the shadow values of land (V_A) and housing stock (V_H).

3.4.2 Output Prices

The local builders are operating on islands of an economy in spatial equilibrium, and therefore demand for the location–and hence, the price of local housing construction–is affected by conditions in other locations. Combining the spatial equilibrium condition expressed as (6) with local optimization (2) and market clearing (3) yields

$$ln(r) = ln(\frac{\gamma}{1+\gamma}) + ln(\frac{1}{1+\gamma}y) + \mu - \bar{U}$$

The price of local housing services depend positively on general preference for housing, local income, and local amenities, and negatively on the reservation utility level in the economy (representing the outside option being higher). Housing consumption will appear in this last term, as described below.

Consumers rent housing each period but builders sell it after construction. The model is dynamic to allow for the potential for forward-looking agents to behave differently in heterogenous spaces. But, I want the model to focus on locational differences in housing supply and structure density in dynamic equilibrium, not a landlord's problem of renting existing stock to consumers or a complicated asset pricing problem. A modeling device is to assume the builder sells the stock to risk neutral middlemen who in turn rent it to residents. The price of this transaction should be determined as the present value of the future stream of rents. For simplicity, I assume risk neutrality and that incomes are a random walk, so the price is the discounted sum of current rents, inclusive of depreciation of the stock.[19]

$$R = \frac{1}{1-\delta\beta} r \tag{12}$$

The state space for the builder consists of his own local income, the amount of land used already, and the economy's utility. These states determine all costs and revenues, so I can now fully specify the builder's problem with laws of motion on state variables.

$$\begin{aligned} V(y, A, \bar{U}) = max_a[\pi(y, A, Y, \bar{U}) + \beta E(y', A', \bar{U}'|y, A, Y\bar{U}; a)] \\ s.t. \\ A' = (1-\delta)A + a \\ y' = f_y(y, \theta_y) \\ \bar{U}' = S(\bar{U}) \end{aligned} \tag{13}$$

where f_y is the law of motion for incomes and S is a function governing the evolution of the economy's level of utility to be determined in equilibrium, which I now describe.

3.5 Dynamic Equilibrium

Given a set of states describing current housing stocks $\{H_j\}_{j=1}^J$, land stock, $\{A_j\}_{j=1}^J$, and local incomes $\{y_j\}_{j=1}^J$, and their laws of motion, the equilibrium is defined by the following.

(i) Consumers are at their optimum bundle in each location j, (2) $\forall j$

(ii) Local housing markets clear in each location, (3) $\forall j$

(iii) There is no spatial arbitrage, (5)

(iv) The entire population is assigned to a location, $\sum_j^J p_j = P$.

(v) The builder in each location is behaving optimally given his states; that is, in each j, the builder is solving (13) by (11b), (11a).

The equilibrium is a complicated object because it involves the joint solution of J dynamic programs, each in its own heterogeneous space. All of these are solved jointly in an economy

[19] A practical issue for empirical implementation is that the builders either need to cover their cash cost of physical materials or have access to a financing market, and the latter would greatly complicate the model without much value added.

under a no-arbitrage spatial equilibrium, so that the level of utility is determined by the incomes, amenities, and housing stocks of all locations in the economy. The equilibrium cannot be characterized analytically. Even its numerical solution is infeasible if one were to solve the fully specified problem because the state space, consisting of all states of all locations simultaneously, is far too large to solve even one dynamic program, much less all J jointly. However, ignoring forward-looking behavior and instead assuming myopia may be consequential. If the model is to be estimated using construction data to infer primitive parameters, as intended here, assuming myopia may result in misspecification bias to a degree possibly heterogeneous across locations. For example, differences in parameters can impose differences in continuation values, so that ignoring forward-looking behavior can overstate parameter heterogeneity. Moreover, in counterfactual scenarios outside the discipline of data, forward-looking and myopic agents will respond differently to changes in the policy environment, also to possibly heterogeneous degrees, so that predictions can depend on the horizon of the agent.

These issues are discussed at length in Mangum (2017), which proposes a method for solving the dynamic equilibrium. The main argument, in brief, is to leverage the no-arbitrage spatial equilibrium condition to yield a feasible approximation to the full solution. Knowledge of local states and the economy's utility is sufficient to solve any of the builders' problems, so the equilibrium model can be solved by finding a rule governing the evolution of the economy's utility, denoted above as S. In a rational expectations equilibrium, the builders all know the rule S, and it correctly predicts their behavior, so that condition *(v)* produces S; formally, $S(V^*) = \sum_j h(V^*(S))$.

Thus, the spatial equilibrium condition dramatically simplifies the problem from each agent predicting the other $J-1$ response functions (as if it were a dynamic game) to one in which the behavior of other agents is summarized into a single law of motion for an endogenous variable, $\bar{U}$. Still, this can be hard to characterize, so a numerical approximation a la Krusell and Smith (1998) is employed, whereby a function S is iterated upon until behavior is consistent with a rational expectations equilibrium. More details and the computational method are available in Mangum (2017).

One more detail is necessary to implement the model empirically. The level of utility in the economy is not defined and must be normalized. There could be multiple ways of normalizing, but one empirically convenient way is specify that all locations deliver utility equivalent to a national average consumer, i.e. one drawing income and consuming housing at the national rate, with a amenity level of zero. This specifies that any single builder solving the dynamic problem track his own local income and land, and for $S()$ the national aggregate housing stock, income, and population. Rents and population shares are then defined in closed form as

$$\begin{aligned} lnr_j =& ln\Big(\frac{\gamma}{1+\gamma}\Big) + \frac{1+\gamma}{\gamma} ln(y_j) - \frac{1}{\gamma} ln(Y) + \frac{1}{\gamma}\mu_j + ln\frac{P}{H_N} && \text{(14a)} \\ \ln\frac{p_j}{P} =& \frac{1}{\gamma} ln\frac{y_j}{Y} + ln\frac{H_j}{H_N} + \frac{1}{\gamma}\mu_j && \text{(14b)} \end{aligned}$$

where j denotes the local value of each variable, Y is the economy's average income, and H_N is economy's aggregate housing stock. Housing services appear in the utility condition via (14a). In equilibrium, the economy's housing services per capita affect rents in aggregate, and then given local deviations in rents, local housing per capita is determined via (14b). For the builder, the shadow prices of land and housing stock, present in the policy functions (11b) and (11a), account for the effect changes to states will have on future costs and output prices.

In summary, the model is designed to study an economy of interconnected heterogeneous markets, and its essential features are described in two sets paired equations, (11a) and (11b), and (14a) and (14b), all of which jointly determine the state of housing and land stock allocations. Heterogeneity in the cost or production function parameters, as shown in the builder policy functions (11b), (11a), allows for differences between locations in land supply elasticity and structure density. A market that offers less income or less amenities can compensate residents by offering more housing. Moreover, there can be economy wide trends in the stock variables, since amount of housing nationally affects the equilibrium level of utility–more supply nationally changes the reservation level of utility, as seen in (14a). If housing stocks expand faster than populations, rents fall everywhere and housing per capita will increase. The effect of the housing stock expansion on population density depends on the structure density. That all of these components are jointly determined allows for a unified treatment of the various margins that might affect carbon emission within cities and the population allocation between them.

4 Estimation

The model is to be estimated on the 50 locations presented in section 2 above. Functional forms were introduced in discussion of the model. This section describes estimation of key parameters and some additional data.[20] Some parameters can be calibrated outside the model, others have simple estimating equations derived from the model, and the remainder rely on solution of the full (approximated) model.

[20]The section closely follows an appendix of Mangum (2017).

4.1 Initial Calibrations

Some preliminaries are calibrated outside the model. The data come at annual frequency, and transition processes are set accordingly. The builders' common discount rate is set to $\beta = 0.95$. Local per capita income is taken from the regional economic accounts data by the Bureau of Economic Analysis. The cities' annual income processes are found to have unit roots, and I take the empirical distribution of annual differences to be the shock distribution.

The depreciation parameter, δ, is calibrated using a regression of the national time series of housing stock and permits $H_{N,t} - i_{n,t-1} = (1-\delta)H_{n,t-1}$; I use the national series because it is longer than any metro (1947 using the Statistical Abstract of the United States) and the permits data are less subject to measurement error. Note that this measures stock depreciation, not value depreciation. I find $1 - \delta = 0.989$, indicating that about one percent of the single family housing stock is destroyed each year.

Materials costs parameters can also be estimated without solution of the model, but I defer discussion to place it in the context of the other city-specific costs.

4.2 Using Model Conditions

The parameters on utility, including the local amenity terms, are characterized by the spatial equilibrium conditions of (14a), (14b). This will require data on local incomes (just described) and rents. The price of housing in each market is derived from sales of newly constructed homes using the Dataquick microdata on deed transactions, which for most counties of the metro areas in the data can be matched to the assessor records using a unique property identifier. Then, I obtain home values per square foot by averaging transaction prices from 2004 and 2005 sales of newly constructed homes as identified by the year built field. In counties for which no transactions data were available, I used median value for homes built 2005 or later in the 2008-2010 American Community Survey (Ruggles et al. (2015)). The values were converted to 2000 dollars using the Consumer Price Index (less housing) and averaged by CBSA using housing units by county as weights. Values over time are pegged to the Federal Housing Finance Administration's (FHFA) all-transactions housing price index for the CBSA.

To connect sales values available to a builder to a flow cost to consumers, I use the commonly known user cost method (see e.g. Poterba (1992)), finding the implied rental rate, $r = uc \cdot v$, where uc is the user cost rate and v is the house value. Poterba (1992) suggests the user cost formula $uc = [(1-\kappa_t)(m+\kappa_p) + \psi + \nu_r] - \nu_g$, where κ_t is the income tax rate, m is the nominal mortgage rate, κ_p is the property tax rate, ψ is maintenance cost and depreciation, ν_r is the risk premium associated with housing, and ν_g is expected inflation. Calibration of these follows Poterba (1992), Albouy (2009), and my own estimate of average appreciation by market.

The model implies in (2) that housing expenditure constitutes a constant fraction of income, so that the utility parameter γ can be calibrated from expenditures on housing. Davis and Ortalo-Magne (2011) have shown using microdata on incomes and rents that average housing expenditure shares are remarkably consistent across metro areas, lending direct support for this type of utility function. However, the utility function has been challenged for its failure to hold across the income distribution of households (Black et al. (2014), Broxterman and Yezer (2015)). Considering the current model focuses on across city differences and uses homogenous agents, the functional specification seems appropriate for the context. Moreover, the analytical expressions for population and rent that Cobb-Douglas utility delivers are convenient, not a trivial consideration for equations that might otherwise need to be solved numerically millions of times in estimation. Using the mean expenditure share of 0.19 (s.d. 0.03), γ is set to 0.23.

With rents, incomes, population and housing stocks and an estimate of γ, I turn to the utility conditions of the model to derive the city-specific amenity value. To recover these, I run the regression

$$\frac{1}{\gamma}ln\frac{y_j}{Y} + ln\frac{h_j}{h_N} = D_j'\mu + \sigma ln\frac{y_j}{y_N}$$

where h_N is the national average housing services per person and D_j is a matrix of indicator variables for each city. Regressions like this are common in the spatial equilibrium literature, though this model has two small points of departure from standard practice. First, this spatial equilibrium condition directly accounts for the amount of housing services per person, and yet rents (prominent in standard Roback (1982) models) enter only indirectly through γ. The second is the use of an extra parameter σ. Without it, the empirical problem with the utility function (or any of the other standard utility functions over housing and numeraire consumption) is its pairing with the spatial equilibrium condition, as the same parameter is asked to govern average expenditure on housing and the elasticity of local population with respect to income. From (14b), this elasticity is $\frac{1}{\gamma}$, which is empirically far too high.[21] In models of this nature with local labor and housing prices, it is the role of the amenity term to rationalize cross sectional populations conditional on observed prices, but unlike most work employing residual amenity terms which deal only in the cross section, this paper needs to rationalize the cross sectional differences over time. Including the σ parameter relaxes this restriction on population elasticity. Intuitively, this means the amenity is interpreted in units of income.

These parameters are necessary for predicting local populations but not of direct interest for the purposes of this paper, so they are relegated to Appendix Table A1, with relative incomes,

[21] This problem, it should be noted, is not particular to the Cobb-Douglas form on utility. A constant elasticity of substitution form on utility, as advocated by Larson et al. (2012) and Larson and Yezer (2015), has similar issues in matching the population dynamics over time. I have tested this using the Larson et al. (2012) and Larson and Yezer (2015) calibration of $u(c,h) = (\beta_1 c^a + \beta_2 h^a)^{\frac{1}{a}}$ to $\beta_1 = 1$ (normalization), $\beta_2 = 0.1056$, and $a = -1/3$, as well as deriving $\beta_2 = 0.0677$, found in my dataset by the average of $\beta_2 = r(\frac{h}{y-rh})^{1-a}$ as implied by the first-order condition for utility maximization.

housing per capita, and rents reported for reference. The term σ is estimated to be 0.987 (s.e. 0.005), so that the average elasticity of population with respect to relative income is 0.056, much smaller than $\frac{1}{\gamma}$. This small adjustment of one additional parameter yields substantial improvements in the ability of the spatial equilibrium condition to fit city populations over time. With the addition of this one parameter, mean squared errors of predicted population fall 95 percent and mean squared errors on predicted rents fall 72 percent. This fit is important for the model to measure how populations would be affected by reallocation of the housing stock.

4.3 Estimation of Housing Supply Parameters

Having obtained parameters general to all locations and the amenity terms, I turn to estimation of the cost parameters for the builders in each metro area. The richness of the data, with the level and density of construction for many cities over a long time series, provides multiple target moments, thus allowing for estimation of multiple city-specific parameters. However, the data are actually insufficient to separately identify all the parameters specified in the general formulation of the model. Estimation will therefore normalize some parameters in such a way as to maintain the major dimensions of heterogeneity between cities.

First consider the relative factor intensities. These are identified by the building densities observed in the data.[22] Equation (9) is the model's density condition, and one would want to map this to observed densities, but there are two challenges in identification. The first challenge is that true capital is not actually observed (or easily defined, for that matter). Data obtained from the construction analytics firm RS Means Company report the costs of material with installation (labor and equipment) cost, on the basis of square foot of living area, annually by city for 1988-2013. This is the bundled capital cost for a square unit of finished housing, i.e. the k component of i, but not units of k. The only observable density condition is (10), housing services per unit of land. Hence, I set $\kappa = 1$ and subtract the capital cost component from the output price, so that (12) becomes $\frac{1}{1-\delta\beta}r - cc$, where cc is the construction cost from RS Means. There is significant spatial heterogeneity in the capital costs, but virtually no evidence that costs fluctuated with the level of building activity. The finding that construction labor and materials are elastically supplied is common (see Gyourko and Saiz (2006), Wheaton and Simonton (2007)), but costs vary spatially and temporally with construction sector wages (the "installation" component).[23] To measure the correlation of construction costs with income, I run a regression of cc_{jt} on y_{jt}, pooling the data across cities and including city fixed effects: $cc_{jt} = cc_{0,j} + cc_y y_{jt}$, for which I find $cc_y = 0.6142$ (s.e. 0.0388). Mean capital costs for each location are reported in Table 3.

The second challenge for factor intensities is that the parameters determining density–the

[22]Technically, everything is jointly identified, but intuitive connections guide the discussion.

[23]Discussions with data engineers at RS Means corroborated this.

TFP ϕ, the elasticity of substitution parameter α, and land cost ρ–are in practice difficult to identify. The functional form imposes the parameters to have slightly different implications for density, which in principle might be recovered, but in city level data, there are simply not enough data points to distinguish these. The important feature to maintain is locational heterogeneity in the structure density, the capital/land ratio. Thus, I elect to set $\rho = 1, \alpha = 0.5$ for all locations, and estimate ϕ. The only restrictive implication of this selection is that densities rise faster with the shadow cost of land in ceteris paribus dense places, but this feature is satisfied in the data.

Next consider the assembly costs. Intuitively, these are identified by level and elasticity of building activity (conditional on the density and construction material costs). The model's condition is (11a). I first need to specify the functional form for $c(A)$, which I set to be $\frac{1}{1+\nu}(c_1 + c_2 \frac{A_{jt}}{\hat{A}_j}) i^{1+\nu}$ subject to the constraints $c_1 \geq 0, c_2 \geq 0$. This form allows but does not impose that the assembly costs increase in the amount of land already in use ($\hat{A}_j$ is the average land stock in the data, a normalization to make the parameters comparable across cities of different physical size). In practice, it was difficult to separately identify the scale and convexity parameters, so I set $\nu = 1$ (making costs quadratic), which allows heterogeneity in elasticity to come through the c_1, c_2 parameters.

After reducing the parameters, the policy functions (11a), (11b) have become

$$\begin{aligned} a =& \frac{(R - cc)\Phi - 1 + \beta(V_H \Phi + V_A)}{c_1 + c_2 \frac{A_{jt}}{\hat{A}_j} \Phi^2} \\ \frac{i}{a} =& \Phi = \phi(1 - \beta V_A)^{0.5} \end{aligned}$$

These form the joint moment conditions in the estimation of ϕ, c_1, c_2 for each city. The objective function is a vector of squared residuals for each city's T observations,

$$M = \sum_{t=1}^{T} \begin{pmatrix} \left((\hat{a}_{jt}(c_1, c_2, \phi, \upsilon) - a_{jt})/\bar{a}_j\right)^2 \\ \left(\left(\frac{\hat{i}_{jt}(c_1,c_2,\phi,\upsilon)}{\hat{a}_{jt}(c_1,c_2,\phi,\upsilon)} - \frac{i_{jt}}{a_{jt}}\right)/\frac{\bar{i}_j}{\bar{a}_j}\right)^2 \end{pmatrix} \tag{16}$$

where the errors are proportional to the city's typical flow. The density moments can be noisy in some locations, so the residuals are weighted by a Gaussian kernel on their distance away from a quadratic trend. The estimates are then $\left(\hat{c_1}, \hat{c_2}, \hat{\phi}\right) = argmin\, M$. In practice, I find the parameters by first conducting a coarse grid search over a wide guess of values, and then a standard simplex-based minimization routine using the grid search outcome as starting values. This reduces concerns over finding local minima. Following Wooldridge (2002) on M-estimators, standard errors are calculated using numerical first and second derivatives to find the score and Hessian matrix for the objective function (16). The estimation begins with a rational expectation rule S as measured from the time series in the data, and then updates internally. In the first iteration, I use a rule S_0 from the data, estimate the parameters for each location, and then simulate the model to general a new aggregate series. I then run a regression to update the parameters of the rule to S_1, re-estimate the model for each city, and repeat until the parameters of the locations' primitives and the rule can no longer be updated. This ensures the recovered parameters deliver an equilibrium that is as close as possible to the data.

5 Baseline Model

Before engaging in the counterfactual simulations, I briefly discuss the structural estimates and fit of the baseline model simulation to the data.

5.1 Estimates

Table 3 reports the estimates for the structural parameters on local costs and productivity. The first column reports average materials costs, which is one dimension of heterogeneity in supply elasticity. All other cost estimates are arrived at conditioning on differences in the cost of materials and installation. Parameters c_1 and c_2 govern the supply elasticity via the convexity in the assembly cost function. Higher values of these parameters indicate a steeper supply function in which new construction is more quickly "choked off" when demand rises. For example, the parameters for San Francisco are considerably higher than for Atlanta in order to reflect less construction (on average and over time) despite higher housing values in San Francisco. On a technical note, the objective function moments are specified in levels (i.e. square feet of construction, not percentage change in stock), so that all else equal, the c_1 parameter rises as the size of the city falls. This is simply because a percentage change in stock in Salt Lake City, for instance, is many less square feet than a percentage change in New York.

The c_2 parameter governs how quickly the land supply function increases in slope as land is used up. If the cost is high, housing supply elasticity is reduced, though the builder can compensate by using less land per unit of housing (higher structure density). Therefore, densities will be higher as more land is used and in periods when the option value of land is deemed to be higher, the mechanisms by which structural densities trend upwards. Some of the most elastically supplied locations (such as cities in Texas), as well as some Rust Belt cities where there is relatively little land being developed, arrive at the zero constraint for this parameter.

Finally, the ϕ parameter governs the productivity of the local builder's housing production function. Essentially, this picks up differences in average structure density across locations that cannot be rationalized by differences in the option value of land. Thus, the parameter is closely correlated with the structural densities in Table 1, but is less disperse since c_2 also helps to determine optimal structure density, a feature of the using the dynamic model.

Table 3: Structural Parameter Estimates

		1	2	3	4	5	6	7
Size	Parameter:	Avg cc	c_1		c_2		ϕ	
Rank	City	($\$sf$)	*est*	*se*	*est*	*se*	*est*	*se*
0	Residual	53.82	50.7	4.8	0.0	4.3	0.036	0.002
1	New York, NY	75.78	391.9	78.0	272.7	21.3	0.051	0.004
2	Los Angeles, CA	62.67	791.6	84.8	404.2	57.7	0.119	0.009
3	Chicago, IL	64.29	1,210.5	77.2	66.4	41.9	0.090	0.007
4	Philadelphia, PA	64.64	1,504.2	363.1	1,080.8	161.8	0.048	0.005
5	Dallas, TX	49.85	867.7	52.3	0.0	16.2	0.109	0.007
6	Miami, FL	48.97	1,938.8	200.7	379.6	129.4	0.131	0.016
7	Washington, DC	53.58	1,908.9	215.8	574.7	114.7	0.043	0.004
8	Houston, TX	49.69	1,030.2	57.6	0.0	20.4	0.129	0.009
9	Detroit, MI	61.01	1,679.7	176.9	195.2	89.1	0.055	0.004
10	Boston, MA	67.08	820.9	259.7	532.7	41.0	0.030	0.003
11	Atlanta, GA	48.79	569.8	96.3	118.7	61.2	0.045	0.004
12	San Francisco, CA	69.98	3,327.9	2,775.3	767.5	1,784.8	0.120	0.081
13	Riverside, CA	63.24	1,119.8	97.3	99.6	66.8	0.137	0.015
14	Phoenix, AZ	51.58	1,069.7	84.8	54.7	24.6	0.157	0.013
15	Seattle, WA	57.21	1,427.7	283.0	685.0	161.2	0.053	0.005
16	Minneapolis, MN	63.88	1,985.2	303.5	514.1	112.4	0.028	0.002
17	San Diego, CA	61.89	3,434.1	380.4	0.0	229.8	0.061	0.013
18	St Louis, MO	57.11	1,849.6	271.9	494.5	129.5	0.040	0.003
19	Baltimore, MD	52.36	5,439.3	1,126.2	1,540.7	560.5	0.046	0.007
20	Pittsburgh, PA	57.92	1,696.6	415.8	617.0	152.0	0.039	0.003
21	Tampa, FL	49.28	1,810.6	178.2	409.9	136.8	0.104	0.009
22	Denver, CO	54.70	3,063.8	264.4	252.0	131.7	0.071	0.006
23	Cleveland, OH	60.92	3,086.8	251.3	0.0	144.9	0.061	0.003
24	Cincinnati, OH	55.33	2,042.0	265.6	270.3	116.6	0.043	0.003
25	Portland, OR	80.26	1,634.2	325.2	388.3	152.2	0.074	0.007
26	Kansas City, MO	57.87	1,827.3	241.1	306.4	144.1	0.082	0.008
27	Sacramento, CA	63.76	1,574.2	306.9	808.8	227.9	0.079	0.010
28	San Jose, CA	69.39	3,619.9	3,397.0	6,497.7	1,777.8	0.073	0.015
29	San Antonio, TX	46.62	2,331.3	175.9	0.0	86.1	0.077	0.006
30	Orlando, FL	49.77	2,053.1	185.2	160.9	131.6	0.103	0.011
31	Columbus, OH	54.93	3,567.1	316.3	0.0	158.0	0.052	0.003
32	Providence, RI	59.41	1,505.2	445.0	518.7	35.3	0.034	0.003
33	Norfolk, VA	49.53	3,662.1	2,579.6	0.0	1,564.2	0.086	0.080
34	Indianapolis, IN	54.93	1,961.4	340.5	517.6	205.0	0.055	0.004
35	Milwaukee, WI	58.55	8,425.3	645.1	0.0	327.9	0.102	0.008
36	Las Vegas, NV	59.97	2,463.7	226.5	198.7	54.9	0.199	0.016
37	Charlotte, NC	45.02	1,546.4	392.1	826.1	298.4	0.054	0.006
38	Nashville, TN	47.99	1,504.6	513.1	508.6	183.7	0.028	0.003
39	Austin, TX	46.10	2,752.6	165.1	0.0	46.2	0.098	0.006
40	Memphis, TN	48.45	3,304.0	566.8	0.0	295.8	0.042	0.004
41	Buffalo, NY	60.61	4,933.5	1,045.4	0.0	560.1	0.035	0.003
42	Louisville, KY	52.66	2,277.0	3,152.6	0.0	1,656.0	0.017	0.004
43	Hartford, CT	60.13	3,657.7	669.8	716.7	280.5	0.035	0.003
44	Jacksonville, FL	47.57	3,683.1	353.0	831.2	218.0	0.087	0.008
45	Richmond, VA	50.73	3,456.8	394.3	0.0	201.2	0.034	0.002
46	Oklahoma City, OK	46.56	3,712.3	424.9	0.0	242.3	0.043	0.003
47	Birmingham, AL	48.38	1,367.0	540.1	518.6	116.4	0.020	0.002
48	Rochester, NY	56.66	7,866.7	1,018.0	0.0	633.8	0.027	0.001
49	Salt Lake City, UT	49.31	10,271.2	469.3	48.1	74.4	0.138	0.005

NOTES: The table reports point estimates and standard errors for the local builders' structural parameters. Standard errors are calculated using numerical first and second derivatives of the objective function (16).

Some of the counterfactual simulations will alter local land use regulations, and so I will return to these estimates later to give a structural interpretation of what a change in land use policy would imply for the primitives the builder is subject to.

5.2 Model Fit

Figure 3 presents a series of scatter plots demonstrating the fit of the baseline simulation. The actual data is plotted on the horizontal axis and the simulated data on the vertical, with each point representing a city average or a city-by-time outcome, and diagonal lines drawn on each. Plot A displays total construction by each city. For the cross sectional variation in construction activity, the model's fit is nearly perfect. Plot B displays the annual construction in each city, introducing the year-over-year variation in the targeted moments. The fit is on average very good, but the dispersion horizontally (actual) is wider than the dispersion vertically (predicted). The model tends to under-predict the peak periods of construction and over-predict in trough periods. The reason is that time varying components of demand in this model–the "fundamentals," local incomes and population–are far less volatile than most housing cycles. Naturally, this is mostly present is cities with substantial housing cycles (e.g. Las Vegas) and less severe in cities with more muted construction cycles. Even so, the errors balance over time to produce the excellent cross-sectional fit in plot A. Thus, the model does well at matching heterogeneity in levels and trends, but is not as well suited to be a higher frequency (e.g. business cycle) model of construction activity.

Plot C displays the structure density, the FAR, for each city. To show the cross-sectional and temporal variation in FAR, there are separate series for each of three decades in all cities. The model is able to match the cross sectional differences in density reported in Table 1 as well as the trends to FAR exhibited in Figure 1. In particular, the heterogeneous trends in structure density across cities are picked up by the changing continuation value of land present in (9).

Plot D displays the population of each city as predicted by the spatial equilibrium condition (14b) for the beginning and end of the simulation. This plot shows that the spatial equilibrium model (with the augmented amenity specification discussed in section 4) does very well at using housing stocks and incomes to predict populations throughout the time period of the data.

Finally, plots E and F display, respectively, the predicted population density and housing per person compared to the data. These moments were not targeted, but by fitting the stocks, densities, and populations (plots A-D), these also match well.

Overall the model is able to fit the key components in the study of carbon impacts of housing and land stock allocation. It matches differences across cities in the level of building activity and its structure density, city populations, and housing and land consumption.

Figure 3: Model Fit

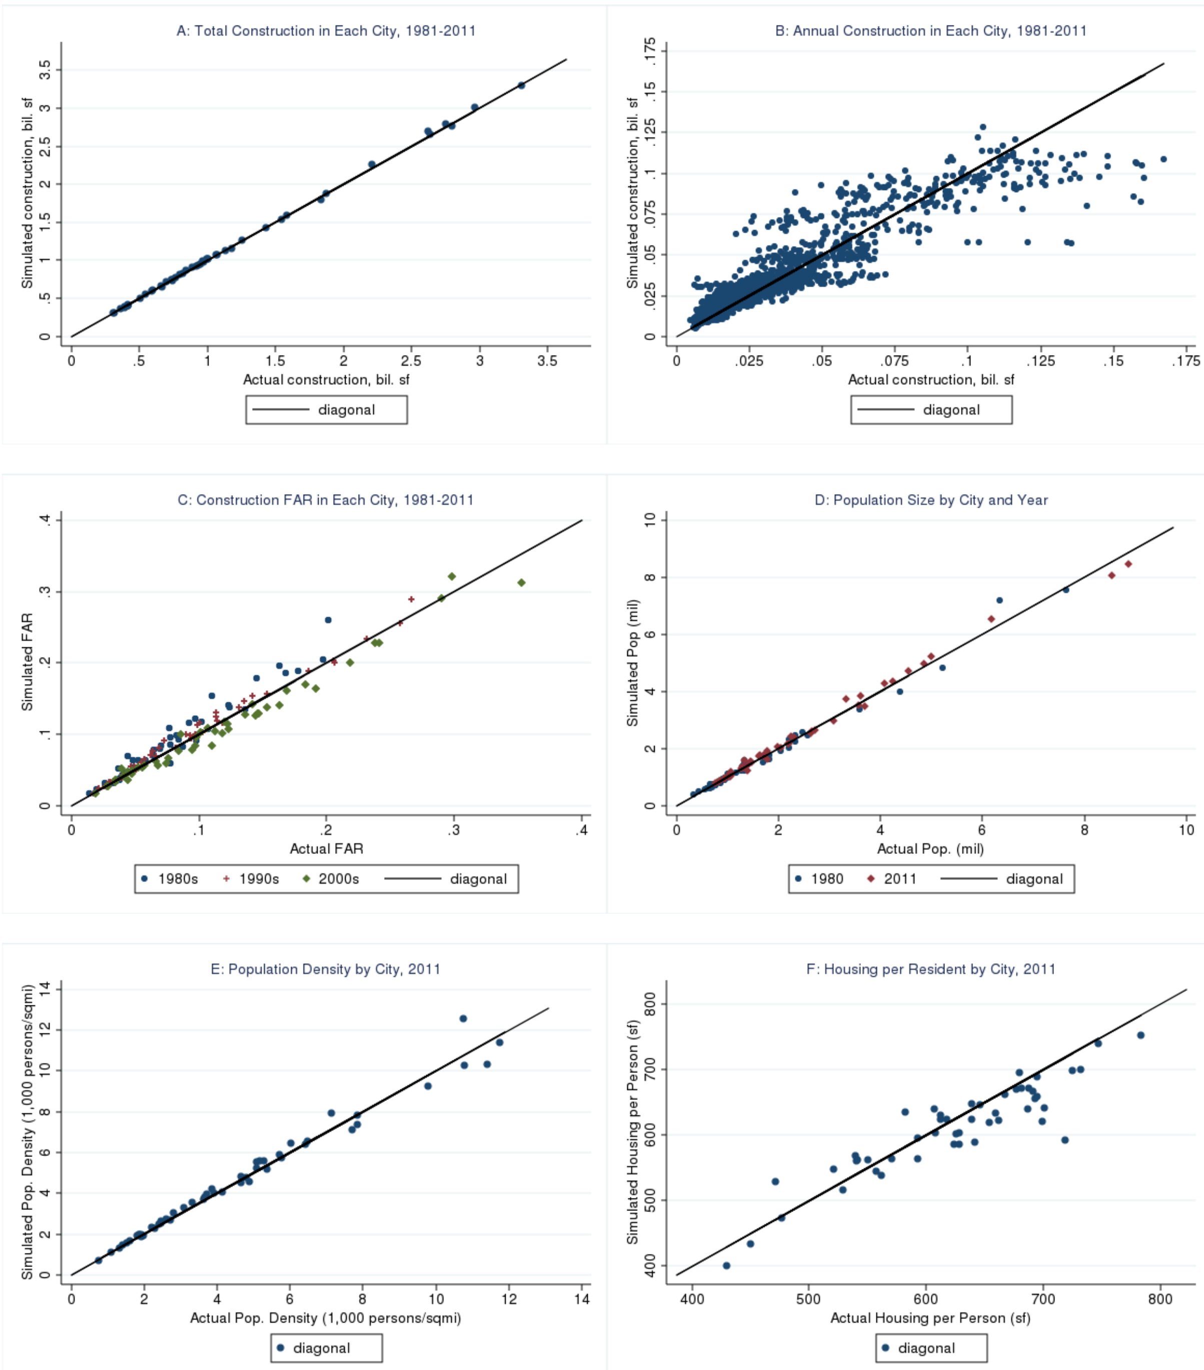

NOTES: The actual data is plotted on the horizontal axis and the simulated data on the vertical, with each point representing a city average or a city-by-year outcome, and diagonal lines drawn on each.

6 Counterfactual Simulations

The main use of a model of this nature is as a laboratory to study alternative policy regimes. This section presents results from counterfactual simulations of the model designed to study the connection between housing allocation, policy, and aggregate carbon emissions. I test three types of policies: (1) a housing demand policy, subsidies affecting the user cost, (2) housing supply policies, local regulations, and (3) the institution of a national carbon tax.

The policies will alter the level and density of construction in heterogeneous ways across markets, which will in turn affect housing and land consumption and the allocation of population across cities. As discussed above, all of these margins interact. Take housing away from a high housing consumption city, for instance, and its attractiveness in equilibrium will be altered, so that population will flow out. Unlike the "accounting" counterfactuals calculated in section 2, the economic model is designed to account for these interactions between housing and land intensities and the equilibrium extensive margin effects. The result is an integrated projection of how policies would affect the economy in equilibrium.

6.1 Projection of Carbon Emissions

First I clarify a few points the procedure for predicting carbon emissions in counterfactual scenarios. The model is designed to simulate housing and land stock evolutions, so to focus on the core endogenous variables, I use projections based on energy usage rates from the data.[24] I admit the possibility that some policies could fundamentally alter the way people use energy, but such concerns are outside the scope of the current study. As described in Section 2, the carbon calculation uses the usage rate per unit of housing by location and vintage, and a prediction of gasoline consumption under counterfactual population density using the elasticity estimate of Karathodorou et al. (2010), which uses worldwide data on metro areas to find an elasticity, $\frac{dln(gas)}{dln(density)} = -0.33$. This is a reduced form estimate inclusive of various mechanisms by which density may affect gasoline consumption.

Like Glaeser and Kahn (2010), my calculations ignore carbon coming from numeraire consumption. This is much harder to measure since it depends on the source and nature of tradables. Moreover, the household's utility function predicts a constant fraction of income be spent on consumption, so there are by assumption no effects of housing policy on carbon from consumption.[25] This assumption is coincidental to the issue described in Section 4 and not essential for

[24] The underlying assumption is that reallocated residents will use energy like incumbents, and that the marginal resident uses like the average. This assumption is reasonable if between city differences in energy usage rates are due to fixed factors, like climate, energy prices, typical building materials, and transportation structure. It would be violated if there were between city sorting on energy preferences–e.g. if Cincinnati was a collection of people who liked using electricity, and would still use at higher rates if relocated.

[25] Some simulations (like the carbon tax) alter the entire budget set, and thus numeraire consumption would move proportionally, magnifying the effects I report. Still, housing stock is the focus of this paper.

the simulations. For a treatment of the effects of carbon from consumption, see Larson and Yezer (2015).

Measurements of carbon will be in percentage changes from the baseline simulation. How big is a percentage change in carbon? For context, consider that the U.S. State Department's 2014 climate action plan set an ambitious goal of a 17 percent reduction in carbon emissions.[26] This target is in one sense larger because it includes all sources, not only personal emissions from urban consumers, but on the other hand, savings would come from improvements in energy efficiency and fuel sources (i.e. changing NERC factor) that I hold fixed in my analysis.

Simulation results are collected in Table 4 and 5. The tables focus on the end of simulation period (2011) in the interest of brevity.[27] All simulations begin with the existing housing stock in 1980 and simulate its evolution as if the policy were instituted at that time. Hence, some portion of each market's stocks were preexisting and unaffected by the policy change. For the aggregate carbon calculation, I also compute a "long run projection" to approximate the economy if the counterfactual policy had always been in place. This is a projection that scales the effect on construction by the share of end-of-simulation housing that was constructed in the sample period.[28] An alternative would be to run the simulation many periods into the future, or even start the simulation many decades before 1980, so that all housing is constructed under the counterfactual policy. However, simulations far outside the sample period (decades, or even centuries) would require strong assumptions about local housing demand which would be difficult to justify and yet have large effects on the results. Instead, I focus the model on the generation of housing in my available data and provide a simple rescaling of effects for a "very long" horizon.

6.2 Housing Demand Policy

The first experiment simulates removal of user costs subsidies to owner occupied housing. Federal income tax deductions for mortgage interest and property tax lower the user cost of housing with the objective of increasing the rate of homeownership. Whether it is successful in doing so (and whether homeownership is a desirable objective) is a matter of debate (Glaeser and Shapiro (2003)). In any case, a potential unintended consequence is that the demand subsidies increase the quantity of housing consumed in addition to affecting tenure choice at the margin because of lower user costs for all consumers. Hilber and Turner (2014) and Hanson (2012) have found that the MID increases the quantity of housing consumed, though the degree varies with market supply elasticity.

[26]See http://www.state.gov/e/oes/rls/rpts/car6/219259.htm [click here].

[27]Cumulative calculations for emissions 1980-2011 are qualitatively similar.

[28]For example, if half of 2011 housing in a city was built since 1980, and construction is one percent lower for that city in the simulation, the long run projection is an effect of two percent ($= \frac{1\%}{0.5}$).

The model of this paper is well suited to study the potentially heterogeneous effects of the user cost subsidies across markets and the resultant impacts on populations, housing intensities, density, and carbon emissions. Operationally, the simulation affects the consumer's budget constraint by increasing the user cost of housing by a factor τ,[29]

$$r_{sim} = (1 + \tau) r_{baseline}$$

From the user cost calibration described in section 4, $\tau = 0.17$. This implementation represents an average effect at the federal level, and a few caveats apply. First, the effect of subsidy removal would vary across U.S. states as some do not tax personal income and others do not allow deduction of mortgage interest or property tax (see Poterba (1992) and Hanson (2012) for more details).[30] This experiment focuses on the removal at the federal level only and studies heterogeneity in impact on housing markets. Second, the government side is ignored. Were the subsidy to be removed in actuality, the extra government revenue presumably would be used to fund some public good or may be rebated through lowering of the income tax rate or increasing a deduction of another sort. How the revenue is treated might have effects on housing demand or otherwise impact carbon emissions. Third, the actual user cost effect for a particular household would depend on its tenure status,[31] household income, whether it itemizes deductions, its property tax and mortgage interest rates, and its place on an amortization schedule. I combine into an average effect partly in the interest of simplicity, but also because many of these complications arise *after* the construction decision has been made.

All discussion of magnitudes is subject to these caveats. More generally, the simulation is interested in how the presence of a user cost subsidy affects the quantity and location of construction.[32] Since the simulation takes a stand on the magnitude of the subsidy, I will treat τ parametrically and examine how the results vary according to the size of the user cost change. A comparison of magnitudes between polices in Table 7 will close the section.

The results of the simulation are summarized in the first column of Table 4. The top panel summarizes aggregate effects on the housing and land stock. With less demand for housing services, construction falls 14.4 percent and comes at a slightly lower density (-1.1 percent). The second panel shows there is considerable heterogeneity in these effects, however, as shown in panel B. In some cities, construction drops by 30 percent or more, and in others, by less than 10 percent. Removing the subsidy has a bigger impact on places that offer a lot of housing or built

[29] The subsidy/tax factor falls out of the spatial equilibrium condition, but affects utility indirectly through rents.

[30] I have accounted for state-by-state tax differences in constructing the data, but the counterfactual is imposed uniformly across locations.

[31] Recall that data are on single family homes, which are the structure most likely to be owner occupied.

[32] There are also implicit user cost subsidies in that owner-occupied housing services and most capital gains are not taxed; see Albouy and Hanson (2014) for discussion.

Table 4: Summary of Simulations

	1 User Cost	2 Lot Size Regulation	3 Regulation “Subsidy”	4 Regulation “Tax”	5 Carbon Tax
A: Housing and Land, in Aggregate (% change)					
Construction	−14.40	−2.39	4.99	−5.84	−19.55
FAR	−1.60	5.88	1.14	1.60	−0.22
Housing pp.	−8.76	−1.06	1.30	−1.92	−11.57
Pop. Density	8.20	4.79	−0.53	2.97	12.23
B: Housing and Land, Statistics Across Cities (% change)					
Construction					
Std Dev	4.08	9.51	9.11	7.07	13.75
Min	−35.61	−8.05	−6.96	−16.68	−74.77
Max	−9.39	43.33	25.03	3.56	4.28
FAR					
Std Dev	5.04	6.67	1.74	1.58	7.45
Min	−16.68	−2.92	−4.70	−1.23	−22.25
Max	3.25	29.14	7.33	7.13	3.52
C: Correlation of Population Change With:					
Housing pp.	−0.34	0.22	−0.62	−0.54	−0.43
Pop. Dens.	0.34	−0.45	0.45	0.26	0.47
Carbon pp.	−0.39	0.19	−0.84	−0.76	−0.72
D: Energy And Carbon (% change)					
Electricity	−10.88	−2.67	−0.72	−4.11	−15.53
Electricity-carbon	−11.03	−2.64	−1.86	−5.05	−16.49
Natural Gas	−9.61	−2.26	0.60	−2.09	−12.22
Oil	−9.40	2.62	1.66	−1.85	−10.31
In-home Carbon	−10.62	−2.30	−1.09	−4.17	−15.18
Gasoline	−1.94	−1.51	0.37	−0.98	−2.90
Total Carbon	−6.22	−1.90	−0.35	−2.55	−8.96
LR Total Carbon	−13.13	−3.23	1.09	−3.28	−21.70
Total Carbon, Δ Pop. Only	−0.15	0.02	−0.70	−0.55	−0.82

NOTES: The table reports the aggregated results from five counterfactual simulations described in the main text. “Percent change” refers to the difference between the counterfactual and baseline simulations

Table 5: Summary of Simulations by City's Carbon Footprint

	1 User Cost		2 Lot Size Regulation		3 Regulation "Subsidy"		4 Regulation "Tax"		5 Carbon Tax	
City Type:	**High**	**Low**	**High**	**Low**	**High**	**Low**	**High**	**Low**	**High**	**Low**
A: Housing and Land, in Aggregate (% change)										
Construction	−13.34	−12.12	−2.41	−2.36	−2.54	15.56	−11.59	2.24	−23.28	−9.26
FAR	−4.67	−8.28	5.79	6.04	0.38	−0.90	0.81	0.06	−5.21	−7.93
Housing Stock	−8.87	−6.95	−1.61	−1.34	−1.79	9.19	−7.95	1.37	−15.39	−5.09
Housing pp.	−8.76	−8.75	−1.06	−1.05	1.68	1.63	−1.60	−1.61	−11.78	−10.67
Pop. Density	5.66	4.28	4.74	4.79	−1.49	−2.59	1.72	1.61	8.47	6.50
B: Housing and Land, Statistics Across Cities (% change)										
Construction										
Mean	−14.44	−12.79	0.24	−2.45	−2.87	14.77	−11.49	2.32	−26.40	−11.33
Std Dev	4.54	3.08	11.73	3.84	1.20	4.24	2.48	0.55	13.05	9.63
Min	−35.61	−21.54	−7.74	−8.05	−6.96	8.42	−16.68	1.76	−74.77	−33.55
Max	−9.52	−9.39	43.33	3.25	−1.85	25.03	−6.23	3.56	−5.74	4.28
FAR										
Mean	−5.89	−8.49	8.45	5.85	−0.04	−1.00	1.00	0.14	−8.25	−8.72
Std Dev	3.99	6.00	7.16	5.65	0.08	2.74	1.95	0.24	8.02	6.45
Min	−13.09	−16.68	−2.92	−1.54	−0.22	−4.70	−1.23	−0.23	−22.25	−20.96
Max	1.26	3.25	29.14	18.27	0.08	7.33	7.13	0.63	3.01	3.52
C: Energy And Carbon (% change)										
Electricity	−10.91	−8.57	−3.30	−2.70	−2.99	8.27	−9.38	0.21	−18.73	−7.55
Electricity-carbon	−10.94	−8.38	−3.15	−2.82	−2.97	8.12	−9.31	0.18	−19.04	−7.91
Natural Gas	−9.77	−7.95	−3.01	−2.37	−2.96	7.37	−8.39	0.16	−15.78	−6.29
Oil	−8.59	−7.56	4.34	0.32	−2.76	8.94	−8.85	0.78	−8.27	−6.85
In-home Carbon	−10.65	−8.16	−2.87	−2.41	−2.96	7.90	−9.14	0.22	−18.10	−7.22
Gasoline	−1.65	0.93	−2.25	−1.54	−2.86	8.41	−7.14	2.44	−6.08	4.59
Total Carbon	−6.49	−2.91	−2.59	−1.91	−2.92	8.19	−8.21	1.50	−12.54	−0.40
D: Energy And Carbon Per Capita (% change)										
Electricity	−10.80	−10.33	−2.76	−2.42	0.44	0.77	−3.13	−2.74	−15.27	−12.99
Electricity-carbon	−10.83	−10.14	−2.60	−2.54	0.46	0.63	−3.06	−2.77	−15.59	−13.33
Natural Gas	−9.66	−9.72	−2.46	−2.09	0.46	−0.07	−2.07	−2.79	−12.19	−11.80
Oil	−8.48	−9.34	4.93	0.61	0.67	1.39	−2.57	−2.19	−4.36	−12.33
In-home Carbon	−10.54	−9.92	−2.32	−2.13	0.46	0.43	−2.87	−2.73	−14.61	−12.68
Gasoline	−1.53	−1.01	−1.70	−1.26	0.57	0.90	−0.73	−0.58	−2.08	−1.56
Total Carbon	−6.37	−4.78	−2.04	−1.63	0.51	0.70	−1.88	−1.49	−8.81	−6.26

NOTES: The table reports the aggregated results from five counterfactual simulations described in the main text. The results are grouped into high carbon (above mean) and low carbon (below mean) cities as measured in Table 2. "Percent change" refers to the difference between the counterfactual and baseline simulations.

a lot of stock in recent years (or both). There is also dispersion in impacts on structure density, with most cities falling but a few cities rising. This is because the lower demand for housing means that land's continuation value declines, but places with higher assembly costs (parameter c_2), where continuation values matter the most, are more affected, so structural densities fall by more in such places.

Returning to the panel A, we see aggregate effects on housing and land intensities. With some preexisting stock surviving, housing per capita falls 8.8 percent nationally. This drop in housing consumption leads to an increase in average population density, even with a slight aggregate decrease in FAR.

As alluded to earlier, removal of the subsidy has heterogeneous impacts across cities, which results in some population reallocation. Panel C reports the correlation of population change (i.e. the difference relative to baseline simulation) with attributes of the cities. Higher housing consumption cities see drops in population, which is reallocated to lower housing consumption cities. On average, this also means lower land consumption (higher population density) cities grow at the expense of lower density places. Effectively, the deductions subsidize locations that offer housing as an outsize portion of their utility bundle. As Table 2 would suggest, this reallocation coincides with the cities' carbon footprints, although not perfectly so. The top five gainers (at about three percent gain each) are San Jose, Boston, New York City, San Francisco, and Philadelphia; the top five losers (at about 4 percent loss each) are Louisville, Richmond, Las Vegas, Riverside, and Austin.

Removal of the subsidy results in both reallocation of people from high carbon cities to low, and lower intensity of housing and land consumption. The bottom panel summarizes the net impacts on the national carbon emissions. Housing per capita has fallen by 8.8 percent, though carbon from in-home sources drops by slightly more (11 percent for electricity, nine percent for other fuels, and 10.6 percent on net) because there are more people consuming in lower energy demand places (with, for instance, milder climates). The increase in population density, coupled with the population reallocation results in slightly less gasoline consumption, though this margin is more difficult to move than were the in-home carbon sources. The net effect is about two percent decline in gasoline. All together, the economy emits 6.2 percent less carbon without the user cost subsidy. Using a long run projection (as if all stock had been built subject to the no-subsidy regime), the carbon effects would be on the order of a 13 percent decline.

The last row of the table disentangles the intensive and extensive margin effects. Here I calculate what the change to aggregate carbon emission would be if only population were reallocated across locations as in the simulation, but the emissions per capita within cities remained as in the baseline. If population across the economy were reweighted by the simulation, aggregate carbon emission would fall by only 0.15 percent. That is, the extensive margin accounts for a mere 2.5 percent of the total effect, and the intensive margin the remainder. Thus, even though population is reallocated away from higher carbon places to lower, the empirical effects on the aggregate are small. This point can also be seen by comparing aggregates of subsets of cities, as I do in Table 5. This table calculates the housing stock changes and carbon emissions separately for higher carbon cities (above the national mean) and lower (below the mean). The table shows that the impacts across city types are broadly similar: housing per capita shrinks and population density rises, and emission rates fall accordingly, in both types of cities. The population shifts cause a slightly larger share of emissions to come from lower carbon places, but the intensive margin effect is present in all locations.

In summary, subsidies to the user cost of housing have the effect of raising carbon emissions nationwide, mostly by incentivizing larger houses and less dense cities. It shifts population towards housing-friendly cities and away from cities offering utility via income/consumption or amenities, though the direct effect of population reallocation on emissions is relatively small.

Before proceeding, note that carbon emission reductions from removal of the subsidy may come at some utility loss. At the calibration of γ, the loss in housing consumption results in a two percent loss in the economy's utility ($\bar{U}$). However, this statement should be interpreted with caution, as it applies only to the utility achieved by the spatial equilibrium. The model has some leakages (most obviously, to government spending) and therefore is not a proper welfare calculation, strictly speaking.[33]

6.3 Housing Supply Policy

There are many forms of housing supply policy, most of which are determined locally. This subsection reviews the carbon impacts of local housing supply regulations.

6.3.1 Measuring Regulations in the Model

In order to do this, the model needs a way to operationalize the impact of changing various land use regulations. To map these into the model, I use the parameter estimates as outcomes in a regression on land use restrictions. The right-hand side variables in the regression are well known measures of local regulatory burden on new construction and the land available for building. For the former, I use the Wharton Residential Land Use Regulatory Index (hereafter, WRI). The WRI is a survey of the practices of local building authorities assembled in Gyourko et al. (2008).[34] It contains several sub-surveys related to types of land use regulation, such as permitting and project approval practices, zoning, lot size restrictions, open space requirements,

[33] A more serious welfare calculation might also account for energy use in the utility function, and to the extent possible, marginal damages from climate change. These are currently beyond the scope of this paper.

[34] There are other similarly designed measures of local land use restrictions, including Malpezzi (1996) and Saks (2008). Gyourko et al. (2008) has wide geographic coverage and provides the underlying microdata, allowing me to form tailored subindices.

and the like. All subindices rate, in one sense or another, the propensity of a community to impede building, but because I am particularly interested in structural densities, I distinguish the density restrictions from the broader regulatory burden index. I build separately a minimum lot size index and an other-regulations index absent the density restrictions using the factor loadings reported by Gyourko et al. (2008).[35] The modified index is rescaled to have a mean of 0 and standard deviation of 1 across communities, and then I take the metro-level average using the community sampling weights provided in Gyourko et al. (2008). The metro level index for the large cities in my sample has a mean near 0 and a standard deviation of 0.59.

To measure naturally occurring barriers to construction, I take a land availability measure published in Saiz (2010), which is based on GIS analysis of a 50 kilometer radius from city centers, quantifying the degree to which a metro area is unsuitable for construction because of coastal and inland water and terrain too steeply sloped for building. Saiz (2010) argued that geographic barriers are important directly and also can be correlated with regulations, affecting the measured impact of regulations on supply elasticities.

The results of the regressions are presented in Table 6, with one panel for each of the three structural parameters. I attempt several specifications to inform the counterfactual calibrations. Column 1 reports coefficients from separate OLS regressions. The regulations index and land share restrictions are positively associated with the log of c_1 assembly cost parameter, indicating that these reduce supply elasticities.[36] The coefficient on lot size restrictions is negative, which is interpreted as more land input being used when lot size minimum is higher, conditional on the other regulatory and geographic barriers. These results are robust to specification. Column 2 uses an inverse covariance weighted OLS regression, where more precisely estimated first stage parameters are given larger weight, and the remaining columns use seemingly unrelated regression to account for cross-equation correlation in errors resulting from the joint structural estimation. Considering these results together, for the counterfactuals I settle on the effect sizes reported in column 6, which are on the conservative side of the confidence intervals. An increase of one standard deviation of the regulatory index will impose a 20 log point increase to a c_1 parameter, and a one acre lot size minimum will impose a 40 log point decrease.

In the next panel, I do the same for the land stock-dependent assembly cost parameter, c_2. This parameter is a bit more complicated because some locations arrive at the zero constraint. These zeroes provide information on the effect of regulation on supply elasticities, but they are numerically a problem because the log of the point estimate approaches negative infinity. One

[35] I also omit the "Local Assembly" subindex (LAI), representing the degree of direct democracy (e.g. town hall meeting votes) in project approval. This is peculiar to New England and may not be well represented in the index. Gyourko et al. (2008) note that this subindex was not part of the survey questions, but information was volunteered by some communities.

[36] The regressions for c_1 also include an unreported control for city size, since this is mechanically related to the land assembly cost which was measured in levels. I use size rank because it is time invariant and closest to linear.

option is to censor the c_2 values at something below the minimum of the nonzero cities. In columns 1 to 4, I impute the value of $ln(c_2) = 0$ for those at the zero bound. The regression coefficients on the regulations index lie between 0.7 and 1.24. Another option is to throw out the cities at $c_2 \approx 0$, which I do in column 5. The coefficient shrinks slightly, but it is notable that regulations are still positively associated with the cost parameters away from the zero boundary. Considering all these coefficients, I will use the effect size of one standard deviation in regulation imposing a 70 log point increase to a c_2 parameter.

Finally, in the last panel, I regress the housing production function total factor productivity parameter, ϕ. The interpretation is simplest for this parameter: lot size rules place a direct restriction on the production function. The magnitudes and precision are somewhat affected by weighting, but the range of estimate suggests that a one acre increase should be associated with a ceteris paribus decline in structure density of 0.02. Note that this is conditional on the control for geographic land availability, itself significantly related to density. The regulations index does not matter for ϕ in any specification.

6.3.2 Density Restrictions

This simulation imagines a nationally instituted reduction in minimum lot size. Specifically, the simulation measures the effects of a one quarter acre decrease to the minimum lot size,[37] or an increase of $-0.02 \cdot -0.25 = 0.005$ to the ϕ and an increase to $ln(c_1)$ of $-0.4 \cdot -0.25 = 0.1$.

Table 4, column 2, reports the results. Panel A shows that the change to parameters causes a 2.4 percent decline in aggregate construction, though it occurs at a higher structure density by 5.9 percent. Builders are not required to use as much land input and construct slightly smaller homes on average. There is heterogeneity in the effects, however, as Panel B shows. The increase in density effect dominates in low ϕ locations, and these construct more, while the land effect dominates in high ϕ locations, and these construct less. In other words, low density places are able to put larger homes on smaller lots, while higher density places use less land and build smaller homes. Structural densities increase almost everywhere. In aggregate, the economy has one percent less housing per capita, and, when combined with a higher average structure density, population density increases by 4.8 percent.

The population effects follow the net changes to construction. Low ϕ locations, building more housing on net, see population increases, at the expense of high density cities. Thus, population changes are negatively correlated with population density. The correlation with housing consumption is positive because many low density places also offer high housing per capita. In total, the correlation to carbon emissions per capita are weakly positive. Table 5 shows that the effects within high and low carbon emitting cities is similar, though the distribution of

[37] The simulation size is just shy of one standard deviation. From the WRI data, in places where minimum lot sizes are known, these average 0.67 acres with a standard deviation of 0.31.

Table 6: Relationship of Costs Estimates to Local Land Use Regulations

	1	2	3	4	5	6
spec:	OLS	OLS (wt'ed)	SUR	SUR (wt'ed)	SUR (wt'ed)	In CF Sims.
				$ln(c_1)$		
Regulations Index	0.2431	0.2507	0.2398	0.2473	0.2720	0.2
	(0.1430)	(0.1227)	(0.1335)	(0.1162)	(0.1369)	
Land Restricted	0.7796	0.8399	0.7496	0.8460	0.6927	
	(0.3691)	(0.3038)	(0.3471)	(0.2878)	(0.3917)	
Lot size	-0.428	-0.605	-0.475	-0.596	-0.776	-0.4
	(0.2496)	(0.1946)	(0.2271)	(0.1842)	(0.2778)	
Cons	6.7528	6.7528	6.8256	6.8702	6.9806	
	(0.2332)	(0.2332)	(0.2139)	(0.1895)	(0.2610)	
				$ln(c_2)$		
Regulations Index	0.7083	1.2446	0.7083	1.2446	0.6279	0.7
	(0.3975)	(0.3286)	(0.3851)	(0.3183)	(0.2354)	
Land Restricted	0.4789	-0.064	0.4789	-0.064	-1.445	
	(1.1762)	(0.9465)	(1.1397)	(0.9170)	(0.6778)	
Cons	4.7468	4.9602	4.7468	4.9602	6.2500	
	(0.3671)	(0.3605)	(0.3557)	(0.3493)	(0.2915)	
				ϕ		
Lot size	-0.002	-0.020	-0.006	-0.017	-0.027	-0.02
	(0.0205)	(0.0120)	(0.0194)	(0.0111)	(0.0103)	
Regulations Index	-0.008	-0.002	-0.007	-0.003	-0.004	0
	(0.0106)	(0.0074)	(0.0101)	(0.0070)	(0.0067)	
Land Restricted	0.0873	0.1172	0.0851	0.1192	0.1288	
	(0.0312)	(0.0217)	(0.0298)	(0.0207)	(0.0213)	
Cons	0.0522	0.0464	0.0548	0.0439	0.0499	
	(0.0182)	(0.0135)	(0.0173)	(0.0127)	(0.0138)	
J	49	49	49	49	34	

NOTES: The table reports coefficient estimates for regressions of the denoted structural parameters on the variables listed. See Table 3 for the structural parameter estimates. The regulations index has mean zero and standard deviation of 0.6, the land restricted is the share (zero to one) of space geographically unavailable for building, and lot size is in minimum acres. The residual location is excluded. Weighted ("wt'ed") specifications use the first stage variance as observation weights. The c_1 regressions include a control for size rank. The regressions on c_2 with 49 observations impute a value of zero for the log of the parameter in cities arriving at the zero constraint. The final column reports the calibration used in sizing the counterfactual policy simulations.

construction changes is more negative for for low carbon cities.

When it comes to carbon accounting, then, there are some countervailing effects, but carbon emissions decrease in aggregate due to the intensive margin effects. With less housing per capita, carbon from in-home sources decrease. (The exception is fuel oil, a minor component of total emissions, which increases because of population growth in New England). With higher population densities, gasoline consumption drops. Together, the effect is a 1.9 percent drop in the economy's emissions. Projecting as if the entire housing and land stock were under the low minimum lot size policy, the effect would be 3.2 percent. Finally, the bottom row of the table shows that all of this effect comes through the intensive margin, since the population reallocation is not in favor of low carbon locations.

Just as user cost subsidies to demand increased carbon output mainly through the intensive margin, lot size restrictions increase carbon emissions national by placing constraints on the amount of housing and land consumed. The relative magnitudes of the policies will be compared in Table 7.

6.3.3 Housing Supply Regulation as Carbon Policy

Next I consider the extent to which local supply regulation can be used to affect national carbon emission. A way of exploring this is the following thought experiment: split the cities into high and low carbon emitters, and then use a "carrot and stick" approach of reducing regulation in low carbon cities and increasing regulation in high carbon cities. Specifically, I apply a change to the high (low) cities' cost parameters associated with a one standard deviation increase (decrease) in the regulation index.[38] The residual location is not altered directly in these simulations. I conduct the "carrot" and "stick" in separate simulations to emphasize their effects. The results continue in Tables 4 and 5, the latter being especially important in these simulations since the policy is designed to apply by carbon footprint.

I begin with the carrot approach of relaxing building restrictions in low carbon places, which operationally lowers these cities' assembly cost parameters c_1, c_2, but has no direct impact on high carbon cities. The aggregate effect is an increase in construction volume and structure density. Table 5 shows the change is composed of a 15.6 percent increase in low carbon cities and a 2.5 percent decrease in high carbon, the latter occurring because of reduced demand for such places in general equilibrium. There are small inverse effects on structure density. In low carbon places, now less regulated, land becomes less dear, and FAR drops by almost one percent. Table 5 panel A compares total stocks and per capita rates. It shows the *total* housing stock increases in low carbon cities and declines in high, but housing consumption *per capita* increases similarly in both because of population flows. This highlights an important point about the aggregate effects of the policy in general equilibrium. By relaxing building restrictions, the supply of

[38] All cost parameters are constrained to be nonnegative.

housing increases and rents fall nationally, raising the reservation utility level through higher average housing services per capita. Thus, an increase in per capita consumption anywhere can affect per capita consumption everywhere in general equilibrium. The result is more housing and land consumption in the aggregate, as shown in Table 4, panel A.[39]

Table 4, panel C, shows the population changes are, by design, negatively correlated with local average housing consumption, land consumption, and carbon emission. (The biggest gainers are San Jose, Sacramento, and Philadelphia, at about 10 percent each. The biggest losers are Birmingham, Louisville, and Nashville at about four percent each.) However, Panel D reports that the effects on total carbon emission are ambiguous, since effects through population reallocation can be overwhelmed by changes in intensive margins. Carbon from electricity drops slightly, but heating fuels increase. Gasoline consumption increases slightly. Table 5 shows the source of these effects. Panel C shows total carbon from low emitting places increases, while total from high emitting places decreases–this is the composition effect due to population reallocation across cities. But *within* each type of city, shown in panel D, emissions increase. The net effect is a mere 0.35 percent drop in aggregate, and the long run projection actually swings positive. The last row of Table 4 shows that population change alone would have caused carbon emission to decrease, but by less than one percent. Population is a "sticky needle," and despite the large differences in emission rates, it is difficult to obtain carbon savings without very large population changes.

The "stick" approach of increasing regulations in high carbon cities is effectively a mirror image of the last simulation. Construction and population shift across cities accordingly, with the inverse general equilibrium effects: by reducing supplies and raising rents economy-wide, housing and land consumption falls. The net result is a decline of emissions both within and between cities, leading to a 2.6 percent drop nationally, up to 3.3 percent in a long horizon. About 0.6 percentage points of the 2.6 come from the reallocation of people across space, and the remainder from the drop in the intensive margins.

The conclusions from these experiments are twofold. First, they highlight some of the incidental effects of local supply restrictions, which, by changing land and housing values, can affect intensive margins. Regulations can play a role in shaping the patterns of consumption that lead to low carbon emission, and are not merely an impediment to moving people to (seemingly) more efficient areas. Second, any housing supply policy with an eye on environmental impacts should consider rebound effects, which in the simulations dominated the impact of reallocation.

[39]Note the national weighted averages, which include the residual outside option location, are slightly different from within city type averages, which exclude the residual city.

6.4 Carbon Tax

The final simulation institutes a carbon tax on the economy. Several nations around the world have implemented various forms of carbon taxes.[40] The following discussion focuses on a tax equivalent to Sweden's at $0.076 per pound of emission, but I close with a parametric treatment of the tax size.

Adding a carbon tax to the model changes the budget constraint of the worker to $y_j = (1+\tau\phi^C)c_j + (r_j + \tau\phi_j^H)h_j + \tau\phi_j^G$, where τ is the tax size and ϕ_j^H, ϕ_j^G refer to, respectively, the city's pass-through rate of carbon per unit of housing and amount of gasoline used per person, and ϕ^C is pass through on the numeraire. The amount of tax owed would depend therefore on the amount of housing consumed as well as the energy used per unit of housing and the carbon content of that energy (which varies by source and location via NERC factor). Because the expenditure on carbon tax depends on endogenous variables, the carbon tax regime requires a new spatial equilibrium condition inclusive the $(\tau, \phi_j^H, \phi_j^H)$ parameters. The new spatial equilibrium condition loses the convenient closed form of (14a), (14b), but can be solved numerically.[41]

The results appear in the final columns of Tables 4 and 5. The tax reduces the budget available for housing, but also is biased against housing goods in places where the carbon pass through rates are high. Thus, demand for housing drops and construction falls (Table 4, panel A), especially in high carbon cities (Table 5, panels A and B). The low demand for housing and the concomitant drop in the value of land actually leads to lower structure densities, though the loss in housing consumption is large enough to increase population densities on net.

This leads to two obvious effects. First, population shifts away from high carbon (and housing and land) locations in favor of lower. The losers are predominantly central and southern cities; Louisville, Oklahoma City, Cincinnati, Orlando, and Nashville are at least 12 percent smaller. The gainers are predominant coastal; The San Francisco Bay Area and Los Angeles, as well as New York, Boston, Philadelphia, are eight to 15 percent larger. Second, it reduces the intensive margins on housing and land. Note that the intensive and extensive margins are linked here;

[40] For a brief overview, see [click here]. Denmark, Finland, and Sweden were early adopters of basic carbon taxes in the 1990s. More recently, Switzerland, Ireland, Japan, France, and South Africa (as of 2016) have implemented carbon taxes, and several other nations tax carbon-emitting fuels at various stages of the supply chain.

[41] The spatial equilibrium condition is:

$$ln\frac{p_j}{P} = \frac{1}{\gamma} ln \frac{y_j - (r_j + \tau\phi_j^H)h_j - \tau\phi_j^G}{y_n - (r_n + \tau\phi_n^H)h_n - \tau\phi_n^G} + ln\frac{H_j}{H_N}$$

with

$$r_j = \gamma\frac{(y_j - \tau\phi_j^G)}{h_j} - \frac{\tau(2\phi_j^H + \phi_j^H\phi_j^C)}{1 + \tau\phi_j^C + \gamma}$$

where ϕ_j^C is the carbon pass-through rate of consumption goods, assumed identical across locations, and all other variables are defined as above.

the average consumption falls both because of the tax on energy *and* because of the shift of population toward lower consumption cities.

These naturally lead to substantial decreases in carbon emissions, especially the more elastic in-home sources. The total effect is about a nine percent drop in carbon, projected to be as large as 22 percent in the long run. Note that less than one percentage point of this comes from the reallocation of population alone. Population is moved from high carbon cities to low, but more important quantitatively is what is occurring within each type of city.

A comment on welfare can be made similar to the user cost simulation. Consumers are unambiguously worse off under the tax, in large part because their consumption of housing goods falls. However, this statement is subject to the caveat that the model ignores the government sector (including the use of tax revenue), and is not intended to give a complete treatment to welfare analyses.

6.5 Summary of Simulations

The simulations above demonstrated the mechanics of the policy changes. The directions of the effects are clear, but I now conclude this section by comparing the relative sizes of each class of policy.

For each type of policy, I conduct a parametric treatment in the neighborhood of the presented simulation. For example, in the user cost experiment, I conducted simulations assuming the removal of the subsidy changed the user cost from one to 33 percent in units of two percentage points. I then calculated the carbon changes for each simulation and measured the slope of the carbon effects with respect to the policy size. This allows a back-of-the-envelope calculation in converting units of one policy to another.

Table 7 compares magnitudes between policies for the sample period and long run projection. The columns under "Carbon Source" report the size of a policy change needed to produce an aggregate one percent decline in carbon emissions from that source. Focusing on the Total column in the data period horizon, we see that a carbon tax of 0.89 cents ($0.0089) would lower aggregate emissions by one percent.[42] To achieve a one percent reduction, the economy would need to remove user cost subsidy of 3.3 percent, decrease lot size minimum by 0.17 acres, or increase regulations in high carbon locations by 0.32 standard deviations. The last column reports the ratios of these relative to a carbon tax. A reduction in user cost subsidy of 3.67 percent would reduce carbon emissions equivalent to a one cent per pound tax, or a 0.2 acre decrease in minus lot size, or a 0.36 standard deviation increase in regulations to high carbon areas. In the long run (second panel), the carbon tax has relatively larger impact than the other policies because it affects all margins of adjustment.

[42] It is worth reminding the reader that my calculations of carbon savings focus on housing and land stock related sources, including personal transportation, not consumption of other goods.

Table 7: Comparison of Policy Magnitudes: Size of Policy Change Needed to Obtain One Percent Reduction in Carbon Emission

Policy Class	Units	Carbon Source						Carbon Tax
		Elec	Nat Gas	Oil	In-home	Gas	Total	Ratio
		In Sample Period						
Carbon Tax	cents/lb	0.49	0.68	0.77	0.53	2.52	0.89	1.00
User Cost of Housing	percent	1.90	2.23	2.15	1.98	8.83	3.26	3.67
Min. Lotsize	acre	0.15	0.21	-0.10	0.18	0.17	0.17	0.20
Reg. Index: High Carbon	SD of WRI	0.16	0.78	0.52	0.21	0.66	0.32	0.36
		Long Run Projection						
Carbon Tax	cents/lb	0.21	0.21	0.21	0.21	1.60	0.37	1.00
User Cost of Housing	percent	0.89	0.85	0.79	0.87	5.52	1.52	4.09
Min. Lotsize	acre	0.09	0.08	-0.09	0.10	0.10	0.10	0.27
Reg. Index: High Carbon	SD of WRI	0.10	0.24	0.17	0.12	0.47	0.19	0.52

NOTES: The table reports the size of the change to a policy required to achieve a one percent aggregate reduction in carbon from the denoted source. The last column uses the ration relative to carbon tax for the total carbon emission.

7 Conclusion

This paper studied the impact of housing and land stock allocation on carbon emissions using a dynamic equilibrium model of housing stock evolution across a system of cities. The main findings are that the intensity of the stock, i.e. housing and land consumption per person, is a major determinant of carbon footprint. Hence, policies that affect housing's intensive margin have implications for carbon emissions, and likewise, carbon policies will affect the nature and allocation of the housing stock. The results indicate that housing-friendly policies like user cost subsidies, density constraints, and lax building regulations may inadvertently contribute to higher carbon emissions. Of course technologies and energy prices will matter for household energy consumption, but this paper highlights the important role for housing policies as well.

On the modeling side, the contribution of the paper is a unified treatment of extensive intensive margins in equilibrium. Part of what makes a high carbon location attractive for households is the large amount of housing consumption it offers; altering its intensive margin affects how many people will want to live there. Moreover, while comparisons between cities are useful, an integrated approach is necessary to measure aggregate effects of policies between connected locations. Locations can be quite heterogeneous, but population flows out of one place will end up somewhere else. By connecting local heterogeneity with regional integration, the model highlights a nuance to the distinction of "intensive" versus "extensive" margins. In equilibrium, intensive margins anywhere can affect extensive margins everywhere by altering the utility available to the economy. The two margins are inextricably linked. The extensive margin is, in a sense, under-measured in that part of population shifts between places manifest as changes to within-city consumption rates with prices moving in general equilibrium. Finally, the model is able to make realistic projections of policy changes by allowing for trends and stock accumulations, something disallowed when doing "putty-putty" models that rebuild cities instantaneously.

I close with some suggested directions for future study. Within the current modeling environment, the most obvious extension would be to include multifamily housing in the empirics. Data on multifamily structure on par with the detailed single family properties were not available, but some close-enough calibration could be possible. Some of the policies examined here would likely be affected by its inclusion, especially density restrictions, and even user cost subsidies that might shift tenure status to single family away from multifamily. The results are likely understated without this channel, though it is interesting how population density can matter even without dramatic "building up" of structures.

An extension of the model could place housing durably *within* a city–centrally, peripherally, etc.–and measure the consequences. Housing in this paper is durable in terms of its across-city location, but fungible within city. Other possible extensions include treatment of land use in commercial, industrial, and infrastructure purposes. Finally, a city's transportation structure could be modeled explicitly. Changes to population density might affect, in lumpy investment style, the construction and use of public transportation, with potentially material effects on carbon per capita.

References

Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr (2016): "Transition to Clean Technology," *Journal of Political Economy*, 124, 52–104.

Ahfeldt, G. M. and D. P. McMillen (2014): "New Estimates of the Elasticity of Substitution Between Land and Capital," Mimeo.

Albouy, D. (2009): "The Unequal Burden of Federal Taxation," *Journal of Political Economy*, 117, 635–667.

Albouy, D. and G. Ehrlich (2016): "Housing Productivity and the Social Cost of Land-Use Restrictions," University of Michigan.

Albouy, D. and A. Hanson (2014): "Are houses too big or in the wrong place? Tax benefits to housing and inefficiencies in location and consumption," *Tax Policy and the Economy*, 28, 63–96.

Aldy, J. E. and R. N. Stavins (2012): "The Promise and Problems of Pricing Carbon: Theory and Experience," *The Journal of Environment and Development*, 21, 152–180.

Battaglini, M. and B. Harstad (2016): "Participation and Duration of Environmental Agreements," *Journal of Political Economy*, 124, 000 – 000.

Baylis, K., D. Fullerton, and D. H. Karney (2014): "Negative Leakage," *Journal of the Association of Environmental and Resource Economists*, 1, 51–73.

Bento, A. M., M. L. Cropper, A. M. Mobarak, and K. Vinha (2005): "The effects of urban spatial structure on travel demand in the United States," *Review of Economics and Statistics*, 87, 466–478.

Black, D. A., N. Kolesnikova, and L. J. Taylor (2014): "Local Labor Markets and the Evolution of Inequality," *Annu. Rev. Econ.*, 6, 605–628.

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

Borck, R. and M. Pflüger (2015): "Green cities? Urbanization, trade and the environment," Tech. rep., IZA Discussion Paper No. 9104.

Brownstone, D. and T. F. Golob (2009): "The Impact of Residential Density on Vehicle Usage and Energy Consumption," *Journal of Urban Economics*, 65, 91–98.

Broxterman, D. A. and A. M. Yezer (2015): "Why does skill intensity vary across cities? The role of housing cost," *Regional Science and Urban Economics*, 55, 14–27.

Combes, P.-P., G. Duranton, and L. Gobillon (2012): "The production function of housing: Evidence from France," *Processed, University of Toronto.*

Davis, M. A. and F. Ortalo-Magne (2011): "Household Expenditures, Wages, and Rents," *Review of Economic Dynamics*, 14, 248–261.

Dietz, S., C. Gollier, and L. Kessler (2015): "The Climate Beta," Toulouse School of Economics.

Eichner, T. and R. Pethig (2011): "Carbon leakage, the green paradox, and perfect future markets," *International Economic Review*, 52, 767–805.

Elliott, J., I. Foster, S. Kortum, T. Munson, F. P. Cervantes, and D. Weisbach (2010): "Trade and carbon taxes," *The American Economic Review*, 100, 465–469.

Epple, D., B. Gordon, and H. Sieg (2010): "A New Approach to Estimating the Production Function for Housing," *American Economic Review*, 100, 905–924.

Fischel, W. (2001): "Municipal Corporations, Homeowners, and the Benefit View of the Property Tax," in *Property Taxation and Local Government Finance: Essays in Honor of C. Lowell Harriss.*, ed. by W. Oates, Cambridge, MA: Lincoln Institute for Land Policy.

Gaigné, C., S. Riou, and J.-F. Thisse (2012): "Are compact cities environmentally friendly?" *Journal of Urban Economics*, 72, 123–136.

Gans, J. S. (2012): "Innovation and Climate Change Policy," *American Economic Journal: Economic Policy*, 4, 125–145.

Glaeser, E. L., J. Gyourko, E. Morales, and C. G. Nathanson (2014): "Housing dynamics: An urban approach," *Journal of Urban Economics*, 81, 45–56.

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

Glaeser, E. L. and J. M. Shapiro (2003): "Benefits of the Home Mortgage Interest Deduction," *Tax Policy and the Economy*, 17, 37.

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014): "Optimal taxes on fossil fuel in general equilibrium," *Econometrica*, 82, 41–88.

Greenstone, M., E. Kopits, and A. Wolverton (2013): "Developing a social cost of carbon for US regulatory analysis: A methodology and interpretation," *Review of Environmental Economics and Policy*, 7, 23–46.

Gyourko, J. and A. Saiz (2006): "Construction Costs and the Supply of Housing Structure," *Journal of Regional Science*, 46, 661–680.

Gyourko, J., A. Saiz, and A. Summers (2008): "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index," *Urban Studies*, 45, 693–729.

Hanson, A. (2012): "Size of home, homeownership, and the mortgage interest deduction," *Journal of Housing Economics*, 21, 195–210.

Hilber, C. A. L. and T. M. Turner (2014): "The Mortgage Interest Deduction and Its Impact on Homeownership Decisions," *Review of Economic and Statistics*.

Karathodorou, N., D. J. Graham, and R. B. Noland (2010): "Estimating the effect of urban density on fuel demand," *Energy Economics*, 32, 86–92.

Krusell, P. and A. A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896.

Larson, W., F. Liu, and A. Yezer (2012): "Energy footprint of the city: Effects of urban land use and transportation policies," *Journal of Urban Economics*, 72, 147–159.

Larson, W. and A. Yezer (2015): "The energy implications of city size and density," *Journal of Urban Economics*, 90, 35–49.

Malpezzi, S. (1996): "Housing prices, externalities, and regulation in US metropolitan areas," *Journal of Housing Research*, 7, 209.

Mangum, K. (2017): "An Approach for Empirical Work in Spatial Dynamics," .

Metcalf, G. E. (2009): "Tax Policies for Low-Carbon Technologies," Working Paper 15054, National Bureau of Economic Research.

Nordhaus, W. (2014): "Estimates of the social cost of carbon: concepts and results from the DICE-2013R model and alternative approaches," *Journal of the Association of Environmental and Resource Economists*, 1, 273–312.

Nordhaus, W. D. and Z. Yang (1996): "A regional dynamic general-equilibrium model of alternative climate-change strategies," *The American Economic Review*, 741–765.

Poterba, J. (1992): "Taxation and Housing: Old Questions, New Answers," *AEA Papers and Proceedings*, 82, 237–242.

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

Ruggles, S., K. Genadek, R. Goeken, J. Grover, and M. Sobek. (2015): "Integrated Public Use Microdata Series: Version 6.0 [Machine-readable database]. Minneapolis: University of Minnesota," .

Saiz, A. (2010): "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 125, 1253–1296.

Saks, R. E. (2008): "Job Creation and Housing Construction: Constraints on Metropolitan Area Employment Growth," *Journal of Urban Economics*, 64, 178–195.

Stern, N. (2008): "The Economics of Climate Change," *American Economic Review*, 98, 1–37.

Stokey, N. L., R. E. Lucas Jr, and E. C. Prescott (1989): *Recursive methods in economic dynamics*, Harvard University Press.

Su, Q. (2011): "The effect of population density, road network density, and congestion on household gasoline consumption in US urban areas," *Energy Economics*, 33, 445–452.

Weitzman, M. L. (2014): "Fat tails and the social cost of carbon," *The American Economic Review*, 104, 544–546.

Wheaton, W. C. and W. E. Simonton (2007): "The Secular and Cyclic Behavior of 'True' Construction Costs," *Journal of Real Estate Research*, 29, 1–25.

Wooldridge, L. W. (2002): *Econometric Analysis of Cross Section and Panel Data*, Cambridge, MA: MIT Press.

A Appendix

Table A1 reports the a summary of incomes, average rent per unit of living area, relative values of income, housing consumption, and rent, and local amenity estimates. For discussion of data and estimation, see, respectively, section 2 and 4.

Table A1: Local Incomes, Rents, Housing Consumption, and the Estimation of Amenity Values

Size Rank	City	Income Per Capita 2011, ($k)	Avg Rent ($/sf)	mean $ln(x_j) - ln(x_{natl})$ Income	Housing Consumption	Rents	Pop Share	μ_j Coef.	se
0	Residual	28.82	7.42	-0.145	-0.028	0.040	0.350	-0.008	0.002
1	New York, NY	43.33	7.93	0.206	0.084	0.018	0.054	0.022	0.002
2	Los Angeles, CA	33.91	7.13	0.050	-0.076	-0.086	0.050	-0.017	0.002
3	Chicago, IL	35.11	7.75	0.081	-0.046	0.027	0.037	-0.009	0.002
4	Philadelphia, PA	37.20	6.68	0.065	-0.304	-0.126	0.031	-0.069	0.002
5	Dallas, TX	33.41	5.97	0.031	0.155	-0.221	0.023	0.036	0.002
6	Miami, FL	32.89	7.03	0.034	-0.039	-0.089	0.019	-0.008	0.002
7	Washington, DC	45.31	8.78	0.254	-0.130	0.132	0.022	-0.027	0.002
8	Houston, TX	36.40	4.96	0.033	0.133	-0.411	0.022	0.031	0.002
9	Detroit, MI	30.55	5.55	0.011	-0.010	-0.310	0.024	-0.002	0.002
10	Boston, MA	44.20	7.85	0.195	0.170	0.005	0.017	0.042	0.002
11	Atlanta, GA	30.32	6.95	-0.009	0.167	-0.069	0.018	0.038	0.002
12	San Francisco, CA	46.88	8.19	0.315	-0.046	0.044	0.018	-0.006	0.002
13	Riverside, CA	22.88	7.29	-0.221	-0.087	-0.050	0.016	-0.023	0.002
14	Phoenix, AZ	28.17	6.20	-0.096	0.048	-0.202	0.014	0.010	0.002
15	Seattle, WA	38.91	7.44	0.119	0.059	-0.039	0.014	0.015	0.002
16	Minneapolis, MN	37.15	7.66	0.112	-0.050	0.012	0.014	-0.010	0.002
17	San Diego, CA	35.74	10.84	0.039	-0.047	0.330	0.012	-0.010	0.002
18	St Louis, MO	32.73	7.63	-0.009	0.020	0.023	0.014	0.005	0.002
19	Baltimore, MD	39.04	8.67	0.067	-0.381	0.128	0.013	-0.087	0.002
20	Pittsburgh, PA	34.34	7.69	-0.042	0.076	0.033	0.014	0.017	0.002
21	Tampa, FL	29.98	6.17	-0.098	0.058	-0.207	0.011	0.012	0.002
22	Denver, CO	37.40	7.89	0.113	0.012	0.042	0.010	0.004	0.002
23	Cleveland, OH	32.35	8.01	0.008	0.082	0.073	0.011	0.019	0.002
24	Cincinnati, OH	31.24	7.43	-0.050	0.082	-0.002	0.010	0.018	0.002
25	Portland, OR	31.53	6.35	-0.020	0.072	-0.198	0.009	0.016	0.002
26	Kansas City, MO	32.88	7.60	-0.011	0.170	0.021	0.010	0.039	0.002
27	Sacramento, CA	31.11	9.20	-0.031	0.022	0.181	0.009	0.005	0.002
28	San Jose, CA	46.64	13.35	0.287	-0.105	0.525	0.008	-0.020	0.002
29	San Antonio, TX	28.08	6.06	-0.205	-0.004	-0.210	0.008	-0.004	0.002
30	Orlando, FL	27.13	6.84	-0.139	0.054	-0.101	0.007	0.011	0.002
31	Columbus, OH	30.68	7.53	-0.058	0.023	0.011	0.008	0.005	0.002
32	Providence, RI	32.98	8.31	-0.068	0.231	0.071	0.005	0.052	0.002
33	Norfolk, VA	32.05	7.92	-0.109	0.059	0.043	0.008	0.012	0.002
34	Indianapolis, IN	30.98	6.11	-0.028	0.103	-0.194	0.008	0.023	0.002
35	Milwaukee, WI	34.06	7.95	0.025	-0.045	0.056	0.007	-0.010	0.002
36	Las Vegas, NV	27.24	10.29	-0.055	-0.003	0.306	0.005	-0.001	0.002
37	Charlotte, NC	30.71	6.19	-0.038	0.128	-0.186	0.007	0.029	0.002
38	Nashville, TN	32.17	6.82	-0.061	0.117	-0.091	0.006	0.026	0.002
39	Austin, TX	30.89	6.14	-0.083	0.129	-0.199	0.005	0.029	0.002
40	Memphis, TN	29.49	7.18	-0.113	0.111	-0.033	0.006	0.024	0.002
41	Buffalo, NY	30.63	7.16	-0.120	0.052	-0.038	0.005	0.011	0.002
42	Louisville, KY	29.81	7.34	-0.104	0.066	-0.018	0.006	0.014	0.002
43	Hartford, CT	40.53	10.08	0.156	0.246	0.292	0.005	0.059	0.002
44	Jacksonville, FL	31.08	6.45	-0.077	0.000	-0.160	0.005	-0.001	0.002
45	Richmond, VA	32.87	7.96	-0.005	0.122	0.059	0.006	0.028	0.002
46	Oklahoma City, OK	30.54	6.74	-0.127	0.096	-0.106	0.006	0.020	0.002
47	Birmingham, AL	31.16	7.56	-0.115	0.094	0.014	0.005	0.020	0.002
48	Rochester, NY	31.83	7.91	-0.034	-0.090	0.063	0.005	-0.021	0.002
49	Salt Lake City, UT	30.23	9.41	-0.136	-0.217	0.213	0.005	-0.052	0.002
	Amenity elasticity σ							0.987	0.005