

The Microgeography of Housing Supply

Nathaniel Baum-Snow*

Lu Han[†]

July 2, 2020

Abstract

Knowledge of microgeographic housing supply elasticities is central for quantitative analysis addressing a wide range of urban questions. Across 50,409 urban census tracts in 306 metros, we find average price elasticities of floorspace, developed land, housing unit and redeveloped unit supply of 3.0, 0.2, 0.5 and 0.15 respectively. With a mean of 87%, floorspace per unit thus accounts for by far the largest fraction of floorspace supply responses to price shocks. Land development, unit and floorspace supply responses all grow with CBD distance, mostly due to the increasing availability of undeveloped land. Tracts with more flat land and less regulation also exhibit more elastic supply. Identification comes from variation in labor demand shocks to commuting destinations, as aggregated using insights from an urban economic geography model. Aggregation of neighborhood level supply elasticities yields metro area supply elasticities that are correlated with but smaller than those in Saiz (2010).

*Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, ON M5S 3E6, Canada, nate.baum.snow@rotman.utoronto.ca.

[†]Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, ON M5S 3E6, Canada, lu.han@rotman.utoronto.ca.

1 Introduction

Housing supply conditions vary considerably both between and within urban areas. While the existing literature documents large differences in housing supply elasticities between cities (Saiz, 2010; Cosman, Davidoff, & Williams, 2018), little empirical evidence exists on how supply elasticities differ *within cities* as a function of distance to the center, land availability, building densities and zoning restrictions. Knowledge of housing supply elasticities at a microgeographic scale is central to understanding spatial variation in booms and busts within housing markets (Glaeser, Gottlieb, & Tobio, 2012; Guerrieri, Hartley, & Hurst, 2013), growth patterns at urban fringes (Glaeser, Gyourko, & Saks, 2005), consequences of neighborhood specific labor or housing demand shocks (Couture, Gaubert, Handbury, & Hurst, 2019) and implications of place-based policy interventions such as targeted neighborhood investment, land use restrictions, and transportation infrastructure investments (Busso, Gregory, & Kline, 2013; Hanson, 2009). This study empirically characterizes housing supply elasticities for all residential neighborhoods in 306 U.S. metro areas. We decompose supply responses of housing services into those for land development, housing units per parcel and floorspace per unit. We also investigate redevelopment in existing developed areas. We demonstrate how housing supply conditions vary by neighborhood location, available land, topography and regulation. We then aggregate neighborhood level housing supply elasticities to the metro area level, highlighting the fact that the aggregation scheme is specific to the nature of the demand shock and to the degree of housing demand substitutability across neighborhoods. Finally, we apply our estimates of neighborhood-level supply and demand function parameters to evaluating the efficacy of “Opportunity Zone” (OZ) provisions in the 2017 “Tax Cuts and Jobs Act”.

We estimate supply elasticities both by units and by floorspace. Across tracts, average unit supply elasticity is about 0.5, with about one-third of the unit supply response coming through new units built on already developed land. Average housing services supply elasticity, which is very close to that for floorspace, is 3.0. With an average estimated land development price elasticity of only 0.2, at least 65% of the floorspace response comes in the form of floorspace per unit rather than total units or new parcel development in over 95% of the neighborhoods in our data.

Housing supply elasticities exhibit substantial within-city variation. Land development, unit

and floorspace supply responses all grow with distance from central business districts (CBDs). At CBDs, they are on average 0, 0.1 and 2.3 respectively, growing to 0.3, 1.0 and 4.8 respectively at urban fringes. The positive CBD distance profile is mainly driven by the fact that the fraction of developed land decreases as one moves away from the CBD. Tracts with more flat land and less stringent regulations also exhibit more elastic supply. Resulting estimates range from 0.1 to 0.9 for the 25th and 75th percentile neighborhoods for unit supply and 2.1 to 3.7 respectively for housing services supply. Conditional on neighborhood topography and land development patterns, we do not find evidence that metro level factors matter for local supply elasticities.

Our supply elasticity estimates are consistent with the literature in several ways. First, housing production function estimates in the literature typically indicate approximately Cobb-Douglas form with a land share of 0.2-0.35 (Ahlfeldt & McMillen, 2014; Combes, Duranton, & Gobillon, 2019; Albouy, Ehrlich, & Shin, 2018), implying housing services supply elasticities of 2-4, which match well with the range of our floorspace estimates. Turning to unit supply, our tract-level elasticities are roughly similar to the municipality-level elasticities in recent work using data from Switzerland and select U.S. cities (von Ehrlich, Schöni, & Büchler, 2018; Orlando & Redfearn, 2018), but mostly smaller than the metro-level supply elasticities in Saiz (2010). One reason is that aggregation from micro to macro elasticities incorporates substitution patterns across neighborhoods in residents' neighborhood system. As neighborhoods become stronger demand substitutes, a shock affecting labor market opportunities in one location affects housing demand in a wider range of areas, as households are more willing to substitute across residential options to take advantage of lower housing prices in some places. This opens up more opportunities for supply elastic neighborhoods to be included and hence results in greater macro supply elasticities than the neighborhood level.

We approach recovery of neighborhood level housing supply elasticities as the fundamentally reduced form problem of identifying coefficients in regressions of changes in tract level housing quantities on changes in a tract level home price index. Our reduced form estimation is micro-founded on a stylized model of neighborhood housing supply. The model provides a natural basis for decomposing the residential floorspace supply elasticity into extensive and intensive margins and for understanding the magnitudes and determinants of our estimated supply elasticities. In doing so, the model highlights the mechanism through which land share, initial development

density, topography and regulation affect housing supply elasticities.

The central challenge in identifying housing supply is to find an exogenous source of variation that shifts neighborhood level housing demand but not local fundamentals including construction costs. This identification challenge is particularly daunting for recovering within-city supply elasticities, as most shocks that impact housing demand in one neighborhood would also affect housing demand for nearby neighborhoods, making it difficult to trace out housing supply in any specific neighborhood. To achieve identification, similar to Severen (2019), we use Bartik (1991) type labor demand shocks to commuting destinations from each residential location as the fundamental source of variation in housing demand shocks, which feed through the commute time matrix to generate exogenous variation in house price growth across residential locations. These labor demand shocks are built using 1990 industry shares in commuting destinations interacted with national industry-specific employment growth rates after year 2000.

One practical challenge is in how to sensibly aggregate these labor demand shocks across all commuting destinations from each residential tract. To do this, we follow Tsivanidis (2018) and nest our reduced-form estimation problem into an urban spatial equilibrium model in which residential demand in neighborhood i depends on “resident market access” (RMA_i), a coherent measure of access to employment from tract i . RMA_i amounts to the commute time discounted sum of employment in each commuting destination from location i . Labor demand shocks in each potential commuting destination are used to generate a simulated counterpart to the change in RMA_i that, conditional on appropriate controls, is purged of shocks to tract housing productivity or changes in other unobserved tract level housing supply factors.

Beyond contributing to the housing production function and supply elasticity literatures, our micro scale estimates provide a supply-side explanation for the recent finding that there was more price growth in the center of metro areas in the 2002-2007 housing boom (Glaeser et al., 2012; Genesove & Han, 2013). In addition, knowledge of neighborhood level housing supply elasticities is central for not only understanding patterns of neighborhood changes (Cutler, Glaeser, & Vigdor, 1999) but also evaluating welfare consequences of affordable housing policies (Favilukis, Mabile, & Van Nieuwerburgh, 2019; Davis, Gregory, & Hartley, 2019). Furthermore, a burgeoning literature examines policies and phenomena that only directly impact a few neighborhoods in cities in the context of general equilibrium urban models (Calabrese, Epple, & Romano, 2011). These stud-

ies typically calibrate micro elasticities to macro estimates from the literature. Our evidence shows that this choice can lead to misleading conclusions about incidence. We thus hope that our supply elasticity estimates are useful for improving quantitative evaluation of various policies that are targeted to particular neighborhoods.

As an example application, we explore the welfare consequences of the Opportunity Zone (OZ) provisions of the 2017 Tax Cuts and Jobs Act. The OZ program targets poorer neighborhoods with reduced capital gains taxes on new real estate investments. The resulting lower cost of capital associated with new construction in these neighborhoods is reflected in reduced marginal costs and downward shifts in neighborhood supply functions. OZ neighborhoods happen to have below average local supply elasticities for their metro areas. As a result, we show that the potential gains in consumer surplus from implementing the same tax incentive in non-OZ neighborhoods is greater by about \$4 million per tract on average. Moreover, these OZ tract gains are overstated by 30% if calculated using regional rather than tract level supply elasticity estimates.

2 Data

We assemble a data set that brings together information from a number of different sources. Using the Zillow's Assessor and Real Estate Database (ZTRAX) data files, we build various housing price and quantity measures, supplemented with aggregate census and American Community Survey (ACS) data from 1990, 2000 and 2008-12. We measure local labor demand conditions using the place of work and journey to work tabulations in the 1990 and 2000 U.S. Censuses of Population and the 2006 and 2010 LODES data plus census tract aggregate data from 1990-2010. Finally, we use remote sensing information on land cover in 2001 to measure baseline tract development intensity and topography and that from 2001-2011 to measure changes in developed land. All data are keyed to 2000 definition census tracts, covering 63,897 tracts in 306 metro areas (with some overlap across metros). Below we describe in more detail how we process each data source.

2.1 Housing Prices

Our primary source for housing data is the Zillow's Assessor and Real Estate Database (ZTRAX) (Zillow, 2017). These come in the form of files for transactions, most recent assessments before

2017 and prior assessments. These data cover more of the U.S. over time from 2000 to 2010, going from coverage of at least part of 267 metro areas and about two-thirds of sample tracts to three-quarters by 2010. Because of incomplete coverage, particularly in year 2000, we supplement Zillow data with decennial Census data, as explained in further detail below.

Transactions information is transcribed from local Recorders of Deeds and includes the sale price, location and some property attributes. To fill out property attributes, we merge in the most recent assessment data. We primarily use the resulting data set to construct price indexes at the census tract level. For the purpose of building home price indexes, we only use arm’s length transactions for resale or new construction. This excludes deed transfers involving non-transactions such as foreclosure by banks or quitclaim deeds. We include all residential units, including single family houses, townhouses and condominiums. We only consider homes that are bought by individual buyers and do not examine institutional buyers. We always exclude homes that sell more than 9 times over our sample period and tract-year combinations with fewer than 10 sales.¹

A well-known challenge for constructing home price indexes is that homes are heterogeneous in observed and unobserved attributes. The goal of the indices is to hold quality constant, eliminating all price variation due to differences in attributes. Leveraging the richness of assessment data on home characteristics, we use census tract-region-year fixed effects a_{irt}^{HI} from the following hedonic regression to build our Hedonic Index (HI).

$$\ln P_{hirtm} = a_{irt}^{HI} + \rho_m^{HI} + X_{hirtm}\beta^{HI} + e_{hirtm}^{HI}$$

Here, h indexes homes in census tract i , region r , year t and month m . X_{hirtm} includes a rich set of characteristics (including unit type, rooms, bedrooms, kitchens, bathrooms, heating and AC, elevator, fireplace, water, sewer, roof type, age and floorspace). Month of sale fixed effects ρ_m^{HI} flexibly account for seasonality in market conditions.

As noted in Section 2.6, we exclude tracts for which Zillow data do not have complete or accurate coverage. To fill out some measure of house prices for these tracts and to facilitate the pre-trends analysis that requires house prices for 1990-2000, we also build a lower quality hedonic price index using self-reported data from the 1990 and 2000 Censuses of Housing and the 2008-

¹This second restriction is intended to eliminate tract-years with selected coverage in the ZTRAX data. It drops 10% of tract-year observations in 2000 and 13% in 2010, with a minor impact on results.

2012 ACS aggregated to the census tract level. These are tract residuals a_{irt}^C from the following cross-sectional regressions estimated separately for 1990, 2000 and 2010:

$$\ln P_{irt}^C = X_{irt}^C \beta_t^C + a_{irt}^C.$$

Here, P_{irt}^C is the average self-reported value of owner-occupied homes in the tract and included in X_{irt}^C are fractions of the tract's owner-occupied units in various building types, with various numbers of bedrooms, and of various vintages. While it is of lower quality, this index covers all residential census tracts in the U.S.

Hedonic indices do not account for unobserved heterogeneity in quality across homes. To account for this, we also use the ZTRAX data set to build a repeat sales index (RS) at the tract-year level. For this index, we exclude any sales fewer than 180 days after the prior sale. Inclusion of home fixed effects α_{hir}^{RS} in the following regression purges individual home heterogeneity that is fixed over time. Tract-year fixed effects a_{irt}^{RS} from this regression form our repeat sales index:

$$\ln P_{hirtm} = a_{irt}^{RS} + \rho_m^{RS} + \alpha_{hir}^{RS} + e_{hirtm}^{RS}.$$

After homes are renovated, we treat them as new homes for the purpose of constructing this index. We recognize that this index may suffer from a less representative sample than the hedonic index and incorporate unwanted capitalization of unobserved home improvements. Nevertheless, it provides a nice complement to the hedonic price index. For further robustness, we compared the Zillow repeat sales index with the tract level Federal Housing Financing Agency (FHFA) repeat sales price index. The correlation between the growth rates of the two indices is about 0.75. The FHFA price index only covers single family house transactions involving conforming and conventional mortgages. For this reason, we focus on Zillow-based measures as they have more complete geographic coverage and include condominium sales.

The top block in Table 1 presents summary statistics about changes in these three home price indexes for the primary estimation sample used in the empirical work. The Zillow hedonic price index growth is 0.62 on average across tracts relative to 0.64 for repeat sales index growth during the 2000-2006 period. For 2000-2010, average growth rates are 0.26 for each, reflecting the 2007-2008 housing market crash. Correlations between the two Zillow indexes is 0.92 for the 2000-2010

period but the correlation with the growth in the Census price index is only about 0.45 for both Zillow based indexes.

2.2 Housing Quantities

We construct five measures of housing quantity changes to cover different aspects of supply responses, ranging from changes in housing stocks, changes in developed land, changes in floorspace, total new construction, new development and redevelopment. As most of the existing literature on housing supply focuses on units in the housing stock, we begin with this measure. Since the Zillow data have incomplete coverage at the beginning of the sample, most of our stock measures rely on 2000 census data for the base year. To gain a conceptually consistent measure of housing supply, we further explicitly measure floorspace, new construction and redevelopment, which the detailed Zillow assessment data facilitate.

To organize our classification of stock measures, we begin by noting that the total residential floorspace in census tract i , S_i , is the amount of developed land L_i times the average number of housing units per developed parcel H_i/L_i times the average floorspace per housing unit S_i/H_i . Differentiating, the growth rate in floorspace over time can be written as follows:

$$d \ln S_i = \underbrace{d \ln L_i}_{d \ln H_i} + \underbrace{d \ln \frac{H_i}{L_i} + d \ln \frac{S_i}{H_i}}_{d \ln A_i} \quad (1)$$

The first two terms on the right hand side of equation (1) also add up to changes in housing units, $d \ln H_i$, which can be decomposed as $\frac{dH_i^R}{H_i} + \frac{dH_i^U}{H_i} + \frac{dH_i^T}{H_i}$, where dH_i^R refers to new units developed on already-developed land (redevelopment), dH_i^U refers to new units developed on land that has not been developed before (new development), and dH_i^T refers to the combination of full depreciation and teardowns and is always negative. In addition, the second and third terms on the right hand side of equation (1) add up to changes in housing services per parcel, $d \ln A_i$, where $A_i = \frac{S_i}{L_i}$ is measured as floorspace per parcel.

2.2.1 Housing Units and Total New Construction

The simplest unit quantity growth measures ($\frac{dH_i}{H_i}$) are the 1990-2000 and 2000-2010 growth rates in tract occupied housing units reported in the 100% count Decennial Census data. We use occupied

units instead of all units to be consistent with Saiz (2010) and because vacant units may be under-reported or not habitable. While it has the best neighborhood coverage, this measure is not an ideal new supply measure as it includes teardowns and depreciation.

We separate out new construction $\frac{dH_i^R + dH_i^U}{H_i}$ using information on building age in the ACS and ZTRAX data sets. In the 2008-2012 ACS tract aggregates, we observe the number of units in the stock built between 2000 and 2009. Following the ZTRAX historical assessment data forward, we record the earliest year built after 2000 for every housing unit in the ZTRAX data set.² Because of incomplete ZTRAX coverage in the earlier part of our sample and to be consistent across measures, we use the occupied housing stock reported in the 2000 census as a base for both measures. As the ACS is based on a 5% sample of occupied units while ZTRAX in principle covers the universe of new construction, the ZTRAX measure is more accurate.

Summary statistics for these three measures are presented in the second block of Table 1. Here we see that for the 2000-2009 calendar years, the average construction rates of new units across tracts in our sample, measured using both the ACS and Zillow data, are very close at about 11 percent growth over the census-defined base in 2000. The average census growth number for the same period, which incorporates teardowns, is 8 percent. Pairwise correlations between these three measures are all over 0.91.

2.2.2 Redevelopment New Construction Units

Redevelopment (dH_i^R) is an important component of housing supply, as it may have a different cost structure than new construction on undeveloped land. Moreover, in cities where building density is already high, builders can only increase housing supply through redevelopment. Urban redevelopment can take many forms, including teardowns and infills, which became widespread during the housing boom of the 2000s. At the peak, the number of demolitions and teardowns in the Chicago metropolitan area approached 40% of sales in 2005 (McMillen & O'Sullivan, 2013). In New York City, annual teardown activity increased almost eight-fold from 1994 to 2004 and peaked in 2005 (Been, Ellen, & Gedal, 2009).

Lacking data on demolition permits or infill construction, we quantify the units built through

²Some rental buildings only report total square footage and do not break out the number of units. In these cases, we impute the number of rental units using the average square footage of units in other rental and condominium buildings of similar size in the tract.

redevelopment by imputing the number of units built on already developed land in the calendar years 2000 through 2009 as follows. We assume that each tract's stock of units reported in the 2010 census is uniformly spatially distributed across the tract's developed area as measured using 2011 satellite land cover information described below. We subtract off the number of ACS reported new construction units 2000-2009 that is imputed to be on newly developed land using this assumption from 2010/2011 about the spatial distribution of housing units in each tract. We infer that the remainder of 2000-2009 new construction is from redevelopment. If this remainder is negative, we assign 0 units to redevelopment. Comparing the new units on already developed land with total new construction units as shown in Table 1, 40% of new construction in an average tract in our sample is redevelopment.

2.2.3 Floorspace

In much of the housing production literature, including the model we develop below, homes are viewed as differing only in the efficiency units of homogenous housing services they provide. If recently built housing units are larger than those built in the early 2000s or some existing housing units have been renovated, using either the difference in stocks or new construction would underestimate the true growth in housing supply, as highlighted in equation (1). To account for changes in housing quality, we first use Zillow assessment data to construct changes in tract floorspace as a way to measure the full housing supply response. We also construct a broader quality-adjusted measure of changes in housing services that incorporates all observed housing attributes, including square footage and the number of different types of rooms, with corresponding hedonic weights. The two measures are highly correlated and generate almost identical empirical results throughout our analysis. Given the possible endogeneity issues associated with constructing the hedonic weights, we focus on the floorspace measure in the analysis hereafter, as it is a more straightforward measure of housing services.

2.3 Satellite Data

We use remote sensing information to measure tract level topography and land development intensity. Land cover information is used to help determine whether new housing is built on previously undeveloped land.

We use three remote sensing data sets for land cover and topography. First, the “Scientific Investigations Map 3085” is derived from the US Geological Survey’s National Elevation Database. This data set uses raster information on slope and elevation range for all 30X30 meter land pixels within a 0.56 km radius (1 sq. km) of each pixel to classify it into one of nine categories that describe how flat or hilly the surrounding area is. After much experimentation with various options, we focus on the fraction of land area surrounded by “flat plains” as our main topographic measure in the empirical work. Flat plains are defined to have a slope of less than 8% in more than half of these nearby pixels and an elevation range of less than 15 meters in this 1 km sq region. On average in our estimation sample defined below, 40% of tract land area is flat plains. We also independently extract elevation range within each tract directly from the US Geological Survey National Elevation 1/9-1/3-1-2 arc second Database.

Development costs not only depend on topographical conditions but also the initial developed state. We construct tract developed fraction from the National Land Cover Database (NLCD) for 2001 and 2011. For each 30X30 meter cell, the NLCD provides one of 4 categories of development (0-19%, 20-49%, 50-79%, 80-100%). We construct the square meters of land in each tract by density of development and aggregate, assigning category medians, to impute average developed fraction for the land area of each census tract. We denote $d \ln L_i$ as the growth rate of tract land that has ever been developed, both as measured and conceptually in the model developed below. The average tract in our estimation sample experienced 8 percent 2001-2011 growth in developed land off of a base of 33% developed in 2001.

We aggregate the resulting tract level data to construct various land unavailability measures for each metro area. To be consistent with Saiz (2010), we calculate the fraction of area within 50 km of the CBD of each region that is undevelopable due to a steep slope (e.g. mountains), water or wetlands (e.g. oceans, lakes, etc), and that is developed. We also build variants of these two measures instead aggregating to the metro area level and within 50% or 100% radii from the CBD to the furthest tract in each metro area. As these measures are highly correlated, our results are not sensitive to the choice of aggregation. Lutz and Sand (2019) construct similar measures of land unavailability for metro areas using some of the same data sources.

Figure 1 Panel A shows kernel densities of fraction flat and fraction developed. Both are bimodal. Fraction flat has modes near the extremes of 0 and 1 while fraction developed is a smoother

distribution with modes near 0 and 0.4. Figure 1 Panel B shows that both decline on average with CBD distance, though fraction developed declines more rapidly.

2.4 WRLURI and FAR

The Wharton Residential Land Use Regulatory Index (WRLURI) is constructed from a battery of survey questions sent to a weighted random sample of municipalities nationwide in the US in year 2005. The index is expressed in population-weighted standard deviation units. While larger urban municipalities were sampled with higher probability, a large number of smaller suburban municipalities were also included in the sample. 261 of the 306 regions in our sample have at least one municipality surveyed. However, the municipality of the CBD is sampled in only 164 of our sample regions. Overall, our data includes 2,373 municipalities and 30,526 tracts with WRLURI information.

We also incorporate separately collected information on Floor Area Ratio (FAR) restrictions on residential development from the municipalities of Atlanta, Boston, Chicago, Denver, Los Angeles, New York, San Francisco, and Washington.³ For each residential land parcel, local zoning maps provide the residential FAR. We use the average of these within each census tract, weighted by parcel area.

Figure 1 Panel B shows that both FAR and the Wharton Index fall with CBD distance out to about 15% of the way to the urban fringe. As the Wharton Index is measured at the municipality level, its decline within this range of CBD distance is fully due to between region variation from increased representation of less dense central cities, which are typically less heavily regulated. The Wharton Index rises steeply beyond that such that land use in municipalities 30% of the way to the urban fringe is on average more heavily regulated than at CBDs. Beyond 30%, we have no FAR data and we do not have sufficient Wharton Index coverage to precisely measure regulation.

2.5 Population, Employment and Commutes

The Census Transportation Planning Package (CTPP) reports tabulations of 1990 and 2000 census data by residential location, work location and commuting flow. The 1990 CTPP geography determines our study regions. The 1990 CTPP assigns microgeographic units the size of census tracts

³Most of these data were generously provided by Ruchi Singh (Brueckner & Singh, 2020).

or smaller to “regions”, which roughly correspond to metropolitan areas. These regions can overlap. Commuting flows and times are reported for pairs of census tracts, traffic analysis zones or block groups within each region only. Employment by place of work, sex and 18 industry groups are reported for these same geographic units. For Connecticut and New Jersey, which are fully contained in one large 1990 CTPP region each, we develop new regions that each have a 25 km radius around each CBD in each state. We map 1990 CTPP geography to 2000 definition census tracts by overlaying their digital maps and using land area as allocation weights. The 2000 CTPP is more spatially comprehensive and thus can be restricted to cover only 1990 region definition geography. The result is a total of 63,896 2000-definition census tracts (comprising 50,410 unique tracts) in 306 regions.

For most regions, central business district (CBD) locations are taken as the centroid of the set of census tracts reported as being in the CBD in the 1982 Economic Census. Remaining CBD assignment is done by eyeballing a location that is near city hall and the most historical bank branches in the region’s largest city.

Empirical implementation requires information on the commute time between each pair of census tracts in each region. Because they are based on only a sample and flows of fewer than 5 sampled workers are suppressed, commutes are not observed between about one-half of tract pairs in 1990 and two-thirds of tract pairs in 2000. To fill in the rest, we develop a forecasting model based on tract relative locations. In particular, we impute origin-destination commute times using out of sample predictions from a regression of log travel time on region fixed effects, log travel distance, log CBD distance of workplace and log CBD distance of residence. See Baum-Snow, Hartley, and Lee (2019) for details.

For 2006 and 2010, we use the LEHD origin destination employment statistics (LODES) data to measure employment by place of work. As this data set does not have commute times, we maintain year 2000 commute times for these later years.

We take census tract aggregates for 1970-2010 from the Neighborhood Change Database supplemented with some Summary Tape File 4 variables from 1980. We use these variables to measure aggregate outcomes and to control for pre-treatment trends in observables.

2.6 Estimation Sample

Our analysis requires reliable information on housing quantities. To this end, we only include census tracts with Zillow unit counts that are close to 2010 100% census counts. In particular, we exclude all tracts from the estimation sample for which our 2010 Zillow unit stock is more than 20% above or below the occupied housing stock reported in the 2010 census. In addition, we exclude tracts in non-disclosure states for which Zillow reports they have incomplete coverage.⁴ For these tracts, we are particularly concerned about under-measurement of 2000-2010 new construction. As a result, our estimation sample is cut from 50,410 tracts (63,896 observations) in 306 partially overlapping regions to 19,985 tracts (25,361 observations) in 170 regions. Primary estimation sample sizes for specific variables are reported in Table 1.

For building instruments and for structural estimation of the model, we need information on labor market opportunities that are relevant to each census tract in the estimation sample. Because the CTPP and LODES fully cover our sample area, these data sets do not introduce any sample constraints.

3 Conceptual Framework

The main object of our analysis is to empirically characterize how housing supply elasticities differ across neighborhoods and regions. While we use the decomposition in (1), developed further below, to study different components of supply, we intentionally impose as little structure as possible on the form of the housing supply function.

As such, we focus on recovering estimates of γ_{ir} in the following reduced form expression. Q_{ir}^s denotes any of the quantity measures listed in (1) in tract i of metropolitan area r and P_{ir} is the observed price per unit of housing services. To accommodate ad-valorem taxes, we can think of the price developers receive per unit of housing services sold as $P_{ir}(1 - t_r)$. Region fixed effects θ_r thus capture potential tax wedges and other region-specific factors that influence construction costs.

$$\ln Q_{ir}^s = \theta_r + \gamma_{ir} \ln P_{ir} + u_{ir} \quad (2)$$

⁴These states are Alaska, Idaho, Kansas, Louisiana, Mississippi, Missouri, Montana, New Mexico, North Dakota, Texas, Utah, and Wyoming.

u_{ir} includes within-metro area supply shifters, both observed and unobserved. We allow γ_{ir} to depend on tract i 's observed heterogeneity, including initial building density, geographic features and distance to the central business district in metropolitan area r . Because of the durability and immobility of housing, (2) is likely to hold with a greater γ_{ir} for price growth than for price declines (Glaeser & Gyourko, 2005). For this reason, despite using price and quantity information for 2000-2010, we rely on the 2000-2006 period for demand shocks. During this time, price growth was positive in over 98 percent of the tracts in our sample, more so than for any other time period in our data.

In this section, we first sketch a simple model of neighborhood housing supply. While stylized, the model delivers a natural decomposition of the residential floorspace supply elasticity into intensive (floorspace per parcel) and extensive (parcel development) margins. It also delivers a theoretical basis for the unit supply elasticities we measure in the empirical work. Finally, it provides rough calibrated quantification of the intensive margin component and microfounds our empirical specifications. We then incorporate the housing supply function into an urban spatial equilibrium model which links neighborhood housing and labor markets in an urban area. This part of the model provides theoretical support for the instruments we use to pin down consistent estimates of parameters used in the construction of γ_{ir} and helps guide our opportunity zone application in Section 6. Altogether, our conceptualization of the data generating process serves as a microfoundation that guides both our empirical strategy for estimating housing supply functions and the interpretation of empirical findings.

3.1 A Theory of Neighborhood Housing Supply

We analyze a static environment in which competitive developers with access to the same technology produce housing on some land parcels in each neighborhood. As the model is static, it is most natural to view it as describing comparisons of housing supply responses across different neighborhoods that are ex-ante identical but experience different exogenous increases in the price of housing services. These are the treatment effects that the empirical work is set up to measure. The key object delivered by the model relevant for the empirical work is a description of relationships between relative housing stocks and relative prices across these ex-ante identical neighborhoods.

We analyze a representative developer. The developer only builds on land parcels with fixed

development costs that are sufficiently low such that the variable profit minus fixed development cost is weakly positive. Each parcel's land value is the variable profit net of the fixed development cost. Conditional on development, the amount of floorspace supplied on each parcel in neighborhood i is A_i . A_i thus represents the intensity of development on any given lot in neighborhood i , incorporating both floorspace per housing unit and units per parcel. It is chosen based on tract housing productivity, parcel size and demand conditions captured by P_i . The price per unit of housing services P_i is the same for all parcels in neighborhood i .

Developers combine land and capital to produce housing services. Each building lot l in neighborhood i has the fixed lot size \bar{M}_i , faces the same variable cost function $C_i(A_i)$ and has the lot-specific fixed development cost g_{il} . The fixed development lot size assumption reflects land assembly frictions that are likely to bind over the 5-10 year time horizon that is the focus of our empirical analysis (Brooks & Lutz, 2016). The fixed cost g_{il} captures development fees and permitting costs plus land preparation costs. Each tract has its own distribution of fixed development costs $F_i(x)$.

A representative developer's profit associated with building on parcel l in neighborhood i is

$$profit_{il} = P_i A_i(P_i) - g_{il} - C_i(A_i) - p_{il},$$

where p_{il} is the endogenous parcel acquisition price. Imposing 0 profits and perfect competition (marginal cost pricing), we have

$$p_{il} = C_i(A_i(P_i)) \left(\frac{d \ln C_i(A_i(P_i))}{d \ln A_i} - 1 \right) - g_{il}.$$

This is the bid-rent function for lot l in neighborhood i . Consistent with a fixed parcel size and Cobb-Douglas housing production in land and capital, the variable cost function exhibits the property that $\frac{d \ln C_i(A_i(P_i))}{d \ln A_i} > 1$. As such, the first term reflects the intuition that more development implies greater variable profits, which get capitalized into a higher parcel price. The second term reflects capitalization of the fixed development cost into the parcel price. Henceforth, consistent with Cobb-Douglas production, we assume that $\frac{d \ln C}{d \ln A} - 1 = \phi > 0$. Normalizing the opportunity cost per unit of land to 0, this means that the fraction of land developed in each tract is $F_i[\phi C_i(A_i(P_i))]$. All derivations are in the Appendix.

The stock of developed land in tract i , $L_i(P_i)$, is $M_i F_i(\phi C_i[A_i(P_i)])$, where M_i is tract i 's total land endowment. Differentiating the developed land supply function yields:

$$\gamma_i^{land} \equiv \frac{d \ln L_i(P_i)}{d \ln P_i} = \frac{f_i(\phi C_i[A_i(P_i)])}{F_i(\phi C_i[A_i(P_i)])} \frac{d \ln A_i(P_i)}{d \ln P_i} \phi P_i A_i(P_i). \quad (3)$$

Expression (3) highlights that tracts with a greater density of parcels available for developing at the fixed cost that equals marginal variable profit, represented by a higher $f_i(\phi C_i[A_i(P_i)])$, exhibit more elastic land supply. Land supply is a component of floorspace and unit supply, which we turn to now.

Starting with the tract-level aggregate housing services supply function $\ln S_i(P_i) = \ln A_i(P_i) + \ln L_i(P_i)$, it follows that

$$\gamma_i^{space} \equiv \frac{d \ln S_i(P_i)}{d \ln P_i} = \frac{d \ln A_i(P_i)}{d \ln P_i} + \frac{f_i(\phi C_i[A_i(P_i)])}{F_i(\phi C_i[A_i(P_i)])} \frac{d \ln A_i(P_i)}{d \ln P_i} \phi P_i A_i(P_i). \quad (4)$$

To get a sense of magnitudes and to connect this framework more closely to the empirical work, we parameterize assuming a Cobb-Douglas production technology with the common land share α and a tract-specific productivity. The resulting parcel-specific housing services supply function is $A_i(P_i) = \rho_i P_i^{\frac{1-\alpha}{\alpha}}$.⁵ We further assume that the fixed cost follows the Frechet distribution, where $F_i(x) = \exp(-\Gamma_i x^{-\lambda})$, with the common dispersion parameter $\lambda > 1$ and the tract-specific scale parameter $\Gamma_i > 0$. With these functional form assumptions, the generalized supply elasticity expression from (4) becomes

$$\frac{d \ln S_i}{d \ln P_i} = \underbrace{\frac{1-\alpha}{\alpha}}_{\frac{d \ln A_i}{d \ln P_i}} + \underbrace{\alpha^{-1-\lambda} \lambda \rho_i^{-\lambda} P_i^{-\frac{\lambda}{\alpha}} \Gamma_i}_{\frac{d \ln L_i}{d \ln P_i}}. \quad (5)$$

Equation (5) highlights how developers respond to positive demand shocks along both intensive and extensive margins. The first term captures the intensive margin. As price rises, developers increase the quantity of housing services supplied per parcel by $\frac{1-\alpha}{\alpha}$, where α is the land share in housing production. With α estimated to be 0.2-0.33 in the literature (Combes et al., 2019), $\frac{1-\alpha}{\alpha}$ reflects a floorspace supply elasticity of 2-4 holding the amount of developed land fixed.

⁵ $\rho_i = \iota^{\frac{\alpha-1}{\alpha}} (1-\alpha)^{\frac{1-\alpha}{\alpha}} \kappa_i^{\frac{1}{\alpha}} \bar{M}_i$ where ι is the cost of capital and κ_i is tract housing productivity.

The second term relates the extensive margin positively to the scale parameter, Γ_i , and negatively to the initial price level, P_i . In the Frechet example, fixed cost distributions in tracts with a higher Γ_i have higher means and variances and hence a thicker right tail. This implies a higher density of land available for development at the marginal variable profit and hence higher γ_i^{land} , which in turn boosts γ_i^{space} . We expect that tracts with lower initial development density, more flat land and less stringent regulation are characterized with fixed cost distribution with a higher scale parameter, Γ_i , and hence more responsive along the extensive supply margin.⁶ In addition, in the Frechet example a higher initial price in tract i implies a thinner right tail of the fixed cost distribution and hence less land available for development at the marginal fixed cost, which in turn causes the extensive margin to be less responsive.

We now turn to the tract-level unit supply function $\ln H_i(P_i) = \ln(H_i/L_i)(P_i) + \ln L_i(P_i)$. It follows that

$$\gamma_i^{unit} \equiv \frac{d \ln H_i}{d \ln P_i} = \frac{d \ln(H_i/L_i)}{d \ln P_i} + \frac{f_i(\phi C_i[A_i(P_i)])}{F_i(\phi C[A_i(P_i)])} \frac{d \ln A_i(P_i)}{d \ln P_i} \phi P_i A_i.$$

Conditional on development, we assume that units per parcel in all areas of a tract (H_i/L_i) is identical. Moreover, the nature of demand shocks are such that the positive intensive margin response $\frac{d \ln A_i}{d \ln P_i}$ is split into two positive components.⁷ As a result, we predict that $\gamma_i^{unit} > \gamma_i^{land}$ and that tract level characteristics affect unit supply elasticities in the same direction as land supply elasticities.

One tract characteristics of particular interest is the floor-area-ratio (FAR) restriction that constrains developers from building beyond some maximum intensity \bar{A}_i . A binding FAR constrains the intensive margin supply response to 0. As a result, the tract supply elasticity comes from the extensive margin only.⁸

$$\gamma_i^{space,FAR} = \gamma_i^{unit,FAR} = \gamma_i^{land,FAR} = \frac{f_i(P_i \bar{A}_i - C_i(\bar{A}_i))}{F_i(P_i \bar{A}_i - C_i(\bar{A}_i))} \bar{A}_i. \quad (6)$$

⁶In the empirical work we recognize that Γ_i and ρ_i may additionally depend on unobserved tract characteristics. We also recognize that these same attributes may be supply shifters.

⁷The equilibrium split of ($A_i \equiv S_i/L_i$) into floorspace per units (S_i/H_i) and units per parcel (H_i/L_i) depends on the composition of housing demand (families vs. single people for example), which neither our data nor identification strategy are well suited to handle. However, we note that single family home construction, or 1 unit per parcel, makes up over 90 percent of 2000-2010 new construction units in our sample area.

⁸See Appendix A.6 for the derivation.

Equation (6) shows two forces through which a less binding FAR (increase in \bar{A}_i) affects the supply elasticity. First, the mechanical effect of being allowed to build more increases supply elasticity by making more parcels viable for development with a marginal price increase. Second, a higher \bar{A}_i attracts more development on available parcels, thereby changing the supply of developable parcels at the marginal fixed cost $P_i\bar{A}_i - C_i(\bar{A}_i)$. This could dampen or increase supply elasticity depending on the form of the fixed cost distribution $f_i(x)$. The net effect of a relaxation of a FAR on supply elasticity is thus an empirical question, which we examine in Section 5.2.

Urban growth takes the form of both development and redevelopment. Our supply model can be further extended to incorporate redevelopment. To see this, consider an environment in which $f_i^L(x)$ is the density of fixed development costs across all parcels in tract i as if they had no prior development and r is an additional fixed redevelopment cost. Then, $f_i^R(x)$ is the left tail of the $f_i^L(x)$ distribution up to ϕC_i but shifted by r to reflect the additional redevelopment cost; $f_i^U(x)$ is its right tail at fixed costs above ϕC_i only and rescaled to integrate to 1. Both have support between fixed costs ϕC_i and $\phi C_i + r$. With this extension, the fixed cost distribution $F_i(x)$ is decomposed into $F_i(x) = \frac{M_i^R}{M_i} F_i^R(x) + \frac{M_i^U}{M_i} F_i^U(x) = F_i^L(x - r)1(x \leq \phi C_i) + F_i^L(x)1(x \geq \phi C_i)$, where M_i^R is the mass of previously developed land parcels in tract i and M_i^U is the analogous object for undeveloped parcels. A corresponding version of equation (3) is then:

$$\gamma_i^{land} \equiv \frac{d \ln L_i(P_i)}{d \ln P_i} = \left(\frac{M_i^R f_i^R(\phi C_i)}{M_i^{R'}} + \frac{M_i^U f_i^U(\phi C_i)}{M_i^{R'}} \right) \frac{d \ln A_i(P_i)}{d \ln P_i} \phi P_i A_i(P_i), \quad (7)$$

where $M_i^{R'} = M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)$ is the amount of land that is newly developed. Equation (7) decomposes the land supply response into land redevelopment and new construction on undeveloped land respectively.

Figure A1 plots $f^L(x)$, $f^R(x)$ and $f^U(x)$ for an example tract. A few implications follow. First, as P_i rises, marginal land parcels are developed left to right in the region of overlapping support of the $f^R(x)$ and $f^U(x)$ distributions. Developers in neighborhoods with greater price growth carry out both additional redevelopment and additional development of previously undeveloped land, relative to developers in neighborhoods with smaller price growth, to supply new housing. Second, the relative magnitudes of the land redevelopment versus new land development elasticities depend on the density of the fixed cost distribution $f_i^L(x)$ at ϕC_i and $\phi C_i + r$ and the relative amounts of previously undeveloped versus developed land in the tract. If a large fraction of tract

land was previously developed, that will tend to boost the redevelopment elasticity as it gives developers greater opportunity to find parcels with relatively low fixed costs. Third, given that $f_i^R(\phi C_i)$ and $f_i^U(\phi C_i)$ both depend on parameters that govern the $F^L(x)$ distribution, tract characteristics affect the redevelopment supply elasticity in the same way as the extensive margin unit supply elasticity in equation (5).⁹

3.2 Housing Demand

We incorporate housing supply conditions that are allowed to differ across locations within cities into a version of the quantitative urban model developed by Ahlfeldt, Redding, Sturm, and Wolf (2015) and extended by Tsivanidis (2018). While tracing out housing supply functions is ultimately about estimating reduced form impacts of housing demand shocks on housing quantities and prices, this part of the theory is helpful in operationalizing this goal in three ways.

First, the model shows how to leverage variation across space within cities in local labor demand shocks to isolate exogenous variation in housing demand shocks across census tracts. The model structure facilitates recovery of causal linkages from labor demand shocks to housing demand shocks, as filtered through the commuting time matrix. We show how housing demand conditions in each census tract i can be summarized through “Resident Market Access” RMA_i , which is the sum of commute time discounted skill prices available to residents of tract i . This object can be readily calculated with available data on the numbers of workers and residents in each tract. Shocks to skill prices in commuting destinations are reflected as shocks to RMA_i .

Second, the model makes clear the conditions required for census tract level “Bartik” shocks to represent a valid source of econometric identification. These shocks are calculated by predicting 2000 to 2006 tract level growth in employment by industry using 1990 tract level employment shares by industry with national industry specific employment growth outside of the metro area in question. Originally proposed by Bartik (1989, 1991), this source of variation has been used at the metro area level in Saiz (2010), Notowidigdo (2013) and Diamond (2015) among many others. This and Baum-Snow et al. (2019) are among the first papers to use this source of variation for identification at the sub-metro level of geography. To help us do this in a sensible way, we explicitly introduce industries k into the model.

⁹Appendix A.5 formally lays out unit and floorspace supply elasticities under redevelopment.

Finally, the model delivers enough structure to aggregate neighborhood housing supply functions to the metro area given an array of shocks to neighborhood fundamentals and to perform welfare analysis of place based policies, as we do in Section 6.

3.2.1 Setup

While our main empirical work uses data for over 150 metropolitan areas, our empirical focus is on within-metro area variation in housing supply elasticities. As such, our model is of a single metro area.

The model features a continuum of ex-ante identical workers indexed by ω who choose residential tract i , work tract j and industry of work k within the metro area. They receive productivity shocks $z_{ijk\omega}$ over commute origin-destination and industry triplets and preference shocks $v_{i\omega}$ over residential locations. The preference shocks are revealed first, leading agents to first choose residential locations anticipating the quality of employment opportunities nearby but before productivity shocks are revealed. Productivity shocks are then revealed and agents choose work locations second. In practice, the shocks primarily allow the model to generate enough variation in tract populations, home prices and employment to rationalize the data.¹⁰

The indirect utility person ω receives from living in tract i , commuting to tract j and working in industry k is

$$v_{ijk\omega} = \frac{v_{i\omega} B_i z_{ijk\omega} w_{jk}}{P_i^{1-\beta} e^{\kappa\tau_{ij}}}, \quad (8)$$

where B_i is a local amenity, w_{jk} is the price of a unit of skill in commuting destination j and industry k , P_i is the price of one unit of housing services in i and $\kappa\tau_{ij}$ is the fraction of time spent commuting for those living in i and working in j . In the data, we observe the price P_i in year 2000 and beyond and the commuting time τ_{ij} in 1990 and 2000. The productivity shock $z_{ijk\omega}$ is drawn from the Frechet distribution with shape parameter ε .

$$F_z(z_{ijk\omega}) = e^{-z_{ijk\omega}^{-\varepsilon}}, \varepsilon > 1 \quad (9)$$

Following Tsivanidis (2018) and Couture et al. (2019), we introduce a nested preference shock

¹⁰The sequencing of shock revelation is for analytical convenience. As our main goal is only to provide enough structure to show how to use Bartik shocks in commuting destinations to generate variation in housing demand in residential locations, this choice is not central for our analysis.

over residential locations $v_{i\omega}$. This shock is also distributed Frechet but with shape parameters η and ψ . This nested structure allows individuals to have different elasticities of substitution in demand between neighborhoods within versus between municipalities, where municipalities are indexed by m and $i(m)$ refers to neighborhood i in municipality m .

$$F_v(v_{i\omega}) = \exp\left[-\sum_m \left[\sum_{i(m)} v_{i\omega}^{-\eta}\right]^{-\frac{\psi}{\eta}}\right], \psi > 1, \eta > 1 \quad (10)$$

Incorporation of this second shock allows the model to generate situations in which people would choose to reside in tracts with lower expected utilities as calculated based on $\frac{B_i z_{ijk\omega} w_{jk}}{P_i^{1-\beta} e^{\kappa\tau_{ij}}}$ only. As a practical matter, it also delivers a convenient expression for mean income net of commuting cost in each tract, as is derived below. If the distribution functions for the two shocks are identical and $\eta = \psi = \varepsilon$, the utility shock becomes redundant and this model reduces to one similar to that in Ahlfeldt et al. (2015).

3.2.2 Resident Market Access

Solving the model backwards, conditional on living in residential location i the probability that work location j provides the highest utility is

$$\pi_{ij|i} = \frac{\sum_k [w_{jk} e^{-\kappa\tau_{kj}}]^\varepsilon}{\sum_k \sum_{j'} [w_{j'k} e^{-\kappa\tau_{ij'}}]^\varepsilon} \equiv \frac{\sum_k [w_{jk} e^{-\kappa\tau_{kj}}]^\varepsilon}{RMA_i}. \quad (11)$$

RMA_i is a convenient summary measure of the access to employment opportunities from residential neighborhood i . In particular, many objects in the model are constant elasticity in RMA_i and it can be readily calculated with available data. Derivations are in the appendix.

Before the productivity shock is revealed, the expected income (wage net of commuting cost) \bar{y}_i associated with residing in tract i is

$$\bar{y}_i = \Gamma\left(1 - \frac{1}{\varepsilon}\right) (RMA_i)^{\frac{1}{\varepsilon}}. \quad (12)$$

As a result, population supply to tract i is given by

$$\pi_i = \mu \left[\sum_{i' \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} RMA_{i'}^{\frac{1}{\varepsilon}} \right)^\eta \right]^{\frac{\psi}{\eta}-1} \left(B_i P_i^{\beta-1} RMA_i^{\frac{1}{\varepsilon}} \right)^\eta. \quad (13)$$

This expression reflects the attractiveness of neighborhood i 's amenities and labor market opportunities, as balanced against its housing cost. This attractiveness is relative to the attractiveness to other neighborhoods in the municipality $m(i)$, captured by the object inside the summation.

Equilibrium commute flows, calculated as $\pi_{ij} = \pi_{ij|i}\pi_i$, follow a standard gravity equation in commute time τ_{ij} .

$$\ln \pi_{ij} = a_i^g + b_j^g - (\kappa\varepsilon)\tau_{ij} \quad (14)$$

That is, a regression of log commute probabilities between each origin-destination pair on origin and destination fixed effects plus commute time τ_{ij} recovers an estimate of the parameter bundle $\kappa\varepsilon$.

Labor supply to tract j is given by

$$L_j = \mu \sum_k \left[w_{jk}^\varepsilon \right] FMA_j, \quad (15)$$

where ‘‘Firm Market Access’’ FMA_j is a measure of the access to workers enjoyed by firms in tract j . Plugging into the definition of RMA_i , we have the following system of equations.

$$FMA_j = \sum_i \frac{e^{-\kappa\varepsilon\tau_{ij}} \pi_i}{RMA_i} \quad (16)$$

$$RMA_i = \sum_j \frac{e^{-\kappa\varepsilon\tau_{ij}} L_j}{FMA_j} \quad (17)$$

Using data on employment L_j , residents π_i , the parameter cluster $\kappa\varepsilon$ and commute times τ_{ij} , we can calculate FMA_j and RMA_i by solving this system. We estimate $\kappa\varepsilon$ using separate flow-weighted commuting gravity regressions like (14) with origin and destination fixed effects in 2000 for each metropolitan region.¹¹ Because we do not observe tract-tract commute times after 2000, we hold commute times constant for later years. We calculate RMA_i using (16) and (17) for 2000, 2006 and 2010.

An individual who lives in i and works in j in industry k has housing demand of $(1 - \beta)^{\frac{1}{\beta}} \frac{1}{P_i}$ from Cobb-Douglas preferences. We assume that all sites in each residential location i are perfect

¹¹Across the 306 regions in our broad sample, the median estimated elasticity of commuting flow with respect to one-way commuting minutes in 2000 is -0.04, the minimum is -0.11 and the maximum is -0.01. Estimates of $\kappa\varepsilon$ are about twice as large in big cities like New York and Los Angeles than in small cities like Bryan-College Station, TX. This reflects the fact that households in bigger cities are willing to travel longer to reach work destinations.

demand substitutes, justifying the uniform price per unit of housing services P_i . Adding up, the log aggregate residential floorspace demand in tract i is thus

$$\ln S_i^d = \ln \rho_{HD} + \frac{1}{\varepsilon} \ln(RMA_i) + \ln \pi_i - \ln P_i. \quad (18)$$

This object is increasing in RMA_i conditional on population π_i because greater RMA_i is associated with greater income for tract residents. Conditional on P_i , equilibrium tract residential population π_i is also increasing in RMA_i , as seen in (13). Thus, shocks to RMA_i result in housing demand shocks. This is the key insight used for identification in the empirical work.

The reduced form empirical work uses the housing supply equation (2) in tandem with the housing demand equation formed by substituting (13) into (18). Credible identifying variation in $\ln P_i$ must come from a component of RMA_i that is cleansed of variation in housing productivities and lot sizes. Section 4 lays out how we isolate such variation using a simulated version of RMA_i based on Bartik type labor demand shocks in commuting destinations for residents of tract i .

3.2.3 Equilibrium

Combining conditions governing population supply to residential tracts (13), labor supply to work tracts (15) and imposing housing market clearing yields conditions describing equilibrium tract population and home prices. Differentiating the population condition over time yields the following structural equation.

$$d \ln \pi_i = \frac{\gamma_i + 1}{\gamma_i + 1 + \eta(1 - \beta)} \frac{\eta}{\varepsilon} \left(1 - \frac{1 - \beta}{\gamma_i + 1}\right) d \ln RMA_i + v_m^\pi + u_i^\pi \quad (19)$$

This equation incorporates an intuitive positive relationship between growth in employment opportunities and tract population. This relationship is stronger if housing supply in tract i is more elastic and/or if there is less dispersion in idiosyncratic preferences over locations (η is larger). v_m^π is a municipality fixed effect that captures common population trends in all tracts in municipality m that come through their correlation in neighborhood choices delivered by the outer nest in preferences over neighborhoods. The error term u_i^π is a function of shocks to amenities and housing productivity in tract i . We use (19) above as a basis for structural estimation of η , recognizing that identifying variation in $d \ln RMA_i$ must be uncorrelated with tract level shocks to amenities and

housing productivity for successful identification.

Details and additional discussion of model equilibrium are in the appendix.

4 Empirical Implementation

Our main estimation equation amounts to the differenced counterpart to the simple tract level supply equation (2).

$$\Delta \ln Q_{ir}^s = \theta_r + X_{ir}\delta + \gamma_{ir}\Delta \ln P_{ir} + \tilde{\rho}_{ir} \quad (20)$$

Observations are for tract i in metro region r . To allow for observed heterogeneity in supply elasticities, we parameterize γ_{ir} to depend on metro region and tract-specific observables Z_r^1 and Z_{ir}^2 .

$$\gamma_{ir} = Z_r^1\gamma_1 + Z_{ir}^2\gamma_2 \quad (21)$$

As is detailed in Section 2, these sources of observed heterogeneity are topography, developed fraction, land use regulation and regulatory burden. Because we do not observe some relevant factors that may differ by CBD distance, we also include CBD distance interactions in some specifications. As our empirical setting only allows us to recover relationships between observed tract attributes and supply elasticities, interaction coefficients also likely incorporate influences of unobserved factors. For example, if tract fraction developed is correlated with unobserved input costs, estimates of the coefficient on the interaction between fraction developed and price growth would in part capture impacts of input cost differences on supply elasticities. This means that while our estimates are well suited for characterizing existing tract level housing supply elasticities, they are less appropriate for making causal predictions about impacts of changing one observed attribute holding all else constant. Instead, our empirical implementation is primarily oriented toward ensuring that variation in price growth across tracts is uncorrelated with unobserved supply shifters, which allows us to recover information about supply elasticities.

The first two terms in (20) are included for identification reasons. Fundamental to our empirical strategy is inclusion of metro region fixed effects θ_r . Their inclusion ensures that we compare different neighborhoods in the same labor market for identification. Robustness checks include these fixed effects interacted with 2-2.5 km CBD distance rings. In tract characteristics X_{ir} , our main specification includes lagged demographic attributes, a quadratic in CBD distance, 2001

tract developed fraction and share flat land, and controls for tract-specific labor demand shocks. Our controls for 1990 and 2000 tract demographic characteristics account for potential influencers of the tract regulatory environment that may be correlated with the instruments we lay out below. CBD distance controls hold constant any potential spatial trends in price growth that are related to costs and are useful given the stronger 2000-2010 labor demand growth in suburban areas. 1990 and 2000 Census rent and price indexes help to account for decadal mean reversion in home price growth. Controls for developed fraction and topography hold constant obvious potential sources of housing supply shocks. Finally, 1990 employment and a 2000-2006 tract-specific Bartik labor demand shock (explained below) ensure that our IV implementation is only using variation from outside of tract ir for identification.

4.1 OLS Results

Table 2 presents basic OLS relationships between various measures of post-2000 house quantity changes and contemporaneous price growth, controlling for metro region fixed effects and the factors described above. Column (1) is for 2000-2006 and the remaining columns focus on the 2000-2010 period. Housing supply responses are measured as Zillow new construction in columns (1)-(2), ACS new construction in column (3), changes in census housing stocks in column (4), changes in developed land in column (5), changes in floorspace in column (6) and redevelopment in column (7). The first row uses the repeat sales index and the second row uses the hedonic price index. We see small positive coefficients of up to 0.12 for housing units and only up to 0.21 for floorspace. Coefficients turn negative if home price growth is measured using the census home price index (unreported).¹²

The implausibly low relationships between housing price growth and quantity growth point to several identification challenges in estimating housing supply. First, neighborhoods that experience stronger demand shocks may follow with unobserved changes in housing regulation in part in order to cope with these demand shocks - a classic endogeneity problem discussed extensively in Davidoff (2016). In particular, if housing development restrictions in the 2000s were loosened in response to a positive demand shock from the 1990s, the observed OLS relationship would trace out the demand curve and become negative, as our census based estimates show. Second,

¹²Ouazad and Ranciere (2019) find similarly small OLS relationships between price growth and quantity growth for the San Francisco metro region.

it is possible that positive productivity shocks outside of the construction sector may simultaneously boost local housing demand through higher household earnings and reduce housing supply through higher construction costs. This would further bias the OLS relationship between price and quantity growth downward. Moreover, our price index measure, while constructed as carefully as possible, is sure to be a noisy measure of the true price of housing services. Mechanical mean reversion in decadal house price growth that could reflect classical measurement error would lead to attenuation bias.

4.2 Instrument Construction

The broad message is the possibility for tract unobservables that predict supply shocks to be correlated with demand shocks, thereby generating a downward bias in OLS. A valid identification strategy must address the classic endogeneity concern of simultaneity in demand and supply by finding variation in local housing demand shocks across neighborhoods that are uncorrelated with shocks to local construction costs or housing productivity. To see where this can come from, consider the following tract level inverse housing demand equation from the model.

$$\begin{aligned} \ln P_i = & \tilde{\rho}_{HD} + \frac{\psi/\eta - 1}{1 + \eta - \eta\beta} \ln \sum_{i' \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} RMA_{i'}^{\frac{1}{\varepsilon}} \right)^\eta + \frac{1}{1 + \eta - \eta\beta} \frac{1 + \eta}{\varepsilon} \ln(RMA_i) \\ & - \frac{1}{1 + \eta - \eta\beta} \ln S_i^d + \frac{\eta}{1 + \eta - \eta\beta} \ln B_i \end{aligned}$$

The fact that the housing price in tract i is increasing in RMA_i through impacts on housing demand is intuitive. Labor demand conditions relevant to neighborhood i , as summarized in RMA_i , represent a key source of variation in house prices. However, any component of RMA_i that is correlated with tract housing productivity or land parcel size is endogenous to housing supply. Indeed, through its codetermination with FMA_i , RMA_i depends structurally on tract population which itself depends on tract housing productivity and parcel size. As such, we develop instruments that pick out components of $d \ln RMA_i$ that are likely orthogonal to levels of and shocks to productivity or other factors that influence local construction costs.

To build instruments, denoted $\Delta \ln \widetilde{RMA}_i$, we start with (16) and (17) as a basis for calculating a simulated version of $d \ln RMA_i$ that plausibly excludes shocks to tract housing productivity and its correlates. This simulated instrument serves a dual purpose. First it is a reduced form

housing demand shock that drives exogenous variation in tract level house price growth, as represented above. Second, it is a predictor of the structural object $d \ln RMA_i$ that is unrelated to tract level shocks to local amenities or housing productivities. The latter use will allow us to recover estimates of η for use in the Opportunity Zone application.

Instead of using actual employment in all commuting destinations in these calculations, we use the employment predicted by national growth rates and initial industry composition in each tract to solve for \widetilde{RMA}_i , after evenly scaling up the residential population of each tract to maintain labor market clearing and allowing us to solve jointly for \widetilde{FMA}_j . For components of instruments, we impose 1990 commute times and initial employment shares by industry and use estimates of $\varepsilon\kappa$ for 2000. We exclude all tracts within 2 km of origins in order to reduce the likelihood that nearby industry composition could be related to trends in tract productivity.

In particular, we calculate the year 2000 component \widetilde{RMA}_i^{2000} of our main instrument as:

$$\widetilde{RMA}_i^{2000} = \sum_j \frac{e^{-\widehat{\varepsilon\kappa}\tau_{ij}^{90}} \mathbf{1}(dis_{ij} > 2km) \sum_k L_{jk}^{90} [E_{r'(j)k}^{2000} / E_{r'(j)k}^{1990}]}{\widetilde{FMA}_j^{2000}} \quad (22)$$

$$\widetilde{FMA}_j^{2000} = \sum_i \frac{e^{-\widehat{\varepsilon\kappa}\tau_{ij}^{90}} \mathbf{1}(dis_{ij} > 2km) \pi_i^{90} \left[\frac{\sum_k L_{jk}^{90} [E_{r'(j)k}^{2000} / E_{r'(j)k}^{1990}]}{\sum_j L_j^{90}} \right]}{\widetilde{RMA}_i^{2000}} \quad (23)$$

In these expressions, τ_{ij}^{90} is the reported or forecast commute time from i to j in the 1990 CTPP. $\widehat{\varepsilon\kappa}$ is estimated separately for each region in year 2000 using gravity regressions of log reported commute flows on commute times and origin plus destination fixed effects using 2000 CTPP data.¹³ Distances from residential to work locations dis_{ij} are calculated using tract centroids. Employment in industry k in work location j , L_{jk}^{90} , is measured from the 1990 CTPP. $E_{r'(j)k}^{2000}$ and $E_{r'(j)k}^{1990}$ are the 2000 and 1990 nationwide employment in industry k excluding the region of tract j , respectively. That is, $\sum_k L_{jk}^{90} [E_{r'(j)k}^{2000} / E_{r'(j)k}^{1990}]$ captures the predicted amount of employment that would exist in tract j if 1990 employment by industry grows at national rates (excluding region r) to year 2000. $\frac{\sum_k L_{jk}^{90} [E_{r'(j)k}^{2000} / E_{r'(j)k}^{1990}]}{\sum_j L_j^{90}}$ is a constant that captures the population growth rate needed to match the aggregate simulated employment in the region in year 2000. The 2006 component of the instrument is calculated analogously, with $E_{r'(j)k}^{2000}$ in (22) and (23) replaced by $E_{r'(j)k}^{2006}$.

¹³Details of these estimates are reported in (Baum-Snow et al., 2019).

The log difference in \widetilde{RMA}_i for 2000-2006, $\Delta \ln \widetilde{RMA}_i$, is our main instrument for $\Delta \ln P_i$ as measured for both the 2000-2006 and 2000-2010 time periods. We build our instrument for the 2000-2006 period only as this is the time period for which first stage predictive power is strongest. The 2006-2010 period experienced employment declines that are not well predicted by Bartik type instruments.

4.3 Instrument Validity

The fundamental sources of identifying variation used are tract level “Bartik” (1991) type shocks in each employment location, written out as follows.

$$Bartik_{jr} = \sum_k \frac{Emp_{jk}^{90}}{\sum_k Emp_{jk}^{90}} [\ln E_{r'(j)k}^{06} - \ln E_{r'(j)k}^{00}] \quad (24)$$

A prerequisite for the spatial aggregation of such shocks into $\Delta \ln \widetilde{RMA}_i$ to successfully predict $\Delta \ln RMA_i$ is for the tract level counterparts to successfully predict tract level employment growth.

Table 3 presents evidence to this effect. It presents regressions of 2000-2006 or 2000-2010 employment growth in tract jr on $Bartik_{jr}$ and controls for 1990 employment level, past demographic composition of tract residents, CBD distance and metro region fixed effects. This tract level variation gets aggregated into the identifying variation for the first stage of our main analysis. All tracts in primary sample regions are included, as they all contribute to measures of \widetilde{RMA}_i in tracts that contribute data to our main estimation exercises. We control for past employment to isolate employment growth due only to variation in industry composition. Lagged demographics and CBD distance controls account for potentially differing labor supply conditions.

Results indicate that we can plausibly isolate labor demand shocks at the tract level. A one percent higher tract Bartik shock predicts 0.4% greater 2000-2006 tract employment growth and 0.7% greater 2000-2010 tract employment growth. Inclusion of 2-2.5 km CBD distance ring fixed effects interacted with metro region (columns 2 and 4) does not affect these conclusions.¹⁴

One challenge we face when estimating the housing supply equation is that home price growth is negatively serially correlated across decades. The first two columns of Table 4 show this pattern. In addition, the positive serial correlation in quantity growth in columns (3) and (4) may reflect

¹⁴Results are also robust to lagging the demographic tract controls by one additional decade.

serially correlated demand or supply shocks. Commensurate with our discussion of the OLS results above, these patterns suggest that there could be local unobserved history that drives both relative price declines and more construction, inducing a downward bias in an OLS estimation of housing supply. One legitimate potential concern is that our instrument $\Delta \ln \widetilde{RMA}_i$ may be correlated with such unobserved history. The results presented in the bottom panel of Table 4 help allay such concerns by showing relationships between the instrument and pre-treatment trends in key endogenous variables. Columns (5) and (7) show that $\Delta \ln \widetilde{RMA}_i$ does not predict 1990-2000 housing price nor quantity growth for our main specification with region fixed effects.¹⁵ With inclusion of region-ring fixed effects, $\Delta \ln \widetilde{RMA}_i$ predicts *negative* 1990-2000 price growth (column 6) and again no quantity growth (column 8). This evidence for pre-trends shows it is unlikely that our instrument is correlated with unobserved local history in a way that will bias our supply elasticity estimates.

Table 5 presents our main first stage estimates. The top two blocks show strong positive relationships between our primary measures of $\Delta \ln P$ and $\Delta \ln \widetilde{RMA}$ that are robust to inclusion of region-ring fixed effects. Slightly smaller first stage coefficients for 2000-2010 relative to 2000-2006 reflects the fact that the 2007-2010 period mostly saw housing market declines. Results at the bottom left show that we only have strong first-stage power predicting the census hedonic index when conditioning on region-ring fixed effects. For this reason, we only use the census index to account for pre-2000 price trends. Results on the bottom right show relationships between $\Delta \ln \widetilde{RMA}$ and $\Delta \ln RMA$ for 2000-2010. It shows significant estimated elasticities of 0.7-0.8 that are robust to the specification used. According to the model, this is the mechanism through which $\Delta \ln \widetilde{RMA}$ predicts $\Delta \ln P$. Tracts that appear in multiple metro regions are weighted equally to tracts that appear in just one. Standard errors are adjusted for spatial autocorrelation out to 25 km using a triangular kernel.¹⁶

Overall, we find that changes in simulated resident market access ($\Delta \ln \widetilde{RMA}$) strongly predict nearby employment growth, home prices, and therefore housing demand growth. Moreover, they are not correlated with pre-trends in home prices conditional on appropriate controls. These find-

¹⁵Goldsmith-Pinkham, Sorkin, and Swift (2018) suggest this sort of pre-trend test for evaluating the validity of Bartik instruments. Their other suggested validity tests use base year industry shares, which are not easily defined in our setting as they are nonlinearly aggregated across all potential commuting destinations into $\Delta \ln \widetilde{RMA}_i$.

¹⁶As should be expected, we also see positive reduced form relationships between our instrument and 2000-2006 and 2000-2010 Zillow new construction. We spatially correct standard errors to 25 km because errors from our main IV regressions in Table 6 are significantly correlated to this distance.

ings provide reassuring support for using $\Delta \ln \widetilde{RMA}$ as a valid instrument in estimating housing supply elasticities.

5 Main Results

5.1 Unified Supply Elasticity Estimates

Table 6 presents unified regressions of housing quantity growth on house price growth using the same specification as the OLS regressions reported in Table 2. In all regressions, $\Delta \ln \widetilde{RMA}$ enters as an instrument for $\Delta \ln P$. We show results using just region fixed effects in Panel A and with region-ring fixed effects in Panel B. We explore seven measures of quantity changes: 2000-2006 or 2000-2010 Zillow new construction units; 2000-2010 ACS new construction units; 2000-2010 change in census housing units; 2001-2011 change in developed land; 2001-2011 change in floorspace; 2000-2010 ACS new construction units through redevelopment only.

Before discussing the estimates, we highlight a few observations. First, unlike small and insignificant estimates from the OLS regressions, the estimated coefficients in Table 6 are positive across all specifications with magnitudes that are similar to or slightly smaller than other supply elasticity estimates from the literature. This is consistent with our narrative that OLS relationships between quantities and prices in part reflect movement along demand rather than supply curves. Second, estimated housing units responses to each price change measure are not significantly different. Third, controlling for CBD distance ring fixed effects interacted with metro area (Panel B) yields similar estimates compared to using metro area fixed effects only. Note that these local average treatment effects are not necessarily indicative of average supply elasticities across urban tracts nationwide. Indeed, we show below that average supply elasticities are somewhat larger.

Results in columns (1)-(4) show units supply elasticity estimates of 0.2 to 0.4, depending on the specification and quantity measure used. Conditional on region-ring fixed effects and based on the Zillow data, the new construction response during 2000-2010 as reported in column (2) is larger than that during 2000-2006 as reported in column (1), likely reflecting construction lags after 2000-2006 price shocks. Focusing on 2000-2010, the estimated elasticity of changes in housing units from the census (column 4) is not significantly different from the estimated elasticities of new construction units reported in columns (2)-(3), even though the former incorporates the negative

impacts of teardowns and full depreciation. This indicates that positive demand shocks have little impact on these two negative margins of response.

In column (5), we report elasticities of observed developed land with respect to home price. These are significantly smaller than unit supply elasticities at 0.1-0.2. Recall that $\gamma_i^{unit} = \gamma_i^{H/L} + \gamma_i^L$, where $\gamma_i^{H/L}$ is the elasticity of average units per ever developed parcel with respect to price. The difference between the unit and developed land supply elasticities shows that average units per ever developed parcel responds positively to price increases. As seen through the lens of the model, part of this likely occurs through more intensive land redevelopment in tracts with larger price increases.

As housing units differ in size, the unit supply elasticity alone does not represent the total housing supply response. To account for differences in size, column (6) reports the floorspace supply elasticity, which ranges between 2.3 and 4 depending on the price growth measures and the fixed effects specification used. We also estimate the supply elasticity using a broader quality-adjusted measure of housing services, resulting in estimates of a similar magnitude.

The difference between floorspace and unit supply elasticities yields the intensive margin of housing supply, as $\gamma_i^{S/H} = \gamma_i^{space} - \gamma_i^{unit}$. Comparing the estimated γ s in column (6) with those in column (2), we obtain an intensive margin housing supply elasticity that ranges between 1.97 and 3.62, consistent with estimated land shares in the housing production literature.¹⁷ Out of total supply response, the response from floorspace per unit accounts for over 80%, which takes the form of renovation and upsizing of existing home units as well as increased unit size in new constructions. The large intensive margin supply elasticity is consistent with the observations that almost half of American homeowners have renovated their homes (Plaut & Plaut, 2010) and that total renovation expenses reached \$326 billion in 2007 (Choi, Hong, & Scheinkman, 2014).

Column (7) reports estimated elasticities for the number of new units constructed 2000-2010 on land that was already developed in 2001. Recall that $\gamma_i^{unit} = \gamma_i^{unit,new} + \gamma_i^{unit,redevelop}$. Comparing the estimated γ s in column (7) with those in columns (2) and (3), we find that about one-third of newly built homes are through redevelopment of already-developed land, consistent with our

¹⁷Other estimates of the land share range from 0.10 for Centre County, PA (Yoshida, 2016) to 0.14 for Allegheny County, PA (Epple, Gordon, & Sieg, 2010) to 0.35 for France (Combes et al., 2019) to 1/3 for the average US housing market (ranging from 0.11 to 0.48 in low to high-value areas (Albouy et al., 2018)). Ahlfeldt and McMillen (2014) provide empirical support for the Cobb-Douglas functional form as a reasonable approximation to the housing production function.

model that developers carry out both new development and redevelopment in response to an increase in demand. The redevelopment supply elasticity further supports the urban literature that neighborhood renewal is largely driven by the deterioration and subsequent redevelopment of the existing housing stock (Rosenthal, 2018; Brueckner & Rosenthal, 2009). In addition, the smaller redevelopment relative to new development elasticity supports the conjecture that redeveloping existing land is typically costlier than developing new land.

5.2 Tract Level Heterogeneity

The local average treatment effect tract-level supply estimates in Table 6 mask substantial variation across neighborhoods. This sub-section shows how tract-level housing supply elasticities vary as a function of distance to CBD, land availability, topographical features, and land use regulations.

5.2.1 Unit Supply

Table 7 repeats the IV regressions in column (2) of Panel A of Table 6 with the addition of a set of interactions between price growth and tract-level factors that may influence supply. Price growth is constructed using the repeat sales index for columns (1)-(4) and the hedonic index for columns (5)-(8). The housing supply response is measured by 2000-2010 Zillow new construction. All specifications include interactions of CBD distance and an indicator for being over two-thirds of the way from the CBD to the region edge with price growth. CBD distance is measured as the fraction of the way from the CBD to the furthest census tract in the region from the CBD, running from 0 to 0.66 only. We use this functional form as our identifying variation is weakest in CBD distance band 0.66-1, where census tracts are typically very large and there is not a lot of 1990 variation in employment composition for identification. Overall, results in Table 7 show that there is substantial within-region variation in local housing supply elasticities. In addition, results across the two measures of price growth are very consistent.

Results in column (1) summarize the average relationship between housing supply elasticity and CBD distance. New construction supply elasticity increases with the distance to the city edge at a marginally decreasing rate. At the CBD, the average supply elasticity is near 0. This number increases to 0.5 at halfway to the city edge and beyond. These microgeographic estimates provide a supply-side explanation for the recent finding of more price growth in the center of metropolitan

areas in the latest boom (Glaeser et al., 2012; Yoshida, 2016). They are also consistent with the observation that a given common increase in demand throughout an urban area leads to relatively smaller price responses and relatively greater quantity responses the further away from the center one gets (Genesove & Han, 2013).

This positive CBD distance profile can be largely explained by the neighborhood level factors that affect the fixed development cost. The model predicts that the extensive margin supply is more responsive in tracts where fixed cost distributions have a fatter right tail. The latter is represented by easier development conditions, which are associated with sparser initial development, flatter land and looser regulation. Panel B of Figure 1 shows that the average tract in our data is almost 60% developed at the CBD but less than 10% developed at the region edge. In addition, the fraction of tract land that is flat declines from 45% to 38% from CBDs to region edges. Furthermore, land use are more regulated at 30% of the way than at the CBD. Thus, the increasing supply elasticity along the distance to the CBD, as observed in column (1), could be a result of the combination of these various factors.

To see how much these tract-level factors matter for the CBD distance effect, column (2) of Table 7 expands the specification in column (1) by adding interactions between $\Delta \ln P$ and a quadratic in the 2001 fraction of land developed in each tract and the 2001 fraction of flat land in each tract.¹⁸ Consistent with the model, we find declining supply elasticity with developed fraction, at a decreasing rate, and a positive impact of flat land on unit supply. Moreover, CBD distance coefficients not only attenuate but also turn negative. The attenuation reflects a significant negative correlation between the fraction of developed land and CBD distance as shown in Figure 1.

Conditional on topography and developed fraction, the negative CBD distance effect in column (2) likely reflects the impact of local regulations, as regulations increase on average with CBD distance within metro regions over the range which is well covered by our data, as seen in Figure 1 Panel B. To explicitly examine the role of regulation, column (3) expands the specification in column (2) by including an interaction with the municipality level Wharton Regulation Index. Its estimated impact on housing supply elasticity is negative, indicating that housing supply is less responsive in more regulated areas. Moreover, including this variable turns the coefficient on

¹⁸As the CBD distance squared interaction is no longer significant, we drop this variable. Neither fraction flat squared nor the interaction of fraction flat and fraction developed price interactions are significant, so we exclude these variables as well.

CBD distance from strongly negative to slightly positive, supporting the conjecture above that the negative CBD distance effect in column (2) is partly due to regulation. However, we lose about half the sample from incomplete Wharton Index coverage. In Column (4), we further restrict the sample to the 1,528 tracts in 8 cities (represented in 12 overlapping regions) for which we have FAR information for residentially zoned parcels. Inclusion of the FAR restriction yields an expected positive interaction coefficient of 0.12. Increasing FAR by 2 (one standard deviation) thus increases the unit supply elasticity by 0.24, again consistent with the model.¹⁹

In columns (5)-(8), we repeat the specifications in columns (1)-(4) but use the Zillow hedonic price index instead. Estimates are consistent across different measures of price growth. Across columns (1)-(8), additional inclusion of tract level elevation range and other topographical features has no significant estimated impact on supply elasticity (unreported). Given the much smaller size and likely endogeneity concerns associated with including the regulation controls, we choose columns (2) and (6) as the primary specifications for predicting tract level unit supply elasticities.

In unreported results, we further investigate whether the estimates from the main specifications in column (2) and (6) are robust to the inclusion of metro-level factors. In particular, we consider the fraction of developed land and the fraction of area that is lost to hills, water and wetland, both measured within the 50% of the maximum radius from the city center or out to 50 km. The effects of the tract-level factors remain consistent and robust with no additional statistically significant effect of the metro-level factors. However, an aggregation exercise below shows how our tract level estimates aggregate up such that these metro level factors do predict metro level housing supply elasticities.

In columns (9) and (10) we focus on redevelopment. We repeat the main specification in columns (2) and (6) but with redeveloped units as the dependent variable. These estimates are simply attenuated versions of those in columns (2) and (6), with this attenuation differing with CBD distance. Conditional on developed fraction and topography, the redevelopment elasticity is about one-third of the full unit supply elasticity on average, with this fraction decreasing with CBD distance. This is consistent with findings in the literature that prime teardowns are near public transportation and traditional village centers in Chicago (Dye & McMillen, 2007) and closer to the CBD and the coast in Miami (Munneke & Womack, 2015), possibly due to the fact that land

¹⁹We have not found plausible instruments for regulatory constraints, though controls for 1990 and 2000 tract level demographic characteristics may account for key determinants of the regulatory environment (Murphy, 2018).

assembly for development is easier with undeveloped parcels in more suburban areas. In addition, the developed fraction and topography affect the redevelopment supply elasticity in the same way as they affect total unit supply, consistent with the model's prediction. Together, the significant within-city heterogeneity, both in new construction and in redevelopment, emphasizes the importance of examining housing supply from a micro-geographic perspective.

5.2.2 Land Development and Floorspace

Table 8 focuses on land development and floorspace supply. Land development results in Columns (1)-(8) very much mirror those for unit supply in columns (1)-(8) of Table 7, though with attenuated estimates. This is consistent with the supply model in Section 3.1, as land development elasticities do not incorporate the densification and redevelopment that are part of the unit supply elasticities in Table 7. As expected, developed fraction, the prevalence of flat land and FAR restrictions all matter for land development elasticities.

Columns (9) and (10) examine determinants of floorspace supply elasticities. Interestingly, CBD distance is not significant for floorspace. As laid out in the supply model, given the relatively small unit and land development supply elasticities, intensive margin responses conditional on parcel development quantitatively dominate extensive margin responses. Because they do not matter, we exclude CBD distance effects in our main specification for floorspace supply, only letting this supply response depend on initial land development and topographical features. As with the unit supply response, floorspace supply is more responsive when there is less developed land. Consistent with Table 6, Table 8 shows that floorspace supply is 7-10 times more price elastic than unit supply, though with large standard errors. Estimates are remarkably similar for a broader housing quantity index that is constructed using additional attributes as well (not reported).

5.3 Distributions of Tract Supply Elasticities

We use the specifications in columns (2), (5), (9) and (10) in Table 7 to predict tract-level total unit and unit redevelopment supply elasticities respectively for each tract. We similarly use specifications in columns (2), (5), (9) and (10) in Table 8 to predict land development and floorspace supply elasticities. While the estimation sample for Tables 7 and 8 is limited, we use coefficient estimates to predict supply elasticities for all census tracts in the 306 metro regions in our data.

Table 9 provides summary statistics of these imputed elasticities plus the decomposition of the components of floorspace elasticity as shown in the table. Imputed elasticities are based on repeat sales price growth in Panel A and hedonic price growth in Panel B. Results are fully consistent across these two measures of price growth.

Across all tracts, the mean unit supply elasticity is 0.5, of which 20-40% is from redevelopment. The average floorspace elasticity is 3, roughly consistent with the recent finding that the estimated land share in housing production is about 0.35 (Combes et al., 2019). These objects have standard deviations of 0.6 and 1.1 respectively, with the most supply elastic tracts having floorspace elasticities of about 4. On average, 87% of floorspace supply responses come from the intensive margin (floorspace per unit), with only 5% from new land development and 8% from units per parcel. However, the interquartile range of intensive margin share is 21 percent. Figure A2 shows kernel density graphs of each of our four supply elasticity measures and their lower and upper 95% confidence bands. Here we see that while there are some negative estimates for all but floorspace elasticity, they are statistically significant in less than 2% of cases. More dispersion in floorspace supply elasticities exists within than between regions. In particular, the standard deviation of the floorspace supply elasticity is 1.1 across tracts, which is also approximately the tract-weighted mean of within metro region standard deviations. In comparison, it is only 0.5 across regions after aggregation across tracts within region, as shown below in Table 10.

Our earlier findings from Tables 7 and 8 indicate the robust pattern that local supply elasticities typically increase with CBD distance. Panel B of Figure 1 shows that on average, the fraction of tract developed land ranges from 60% at the CBD to less than 10% at 90 percent of the way to the city edge. This pattern is consistent with the prediction from a monocentric model of urban land use that the density of construction declines as one moves away from the CBD (Duranton & Puga, 2015). Areas further away from CBDs effectively have more land with which to deliver housing, facilitated mostly by initial development density. Average topographical conditions are quite constant with CBD distance, with fraction flat land declining only slightly. It is therefore tempting to conjecture that the CBD distance pattern in supply elasticities is mostly explained by the increasing availability of developable land as one moves away from the center.

Figure 2 shows how much this is the case. Each panel shows how a different predicted supply elasticity (unit new construction, land development, unit redevelopment and floorspace) changes

with CBD distance for the average tract. In each panel, the solid line represents the imputed supply elasticity, which is almost indistinguishable from the long dash-dot line that represents the imputed supply elasticity holding the fraction of flat land constant with CBD distance. Each supply elasticity except for floorspace increases with CBD distance, then declines some in the suburbs, before increasing again toward the urban fringe. The floorspace elasticity increases monotonically on average from CBDs to metro edges.

The remaining two lines in each panel represent the imputed supply elasticities when holding the initial developed fraction at each metro region's mean and when holding both developed fraction and topography at the metro mean, respectively. The two lines coincide with each other, indicating that, on average, topography alone does not play a big role in explaining the CBD distance pattern in local supply elasticities. Under both counterfactuals, we see supply elasticity falling over most of the range of CBD distance for unit, land and redeveloped unit supply elasticities. Mechanically, this is because of the negative CBD distance coefficients in Tables 7 and 8, which appear to capture regulation. At CBD distance 0.66 and beyond, the specification in Tables 7 and 8 impose only a level effect, reversing this downward trend as developed fraction continues to fall.

We plot similar figures for unit and floorspace supply elasticities in select metropolitan areas (Los Angeles, New York, Madison and Pittsburgh) in Figures 3 and 4. Regardless of whether a metropolitan area is considered elastic as a whole, its within-city variation in supply elasticities is always amplified by spatial variation in the fraction developed land and topography. While there is a tendency for supply elasticity to increase with CBD distance, this effect gets mitigated in Pittsburgh, Los Angeles and New York by the hilly topography and variable development density at the edges of these metro regions. In contrast, in Madison there is true monotonicity in supply elasticity with respect to CBD distance. Holding only topography constant does not affect supply elasticities much, though it does smooth out some local variation, which is especially evident in the top-left graph for Los Angeles. Fixing development fraction instead leads to declining supply elasticity with CBD distance, as does fixing both of these determining factors.

6 Aggregation

Much of the existing evidence on housing supply elasticities uses metro areas as the unit of analysis. In order to connect our estimates to these metro level estimates, in this section we aggregate estimated tract-level housing supply elasticities to the metro area level. Aggregation brings up a number of conceptual and practical challenges. The typical approach, for example in Saiz (2010), has been to use metro level labor demand and/or population supply shocks to deliver housing demand shocks. However, metro level studies may find supply elasticities that weight certain types of neighborhoods above their share of metro populations. As neighborhoods are linked in the residential demand system, metro level demand shocks of the same size but aggregated from different combinations of changes in neighborhood fundamentals can imply different metro level housing supply elasticities. Because of this variance to setting, here we provide a few examples of the macro supply elasticities implied from some simple broad-based neighborhood-specific shocks. Context matters and neighborhood level supply elasticities must be aggregated as appropriate to the application at hand.

To get a handle on aggregation, we first note that the tract level supply elasticity for supply measure Q , γ_{ir}^Q , generically aggregates to a metro region level elasticity, γ_r^Q , as follows. Aggregating tract level supply growth to the metro level means taking a sum weighted by initial neighborhood shares of the housing stock:

$$d \ln Q_r = \sum_i \frac{Q_{ir}}{Q_r} d \ln Q_{ir} = \sum_i \frac{Q_{ir}}{Q_r} \gamma_{ir}^Q d \ln P_{ir} = \gamma_r^Q d \ln P_r$$

Solving out, by definition the region level elasticity is given by

$$\gamma_r^Q \equiv \left[\sum_i \frac{Q_{ir}}{Q_r} \gamma_{ir}^Q d \ln P_{ir} \right] / \left[\sum_i \frac{Q_{ir}}{Q_r} d \ln P_{ir} \right]. \quad (25)$$

Here, we see that the metro level elasticity depends on the mix of neighborhoods experiencing price growth that has been spurred by demand shocks. As neighborhoods are linked in spatial equilibrium, it is difficult to imagine how price changes in multiple neighborhoods may occur in mutual isolation. The following sub-section provides two examples that impose different strong assumptions about the form of demand linkages across neighborhoods.²⁰

²⁰In our data, we can only accurately measure year 2000 tract level quantities as housing units. As such, for the

6.1 Aggregation of Common Neighborhood Demand Shocks

In this sub-section we first consider the case in which all neighborhoods simultaneously experience identical housing demand shocks. Because of differing housing supply elasticities, these shocks manifest themselves as different combinations of housing price and quantity changes, depending on the neighborhood. In the context of the model from Section 3.2, the tract level expenditure share on housing is $\frac{S_{ir}P_{ir}}{y_{ir}\pi_{ir}} = 1 - \beta$. Therefore, if the demand shock changes aggregate expected income by the same percentage in every neighborhood, $d \ln S_{ir} + d \ln P_{ir}$ is a constant, call it x . This would happen if the outside option $\ln \mu$ in (13) changes and leaves little scope for households to substitute across neighborhoods in response to this shock.

The resulting metro housing supply elasticity is a weighted average of tract-level elasticities, where the weight is the initial housing share adjusted for neighborhood supply elasticity. Using $d \ln P_{ir} = x / (1 + \gamma_{ir})$ and (25) to aggregate over tracts, we have

$$\frac{d \ln S_r}{d \ln P_r} = \gamma_r^1 = \frac{\sum_i H_{ir} \frac{\gamma_{ir}}{1 + \gamma_{ir}}}{\sum_i H_{ir} \frac{1}{1 + \gamma_{ir}}}.$$

This expression reflects the fact that tracts with more elastic supply will receive lower weight in aggregation because price growth is lower in these locations for a given demand shock. We apply this same expression to aggregate unit supply elasticities, recognizing that this requires tract per-unit price growth to match that for floorspace.

Second, we consider the alternative case in which aggregate housing demand shifts out in the city but agents get redistributed across neighborhoods in a way that maintains the same relative home prices across neighborhoods. This environment can be justified from the spatial equilibrium condition of a simpler model than that in Section 3.2, as in Roback (1982), in which neighborhoods differ in amenities but get hit with the same per-capita potential income shock, thereby driving the same price growth rate but different tract population changes depending on the local supply elasticity. In this setting, conditional on amenities and wages net of commuting costs, neighborhoods

purposes of aggregation $\frac{Q_{ir}}{Q_r}$ is calculated using units as $\frac{H_{ir}}{H_r}$ for all supply measures investigated.

are perfect substitutes. In this case, the metro-level supply elasticity is

$$\frac{d \ln S_r}{d \ln P_r} = \gamma_r^2 = \sum_i \left[\frac{H_{ir}}{H_r} \gamma_{ir} \right].$$

Tracts with higher initial housing stocks are typically associated with smaller housing supply response, but these tracts are weighted more in γ_r^2 .

Table 10 presents summary statistics for four versions of our two aggregate supply elasticity measures γ_r^1 and γ_r^2 . Across metro regions, the mean repeat sales based unit supply elasticity is 0.6 for γ_r^1 and 0.7 for γ_r^2 , with standard deviations of 1.2 and 0.3 respectively. Our second aggregation scheme reduces dispersion because it incorporates population flows between neighborhoods. Analogous floorspace elasticities are 3.5 and 3.8 respectively with a standard deviation across metros of 0.5. Compared to Saiz (2010) estimates, these unit elasticity estimates are smaller. Rank correlations between these metro level supply elasticity measures and those found by Saiz (2010) are between 0.35 and 0.42.

In Table 11, we examine relationships between metro-level conditions and our imputed metro-level housing supply elasticities. Following Saiz (2010), we construct three measures of metro-level factors that may influence supply elasticities: the fraction of land that is developed within some CBD radius, the fraction of area that is lost to hills, water and wetland within the same radius, and the metropolitan area level Wharton Regulation Index. We construct the first two variables within the 10 km, 20 km, 50 km, 10%, 50% and 100% of the maximum radius from the CBD and for the entire metropolitan area. Table 11 reports regressions of our metro supply elasticities on these three variables, using the 50% of the way to metro edge version, though all results discussed below are consistent across these different radii.²¹

We find that the metro-level developed fraction is predictive of metro supply elasticities, with expected negative signs in all cases. Conditional on initial development density, the fraction of a metro area that is lost to hills, waters and wetland also has a negative relationship with unit supply elasticities, but, if anything, positive relationships with floorspace supply elasticities. This indicates that in metros with the same initial development footprint but less land available due to difficult topographical features, average housing units are expanded more even as fewer new

²¹Since there is large variation in metro size, measures that use percentages rather than absolute distances generate more precise estimates.

units get built. Finally, there are significant negative relationships between regulation and metro housing supply elasticities, as expected. Note that we did not use regulation to predict the tract level supply elasticities before aggregation. Therefore, regulation is correlated with some combination of developed fraction and land unavailable for development.

6.2 Recovering the Neighborhood Demand System

In anticipation of our evaluation of the Opportunity Zone program, we first recover estimates of parameters that govern demand substitution patterns across neighborhoods in the model from Section 3.2. From (13) and (18), aggregate housing demand in each tract is given by

$$\ln S_i^d = \ln \tilde{\rho}_{HD} + \left(\frac{\psi}{\eta} - 1\right) \ln \sum_{i' \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} RMA_{i'}^{\frac{1}{\varepsilon}} \right)^\eta + \frac{1+\eta}{\varepsilon} \ln(RMA_i) - (\eta(1-\beta) + 1) \ln P_i + \eta \ln B_i.$$

Housing demand becomes more elastic as η and $\frac{\psi}{\eta}$ grow, as these objects reflect how fundamentally substitutable neighborhoods are for each other by residents. Therefore, understanding how much a given exogenous demand shock (through changes in B_i for example) affects prices versus quantities requires knowledge of these parameters.

Using the housing demand equation above, the floorspace demand elasticity is

$$\frac{d \ln S_i^d}{d \ln P_i} = (\beta - 1)[\eta + s_i(\psi - \eta)] - 1, \quad (26)$$

where s_i is the share of municipality $m(i)$'s population that is in tract i . The implied elasticity of substitution between two neighborhoods in the same municipality is $1 + \eta(1 - \beta)$ and that between neighborhoods in different municipalities $1 + (1 - \beta)[\eta + (s_i + s_{i'})((\psi - \eta))]$. We expect neighborhoods within the same municipality to be closer mutual substitutes than those in different municipalities, or that $\psi < \eta$. Because of Cobb-Douglas preferences, the analogous population and units elasticities of substitution are $\eta(1 - \beta)$ between two neighborhoods in the same municipality and $(1 - \beta)[\eta + (s_i + s_{i'})((\psi - \eta))]$ between neighborhoods in different municipalities.

We develop an “estimation” strategy for η using the structural equation (19) from the model. Using estimates of γ_{ir}^{space} discussed above, estimates of $(\kappa\varepsilon)_r$ from gravity regressions and calibrated values of 0.8 for β and 0.01 for κ as inputs,²² we estimate various specifications of equation

²² $\kappa = 0.01$ implies that 1 minute of commuting in one direction reduces full income by 1%. We also tried estimating

(19) with GMM. This includes a metro-region fixed effects specification and a municipality fixed effects specification where each metro area is divided into 5 municipalities: one for the central city and one each for suburbs in north, south, east and westerly directions. We impose that the error term in (19) is orthogonal to our main instrument $\Delta \ln \widetilde{RMA}_i$. In theory, we can adopt the same strategy to structurally estimate ψ . In practice, we do not have sufficient statistical power to precisely estimate ψ using a municipality level estimation equation. Instead, we infer the importance of municipalities by estimating (19) with and without municipality fixed effects. We use the same sample as in Tables 5-7 and the same tract weights, resulting in an estimation sample of 17,646 tracts.

As our fundamental source of identifying variation comes through labor demand shocks that lead to housing demand (rather than supply) shocks, we lean heavily on the model's structure to recover these neighborhood demand estimates. In particular, the model delivers how much housing prices must change in a neighborhood for a given exogenous RMA_i shock holding population constant. Then, observations about population changes are informative about η and ψ because these parameters govern elasticities of substitution in demand across neighborhoods. By model construction, they also govern own price demand elasticities. While not ideal, this procedure provides us a rough estimate to work with below.

The version with municipality fixed effects yields an estimate of η of 8.8 (se=0.9). The version with only metro fixed effects yields an estimate of η^R of 3.3 (se=0.1). The lower estimate without municipality fixed effects reflects less substitutability between neighborhoods in different municipalities than between neighborhoods in the same municipality. The resulting implied average elasticity of demand for floorspace in each neighborhood is $-0.2 \cdot 3.3 - 1 = -1.7$ while that for units is -0.7 . These estimates are similar to estimates in Hanushek and Quigley (1980).

6.3 Opportunity Zones

As an application, we use our tract-level supply elasticity estimates to provide a rough evaluation of the efficacy of the Opportunity Zone (OZ) provisions in the U.S. federal "Tax Cuts and Jobs Act" of 2017. The OZ program was created to incentivize investment in economically distressed communities. Among other incentives, the program provides preferential tax treatment of capi-

κ jointly with neighborhood demand parameters, but this yielded implied values of ϵ that were too low, leading us to our "estimation" strategy.

tal gains for new real estate investments within designated low-income census tracts. This may affect local housing markets in two ways. First, the reduction in the capital gains tax liability for investors reduces the financing costs of real estate development in these areas, which we model as a reduction in the marginal cost of building housing by Δs log points relative to other tracts. Second, the tax abatement encourages investors to divert to OZ neighborhoods from other similar neighborhoods. The investment incentive is further reinforced if OZ status spurs local governments to invest in tract amenities. We treat this as an outward shift in housing demand by Δd log points in terms of quantities relative to observationally equivalent tracts that do not attain OZ designation. Comparing OZ to similar looking non-OZ tracts, Chen, Glaeser, and Wessel (2019) find that the OZ program had little effect on home prices. However, the program may have precipitated both demand and supply shocks for these neighborhoods, resulting in greater quantities of housing. Whatever the true benefits of the OZ program, the primary goal of this treatment is to use the OZ program as an example to demonstrate that paying attention to variation in housing supply elasticities within metro regions can matter for policymaking.

Following the model from Section 3, we assume that the demand and supply for floorspace and housing units have constant elasticity forms. Generically, these equations are:

$$\ln Q_i^d = d_i + \varepsilon_D \ln P_i$$

$$\ln Q_i^s = s_i \varepsilon_S + \varepsilon_S \ln P_i$$

Shocks of Δd to demand and Δs to supply yield $\Delta \ln P_i = \frac{\Delta d_i - \varepsilon_{S_i} \Delta s_i}{\varepsilon_{S_i} - \varepsilon_{D_i}}$. To avoid relying very strongly on function form assumptions, this setup is most sensible for analyzing welfare impacts of small supply shocks on consumers and small demand shocks on producers. The dollar changes in CS and PS for small supply and demand shocks respectively are:

$$\Delta CS_i = -\Delta P_i H_i - \frac{1}{2} \Delta P_i \Delta H_i$$

$$\Delta PS = \Delta P_i H_i + \frac{1}{2} \Delta P_i \Delta H_i.$$

We measure base year prices using the 2016 repeat sales index and base year quantities of units and floorspace in each tract using Zillow data from 2016. All values are in 2010 dollars. We impose

tract unit and floorspace demand elasticities of -0.7 and -1.7 respectively, as calculated in Section 6.2. We predict out tract supply elasticities using the results in Tables 7 and 8 and tract developed fraction from 2011. To get a sense of the importance of local heterogeneity, we compare results using tract level supply elasticities to those using our second version of aggregate region level supply elasticities γ_r^2 . Using these assumptions, we calculate changes in CS and PS for all census tracts for which we have reliable housing stock and price index information in 2016, including 3,473 OZ designated tracts, given either $\Delta s = 0.05$ or $\Delta d = 0.05$. Assuming a capital gain of 25% on an average real estate property and a savings of 20% in capital gains tax from the OZ program, using $\Delta s = 0.05$ seems reasonable.²³

Table 12 presents the results, in which Panel A examines the market for housing units and Panel B examines the market for floorspace. Opportunity zones are primarily located in more urban and developed areas. As a result, supply elasticities in these areas are lower than in the average neighborhood. The average OZ tract has a unit supply elasticity of 0.4 and a floorspace supply elasticity of 2.6 relative to other tracts, with elasticities of 0.5 and 2.9 respectively. As a result, the average opportunity zone tract has a smaller scope for gains in consumer surplus through the OZ program relative to other tracts. Moreover, using metro region level supply elasticities overstates consumer surplus gains from the program and understates the dispersion in these gains.

Results in the first two columns of Panel A show that in the market for housing units our assumed supply shock would increase CS by an estimated \$5.5 million on average in OZ tracts and of \$9.2 million on average if it were imposed in other tracts. Imposing metro level supply elasticities would instead yield implies CS changes of \$7.1 million in OZ tracts, with a smaller standard deviation. Results in the third and fourth columns show that PS would also increase more in non-OZ tracts from demand shocks of 5%. The average PS increase of \$19.5 million would be understated by \$2 million using metro area level supply elasticities.

Results for the floorspace market in Panel B exhibit a similar pattern. In particular, we calculate that the average OZ tract would experience an increase in CS of \$11.9 million relative to \$17.0 million for a typical non-OZ tract. These numbers are slightly larger using region level supply elasticities. For PS, again our evidence indicates that developers in the average non-OZ tract would benefit more from a demand shock of this type than in OZ tracts. The fact that floorspace

²³Long-term capital gains are taxed at either 0, 15, or 20 percent depending on the taxpayer's income.

is relatively elastically supplied justifies the greater magnitudes of CS than PS changes in the floorspace analysis.

7 Conclusion

Since DiPasquale (1999)'s lament on the limited work on the supply side of housing, a number of studies have identified regulation and topographical conditions as determinants of supply elasticities. Saiz (2010) in his seminal work estimates housing supply elasticities at the metropolitan area level and characterizes them as a function of both physical and regulatory constraints. In this paper, we follow his insights and present the first set of estimates of tract-level supply elasticities. Knowledge of local housing supply elasticities at a microgeographic scale is not only central to understanding within-city house price dynamics (Glaeser et al., 2012; Guerrieri et al., 2013), but also important for evaluating place-based policy interventions (Busso et al., 2013; Hanson, 2009).

We find that housing supply becomes more elastic further out from urban centers and that there is more variation within than between metro areas in housing supply elasticities. This pattern is in part but not entirely due to a larger fraction of land available for development. Initial development density, availability of flat land and zoning regimes are all important determinants of local housing supply.

Acknowledgements: We thank Camilo Acosta, Rolando Campusano, Ramin Forouzandeh and Ian Herzog for excellent research assistance. We also thank Gilles Duranton, Chris Foote, Laurent Gobillon, Dan McMillen and Lindsay Relihan for helpful comments. We gratefully acknowledge financial support from the Social Sciences and Humanities Research Council of Canada. All errors are our own.

References

- Ahlfeldt, G. M., & McMillen, D. P. (2014). New estimates of the elasticity of substitution of land for capital. *Louvain-la-Neuve: European Regional Science Association (ERSA)*.
- Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., & Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6), 2127–2189.
- Albouy, D., Ehrlich, G., & Shin, M. (2018). Metropolitan land values. *Review of Economics and Statistics*, 100(3), 454–466.
- Bartik, T. J. (1991). Who benefits from state and local economic development policies?
- Baum-Snow, N., Hartley, D. A., & Lee, K. O. (2019). The long-run effects of neighborhood change on incumbent families.
- Been, V., Ellen, I., & Gedal, M. (2009). Teardowns and land values in new york city. In *Lincoln institute working paper*.
- Brooks, L., & Lutz, B. (2016). From today's city to tomorrow's city: An empirical investigation of urban land assembly. *American Economic Journal: Economic Policy*, 8(3), 69–105.
- Brueckner, J. K., & Rosenthal, S. S. (2009). Gentrification and neighborhood housing cycles: Will america's future downtowns be rich? *The Review of Economics and Statistics*, 91(4), 725–743.
- Brueckner, J. K., & Singh, R. (2020). Stringency of land-use regulation: Building heights in us cities. *Journal of Urban Economics*.
- Busso, M., Gregory, J., & Kline, P. (2013). Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review*, 103(2), 897–947.
- Calabrese, S. M., Epple, D. N., & Romano, R. E. (2011). Inefficiencies from metropolitan political and fiscal decentralization: Failures of tiebout competition. *The Review of Economic Studies*, 79(3), 1081–1111.

- Chen, J., Glaeser, E. L., & Wessel, D. (2019). *The (non-) effect of opportunity zones on housing prices* (Tech. Rep.). National Bureau of Economic Research.
- Choi, H.-S., Hong, H., & Scheinkman, J. (2014). Speculating on home improvements. *Journal of Financial Economics*, 111(3), 609–624.
- Combes, P.-P., Duranton, G., & Gobillon, L. (2019). The production function for housing: Evidence from France.
- Cosman, J., Davidoff, T., & Williams, J. (2018). Housing appreciation and marginal land supply in monocentric cities with topography.
- Couture, V., Gaubert, C., Handbury, J., & Hurst, E. (2019). *Income growth and the distributional effects of urban spatial sorting* (Tech. Rep.).
- Cutler, D., Glaeser, E., & Vigdor, J. (1999). The rise and decline of the American ghetto. *Journal of Political Economy*, 107(3), 455–506.
- Davidoff, T. (2016). Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors. *Critical Finance Review*, 5(2), 177–206.
- Davis, M., Gregory, J., & Hartley, D. (2019). The long-run effects of low-income housing on neighborhood composition.
- DiPasquale, D. (1999). Why don't we know more about housing supply? *The Journal of Real Estate Finance and Economics*, 18(1), 9–23.
- Duranton, G., & Puga, D. (2015). Urban land use. In *Handbook of regional and urban economics* (Vol. 5, pp. 467–560). Elsevier.
- Dye, R. F., & McMillen, D. P. (2007). Teardowns and land values in the Chicago metropolitan area. *Journal of urban economics*, 61(1), 45–63.
- Epple, D., Gordon, B., & Sieg, H. (2010). A new approach to estimating the production function for housing. *American Economic Review*, 100(3), 905–24.
- Favilukis, J., Mabilie, P., & Van Nieuwerburgh, S. (2019). Affordable housing and city welfare.
- Genesove, D., & Han, L. (2013). A spatial look at housing boom and bust cycles. In *Housing and the financial crisis* (pp. 105–141). University of Chicago Press.
- Glaeser, E. L., Gottlieb, J. D., & Tobio, K. (2012). Housing booms and city centers. *American Economic Review*, 102(3), 127–33.
- Glaeser, E. L., & Gyourko, J. (2005). Urban decline and durable housing. *Journal of political economy*,

113(2), 345–375.

- Glaeser, E. L., Gyourko, J., & Saks, R. E. (2005). Urban growth and housing supply. *Journal of economic geography*, 6(1), 71–89.
- Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2018). *Bartik instruments: What, when, why, and how* (Tech. Rep.). National Bureau of Economic Research.
- Guerrieri, V., Hartley, D., & Hurst, E. (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics*, 100, 45–60.
- Hanson, A. (2009). Local employment, poverty, and property value effects of geographically-targeted tax incentives: An instrumental variables approach. *Regional Science and Urban Economics*, 39(6), 721–731.
- Hanushek, E. A., & Quigley, J. M. (1980). What is the price elasticity of housing demand? *The Review of Economics and Statistics*, 449–454.
- Lutz, C., & Sand, B. (2019). Highly disaggregated land unavailability. Available at SSRN 3478900.
- McMillen, D., & O’Sullivan, A. (2013). Option value and the price of teardown properties. *Journal of Urban Economics*, 74, 71–82.
- Munneke, H. J., & Womack, K. S. (2015). Neighborhood renewal: the decision to renovate or tear down. *Regional Science and Urban Economics*, 54, 99–115.
- Murphy, A. (2018). A dynamic model of housing supply. *American Economic Journal: Economic Policy*, 10(4), 243–67.
- Orlando, A., & Redfearn, C. (2018). What’s lost in the aggregate: Lessons from a local index of housing supply restrictions. *Unpublished Manuscript*.
- Ouazad, A., & Ranciere, R. (2019). City equilibrium with borrowing constraints: Structural estimation and general equilibrium effects. *International Economic Review*, 60(2), 721–749.
- Plaut, P., & Plaut, S. (2010). Decisions to renovate and to move. *Journal of Real Estate Research*, 32(4), 461–484.
- Roback, J. (1982). Wages, rents, and the quality of life. *Journal of political Economy*, 90(6), 1257–1278.
- Rosenthal, S. S. (2018). Owned now rented later? housing stock transitions and market dynamics.
- Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3), 1253–1296.
- Severen, C. (2019). *Commuting, labor, and housing market effects of mass transportation: Welfare and*

identification (Tech. Rep.).

Tsivanidis, N. (2018). The aggregate and distributional effects of urban transit infrastructure: Evidence from bogotas transmilenio. *Unpublished manuscript*.

von Ehrlich, M., Schöni, O., & Büchler, S. (2018). On the responsiveness of housing development to rent and price changes: Evidence from switzerland. *Unpublished Manuscript*.

Yoshida, J. (2016). Structure depreciation and the production of real estate services. *Available at SSRN 2734555*.

Zillow. (2017). *Ztrax: Zillow transaction and assessor dataset, 2017-q4*. (Zillow Group, Inc. <http://www.zillow.com/ztrax/>)

Table 1: Summary Statistics

	Mean	St Dev	Obs	Tracts
Tract Home Price Changes, Estimation Sample				
Repeat Sales Index, 2000-2006	0.64	0.35	24,266	19,126
Hedonic Index, 2000-2006	0.62	0.34	24,192	19,001
Repeat Sales Index, 2000-2010	0.26	0.36	23,918	18,978
Hedonic Index, 2000-2010	0.26	0.34	24,447	19,235
Census Index, 2000-2010	0.54	0.27	25,267	19,909
Tract Housing Quantity Changes, Estimation Sample				
Stock of Housing Units, Census, 2000-2010	0.08	0.23	25,361	19,985
New Units, ACS, 2000-2009	0.12	0.21	25,353	19,977
New Units, Zillow, 2000-2006	0.09	0.20	25,361	19,985
New Units, Zillow, 2000-2009	0.11	0.22	25,361	19,985
Floorspace, 2001-2011	0.34	0.83	24,044	18,728
New Units on Developed Land, ACS, 2000-2009	0.05	0.12	25,361	19,985
Developed Land, 2001-2011	0.08	0.13	25,361	19,985
Tract Employment and Population Variables				
Tract Employment, 2000-2010, Regions in Est. Sample	-0.15	0.89	55,981	42,693
Tract Level Bartik Instrument, 2000-2006, Regions in Est. Sample	0.09	0.05	56,238	42,866
RMA, 2000-2010, Estimation Sample	0.04	0.05	25,361	19,985
Simulated RMA, 2000-2006, Estimation Sample	0.05	0.01	25,361	19,985
Tract Level Supply Influencers, Estimation Sample				
Fraction of Land Area Developed, 2001	0.33	0.21	25,361	19,985
Fraction of Land Area Flat	0.40	0.43	25,361	19,985
Wharton Real Estate Index (municipality level variation)	0.35	1.00	10,309	8,180
Residential Floor Area Ratio (8 cities)	1.74	1.31	2,001	1,382
Fraction of Way from CBD to Metro Edge	0.27	0.20	25,361	19,985
Tract Level Supply Influencers, Non-Estimation Sample				
Fraction of Land Area Developed, 2001	0.37	0.24	38,535	30,424
Fraction of Land Area Flat	0.45	0.43	38,535	30,424
Wharton Real Estate Index (municipality level variation)	0.02	0.96	20,214	14,699
Residential Floor Area Ratio (8 cities)	2.26	1.88	6,706	3,385
Fraction of Way from CBD to Metro Edge	0.27	0.23	38,536	30,425

All changes are in percentage terms. The full sample includes 50,410 unique census tracts in 306 partially overlapping metro regions with 1990 information on employment. The estimation sample includes 19,985 equally weighted unique tracts in 170 regions with at least 10 housing market transactions in 2000 and 2010 in the ZTRAX data. It excludes tracts for which the 2010 housing unit counts from Zillow is more than 20% different from the analogous 2010 census count or both the 2000 census and Zillow counts are more than 20% apart and the tract is outside a state for which Zillow collects complete assessment data.

Table 2: OLS Housing Supply Elasticity Estimates

Quantity Measure Source Time Period	$\Delta \ln S(\text{supply})$						
	(1) New Units Zillow 2000-2006	(2) New Units Zillow 2000-2010	(3) New Units ACS 2000-2010	(4) Units Census 2000-2010	(5) Develop. Land USGS 2001-2011	(6) Floorspace Zillow 2001-2011	(7) Redev Units Zillow & USGS 2000-2010
$\Delta \ln P$ (Repeat Sales)	0.06*** (0.01)	0.06*** (0.01)	0.06*** (0.01)	0.09*** (0.01)	0.02*** (0.00)	0.17*** (0.02)	0.03*** (0.00)
Obs	24,265	23,917	23,909	23,917	23,917	22,654	23,917
$\Delta \ln P$ (Hedonic)	0.07*** (0.01)	0.08*** (0.01)	0.07*** (0.01)	0.12*** (0.01)	0.03*** (0.00)	0.21*** (0.02)	0.03*** (0.00)
Obs	24,189	24,442	24,434	24,442	24,442	23,447	24,442

The price growth measures are constructed from the Zillow. Each cell represents a separate regression. Regressions include metro region fixed effects, a quadratic in fraction of the way from the CBD to the metro edge, an indicator for being beyond 66% of the way to the edge, fraction tract developed in 2001, fraction of tract land flat, log 1990 tract employment, the 2000-2006 Bartik shock for the tract and the following tract attributes measured in 1990 and 2000: census home price index, rent index, log population, log avg hh income, share black, share white and share college. The estimation sample for the repeat sales index uses data from 161 metro regions while that for the hedonic index uses data from 164 regions. These samples are reduced by 1 region for the repeat sales index and 3 regions for the hedonic index in Column 1 and by 4 regions in column 6 for both price measures due to missing data on the outcome. The hedonic index sample excludes tracts in the repeat sales sample containing homes for which age and/or floorspace are not observed in 2000 or 2010. The repeat sales sample excludes tracts in the hedonic sample containing homes that only transacted once. Robust standard errors are in parentheses. While some are duplicated across regions, each census tract receives equal weight.

Table 3: Regressions of Emp. Growth on Bartik Shocks

	(1)	(2)	(3)	(4)
	Change in log Employment			
	2000-2006		2000-2010	
Tract Bartik Shock, 2000-2006	0.38***	0.27**	0.73***	0.68***
	(0.12)	(0.12)	(0.12)	(0.12)
No 1990 Emp. Info.	0.03	0.01	0.07***	0.06**
	(0.03)	(0.03)	(0.03)	(0.03)
Observations	53,732	53,444	55,981	55,681
R-squared	0.05	0.05	0.05	0.05
Number of Regions	162		170	
Number of Region_Rings		2,453		2,573

Regressions also include log 1990 tract employment and the following tract attributes from 1990 and 2000: census home price index, rent index, log pop, log avg hh income, fraction black, fraction white, fraction college. Sample includes all tracts in metro regions that are in the primary sample. Each tract receives equal weight. Robust standard errors in parentheses.

Table 4: Tract-Level Housing Market Dynamics

	(1)	(2)	(3)	(4)
	$\Delta \ln$ House Price 2000-2010		$\Delta \ln$ House Quant. 2000-2010	
$\Delta \ln$ House Price, 1990-2000	-0.24*** (0.01)	-0.25*** (0.01)	0.02 (0.01)	0.01 (0.02)
$\Delta \ln$ House Quantity, 1990-2000	0.01 (0.03)	0.01 (0.03)	0.22*** (0.04)	0.17*** (0.04)
Region All FE	Yes	No	Yes	No
Region-Ring FE	No	Yes	No	Yes
	(5)	(6)	(7)	(8)
	$\Delta \ln$ House Price 1990-2000		$\Delta \ln$ House Quant. 1990-2000	
$\Delta \ln$ Sim. RMA, 2000-2006	-0.58 (1.17)	-2.00** (0.95)	0.22 (0.34)	0.03 (0.35)
Region All FE	Yes	No	Yes	No
Region-Ring FE	No	Yes	No	Yes

Price changes and housing quantity changes are constructed based on the Census data, as detailed in Section 3. Each entry is from a separate regression of the variable at top on the variable at left with the indicated fixed effects. Lagged demographic controls are the same as in Table 2, excluding the 1990 and 2000 census house price and rent indexes. Standard errors are corrected for spatial autocorrelation up to 25km.

Table 5: First Stage Results

	$\Delta \ln$ House Price (Zillow Repeat Sales)			
	2000-2006		2000-2010	
$\Delta \ln$ Simulated RMA, 2000-2006	6.76*** (1.81)	7.83*** (1.57)	6.11*** (1.44)	4.25*** (1.23)
Observations	24,266	24,266	23,918	23,918
R-squared	0.14	0.11	0.21	0.17
Number of Fixed Effects	160	1,793	161	1,800
	$\Delta \ln$ House Price (Zillow Hedonic)			
	2000-2006		2000-2010	
$\Delta \ln$ Simulated RMA, 2000-2006	6.66*** (1.84)	7.46*** (1.66)	6.06*** (1.42)	3.86*** (1.21)
Observations	24,192	24,192	24,447	24,447
R-squared	0.16	0.13	0.20	0.16
Number of Fixed Effects	161	1,762	164	1,784
	$\Delta \ln$ House Price (Census Hedonic)		RMA	
	2000-2010		2000-2010	
$\Delta \ln$ Simulated RMA, 2000-2006	0.17 (1.25)	2.35** (1.15)	0.76*** (0.18)	0.71*** (0.17)
Observations	25,267	25,267	25,361	25,361
R-squared	0.08	0.07	0.04	0.02
Number of Fixed Effects	166	1,858	166	1,860
Region All FE	Yes	No	Yes	No
Region-Ring FE	No	Yes	No	Yes

Regressions include the same controls as in Table 2. Tracts are equally weighted, even if they appear in multiple metro regions. Standard errors are corrected for spatial autocorrelation up to 25 km using a Bartlett kernel.

Table 6: Unified IV Results for Housing Supply

Quantity Measure Source Time Period	$\Delta \ln S(\text{supply})$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	New Units Zillow 2000-2006	New Units Zillow 2000-2010	New Units ACS 2000-2010	Units Census 2000-2010	Develop. Land USGS 2001-2011	Floorspace Zillow 2001-2011	Redev Units Zillow & USGS 2000-2010
Panel A: Controls for Region Fixed Effects							
$\Delta \ln P$, Repeat Sales	0.26** (0.11)	0.31** (0.12)	0.24** (0.11)	0.35*** (0.13)	0.13 (0.08)	2.28** (0.91)	0.11* (0.06)
$\Delta \ln P$, Hedonic	0.25** (0.11)	0.30** (0.12)	0.22** (0.11)	0.33** (0.13)	0.13 (0.08)	2.59*** (0.92)	0.09 (0.06)
Panel B: Controls for Region-Ring Fixed Effects							
$\Delta \ln P$, Repeat Sales	0.18** (0.09)	0.40** (0.20)	0.28* (0.16)	0.30* (0.17)	0.19 (0.12)	3.73*** (1.41)	0.11 (0.08)
$\Delta \ln P$, Hedonic	0.19** (0.09)	0.42* (0.22)	0.26 (0.17)	0.26 (0.18)	0.17 (0.13)	4.04*** (1.47)	0.10 (0.09)

The price growth measures are constructed from the Zillow. Each cell represents a separate regression. Specifications are the same as in Table 2, except that the 2000-2006 change in \ln simulated RMA instruments for changes in prices. See the notes to Table 2 for details about specification and sample sizes. Standard errors are corrected for spatial autocorrelation to 25 km. First stage F-statistics can be determined from results in Table 5.

Table 7: Heterogeneity in Unit Supply Elasticities by CBD Distance and Tract Conditions

	$\Delta \ln S(\text{supply})$									
	New Units Construction								Units Redevelopment	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\Delta \ln P$	-0.01 (0.14)	2.00*** (0.72)	1.24*** (0.26)	-0.90 (0.67)	-0.00 (0.13)	1.91*** (0.71)	1.26*** (0.29)	-1.28 (0.89)	0.84** (0.33)	0.85** (0.34)
$\Delta \ln P \times (\text{CBD Dis})$	1.60*** (0.55)	-1.19** (0.61)	0.23 (0.57)	-1.08 (1.66)	1.25*** (0.48)	-1.46** (0.69)	-0.02 (0.65)	-0.94 (1.62)	-0.66** (0.28)	-0.83** (0.33)
$\Delta \ln P \times (\text{CBD Dis})^2$	-1.28** (0.54)				-1.14** (0.50)					
$\Delta \ln P \times \text{Edge}$	0.51*** (0.18)	-0.96** (0.44)	-0.04 (0.25)	1.39 (1.31)	0.34** (0.14)	-1.05** (0.49)	-0.15 (0.23)	-1.43 (2.67)	-0.53** (0.21)	-0.58** (0.24)
$\Delta \ln P \times \% \text{Dev}$		-6.52*** (1.83)	-3.67*** (0.68)	0.37 (0.60)		-5.92*** (1.86)	-3.85*** (0.75)	0.86 (0.71)	-2.82*** (0.90)	-2.80*** (0.96)
$\Delta \ln P \times \% \text{Dev}^2$		5.00*** (1.31)	2.80*** (0.71)	-0.33 (0.50)		4.24*** (1.27)	2.99*** (0.80)	-0.70 (0.55)	2.26*** (0.68)	2.11*** (0.69)
$\Delta \ln P \times \% \text{Flat}$		0.55* (0.30)	0.10 (0.14)	0.62 (0.60)		0.62* (0.36)	0.07 (0.12)	0.95 (0.81)	0.25* (0.13)	0.29* (0.16)
$\Delta \ln P \times \text{WRLURI}$			-0.03** (0.02)				-0.02 (0.02)			
$\Delta \ln P \times (\text{Res. FAR})$				0.12*** (0.04)				0.10** (0.04)		
Observations	23,917	23,917	9,431	1,528	24,442	24,442	9,887	1,952	23,917	24,442
$\Delta \ln P$ Measure	RS	RS	RS	RS	HI	HI	HI	HI	RS	HI

Regressions are the same specification as in column (2) of Panel A in Table 6 with the addition of indicated interaction terms, though WRLURI and FAR main effects are excluded in Columns 3, 4, 7 and 8 to maintain statistical power. SE adjusted for spatial autocorrelation to 25km. If included where omitted, coefficients on $\Delta \ln P \times (\text{CBD Dis})^2$ would be insignificant.

Table 8: Heterogeneity in Land & Floorspace Supply Elasticities by CBD Dis. & Tract Conditions

	$\Delta \ln S(\text{supply})$									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Developed Land								Floorspace	
$\Delta \ln P$	0.02	0.73**	0.81***	0.03	0.01	0.75**	0.80***	-0.05	5.43**	4.83***
	(0.08)	(0.35)	(0.18)	(0.21)	(0.09)	(0.34)	(0.21)	(0.27)	(2.71)	(1.85)
$\Delta \ln P \times (\text{CBD Dis})$	0.46*	-0.43	0.17	0.04	0.40	-0.51*	-0.09	-0.19		
	(0.27)	(0.28)	(0.42)	(0.44)	(0.25)	(0.31)	(0.51)	(0.46)		
$\Delta \ln P \times (\text{CBD Dis})^2$	-0.29				-0.25					
	(0.28)				(0.28)					
$\Delta \ln P \times \text{Edge}$	0.15	-0.38*	-0.05	0.04	0.13	-0.41*	-0.12	0.16		
	(0.10)	(0.19)	(0.18)	(0.16)	(0.08)	(0.22)	(0.18)	(0.58)		
$\Delta \ln P \times \% \text{Dev}$		-2.32***	-2.49***	-0.42**		-2.35***	-2.66***	-0.37	-14.28***	-11.90***
		(0.81)	(0.48)	(0.20)		(0.83)	(0.53)	(0.23)	(5.23)	(3.44)
$\Delta \ln P \times \% \text{Dev}^2$		1.65***	1.76***	0.23		1.62***	1.98***	0.23	14.30***	12.84***
		(0.59)	(0.52)	(0.19)		(0.59)	(0.58)	(0.19)	(5.38)	(4.48)
$\Delta \ln P \times \% \text{Flat}$		0.24*	0.13	0.10		0.27*	0.11	0.17	0.24	0.03
		(0.13)	(0.11)	(0.15)		(0.16)	(0.10)	(0.23)	(0.51)	(0.54)
$\Delta \ln P \times \text{WRLURI}$			-0.01				-0.01			
			(0.01)				(0.01)			
$\Delta \ln P \times (\text{Res. FAR})$				0.03***				0.02*		
				(0.01)				(0.01)		
Observations	23,917	23,917	9,431	1,528	24,442	24,442	9,887	1,952	22,654	23,447
$\Delta \ln P$ Measure	RS	RS	RS	RS	HI	HI	HI	HI	RS	HI

Regressions are the same specification as in Table 6 with the addition of indicated interaction terms, though WRLURI and FAR main effects are excluded in Columns 3, 4, 7 and 8 to maintain statistical power. SE adjusted for spatial autocorr. to 25km. If included where omitted, coefficients on $\Delta \ln P \times (\text{CBD Dis})^2$ would be insignificant.

Table 9: Summary of Tract Level Supply Elasticities (γ_i)

	25th Pctile	Median	75th Pctile	Mean	St. Dev.
Panel A: Based on Repeat Sales Price Growth					
		$\gamma_i^{unit} = \gamma_i^{redevelop} + \gamma_i^{new\ develop}$			
Units	0.1	0.4	0.9	0.5	0.6
Redevelopment	0.0	0.2	0.3	0.2	0.2
		$\gamma_i^{space} = \gamma_i^{land} + \gamma_i^{unit/land} + \gamma_i^{space/unit}$			
Floorspace	2.1	2.6	3.7	3.0	1.1
Land Development	0.0	0.2	0.3	0.2	0.2
	2%	6%	10%	5%	6%
Units per Land	0.1	0.2	0.6	0.3	0.4
	3%	10%	15%	8%	9%
Floorspace per Unit	2.0	2.3	3.0	2.5	0.7
	75%	84%	96%	87%	15%
Panel B: Based on Hedonic Price Growth					
		$\gamma_i^{unit} = \gamma_i^{redevelop} + \gamma_i^{new\ develop}$			
Units	0.1	0.4	0.8	0.5	0.6
Redevelopment	0.0	0.1	0.3	0.1	0.2
		$\gamma_i^{space} = \gamma_i^{land} + \gamma_i^{unit/land} + \gamma_i^{space/unit}$			
Floorspace	2.2	2.6	3.5	2.9	0.9
Land Development	0.0	0.2	0.3	0.2	0.2
	1%	7%	11%	6%	7%
Units per Land	0.0	0.2	0.5	0.3	0.3
	2%	9%	15%	8%	10%
Floorspace per Unit	2.0	2.3	2.9	2.5	0.6
	74%	84%	97%	86%	17%

50,409 unique tracts are equally weighted across 63,896 observations. Percentages capture the distributions of fractions of floorspace elasticities accounted for by indicated components. To impute percentages, we first impute the fraction of a given component in each tract and then average the fractions across tracts.

Table 10: Summary Statistics of Metro Level Supply Elasticities (γ_r)

	25th Pctile	Median	75th Pctile	Mean	St Dev	Rank Corr w/ Saiz
γ_r^1 : Less Substitution Across Neighborhoods (306 Metros)						
Units (RS)	0.4	0.6	0.8	0.6	1.2	0.39
Units (HI)	0.3	0.6	0.8	0.5	0.5	0.35
Floorspace (RS)	2.8	3.1	3.5	3.2	0.5	0.39
Floorspace (HI)	2.7	3.0	3.3	3.0	0.4	0.37
γ_r^2 : More Substitution Across Neighborhoods (306 Metros)						
Units (RS)	0.6	0.8	1.0	0.7	0.3	0.40
Units (HI)	0.5	0.7	0.9	0.7	0.3	0.37
Floorspace (RS)	3.0	3.4	3.8	3.5	0.5	0.42
Floorspace (HI)	2.9	3.2	3.5	3.2	0.4	0.39
Saiz (2010) Unit Elasticities (236 Metros)						
	1.6	2.3	3.4	2.6	1.5	1

Table 11: Heterogeneity in Metro-Level Supply Elasticities

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Less Sub. Across Neighborhoods (γ_r^1)				More Sub. Across Neighborhoods (γ_r^2)			
	Units	Units	Floorspace	Floorspace	Units	Units	Floorspace	Floorspace
	(RS)	(HI)	(RS)	(HI)	(RS)	(HI)	(RS)	(HI)
Frac. Devel w/in 50% of Edge	-4.82***	-5.04***	-3.88***	-2.81***	-3.66***	-3.72***	-4.64***	-3.29***
	(1.42)	(0.58)	(0.51)	(0.41)	(0.28)	(0.28)	(0.53)	(0.43)
Frac. Unavail for Dev. w/in 50% of Edge	0.36	-0.35**	0.24*	0.29***	-0.15**	-0.21***	0.22	0.28**
	(0.38)	(0.16)	(0.14)	(0.11)	(0.08)	(0.08)	(0.14)	(0.11)
Metro Wharton Regulation Index	-0.09	-0.12***	-0.13***	-0.08***	-0.09***	-0.09***	-0.14***	-0.09***
	(0.09)	(0.04)	(0.03)	(0.03)	(0.02)	(0.02)	(0.03)	(0.03)
Constant	0.89***	0.92***	3.39***	3.16***	1.02***	0.98***	3.72***	3.37***
	(0.15)	(0.06)	(0.06)	(0.05)	(0.03)	(0.03)	(0.06)	(0.05)
Observations	306	306	306	306	306	306	306	306
R-squared	0.06	0.26	0.25	0.23	0.45	0.45	0.30	0.26

The dependent variable is the indicated aggregated metro level supply elasticity.

Results are similar when using analogous variables measured within 50 km of CBDs instead, as in Saiz (2010).


Table 12: Welfare Consequences of the OZ Program (millions of \$2016)

Opportunity Zone	Yes	No	Yes	No
Δd	0	0	0.05	0.05
Δs	0.05	0.05	0	0
Welfare Measure	CS	CS	PS	PS
Panel A: Market for Housing Units				
Tract Supply Elasticity	0.4	0.5	0.4	0.5
	(0.4)	(0.5)	(0.4)	(0.5)
Calc. Using Tract Supply Elasticities	5.5	9.2	19.5	25.0
	(9.1)	(12.4)	(17.1)	(23.9)
Calc. Using Region Supply Elasticities	7.1	9.3	17.3	24.8
	(8.5)	(10.9)	(18.8)	(25.7)
Panel B: Market for Floorspace				
Tract Supply Elasticity	2.6	2.9	2.6	2.9
	(0.8)	(1.1)	(0.8)	(1.1)
Calc. Using Tract Supply Elasticities	11.9	17.0	4.4	5.9
	(11.8)	(15.9)	(3.6)	(5.0)
Calc. Using Region Supply Elasticities	12.4	17.2	4.1	5.8
	(11.2)	(15.2)	(3.7)	(5.1)


Tract means with standard deviations in parentheses. Entries are calculated using assumptions about demand and supply shocks indicated in column headers. Analysis is for 3,473 opportunity zone tracts and 35,472 other tracts for which we could construct house price and quantity information in 2016. Estimates use a floorspace demand elasticity of -1.7 and a units demand elasticity of -0.7. Regional supply elasticities are 0.5 for units and 3.0 for floorspace on average for OZ and non-OZ tracts alike. In the floorspace market, indicated increases amount to 2.1 percent of CS and 4.2 percent of PS using tract supply elasticities for both OZ and non-OZ tracts. These percentages are undefined in the market for housing units because inelastic demand makes CS infinite.

Figure 1: Development, Topography and Regulation

Panel A: Land Attribute Distributions



Panel B: Average Tract Attributes by CBD Distance



Wharton Index and FAR are only shown for CBD distances with good coverage.

Figure 2: Supply Elasticities by CBD Distance





Figure 3: Units Supply Elasticities by CBD Distance


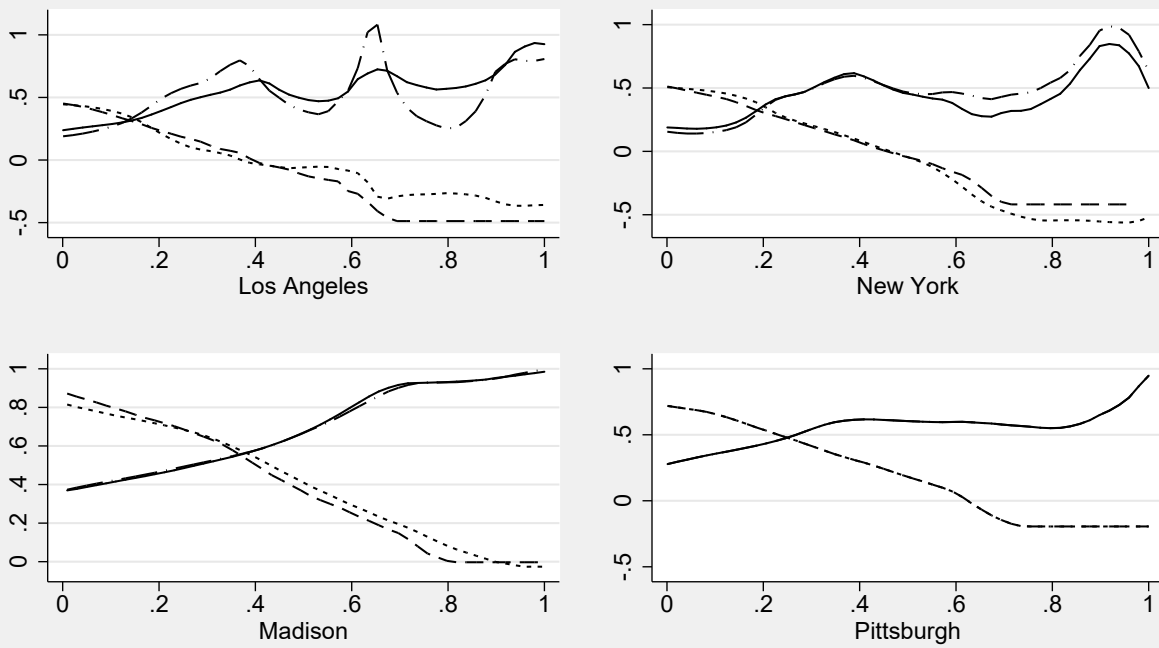


Figure 4: Floorspace Supply Elasticities by CBD Distance

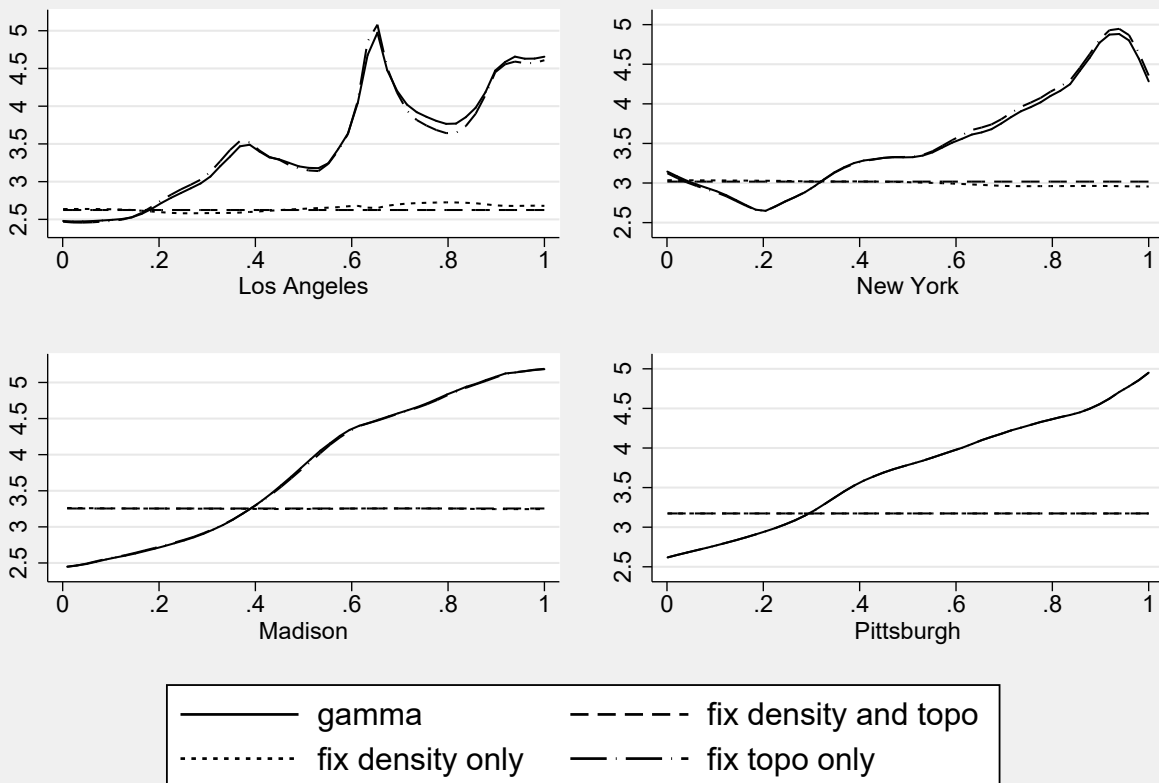


Figure A1: Example Fixed Cost Distributions


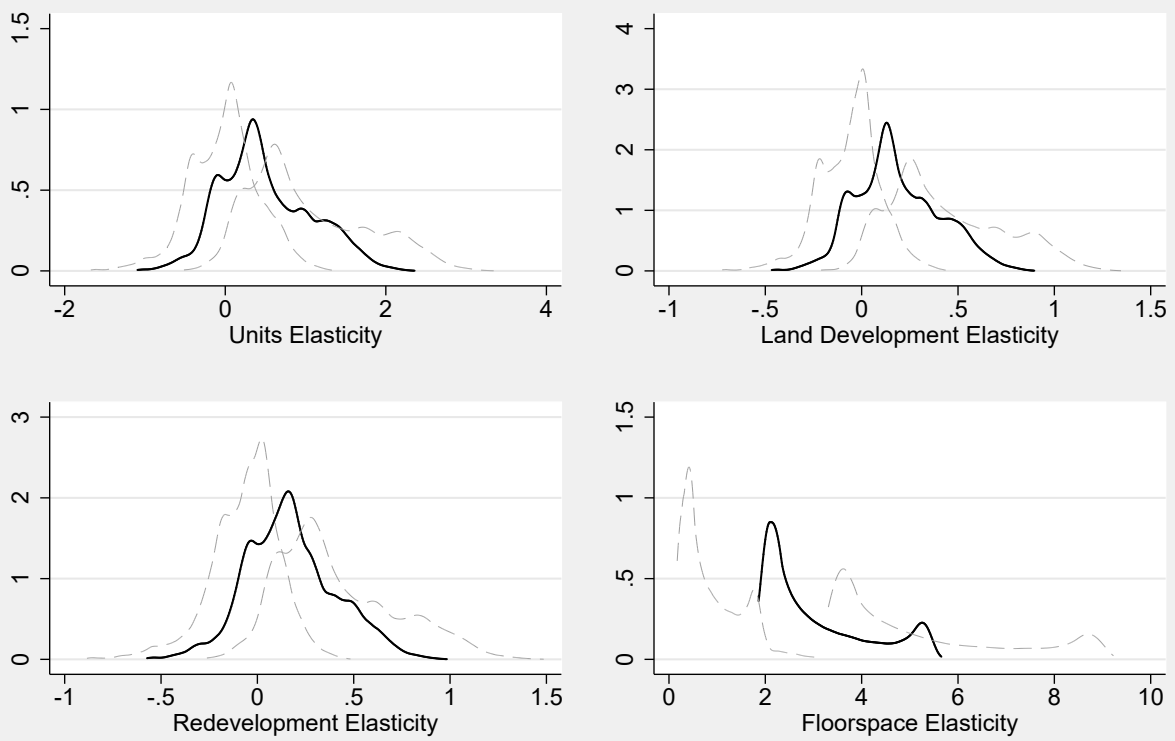


Figure A2: Kernel Dens. of Supply Elas. Components



Dashed lines indicate 95% confidence intervals of RS elasticities.

A Housing Supply

A.1 Setup

The developer's profit consists of revenue net of the fixed development cost for plot il , the variable cost and the land acquisition cost, respectively.

$$\pi_{il} = P_i A_i - g_{il} - C_i(A_i) - p_{il}$$

The developer is subject to marginal cost pricing $P_i = \frac{dC_i(A_i)}{dA_i}$ and chooses the optimal amount of floorspace $A_i^* = A_i(P_i)$ accordingly. The 0 profit condition requires that the bid for each plot of undeveloped land is, after substituting $\frac{dC_i(A_i)}{dA_i}$ for P_i :

$$p_{il} = C_i(A_i) \left(\frac{d \ln C_i(A_i^*)}{d \ln A_i} - 1 \right) - g_{il}.$$

This is the bid rent function for plot of land il . Henceforth, assume that $\frac{d \ln C_i(A_i)}{d \ln A_i} - 1 = \phi > 0$. This is consistent with Cobb-Douglas production, as described below.

Each tract has a distribution of the fixed costs of development $F_i(x)$, with this distribution depending on some tract-specific parameter. Normalizing the opportunity cost per unit of land of 0, this means that the fraction of land developed in each tract is $F_i[\phi C_i(A_i^*)]$.

A.2 Tract Housing Services Supply

The amount of developed land in tract i is $M_i F_i(\phi C_i(A_i^*))$, where M_i is the number of plots of land in tract i . The implied tract aggregate housing services (floorspace) supply function $S_i(P_i)$ is [housing services per parcel]x[number of parcels of land]x[fraction of plots developed]. Taking logs, we have

$$\ln S_i(P_i) = \ln A_i(P_i) + \ln M_i + \ln F_i(\phi C_i[A_i(P_i)]).$$

Differentiating, the supply elasticity decomposes as

$$\begin{aligned} \frac{d \ln S_i}{d \ln P_i} &= \frac{d \ln A_i(P_i)}{d \ln P_i} + \frac{f_i(\phi C_i[A_i^*])}{F_i(\phi C_i[A_i^*])} \phi MC(A_i^*) P_i \frac{d A_i(P_i)}{d P_i} \\ &= \frac{d \ln A_i(P_i)}{d \ln P_i} + \frac{f_i(\phi C_i[A_i^*])}{F_i(\phi C_i[A_i^*])} \frac{d \ln A_i(P_i)}{d \ln P_i} \phi P_i A_i^* \end{aligned}$$

This expression reflects intensive and extensive margin responses respectively.

A.3 Parameterization with Cobb-Douglas Production

The production function is $A_i = \kappa_i \bar{M}_i^\alpha K_i^{1-\alpha}$, where \bar{M}_i is the exogenous parcel size and K_i is the only variable factor. Going through profit maximization, as above, yields the following factor demand, where ι is the cost of capital. As is standard in the literature, we assume that ι does not vary by location.

$$K_i^* = \left(\frac{1-\alpha}{\iota}\right)^{\frac{1}{\alpha}} \kappa_i^{\frac{1}{\alpha}} P_i^{\frac{1}{\alpha}} \bar{M}_i$$

Floorspace per parcel is

$$A_i^* = \left(\frac{1-\alpha}{\iota}\right)^{\frac{1-\alpha}{\alpha}} \bar{M}_i \kappa_i^{\frac{1}{\alpha}} P_i^{\frac{1-\alpha}{\alpha}}$$

and variable cost is

$$C_i = \iota \left[\frac{A_i}{\kappa_i \bar{M}_i^\alpha} \right]^{\frac{1}{1-\alpha}} = (1-\alpha)^{\frac{1}{\alpha}} \iota^{1-\frac{1}{\alpha}} \kappa_i^{\frac{1}{\alpha}} P_i^{\frac{1}{\alpha}} \bar{M}_i.$$

From these objects, note that $\frac{d \ln A_i}{d \ln P_i} = \frac{1-\alpha}{\alpha}$, $MC(A_i) = \frac{dC_i}{dA_i} = \frac{\iota}{1-\alpha} \left[\frac{1}{\kappa_i \bar{M}_i^\alpha} \right]^{\frac{1}{1-\alpha}} A_i^{\frac{\alpha}{1-\alpha}}$, and the developer's revenue from one parcel is $\bar{M}_i \left(\frac{1-\alpha}{\iota}\right)^{\frac{1-\alpha}{\alpha}} \kappa_i^{\frac{1}{\alpha}} P_i^{\frac{1}{\alpha}}$. As a result, $\frac{\text{Revenue}_i}{\text{Variable Cost}_i} = \frac{1}{1-\alpha}$ and $\phi = \frac{\alpha}{1-\alpha}$.

Plugging in, the resulting elasticity expression is

$$\frac{d \ln S_i}{d \ln P_i} = \frac{1-\alpha}{\alpha} + \frac{f_i[\alpha(\frac{1-\alpha}{\iota})^{\frac{1-\alpha}{\alpha}} \bar{M}_i \kappa_i^{\frac{1}{\alpha}} P_i^{\frac{1}{\alpha}}]}{F_i[\alpha(\frac{1-\alpha}{\iota})^{\frac{1-\alpha}{\alpha}} \bar{M}_i \kappa_i^{\frac{1}{\alpha}} P_i^{\frac{1}{\alpha}}]} \left(\frac{1-\alpha}{\iota}\right)^{\frac{1-\alpha}{\alpha}} \bar{M}_i \kappa_i^{\frac{1}{\alpha}} P_i^{\frac{1}{\alpha}}.$$

A.4 Parameterization with a Frechet fixed Cost Distribution

We consider Frechet fixed cost distributions with the common dispersion parameter λ and tract-specific scale parameter Γ_i . We express the CDF as $F_i(x) = \exp[-\Gamma_i x^{-\lambda}]$ and the associated PDF as $f_i(x) = \lambda \Gamma_i x^{-1-\lambda} \exp[-\Gamma_i x^{-\lambda}]$. Therefore, $\frac{f_i(x)}{F_i(x)} = \lambda \Gamma_i x^{-1-\lambda}$.

Plugging into the expression above, the overall supply elasticity is

$$\frac{d \ln S_i}{d \ln P_i} = \left[\frac{d \ln A(P_i)}{d \ln P} + \frac{\lambda \Gamma_i}{(\phi C_i)^{1+\lambda}} \frac{d \ln A(P_i)}{d \ln P_i} \phi P_i A_i^* \right].$$

Under Cobb-Douglas production,

$$\frac{d \ln S_i}{d \ln P_i} = \frac{1 - \alpha}{\alpha} + \alpha^{-1-\lambda} \left(\frac{1 - \alpha}{\iota} \right)^{-\lambda \frac{1-\alpha}{\alpha}} \lambda \bar{M}_i^{-\lambda} \kappa_i^{-\frac{\lambda}{\alpha}} P_i^{-\frac{\lambda}{\alpha}} \Gamma_i.$$

Defining $\rho_i = \left(\frac{1-\alpha}{\iota} \right)^{\frac{1-\alpha}{\alpha}} \bar{M}_i \kappa_i^{\frac{1}{\alpha}}$, the extensive margin component can be written as $\alpha^{-1-\lambda} \lambda \rho_i^{-\lambda} P_i^{-\frac{\lambda}{\alpha}} \Gamma_i$.

A.5 Land Redevelopment

Above we treated $f_i(x)$ as the density of fixed costs of land preparation for development, whether or not parcels had been developed prior. Denote as $f_i^L(x)$ the density of fixed costs of land preparation across all parcels in a tract in an environment in which it had not been developed prior.

Fraction $\frac{M_i^R}{M_i}$ of parcels up to fixed cost \bar{g}_i has been previously developed. We suppose that previously developed parcels become r dollars more expensive to redevelop relative to the original fixed development cost. Then the new density of redevelopment fixed costs is $f_i^R(z) = \frac{f_i^L(x-r)}{M_i^R/M_i} \mathbf{1}(z \leq \bar{g}_i + r)$. This is the left tail of the $f_i^L(x)$ distribution, shifted by the additional parcel preparation cost r and rescaled to integrate to 1, where $M_i^R/M_i = F^L(\phi C_i)$. The new distribution of fixed costs of developing undeveloped land is $f_i^U(z) = \frac{f_i^L(x)}{M_i^U/M} \mathbf{1}(z \geq \bar{g}_i)$. This is the right tail of the $f_i^L(x)$ distribution. That is, both have support between fixed costs of \bar{g}_i and $\bar{g}_i + r$. Using these ideas, we decompose the fixed cost distribution $F_i(z)$ from Section A.1 into $F_i(z) = \frac{M_i^R}{M_i} F_i^R(z) + \frac{M_i^U}{M_i} F_i^U(z) = F_i^L(x-r) \mathbf{1}(x \leq \bar{g}_i) + F_i^L(x) \mathbf{1}(x \geq \bar{g}_i)$. Figure A1 shows plots of $f^L(x)$, $f^R(x)$ and $f^U(x)$ for an example tract. As P_i rises, marginal land parcels are developed left to right in the region of overlapping support of the $f^R(x)$ and $f^U(x)$ distributions.

With these definitions of $F_i^R(x)$ and $F_i^U(x)$ established, the generic revised aggregate tract floorspace supply function is

$$S_i(P_i) = A_i(P_i) [M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)].$$

Relative to an initial baseline, prices rise such that $\bar{g}_i < \phi C_i < \bar{g}_i + r$. As a result, developers draw

from both the previously developed and undeveloped land for their new developments.

$$\begin{aligned}\frac{d \ln S_i}{d \ln P_i} &= \frac{d \ln A_i(P_i)}{d \ln P_i} + \frac{M_i^R f_i^R(\phi C_i) + M_i^U f_i^U(\phi C_i)}{M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)} \frac{d \ln A_i(P_i)}{d \ln P} \phi P_i A_i^* \\ &= \frac{d \ln A_i(P_i)}{d \ln P_i} \left[1 + \frac{M_i^R f_i^R(\phi C_i)}{M_i^{R'}} + \frac{M_i^U f_i^U(\phi C_i)}{M_i^{R'}} \right] \phi P_i A_i^*.\end{aligned}$$

Here $M_i^{R'}$ denotes the current mass of developed land.

We want the unit supply elasticity γ_i^{unit} to decompose as $\gamma_i^{unit,R} + \gamma_i^{unit,U} = \frac{dH_i^R/H_i}{d \ln P_i} + \frac{dH_i^U/H_i}{d \ln P_i}$.

As a baseline, observe that

$$H_i = H_i^R + H_i^U = \frac{H_i}{L_i} [M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)].$$

Differentiating,

$$\begin{aligned}\frac{dH_i^R}{H_i} &= \frac{\frac{H}{L} d[M_i^R F_i^R(\phi C_i)] + M_i^R F_i^R(\phi C_i) d\frac{H}{L}}{\frac{H}{L} [M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)]} \\ &= \frac{d[M_i^R F_i^R(\phi C_i)]}{M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)} + \frac{M_i^R F_i^R(\phi C_i)}{M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)} d \ln \frac{H}{L} \\ &= \frac{M_i^R F_i^R(\phi C_i)}{M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)} d \ln [M_i^R F_i^R(\phi C_i)] + \frac{M_i^R F_i^R(\phi C_i)}{M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)} d \ln \frac{H}{L}\end{aligned}$$

Therefore,

$$\gamma_i^{unit,R} = \frac{M_i^R F_i^R(\phi C_i)}{M_i^{R'}} \left[\frac{d \ln H_i / L_i}{d \ln P_i} + \frac{f_i^R(\phi C_i)}{F_i^R(\phi C_i)} \phi P_i A_i^* \right].$$

Analogously,

$$\gamma_i^{unit,U} = \frac{M_i^U F_i^U(\phi C_i)}{M_i^{R'}} \left[\frac{d \ln H_i / L_i}{d \ln P_i} + \frac{f_i^U(\phi C_i)}{F_i^U(\phi C_i)} \phi P_i A_i^* \right].$$

A.6 FAR Restriction

The tract level FAR constraint is $\frac{\text{Floorspace}}{\text{Lot Size}} = D_i$. Suppose that the price is sufficiently high such that the developer builds up to the point that $\bar{A}_i = D_i \bar{M}_i$, and is constrained at this point. Profits are now

$$\pi_{il} = P_i \bar{A}_i - g_{il} - C_i(\bar{A}_i) - p_{il},$$

directly pinning down the parcel price through the 0-profit condition. Therefore, fraction $F_i(P_i\bar{A}_i - C_i(\bar{A}_i))$ of available land is developed.

With the tract developed at the maximum allowed floorspace on each parcel, the supply function is

$$\ln S_i^{FAR} = \ln(\bar{A}_i) + \ln M_i + \ln F_i(P_i\bar{A}_i - C_i(\bar{A}_i)).$$

Therefore, the tract supply elasticity comes from the extensive margin only. It is

$$\frac{d \ln S_i^{FAR}}{d \ln P_i} = \frac{d \ln L_i^{FAR}}{d \ln P_i} = \frac{f_i(P_i\bar{A}_i - C_i(\bar{A}_i))}{F_i(P_i\bar{A}_i - C_i(\bar{A}_i))} \bar{A}_i.$$

Depending on the shape of f , this could mean the relaxation of a FAR results in a greater or smaller extensive margin supply elasticity. The selection effect of increasing variable profit from $P_i\bar{A}_i - C_i(\bar{A}_i)$ to $\phi C_i(A_i)$ may go in the opposite direction of the incentive effect of increasing marginal revenue from \bar{A}_i to $P_i A_i^*$.

B Housing Demand Model

The fraction of residents of tract i that work in j $Pr(\frac{v_{i\omega} B_i z_{ijk\omega} w_{jk}}{P_i^{1-\beta} e^{\kappa\tau_{ij}}} >= \max_{j',k'} \frac{v_{i\omega} B_i z_{ij'k'\omega} w_{j'k'}}{P_i^{1-\beta} e^{\kappa\tau_{ij'}}})$ can be determined using the properties of the Frechet draws $z_{ijk\omega}$.

$$\pi_{ij|i} = \frac{\sum_k [w_{jk} e^{-\kappa\tau_{ij}}]^\varepsilon}{\sum_k \sum_{j'} [w_{j'k} e^{-\kappa\tau_{ij'}}]^\varepsilon} \equiv \frac{\sum_k [w_{jk} e^{-\kappa\tau_{ij}}]^\varepsilon}{RMA_i} \quad (27)$$

We write this expression as a function of resident market access $RMA_i \equiv \sum_k \sum_{j'} [w_{j'k} e^{-\kappa\tau_{ij'}}]^\varepsilon$, which is a summary measure of the access to employment opportunities from residential neighborhood i .

Before the productivity shock is revealed, the expected income (wage net of commuting cost) \bar{y}_i associated with residing in tract i is $E(\max_{j,k} \frac{w_{jk} z_{ijk\omega}}{e^{\kappa\tau_{ij}}})$. Solving this through,

$$\bar{y}_i = \Gamma(1 - \frac{1}{\varepsilon})(RMA_i)^{\frac{1}{\varepsilon}} \quad (28)$$

This object is increasing in RMA_i and declining in ε . As ε increases, there is a smaller dispersion in skill prices across locations, reducing the probability of individuals receiving high wage offers

in any location.

The probability that i is the highest utility residential location is the probability that the inclusive value of municipality r is the highest times the probability that neighborhood i is the highest utility neighborhood in municipality m . Using properties of the Fréchet distribution, this second object is $\frac{(B_i P_i^{\beta-1} \bar{y}_i)^\eta}{\sum_{i'} (P_{i'}^{\beta-1} B_{i'} \bar{y}_{i'})^\eta}$. The second object is $\frac{\sum_{i' \in m(i)} (B_{i'} P_{i'}^{\beta-1} \bar{y}_{i'})^\eta \frac{\psi}{\eta}}{\sum_m [\sum_{i' \in m(i)} (B_{i'} P_{i'}^{\beta-1} \bar{y}_{i'})^\eta \frac{\psi}{\eta}]}$. Plugging in for \bar{y}_i gives the population supply function in the text.

The probability that j is the highest utility work location for a resident of any given tract i is $\pi_{ij|i} = \frac{\sum_k [w_{jk} e^{-\kappa \tau_{ij}}]^\varepsilon}{\sum_k \sum_{j'} [w_{j'k} e^{-\kappa \tau_{ij}}]^\varepsilon}$. Summing over the probability of living in i , we recover the labor supply to tract j in the text.

Substituting for (19) in housing demand (18), setting it equal to housing supply (2), solving for price and differentiating yields the growth rate in tract house price, expressed as follows.

$$d \ln P_i = \frac{1}{\varepsilon(\gamma_i + 1)} + \frac{1}{\gamma_i + 1 + \eta(1 - \beta)} \frac{\eta}{\varepsilon} \left(1 - \frac{1 - \beta}{\gamma_i + 1}\right) d \ln RMA_i + v_m^P + u_i^P \quad (29)$$

This equation shows, as is intuitive, that positive shocks to employment opportunities get capitalized into home prices. The amount of capitalization is of course decreasing in the floorspace supply elasticity γ_i and in the dispersion of amenity draws within municipality m , an object which is negatively related to η . Changes in housing productivity, average lot size and the local amenity B_i show up in the error term u_i^P . Because RMA_i itself depends on quantities and prices throughout the region, it is also a function of these objects.

Finally, the model delivers the following implicit equation which describes the relationship between change in $d \ln RMA_i$ and municipality level aggregates of tract population growth $d \ln \pi_i$.

$$\sum_{i(m)} s_i \left[1 + \frac{\psi(1 - \beta)}{1 + \gamma_i}\right] d \ln \pi_i - \Lambda - \frac{\psi}{\varepsilon} \sum_{i(m)} \left[s_i \left(1 - \frac{1 - \beta}{1 + \gamma_i}\right)\right] d \ln RMA_i = u^m \quad (30)$$

In this equation, s_i is the base year share of municipality r 's population living in tract i . As ψ rises, dispersion in preferences across municipalities falls. As a result, positive shocks to RMA in any neighborhoods within m result in more rapid population growth in this municipality.