

The Microgeography of Housing Supply*

Nathaniel Baum-Snow, University of Toronto

Lu Han, University of Toronto

September 8, 2019

Abstract

This paper reports estimates of housing supply elasticities for most urban neighborhoods in the U.S. Consistent with housing production function estimates from the literature, we find that the average neighborhood level elasticity of the supply of housing services with respect to price is about 4. The average price elasticity of quantity of new units supplied is only about 0.4, with about one-third of this response due to redevelopment. Housing supply increases with CBD distance, in part accounted for by increasing neighborhood land availability from reductions in the fraction already developed. Tracts with flat and gently sloping land also exhibit more elastic supply. Conditional on census tract level measures of land unavailability, analogous metro area measures do not significantly influence tract level supply estimates. Identification comes from variation in labor demand shocks to commuting destinations, as aggregated using an urban economic geography model. Aggregation of neighborhood level supply elasticities yields metro area supply elasticities that are correlated with but smaller than those reported in Saiz (2010).

*We thank Camilo Acosta, Rolando Campusano and Ian Herzog for excellent research assistance.

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 (Calabrese, Epple, & Romano, 2011; 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). The main aim of this study is to empirically characterize housing supply elasticities for all residential neighborhoods in 306 U.S. metro areas. We also decompose supply responses of total housing services into units and quality components. Within the units component, we further decompose into redevelopment of existing developed areas versus newly developed land. We demonstrate how housing supply conditions vary by neighborhood location, available land, topography and regulation. Finally, we explore various ways of aggregating neighborhood housing supply elasticities to the metro area level.

Our estimates of supply elasticities are about 0.5 for housing units and 4 for housing services on average across neighborhoods in our data, with one-quarter to one-third of the unit supply response coming through land redevelopment. These supply responses grow with distance from central business districts (CBDs) such that unit and housing services supply elasticities at CBDs are on average 0.3 and 2 respectively, growing to 1.7 and 5 respectively at urban fringes. In addition, supply responds in expected directions to the fraction of neighborhood land area that is already developed and has steep slopes. Tracts in areas with more stringent regulation, as measured by Floor-Area-Ratio restrictions (Brueckner & Singh, 2018), are less responsive in supply. Resulting estimates range from 0.2 to 1.6 for the 5th and 95th percentile neighborhoods for unit supply and 3.5 to 4.8 respectively for housing services supply. While we face some statistical power constraints for isolating variation in supply elasticities between metros as functions of land developed frac-

tion, topography, and regulation, we do not find evidence that metro level factors matter for local supply elasticities conditional on neighborhood topography and land development patterns. Using data from the US and Switzerland respectively, Orlando and Redfearn (2018) and von Ehrlich, Schöni, and Büchler (2018) find even smaller unit supply elasticities at the municipal level, though these papers also find that supply elasticities increase with CBD distance.

Supply elasticity estimates for housing services are quantitatively consistent with housing production function estimates in the literature. These indicate approximately Cobb-Douglas form with a land share of 0.2-0.3 (Ahlfeldt & McMillen, 2014; Combes, Duranton, & Gobillon, 2017; Albouy & Ehrlich, 2012), implying housing services supply elasticities of 2.3-4. Our tract level unit supply estimates are mostly smaller than those found in Saiz (2010) for the metro level. One reason is that aggregation from micro to macro elasticities incorporates substitution patterns across neighborhoods in residents' neighborhood demand system. As neighborhoods become stronger demand substitutes, a shock affecting (for example) 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. The result is aggregation to macro supply elasticities that are greater than tract level elasticities. However, we note that macro elasticities are specific to the nature of the demand shock. This motivates us to explore aggregation using estimated demand parameters and shocks generated by a specific place based policy in an urban market access model similar to that in Tsivanidis (2018) (in progress).

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 house price index. The central challenge 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 changes in house prices in residential locations. These labor demand shocks depend on 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 in 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, we hope our supply estimates will be useful both for policy evaluation and to help improve understanding of patterns of neighborhood change. For example, 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). On the other hand, the greater house price fluctuations at metro area edges where housing is the most land intensive (Zhou & Haurin, 2010) seem most likely driven by demand fluctuations. A burgeoning literature examines policies and phenomena that only directly impact a few neighborhoods in cities in the context of general equilibrium urban models. These studies typically calibrate micro elasticities to macro estimates from the literature. Our evidence shows that this choice can lead to misleading conclusions about incidence. In particular, landowners in the supply inelastic neighborhoods nearer to CBDs may enjoy most of the gains or bear most of the burdens even if at the metro area level housing supply is elastic. We thus hope that our supply elasticity estimates are useful for improving quantitative evaluation of various policies that are targeted to particular neighborhoods.

2 Data

For this analysis, we construct a data set that brings together information from a number of different sources. Using the Zillow ZTRAX transactions, assessments and historical assessments data files, we build quality-adjusted house price growth between 2000 and 2010 for census tracts in 306 metro areas using both repeat sales and hedonic approaches. We use the same data sources to construct a rich set of housing supply measures including changes in housing stocks, new construction, quality-adjusted changes in housing stocks, redevelopment through both renovation and teardowns, all at the tract level. Local labor demand conditions are measured 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, American Community Survey (ACS) data from 1990-2015. Finally, we use remote sensing information on land cover in 2001 to measure baseline tract development intensity, topography and prevalence of wetlands. 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 ZTRAX data sets (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 the full study area by 2010, with a few exceptions noted below. Because of incomplete coverage, particularly in year 2000, we supplement Zillow data with decennial Census data, as is explained in further detail below.

Transactions information is constructed from information in local Recorders of Deeds and includes the sale price, location and some property attributes. To fill out property attributes, we merge it with the most recent assessment data. We primarily use this 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, semi-detached and condominiums. We consider only homes that are bought by individual buyers and do not examine institutional buyers. We exclude all census

tracts with fewer than 10 transactions over the full sample period. Our main use of the ZTRAX data is to build home price indexes at the census tract level.

A well-known challenge for constructing home price indexes is that homes are heterogeneous in observed and unobserved attributes. The goal of the indexes 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).

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

Here, h indexes homes in census tract i , region r , year t and month m . X_{hirtm} includes a rich set of characteristics (including rooms, bedrooms, kitchens, bathrooms, heating and AC, elevator, fireplace, water, sewer, year built, year renovated, floorspace and lot size). Month of sale fixed effects ρ_m^{HI} flexibly account for seasonality in market conditions. Differing coefficients on housing attributes β_r^{HI} accounts for the possibility that (for example) heating and air conditioning may be valued differently in different regions. Our inclusion of controls for region fixed effects throughout the main empirical analysis below ensures that comparisons are always made across locations within metro regions. Due to incomplete coverage in the ZTRAX data set and our inclusion of tracts with at least 10 valid home sales during our sample period, this index covers only about two-thirds of our full tract sample.

To fill out the tract sample with some measure of house prices for 2000 and 2010 and to facilitate looking at pre-trends, we also build a lower quality hedonic price index using self-reported data from the 1990 and 2000 Censuses of Housing and the 2008-2012 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 census tracts in our sample regions.

Hedonic indexes 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 at the tract-year level. For this index, we exclude homes that sold only once or more than 9 times and 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 are 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 more biased sample than the hedonic index and incorporate unwanted capitalization of unobserved home improvements.

The top block in Table 1 presents summary statistics about changes in these three home price indexes for the primary estimation sample used for the empirical work. The Zillow hedonic price index growth is 0.89 on average across tracts relative to 0.77 for repeat sales index growth during the 2000-2006 periods. For 2000-2010, average growth rates are 0.37 and 0.45 respectively, reflecting the 2007-2008 housing market crash. Correlations between the two Zillow indexes is 0.63 for the 2000-2010 period but the correlation with the growth in the Census price index is only about 0.10 for both Zillow based indexes. Nevertheless, we will see below that all three price indexes generate similar housing supply elasticity estimates.

2.2 Satellite Data

We use remote sensing information to measure tract level topography and land development intensity. Topography is one of the most important factors influencing construction costs and ultimately housing supply. Land cover information is used to help determine whether new housing is built on previously undeveloped land.

We use three remote sensing data sets to obtain land cover and slope. First, the “Scientific Investigations Map 3085” is derived from the US Geological Survey’s National Elevation Database. The Map 3085 data 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. We aggregate these categories into “flat plains”,

“flat non-plains” and “hilly”. Flat plains 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 41% of tract land area is flat plains. Flat non-plains have a slope of less than 8% in more than half of pixels within 1 sq km and a larger elevation range. On average, 45% of tract land area is flat non-plains. Remaining land is hilly. We take elevation range within 1 km of each pixel 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. In particular, the NLCD provides for each 30X30 meter cell 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 to impute developed fraction for the land area of each census tract. The average tract in our sample had 38% of land area 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 the metro area. As these measures are highly correlated, our estimation results are not sensitive to the choice of aggregation.

2.3 Housing Quantities

We construct six measures of housing quantity changes to cover different aspects of supply responses. As most of the existing literature on housing supply focuses on units in the housing stock, we begin with this measure. However, we also aim to explicitly measure new construction and to quality adjust the stock, which the detailed Zillow assessment data allows for. However, since the Zillow data has incomplete coverage at the beginning of the sample, most of our stock measures rely on 2000 census data for the base year.

To organize our classification of stock measures, we begin by noting that the total quantity of efficiency units of housing services in census tract i , H_i^{ef} , is the number of housing units H_i times the average quantity of services provided per housing unit I_i . Differentiating, the growth rate in

the quantity of housing services over time can be written as follows:

$$d \ln H_i^{ef} = \frac{dH_i^R}{H_i} + \frac{dH_i^U}{H_i} + \frac{dH_i^T}{H_i} + d \ln I_i \quad (1)$$

We separate out our measures of the total change in housing services in tract i into these three components. New construction is broken into that on existing developed land $\frac{dH_i^R}{H_i}$ and on undeveloped land $\frac{dH_i^U}{H_i}$. We call the former “redevelopment”. The third term is negative and captures teardowns and the housing units that fully depreciate away. Finally, the intensive margin component $d \ln I_i$ is a residual that includes partial depreciation, remodeling and the quality of new construction units relative to that for units in the base period.

The simplest quantity measures are the 1990-2000 and 2000-2010 growth rates in tract level housing occupied unit stock constructed using the 100% count Decennial Census data $\frac{dH_i}{H_i}$.¹ While most complete, 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 number of new construction units across tracts in our sample, counted using both the ACS and Zillow data, are very close at about 235, or about 11 percent growth over the base in 2000. The average census growth number for the same period, which incorporates fully depreciated or torn down units, is 181 or 7 percent growth. Pairwise correlations between these three measures are all over 0.91.

¹We use occupied units instead of all units to be consistent with Saiz (2010) and because vacant units may be under-reported or not habitable.

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 (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 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 the 2011 satellite data. 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. Our measure indicates that about 40% of new construction in an average tract in our sample is redevelopment.

In the context of the model developed below, the tract efficiency units of housing services H_i^{ef} , including the intensive margin of supply, is the most conceptually relevant measure. If recently built housing units are larger than those built in early 2000s, then using either the difference in stocks or new construction would underestimate the true housing supply response. Moreover, we wish to have some way of capturing upgrades to the existing stock.³ To account for quality changes, we use the Zillow data to construct a tract level quality-adjusted housing quantity index. We start by constructing weights for each housing attribute using transactions data from 2006 only. We choose 2006 because this is the first year with near universal coverage in our sample

³The challenge of separating between the price of housing per unit and the quality-adjusted amount of housing is well-established in the housing production literature Ahlfeldt and McMillen (2014); Combes et al. (2017); “Housing Productivity and the Social Cost of Land-Use Restrictions” (n.d.). As noted in Epple, Gordon, and Sieg (2010), “houses are viewed as differing only in the quantity of service they provide, with housing service being homogenous and indivisible.”

area. For all census tracts with at least 10 transactions, we estimate coefficients in the following hedonic regression:

$$\ln P_{hir} = X_{hir}\beta^W + \rho_m^W + \epsilon_{ir} + \eta_{hirm}$$

.Here, X_{hir} includes property age, age squared, floorspace, floorspace squared and home type dummies and ϵ_{ir} are census tract fixed effects. β^W is a vector of hedonic weights that indicate the amount of housing services provided by each attribute in 2006. These attribute prices that are common across locations and are applied in subsequent years.⁴ We take the universe of (assessed) housing units in each tract and year and calculate $H_{hirt}^{ef} = \exp(X_{hirt}\hat{\beta}^W)$. The stock of housing services in tract i at time t is then $H_{irt}^{ef} = \sum_h H_{hirt}^{ef}$.

Our final measure of housing quantity change is the (residual) intensive margin of housing supply holding the number of units fixed $d \ln I_i$. Renovation is an important supply substitute for new construction. Choi, Hong, and Scheinkman (2014) report that American home improvement expenditures on, for instance, a new bathroom or a new deck jumped from around 1% of gross domestic product (GDP) (\$229 billion) in 2003 to 2% of GDP (\$326 billion) in 2007.⁵ Despite its importance, renovation has often been ignored in the literature on housing supply (DiPasquale, 1999). We capture this important dimension (both renovation of existing units and upsizing in newly built units) by taking the difference between the quality-adjusted housing quantity change constructed using ZTRAX data and the census-measured change in housing units.

2.4 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. 1990 CTPP geography determines our study regions. The 1990 CTPP assigns microgeographic units the size of census tracts 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

⁴Because of the limited number of transactions in 2006, we found it better to maintain common weights across regions for the purpose of constructing this index.

⁵According to Bendimerad (2007), Americans spent \$280 billion on home remodeling in 2005, and this number was projected to increase at 3.7% in real terms by 2015. Plaut and Plaut (2010) further report that almost half of American home owners made some renovations during 2003-2004.

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 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 commuting 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 predict 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.

2.5 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 incorporate data 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.⁶ 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.

2.6 Estimation Sample

Our analysis requires reliable information on housing quantities and prices. Our Zillow price indexes are built only for tracts with at least 10 valid home transactions in our sample period. For quantities, 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. Unless the 2000 Zillow stock is within 20% of its 2000 census counterpart, we also exclude tracts in non-disclosure states for which Zillow reports they have incomplete coverage.⁷ For these tracts, we are particularly concerned about under-measurement of 2000-2010 new construction. As a result, our sample is cut in about half, from 63,896 tracts in 306 regions to 31,242 tracts in 275 regions. Beyond this initial sample restriction, we lose about 7,000 additional observations because we lack ZTRAX based price index data for year 2000. Primary estimation sample sizes are reported in the top two blocks of Table 1.

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

⁶Most of these data were generously provided by Ruchi Singh (Brueckner & Singh, 2018).

⁷These states are Alaska, Idaho, Kansas, Louisiana, Mississippi, Missouri, Montana, New Mexico, North Dakota, Texas, Utah, and Wyoming.

3 Conceptual Framework

The main object of our analysis is to estimate housing supply elasticities that are allowed to differ flexibly 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 estimating γ_i in the following reduced form expression. Here, H_i^s denotes any of the quantity measures listed in (1).

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

In the empirical work, u_{ir} includes supply shifters and the error and we allow γ_{ir} to depend on tracts' observed heterogeneity, including initial building density, geographic features and distance to the central business district. 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 and Gyourko (2005). Largely for this reason, our analysis focuses on demand shocks in the 2000-2006 period. 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 lay out relevant microfoundations for housing supply in our context and then we lay out the theory behind the procedure we use to construct demand shocks necessary to pin down consistent estimates of γ_{ir} .

3.1 A Theory of Neighborhood Housing Supply

While microfoundations are not required to use our supply elasticity results, they do help to understand magnitudes of our empirical results reported below. Indeed, conceptual development beyond simple aggregation of individual housing production functions is needed to justify our differing price elasticity estimates for different components of supply. To fix ideas about what our different elasticity measures deliver and how to combine them, here we sketch a model of housing supply to neighborhood i .

We begin with market structure and technology and then add land availability and redevelopment constraints. We assume that the amount of housing services on each parcel S_i is supplied by a developer in a perfectly competitive market. Developers combine land and capital with pro-

ductivity that can differ across neighborhoods. Each building site s in neighborhood i has a fixed lot size \bar{M}_i and a heterogeneous fixed development cost of f_{is} . f_{is} captures development fees and permitting costs plus land preparation costs for development, be it demolition and environmental remediation of existing construction or grading and draining of undeveloped land. 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). Demand conditions and the cost of capital determine the intensity of new development on any given plot in a neighborhood. The price per unit of housing services P_i is the same for all parcels in neighborhood i .

Following is a representative developer's profit associated with building on parcel s in neighborhood i .

$$profit_{is} = P_i S_i(P_i) - f_{is} - C(S_i) - p(f_{is})$$

Here $C(S_i)$ is the variable cost and $p(f_{is})$ is the endogenous parcel acquisition price. Imposing 0 profits and perfect competition (marginal cost pricing), we have

$$p(f_{is}) = C(S_i) \left(\frac{d \ln C(S_i)}{d \ln S} - 1 \right) - f_{is}.$$

Because of the fixed lot size, the variable cost function is likely convex and so $\frac{d \ln C}{d \ln S} > 1$. The first term of this expression reflects the intuition that higher housing prices imply higher marginal costs and greater profits absent land costs. The second term simply reflects capitalization of fixed development cost into the parcel price.

Consistent with housing production function estimates in the literature, we proceed assuming Cobb-Douglas production with a land share of α and productivity ρ_i . This implies a parcel-specific housing services supply function of $S_i(P_i) = \rho_i P_i^{\frac{1-\alpha}{\alpha}}$ and that $p(f_{is}) = \alpha \rho_i P_i^{\frac{1}{\alpha}} - f_{is}$.⁸

Each tract has a CDF of the fixed costs of development $F_i(x)$ that depends on tract characteristics. Normalizing the opportunity cost per unit of land to 0, this means that the fraction of land developed in each tract is $F_i(\alpha \rho_i P_i^{\frac{1}{\alpha}})$. The resulting aggregate housing supply function in tract i is:

$$\ln H_i^{ef,s} = \rho_i P_i^{\frac{1-\alpha}{\alpha}} \bar{M}_i F_i(\alpha \rho_i P_i^{\frac{1}{\alpha}}) \quad (3)$$

⁸ $\rho_i = \iota^{\frac{\alpha-1}{\alpha}} (1-\alpha)^{\frac{1-\alpha}{\alpha}} \kappa_i^{\frac{1}{\alpha}} \bar{M}_i$ where ι is the cost of capital and κ_i is tract housing productivity.

This function reflects the intensive and extensive margin of supply. The first part, $\rho_i P_i^{\frac{1-\alpha}{\alpha}}$, is the intensity of development on each developed plot. The second part, $\bar{M}_i F_i(\alpha \rho_i P_i^{\frac{1}{\alpha}})$, is the total developed land area in the tract. Differentiating, the implied housing supply elasticity is

$$\frac{d \ln H_i^{ef,s}}{d \ln P_i} = \frac{1-\alpha}{\alpha} + \frac{f_i(\alpha \rho_i P_i^{\frac{1}{\alpha}})}{F_i(\alpha \rho_i P_i^{\frac{1}{\alpha}})} \rho_i P_i^{\frac{1}{\alpha}}, \quad (4)$$

which quantifies the intensive and extensive margin responses respectively. With α estimated to be about 0.2 in the literature (Combes et al., 2017), this expression reflects a baseline supply elasticity of about 4 holding the amount of developed land fixed and an additional component reflecting the extensive margin associated with developing additional land or land redevelopment. For tracts in which the fixed cost of developing marginal land parcels is higher (the CDF $F_i(x)$ rises more slowly with x), the extensive margin component of the supply elasticity will be smaller.⁹

Now that we have laid out the basis for intensive and extensive margin responses, we adjust the model to incorporate all of the margins of response we observe in the data. In particular, we explicitly break the extensive margin component into redevelopment and development of undeveloped land. In addition, we want to incorporate teardowns and depreciation. Finally, we would like to capture floor-area ratio (FAR) restrictions that constrain the quantity of housing services developed per parcel to be below the preferred amount of $\rho_i P_i^{\frac{1-\alpha}{\alpha}}$.

To capture the idea that redevelopment may be costlier than development of new land, we decompose $F_i(x)$ into a component for redevelopment and a component for undeveloped land: $F_i(x) = \frac{M_i^R}{\bar{M}_i} F_i^R(x) + \frac{M_i^U}{\bar{M}_i} F_i^U(x)$, in which we think of $F_i^R(x)$ as first-order stochastically dominating $F_i^U(x)$. Second, we impose a cap on development intensity \bar{D} such that if a tract is at its development cap, the supply elasticity reflects the extensive margin only. Since in any initial period, all housing units are on developed land by definition, using the notation in (1) we have $\frac{dH_i^{R,s} + dH_i^{T,s}}{H_i} = \frac{f_i^R(\alpha \rho_i P_i^{\frac{1}{\alpha}})}{F_i^R(\alpha \rho_i P_i^{\frac{1}{\alpha}})} \rho_i P_i^{\frac{1}{\alpha}}$ and $\frac{dH_i^{U,s}}{H_i} = \frac{f_i^U(\alpha \rho_i P_i^{\frac{1}{\alpha}})}{F_i^U(\alpha \rho_i P_i^{\frac{1}{\alpha}})} \rho_i P_i^{\frac{1}{\alpha}}$.

Using this expression to extend (4) and the identity (1), we have the following unified decomposition that forms the basis for our reduced form empirical work.

⁹If $F_i(x)$ is distributed Frechet with tract-specific dispersion parameter T_i , $\frac{f_i(\alpha \rho_i P_i^{\frac{1}{\alpha}})}{F_i(\alpha \rho_i P_i^{\frac{1}{\alpha}})} = T_i (\alpha \rho_i P_i^{\frac{1}{\alpha}})^{-1-T_i}$ and the extensive margin component of the supply elasticity is $T_i \alpha^{-1-T_i} \rho_i^{-T_i} P_i^{-\frac{T_i}{\alpha}}$. Here we see explicitly how the extensive margin of supply may depend on tract topography and developed fraction if T_i depends on these objects.

$$\frac{d \ln H_i^{ef,s}}{d \ln P_i} = \frac{d \ln I_i}{d \ln P_i} \mathbf{1}(I_i < \bar{D}) + \frac{dH_i^{R,s}/H_i}{d \ln P_i} + \frac{dH_i^{T,s}/H_i}{d \ln P_i} + \frac{dH_i^{U,s}/H_i}{d \ln P_i} \quad (5)$$

If the FAR is binding, the first term is 0. The second and third terms reflect new unit construction on redeveloped land and teardowns. The final term is new unit construction on previously undeveloped land. The redevelopment elasticity $\frac{dH_i^{R,s}/H_i}{d \ln P_i}$ may be smaller than the remaining new units elasticity $\frac{dH_i^{U,s}/H_i}{d \ln P_i}$ because of additional teardown and land remediation costs. We assume that variation in depreciation of existing units across tracts in a metro area is unrelated to changes in prices.

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 provides a way to convert the quantity of housing services, which we observe only noisily, to residential population, which we observe accurately. The cost of relying on the model to make this conversion for us is that it depends on assumptions about preferences over housing and requires joint estimation of model parameters that govern labor and population supply conditions to residential and work locations respectively. For these reasons, we show results both from the housing market portion of the model, for which we measure housing quantities directly, and from more complete estimation. Both of these sets of results will use the same sources

of identifying variation. However, even the more isolated housing market estimation results rely on using the structure of the model to partial out the equilibrium impacts of labor demand shocks on local prices. Because these shocks impact housing demand conditions in multiple locations simultaneously, the model is essential for delivering a scaling of relative housing demand shifts across tracts that result from these shocks. This motivation is line with Monte, Redding, and Rossi-Hansberg (2018)’s observation that two identical productivity shocks hitting different cities may impart different treatment effects on city level outcomes because of differences in inter-city commuting and trade linkages.

Third, 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.

Finally, the model provides structure that can be used 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.

3.2.1 Setup

While our main empirical work uses data for 275 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 specified with an eye toward exploiting within metro area variation in labor demand shocks.

The model features a continuum of 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. We think of them as first choosing their residential location, anticipating the distribution of productivity shocks from which they draw, with these shocks only revealed after

committing to a residential location. In practice, the shocks primarily serve to smooth things over such that the model can generate data in which people with different utility realizations can have the same expected utility.

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

$$v_{ijk\omega} = \frac{v_{i\omega} B_i z_{ijk\omega} w_{jk}}{P_i^{1-\beta} e^{\kappa\tau_{ij}}} \quad (6)$$

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 $e^{\kappa\tau_{ij}}$ 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 .

$$F_z(z_{ijk\omega}) = e^{-z_{ijk\omega}^{-\varepsilon}}, \varepsilon > 1 \quad (7)$$

Following Tsivanidis (2018) and Couture et al. (2019), we also introduce the nested preference shock over residential locations $v_{i\omega}$. This shock is also distributed Frechet but with shape parameters η and ϕ . 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 .

$$F_v(v_{i\omega}) = \exp\left[-\sum_m \left[\sum_{i(m)} v_{i\omega}^{-\eta}\right]^{-\frac{\phi}{\eta}}\right], \phi > 1, \eta > 1 \quad (8)$$

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 $\eta = \phi = \varepsilon$, the utility shock becomes redundant and this model reduces to one more similar to that in Ahlfeldt et al. (2015).

3.2.2 Resident Market Access

The Appendix walks through the model in more detail to show that $RMA_i \equiv \sum_k \sum_j [w_{jk} e^{-\kappa\tau_{ij}}]^\varepsilon$ 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.

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

$$\bar{y}_i = \Gamma\left(1 - \frac{1}{\varepsilon}\right)(RMA_i)^{\frac{1}{\varepsilon}}. \quad (9)$$

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

$$\pi_i = \lambda \left[\sum_{i' \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} RMA_{i'}^{\frac{1}{\varepsilon}} \right)^\eta \right]^{\frac{\phi}{\eta}-1} \left(B_i P_i^{\beta-1} RMA_i^{\frac{1}{\varepsilon}} \right)^\eta. \quad (10)$$

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, and the overall attractiveness of all neighborhoods in the metro region $1/\lambda$.

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

$$\ln \pi_{ij} = a_i^g + b_j^g - (\kappa\varepsilon)\tau_{ij} \quad (11)$$

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

$$L_j = \lambda \sum_k \left[w_{jk}^\varepsilon \right] FMA_j, \quad (12)$$

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.

$$FMA_j = \sum_i \frac{e^{-\kappa\varepsilon\tau_{ij}} \pi_i}{RMA_i} \quad (13)$$

$$RMA_i = \sum_j \frac{e^{-\varepsilon\kappa\tau_{ij}} L_j}{FMA_j} \quad (14)$$

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.

RMA_i can be readily calculated (jointly with FMA_i) using our tract level data on employment, population, commuting flows and commuting time. We observe spatial distributions of employment and population in 1990, 2000, 2006 and 2010 and commute flows and times in 1990 and 2000. We estimate $\hat{\epsilon}\kappa$ using separate flow-weighted commuting gravity regressions like (11) with origin and destination fixed effects in 2000 for each metropolitan region.¹⁰ Because we do not observe tract-tract commute times after 2000, we hold commute times constant at 2000 times for the later years. This yields direct measures of RMA_i as in (14) for 2000, 2006 and 2010.

An individual who lives in i and works in j in industry k has housing demand of $(1 - \beta)^{\frac{y_i}{P_i}}$ from Cobb-Douglas preferences. We assume that all sites in each residential location i are perfect demand substitutes, justifying the uniform price per unit of housing services P_i . Adding up, the log aggregate housing demand in tract i is thus

$$\ln H_i^{ef,d} = \ln \rho_{HD} + \frac{1}{\epsilon} \ln(RMA_i) + \ln \pi_i - \ln P_i. \quad (15)$$

This object is expressed in terms of units of housing services. It is increasing in RMA 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 (10). 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 (10) into (15). 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 3.4 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 .

¹⁰Across the 306 regions in our 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. In addition, estimates of $\epsilon\kappa$ 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.

3.3 Equilibrium

We write down equilibrium conditions primarily with an eye toward showing how to use local housing demand shocks to identify of tract-specific housing supply elasticities γ_i . For aggregation of these supply elasticities to levels above the census tract, however, we need estimates of remaining parameters. The following equilibrium conditions are useful in this regard as well.

Combining conditions governing population supply to residential tracts (10), labor supply to work tracts (12) and imposing housing market clearing yields conditions describing equilibrium tract population and house prices. Differentiating the population condition over time and yields the following structural equation.

$$d \ln \pi_i = \frac{\gamma_i + 1}{\gamma_i + 1 + \eta(1 - \beta)} \frac{\eta}{\varepsilon} \left(1 - \frac{1 - \beta}{\gamma_i + 1}\right) d \ln RMA_i + v_m^\pi + u_i^\pi \quad (16)$$

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). In this equation, 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 . Mathematical details are in the appendix. We use (16) below 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.

Substituting for (16) in housing demand (15), setting it equal to housing supply (2), solving for price and differentiating yields the growth rate in tract house price, expressed as follows.

$$d \ln P_i = \frac{1}{\varepsilon(\gamma_i + 1)} + \frac{1}{\gamma_i + 1 + \eta(1 - \beta)} \frac{\eta}{\varepsilon} \left(1 - \frac{1 - \beta}{\gamma_i + 1}\right) d \ln RMA_i + v_m^P + u_i^P \quad (17)$$

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 housing 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 ρ_i , average lot size \overline{M}_i and the local

amenity B_i both 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. So, to restate, our empirical objective must be to pick out variation in $d \ln RMA_i$ that are uncorrelated with local innovations to housing productivity (costs) and lot sizes.

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 $d \ln \pi_i$.

$$\sum_{i(m)} s_i \left[1 + \frac{\phi(1-\beta)}{1+\gamma_i} \right] d \ln \pi_i - \Lambda - \frac{\phi}{\varepsilon} \sum_{i(m)} \left[s_i \left(1 - \frac{1-\beta}{1+\gamma_i} \right) \right] d \ln RMA_i = u^m \quad (18)$$

In this equation, s_i is the base year share of municipality m 's population living in tract i . As ϕ , rises dispersion in preferences across municipalities fall and so positive shocks to RMA in any neighborhoods within m result in more rapid population growth in this municipality. This is the estimation equation we use below to recover estimates of the preference parameter ϕ .

4 Empirical Implementation

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

$$\Delta \ln H_{ir}^s = \theta_r + X_{ir} \delta + \gamma_{ir} \Delta \ln P_{ir} + \tilde{\rho}_{ir} \quad (19)$$

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 .

$$\gamma_{ir} = Z_r^1 \gamma_1 + Z_{ir}^2 \gamma_2 \quad (20)$$

As is detailed in Section 2, these sources of heterogeneity are topography, land use regulation and regulatory burden.

The first two terms in (19) 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 and controls for tract-specific labor demand conditions that we worry may spill over into housing supply

factors. 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. The CBD distance controls holds 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. Finally, 1990 employment and a tract-specific Bartik labor demand shock for the 2000-2006 (explained below) ensures 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 our various measures of post-2000 home quantity changes and contemporaneous price growth. Given that the 2007-2010 period saw mostly declining housing demand and therefore may not be not a good setting for estimating housing supply, we include a specification that focuses on the 2000-2006 period. In the first row, we see negative relationships between Census Hedonic price growth and quantity growth. Remaining rows, which are for our two Zillow price indexes, show small positive coefficients of up to 0.02 for quantity measures in housing units and only up to 0.11 for our quality-adjusted quantity measure. Such small (or negative) housing supply elasticities seem implausible and help to motivate our search for more convincing identifying variation in prices.¹¹

The implausibly low relationships between the housing price growth and quantity growth point to two 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 classical endogeneity problem discussed extensively in Davidoff (2016). To the extent that this phenomenon happens dynamically, it could generate the negative observed relationship between price and quantity growth. If endogenous positive supply shifts after 2000 loosen housing development restrictions in response to demand shocks from the 1990s, observed OLS relationships will trace out demand curves and be negative. Such a story would justify the negative serial correlation in decadal price growth seen in the

¹¹Ouazad and Ranciere (2019) find similarly small OLS relationships between price growth and quantity growth for the San Francisco metro region.

data at the tract level One might also be concerned that positive productivity shocks outside of the construction sector boost local housing demand through higher household earnings and construction costs at the same time. The result is inward shifts in housing supply that are correlated with outward shifts in housing demand. 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 such measurement error would lead to a classic attenuation bias problem.

The broad message is the possibility for a local unobserved history to drive both relative price declines and more construction, inducing a downward bias in OLS. Thus a valid identification strategy needs deal with the classical endogeneity concern of simultaneity in demand and supply by comparing neighborhoods that are similar ex-ante. To resolve this issue, we must isolate a source of tract-specific housing demand shocks which are uncorrelated with supply factors. This is the role of our instrument, which we develop next.

4.2 Instrument Construction

We need a source of identifying variation in home price growth that driven by tract level housing demand shocks and is unrelated to 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.

$$\ln P_i = \tilde{\rho}_{HD} + \frac{\phi/\eta - 1}{1 + \eta - \eta\beta} \ln \sum_{i' \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} RMA_{i'}^{\frac{1}{\varepsilon}} \right)^\eta + \frac{1}{1 + \eta - \eta\beta} \frac{1 + \eta}{\varepsilon} \ln(RMA_i) - \frac{1}{1 + \eta - \eta\beta} \ln H_i^{d,ef} + \frac{\eta}{1 + \eta}$$

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 housing productivity ρ_i is endogenous to housing supply. Indeed, through its code-termination with FMA_i , RMA_i depends structurally on tract population which itself depends on tract housing productivity. As such, we develop instruments that pick out components of $d \ln RMA_i$ that are likely orthogonal to shocks to ρ_i or other factors that influence local construction costs.

To build instruments, instead of using actual employment in all commuting destinations to calculate RMA_i , we instead use the employment predicted by national growth rates and initial

industry composition in each tract to solve for RMA_i , after evenly scaling up the residential population of each tract to maintain labor market clearing and allowing us to solve jointly for FMA_j . For components of instruments, we impose 1990 commute times and initial employment shares by industry and use estimates of $\epsilon\kappa$ 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. For example, tracts near a lot of construction employment may be subject to secular changes in shocks to construction productivity that would show up as part of the instrument.

We measure the 2000 component of this instrument as:

$$\widetilde{RMA}_i^{2000} = \sum_j \frac{e^{-\epsilon\kappa\tau_{ij}^{90}} \mathbf{1}(dis_{ij} > 2km) \sum_k L_{jk}^{90} [E_k^{2000} / E_k^{1990}]}{\widetilde{FMA}_j^{2000}} \quad (21)$$

$$\widetilde{FMA}_j^{2000} = \sum_i \frac{e^{-\epsilon\kappa\tau_{ij}^{90}} \mathbf{1}(dis_{ij} > 2km) L_i^{90} \left[\frac{\sum_k \sum_j L_{jk}^{90} [E_k^{2000} / E_k^{1990}]}{\sum_j L_j^{90}} \right]}{\widetilde{RMA}_i^{2000}} \quad (22)$$

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, as we show below. This is a simulated instrument that serves the dual purpose of being a reduced form demand shock driven exogenous predictor of tract level house price growth and being a predictor of the structural object $\Delta \ln RMA_i$ that is unrelated to tract level shocks to local amenities or housing productivities. The latter use will allow us to recover estimates of structural parameters.

4.3 Instrument Validity

The fundamental source of identifying variation used is the standard tract level ‘‘Bartik’’ (1991) type shocks in each employment location, written out as follows.

$$Bartik_{jr} = \sum_k \frac{Emp_{jk}^{90}}{\sum_k Emp_{jk}^{90}} [\ln E_{r'(j)k}^{06} - \ln E_{r'(j)k}^{00}] \quad (23)$$

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.

Note that the shock component is national employment growth outside of region r in each industry k between 2000 and 2006. We choose this time interval rather than 2000-2010 because of the stronger available identification during this boom period. The 2006-2010 period was associated with employment declines that are not well predicted by Bartik instruments.

Table 3 presents evidence that our underlying identifying variation can successfully predict tract level employment growth. It presents regressions of 2000-2006 or 2000-2010 tract employment growth in tract jr on $Bartik_{jr}$ and controls for 1990 employment level, past demographic composition of tract residence, CBD distance and metro region fixed effects. This is a sort of tract level first stage regression which gets aggregated in the first stage of our main analysis. We control for past employment to isolate employment growth due only to variation in industry composition. Past demographic 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-percentage point increase in the Bartik shock predicts a 0.2% increase in 2000-2006 tract employment and a 0.6% increase in 2000-2010 tract employment. Inclusion of 2-2.5 km CBD distance ring fixed effects interacted with metro region does not affect these conclusions.¹²

One challenge we face when estimating the housing supply equation is that house price growth is negatively serially correlated across decades. The first two columns of Table 4 show this pattern. Also note the positive serial correlation in quantity growth, which likely reflects serially correlated positive demand shocks. Finally, note that quantity growth begets subsequent price declines. 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. We hope that results in the final two columns of Table 4 allay such concerns. They show that 1990-2000 price and quantity growth cannot predict subsequent growth in simulated RMA (our instrument). This suggests that IV estimates are unlikely to suffer from the sort of selection bias that we discuss above. Moreover, results in Table A1 show that our instrument does not predict 1990-2000 housing price growth for four sets of control variables and it does not predict 1990-2000 quantity growth for our primary specification.

¹²Results are also robust to lagging the demographic tract controls by one additional decade.

The top left panel in Table 5 presents the relationship between $\Delta \ln \widetilde{RMA}$ and $\Delta \ln RMA$ for 2000-2006. It shows significant estimated elasticities of 0.5-0.7 that are robust to the specification used. This is a necessary prerequisite for $\Delta \ln \widetilde{RMA}$ to predict $\Delta \ln P$ in (19) as an instrument. The next three blocks show first stage relationships between $\Delta \ln \widetilde{RMA}$ and our three primary measures of $\Delta \ln P$. First stage coefficients are strong and robust to the different specifications used. With metro area fixed effects, first stage coefficients are not significantly affected by inclusion or exclusion of lagged demographic and housing market controls, corroborating evidence from Table 4 that they are uncorrelated with these variables. Smaller first stage coefficients for the 2000-2010 relative to the 2000-2006 period reflects the fact that the 2007-2010 period mostly saw housing market declines.

The bottom two blocks in Table 5 show strongly positive reduced form relationships between our instrument and 2000-2006 and 2000-2010 Zillow new construction. We note that while these reduced form quantity responses are slightly larger over this longer time period, the two are not significantly different. We retain the longer 2000-2010 time horizon for two reason. First, we have higher quality information about some housing quantity measures for 2010 than for 2006. Second, as discussed further below, the intensive margin responses do change from 2006-2010 and we find it important to incorporate these into the analysis. Ratios of reduced form to first stage coefficients from Table 5 preview the full IV results presented in the following section.

Overall, we find that changes in simulated resident market access ($\Delta \ln \widetilde{RMA}$) strongly predict nearby employment growth, home prices, and thereby housing demand growth. Moreover, they are not correlated with pre-trends in house prices conditional on appropriate controls. These findings provide reassuring support for the use of $\Delta \ln \widetilde{RMA}$ as a valid instrument in estimating housing supply elasticities.

5 Main Results

In this section, we present our main reduced form housing supply estimates, estimating (19) by IV, and explore their heterogeneity as a function of CBD distance, initial year development intensity, topography, and regulation.

5.1 Unified Supply Elasticity Estimates

To start, Table 6 presents unified regressions of housing quantity growth on house price growth, with $\Delta \ln \widetilde{RMA}$ as an instrument for $\Delta \ln P$ using the same specification as the OLS regressions reported in Table 2. To ensure the robustness of the results, we use three measures of house price growth: hedonic price growth based on the Census data; hedonic price growth based on the Zillow data; and repeat sales price growth based on the Zillow data. And we show results using just region fixed effects in Panel A and with region-ring fixed effects in Panel B. As discussed further in Section 2.3, we also explore seven measures of house quantity changes: increase in housing stocks between 2000-2010 based on the Census data, new construction during 2000-2010 based on the ACS, new construction during 2000-2010 based on the Zillow data; new construction between 2000-2006 based on the Zillow data; increase in quality-adjusted housing units during 2000-2010 based on the Zillow data; increase in housing supply through renovation and upgrading of existing units between 2000-2010; and increase in housing supply through redevelopment only between 2000-2010. The top panel controls for the metropolitan area fixed effects, while the bottom panel further controls for the CBD distance rings within metropolitan area fixed effects. The combination of different price and quantity measures, along with different fixed effects, yields 40 supply elasticity estimates in total. Equation (5) shows how to decompose the quality-adjusted elasticity into components.

Before discussing the estimates, we highlight a few observations. First, unlike small and insignificant estimates from the OLS regressions, the estimated coefficients γ in Table 6 are positive across all specifications with magnitudes that are in line with 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, while estimated quantity responses to each measure of price changes are not significantly different, the most accurate estimates come from the Zillow hedonic price index. This is particularly so when we measure supply with changes in quality-adjusted housing units, which is also based on a hedonic approach. As such, our discussion focuses on results using this index. This is also the measure for which we have the strongest first stage F when estimating the price elasticity of our housing quantity index, reported in Column 5. Third, controlling for metropolitan area or metropolitan ring fixed effects yields similar estimates, indicating that these supply estimates are robust not

only to variation across metropolitan areas but also across different rings within each metropolitan area.

Comparisons of results in Columns 1 and 2 show the utility of allowing for longer adjustment periods. Over the longer time period, our estimated elasticity of new unit supply with respect to price more than doubles from 0.2 to 0.5. As seen in Table 5, this is primarily because of the more attenuated longer run price response, which reflects the 2007-2009 housing bust. Across the Zillow and ACS measures of unit construction, housing supply elasticity estimates are very stable (columns 2 and 3). The estimated elasticity for the census change (column 4) is not significantly different from the other two, even though this measure incorporates the negative impacts of teardowns and full depreciation. This indicates that price shocks have little impact on these two negative margins of response.

In column 5, we report quality-adjusted housing supply elasticities for the 2000-2010 period. These estimated elasticities are much larger than the unit supply elasticities shown in columns 2-4, at 3.7-4. Based on these estimates, we conclude that housing quality is an important margin of response to price changes. This can come through construction of higher quality (larger) new units or the renovation of existing units. Indeed, it's clear that treating housing units as homogenous misses an important segment of supply. A supply elasticity of 4 is exactly what would be implied by an approximately Cobb-Douglas housing production function with a land share of 0.2 as found in Combes et al. (2017). Thus our estimate of housing supply for the average U.S. tract-level housing market is consistent with the recent micro-founded housing production literature.¹³

In column 6, we focus only the quality dimension of housing supply, measured as the difference between changes in quality adjusted housing stocks and changes in the housing stock used as outcomes in Columns 5 and 4. Using Zillow-based price growth measures, we find that renovations and repairs are quite elastic with respect to price, contrary to DiPasquale (1999). Comparing the estimated γ s in column 6 with those in column 5, we find that renovation and upsizing account

¹³Empirical work on the housing production function dates back to Muth (1964). The estimated land share in the recent literature ranges from 0.10 for Centre County, PA (Yoshida, 2016), to 0.14 for Alleghany County, PA (Epple, Gordon and Sieg, 2010), to 0.2 for France (Combes *et al.*, 2016), to 1/3 for the U.S. average housing market (ranging from 0.11 to 0.48 in low to high-value areas, as reported in Albouy and Ehrlich, 2012). Ahlfeldt and McMillen (2014) provide convincing empirical support for the Cobb-Douglas functional form as a reasonable approximation to the housing production function.

for over 80% of the total supply increase. While these estimates may seem surprisingly large, they are consistent with the observations that almost half of American homeowners renovated their homes Plaut and Plaut (2010) and that total renovation expenses reached \$326 billion in 2007 Choi et al. (2014). These estimates again confirm the importance of adjusting for quality changes in measuring supply elasticity.

Column 7 reports estimated elasticities for the number of new units constructed 2000-2010 on land that was already developed in 2001. Metropolitan areas undergo a continual process of building and rebuilding. This is characterized by equation (1), which decomposes the total new construction into redevelopment of existing developed land and of undeveloped land. Total new construction is measured by the dependent variable in columns 2-3 while redevelopment is measured by the dependent variable in column 7. Thus the difference between the two captures new development. Comparing results in these columns, we find that at least one-third of newly built homes are through redeveloping the already-developed land, consistent with the findings in the urban literature that neighborhood renewal is largely driven by the deterioration and subsequent redevelopment of the existing housing stock (e.g. Rosenthal (2018); Brueckner and Rosenthal (2009)). The smaller estimated redevelopment elasticity supports the conjecture that redevelopment is costlier than development of new land.

To summarize, we find that the quality-adjusted housing supply elasticity at the tract level is about 4, consistent with an approximately Cobb-Douglas housing production function with a land share of about 0.2 as found in the recent literature. Unadjusted for quality changes, the average housing unit supply elasticity is about 0.5, one-third of which coming from redevelopment and the remaining from the development of new land.

5.2 Tract Level Heterogeneity

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

Columns 1 through 6 in Table 7 repeat the IV regressions in column 2 of Table 6 with the addition of a rich set of interactions between the price growth and tract-level factors that may influence supply. Price growth is constructed using the repeat sales index in the top panel and

the hedonic index in the bottom panel. Overall, results in Table 7 show that there is substantial within-city variation in local housing supply elasticities. The results across two measures of price growth are very consistent. As a lot of our investigation is about variation as a function of CBD distance, we focus on the specification that controls for region but not region-ring fixed effects.

In column 1, we expand the baseline supply equation by interacting $\Delta \ln P$ with the distance to CBD and its square, where the distance is measured as the percentage of way to the edge of the metropolitan region. The estimated housing supply elasticities exhibit a monotonically increasing trend with the distance to the city edge at a marginally decreasing rate. Starting from a city center, on average, a one-percentage point increase in house price increases the quantity of housing supplied by only 0.35%. This number increases to 0.83% at half of way to the city edge, and 0.89% at 90% of the way to the city edge. These micro geography level 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 a relatively smaller price response and relative greater quantity response the further away from the center one gets Genesove and Han (2013).

We now evaluate the extent to which initial land development density and local topographical features drive housing supply and whether they explain the observed CBD distance effect. In column 2, we expand the specification by adding an interaction between $\Delta \ln P$ and the 2001 fraction of land developed in each tract. All else equal, a one-percentage point increase in house price increases the quantity of housing supplied by 0.92% in tracts with no developed land, but only 0.48% in tracts where half of land has already been developed (which is half a standard deviation above the mean developed fraction of 0.38). In column 3, we further include interactions between $\Delta \ln P$ and the fraction of tract land that is flat and plains. Tracts with more of these types of topography have significantly more elastic housing supply, consistent with their easier-to-build conditions. In both columns 2 and 3, the CBD distance trend continues to be monotonically increasing at a decreasing rate, and its slope is not significantly affected by including these tract-level features. Conditional on the interactions in column 3, elevation range has no significant estimated impact on supply elasticity.¹⁴

¹⁴Exploration of additional interactions yielded low power and no significant coefficients.

Moving to columns 4 and 5, we examine the role of regulation. In column 4, we restrict the sample to the portions of 8 metropolitan areas in which we observe the average parcel size-weighted floor-area-ratio restriction for residential development at the tract level. The positive and significant estimate indicates that tracts with higher floor-area-ratios have more elastic supply. In column 5, we further include the WRLURI index measured for the municipality of the tract. The negative (but insignificant) coefficient suggests that tracts in more regulated municipalities have less elastic housing supply. These two estimates, although sensible in their direction, suffer from the usual endogeneity concern associated with regulation. Thus we do not include these two controls when imputing the tract-level supply elasticities.¹⁵

In column 6, we examine whether the estimates from the main specification in column 3 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. With the tract-level factors controlled, the metro-level supply conditions appear insignificant when we measure supply response using the repeat sales price growth (top panel) and implausibly positive using the hedonic adjusted price growth (bottom panel). However the effects of the tract-level factors are remarkably consistent and robust. We investigate effects of these metro-level factors further in the following subsection.

In column 7, we focus on redevelopment. In particular, we repeat the main specification in column 3 but with redeveloped units as the dependent variable. These estimates are simply attenuated versions of those in column 3, with this attenuation differing with CBD distance. Conditional on developed fraction and topography, the redevelopment elasticity is 31 percent of the full unit supply elasticity; at 50% of the way to the metro edge this ratio falls to 0.27; at the edge, it is only one-quarter. Once taking into account that the developed fraction increases and average land slopes increase when moving away from CBDs, this gradient becomes even steeper.

This is consistent with findings in the existing literature that prime teardowns are near public transportation and traditional village centers in Chicago Dye and McMillen (2007) and closer to the CBD and the coast in Miami Munneke and Womack (2015). Note that households attracted to redeveloped suburban housing can be quite different from central-city gentrifiers in that the former are likely to value high-quality public schools while the latter value new, larger housing

¹⁵We 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)).

relatively close to the CBD Charles (2013). In addition, land assembly for development may be easier with undeveloped parcels in more suburban areas. Together, the significant CBD distance effects, both in new construction and in redevelopment, emphasize the importance of examining housing supply through the lens of micro-geographical perspectives.

Columns 8-10 focus on quality-adjusted housing supply responses. Interestingly, the CBD distance patterns are no longer evident, as the estimates are either insignificant based on the repeat sales price growth (top panel) or noisy based on the hedonic estimates (bottom panel). Intuitively, intensive margin improvements are mainly capital rather than land intensive, reducing dependence on land availability conditions and hence CBD distance. For this reason, we treat column 10 as the main specification for quality supply where we let the supply response depend on initial land development and topographical features only. As with the unit supply response, the quality adjusted supply is more responsive when there is less developed land and when there is more flat and plain land in a given tract. These effects represent the aggregation of their respective effects on the unit supply (column 3) and on the quality dimension response holding the number of units fixed (column 11). As noted above, quality-adjusted supply is 7-10 times more price elastic than is unit supply, though we have large standard errors on our estimates of this type of supply elasticity.

5.3 Metro Level Heterogeneity

In Table 7, we examine the role of metro-level supply conditions on local housing supply elasticities. Following Saiz (2010), we construct three measures of metro-level factors: the fraction of developed land within the 50 km radius from the city centre; the fraction of area that is lost to hills, water and wetland within the 50 km radius from the city centre; the metropolitan area level WRLURI indices. To ensure the robustness of our results, we also reconstruct the first two variables within the 10 km, 20 km, 10%, 50% and 100% of the maximum radius from the CBD and for the entire metropolitan area. The results are consistent across these different measures.

Column 1 repeats the baseline CBD distance specification in Table 6. In column 2, we add an interaction between $\Delta \ln P$ and the fraction of developed land within the 50 km radius from the city centre. This coefficient is imprecisely estimated; moreover, the distance coefficients become more noisy. Depending on the size of a metropolitan area, 50 km might be far outside of the

metropolitan boundary or only cover part of the area. In light of this, Column 3 reports estimates from a regression that includes an interaction between $\Delta \ln P$ and the fraction of developed land within 50% of the maximum radius from the CBD. This coefficient is again insignificant, although including this variable does not significantly affect the CBD distance pattern in local supply elasticity. In columns 4 and 5, we run alternative regressions in which $\Delta \ln P$ is interacted with the fraction of area that is lost to hills, water and wetland within the 50 km radius from the city centre. In columns 6 and 7, we include both the metro-level initial development condition and metro-level topographical conditions. Across these specifications, none of the metro-level factors appear statistically significant in explaining variations in local supply response. In column 8, we control for an interaction between $\Delta \ln P$ and the metro-level WRLURI. The coefficient is negative and marginally significant in the top panel, consistent with the expected effects of local regulation (Saiz, 2010). However, the usual endogeneity concern applies here. In column 9, we include all the metro-level supply conditions in addition to the CBD distance variables. None of the metro-level factors appear significant for local supply elasticity.

While the previous housing supply literature has established the importance of the metro-level topographical features and regulations on the metro-level housing supply, we cannot find evidence that they explain variations in neighborhood level supply elasticities across metros of different types. We caution that our empirical setting is not well suited for this sort of analysis, as our identification strategy explicitly compares tracts in the same metro region that receive different housing demand shocks. Interactions with metro level factors use between-metro variation in housing demand shocks for identification, which our empirical setup is not set up to deliver. As such, it is not surprising that standard errors in Table 7 are large and estimates are imprecise. We do not conclude that metro level land availability conditions do not matter for housing supply in aggregate, as metros with more unavailable land have lower tract level supply elasticities which aggregate to more inelastic metro level housing supply. We demonstrate this carefully in the following section. The point remains, however, that when analyzing impacts of a local housing demand shock or a neighborhood-targeted policy, it is insufficient to use the metro-level supply elasticity. Instead, neighborhood level supply estimates are needed.

5.4 Distribution of Tract-Level Supply Elasticities

We now return to Table 6 and use the specifications in columns 3 and 10 to generate the tract-level unit supply elasticities and quality-adjusted supply elasticities, respectively. While the estimation sample for Table 6 is limited, we use coefficient estimates to predict supply elasticities for all census tracts in our data.

Table 8 provides summary statistics of these imputed elasticities. Across all tracts, mean unit supply elasticities are 0.8 for repeat-sales price growth and 1.3 for hedonic price growth. For unit supply, analogous elasticities are 4.1 and 4.4 respectively. These objects have standard deviations of 0.4-0.6 with the most supply elastic tracts having quality-adjusted elasticities of about 5. This dispersion is of similar magnitudes within metro areas. With the standard deviation of across region mean supply elasticities of only 0.1-0.3, more dispersion in supply elasticities exists within rather than between metro areas. Therefore, understanding such within region elasticity is really important.

The three metropolitan areas with lowest average tract-level hedonic index based unit supply elasticities are Jersey City (0.9), Newark (1.0) and Los Angeles (1.0), respectively. On the other end, the three metropolitan areas with highest average tract-level supply elasticities are Dothan, AL, Gainesville, FL and Ocala, FL (1.7 each), respectively. For quality-adjusted supply, Newark (3.6), Jersey City (3.8) and San Diego (3.9) have the lowest average elasticities, with Los Angeles ranked sixth, while Jacksonville, NC, Florence and Ocala (5.3 each) have the highest. There is large variation in supply elasticities within metro areas. Within Los Angeles, the unit supply elasticity ranges between 0.4 and 2.0 and quality-adjusted supply elasticities range from 3.0 to 5.7. The main cross-region source of variation in dispersion in supply elasticities comes from the bottom end of the distribution. In Ocala, unit supply elasticities range from 1.0 and 2.0 with the quality-adjusted counterpart ranging from 4.4 to 5.6.

Our results indicate the remarkably consistent pattern that local supply elasticities increase with CBD distance, some of which is attributed to initial development density and topography. Figure 1 shows that the closer to city centers, there is a higher fraction of flat, plain and developed land, thereby contributing to the CBD distance gradient for supply elasticity. On average, the fraction of tract-level 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 and Puga (2015)). Thus areas that are further away from the CBD effectively have more land in which to deliver housing, permitted both by initial development density and by topographical conditions. It is therefore tempting to conjecture that the CBD distance pattern in supply elasticities is mostly explained by the increasing availability of land as one moves away from the center.

Figure 2 shows how much this is the case. The left panel presents four measures of unit supply and the right panel presents four measures of quality-adjusted supply based on repeat sales price growth. In the left panel, the flatter curves show the imputed γ when holding the tract-level fraction of developed land and fraction of flat and plain land at each tract's respective metropolitan area mean level. It is clear that even without spatial variation in initial density and topography, there is already a significant increasing CBD distance effect in local supply elasticities. This effect may be explained by local regulations which we do not control for in our main specification due to endogeneity concerns and associated required sample restrictions. Closer to city centers, building permits and zoning restrictions may be more stringent, making supply less elastic. Almost perfectly coinciding is imputed γ when holding only the fraction of developed land at its metro-mean level, allowing for variation in topography across tracts. Local topographical features do not add much to spatial variation in supply elasticities within cities. The steeper curves plot the imputed γ when fixing only the fraction of flat and plain land at its metro-mean level, allowing initial development density to vary across tracts. These curves are much steeper and almost overlay the imputed overall γ that we discussed in the context of Table 8 (marked by the blue curve). Thus, the monotonically declining development density with CBD distance, as shown in Figure 1, explains almost all of the additional relationship between supply elasticity and CBD distance that exists beyond that implied by the regression results in Table 6 Column 3. In the right panel of Figure 2, we see more of a role for both topography and development density in explaining the increasing profile of quality-adjusted supply elasticities with respect to CBD distance.

We plot similar figures for select metropolitan areas (Log Angeles, New York, Madison and Pittsburgh) in Figure 3. Regardless 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 of developed land and, for quality-adjusted supply, topography. While there is a tendency for

supply elasticity to increase with CBD distance, this effect gets mitigated in Pittsburgh and Los Angeles by the hilly topography at the edges of these metro regions.

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 the 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, as we show above in Section 3, positive labor demand shocks tend to be disproportionately oriented toward suburban areas where housing supply is relatively elastic. This means that metro level studies may find supply elasticities that weight the suburbs above their share of metro populations (though not necessarily land areas). Indeed, 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 γ_{ir} generically aggregates to a metro region level elasticity γ_r as follows. Aggregating tract level supply growth to the metro level means taking a sum weighted by neighborhood shares of the housing stock:

$$\sum_i \frac{H_{ir}}{H_r} [d \ln H_{ir} = \gamma_{ir} d \ln P_{ir}]$$

Comparison with its region level counterpart $d \ln H_r = \gamma_r d \ln P_r$, by definition the region level elasticity is given by

$$\gamma_r \equiv \left[\sum_i \frac{H_{ir}}{H_r} \gamma_{ir} d \ln P_{ir} \right] / \left[\sum_i \frac{H_{ir}}{H_r} d \ln P_{ir} \right] \quad (24)$$

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 think of prices changes in any two neighborhoods in isolation of each other.

If demand shocks were to generate the same price growth rate in every neighborhood, the macro level elasticity would simply be the average of micro elasticities weighted by shares. However, exactly because of differing supply elasticities across locations, the same demand shock in each location would typically generate different price changes and aggregation is thus not straightforward. The following sub-section provides some examples.

6.1 Aggregation of Common Neighborhood Demand Shocks

In this sub-section we consider the case in which all neighborhoods simultaneously experience 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. With Cobb-Douglas preferences, we have that tract level expenditure share on housing is $\frac{H_{ir}P_{ir}}{Y_{ir}} = 1 - \beta$, where $Y_{ir} \equiv \bar{y}_{ir}\pi_{ir}$. Therefore, if the demand shock changes expected income by the same percentage in every neighborhood, $d \ln H_{ir} + d \ln P_{ir}$ is a constant, call it x . This could happen, for example, if a city experiences a change in the individual productivity dispersion parameter ε .

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 using (24) to aggregate over tracts, we have

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

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.

In reality, demand shocks are likely to hit some neighborhoods more than others as households move across tracts in search of lower housing costs. If this is the case, γ_r^1 would understate the true aggregate elasticity. As such, we now turn to the case in which we assume that aggregate housing demand shifts out in the city, but that people get distributed to neighborhoods only based on neighborhood-specific price increases. We further assume that that demand shock occurs in a way such that each tract has the same percentage change in price. This assumption can be justified

in the context of a spatial equilibrium condition as in Roback (1982) in which each neighborhood differs in amenities but gets hit with the same income shock. In this case, the metro-level supply elasticity is the following:

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

Tracts with higher initial housing stock are typically associated with smaller housing supply response, but these tracts are weighted more in γ_r^2 . Thus the resulting metro-level elasticity is likely to underestimate the true elasticity.

Table 9 presents summary statistics about four versions of our two aggregate supply elasticity measures γ_r^1 and γ_r^2 . Because our second aggregation scheme allows for some of residential substitution across tracts, the average of metro supply elasticities are slightly above grand tract-level averages. All region level aggregate measures of supply elasticities are significantly lower in areas with a high fraction of land developed and most are also lower in regions with more unavailable land, with the hedonic price index based supply elasticities are more responsive in expected directions to land unavailability (unreported). The key observation here is that tract-level supply influencers also aggregate up to the metro level. Correlations between our metro level supply elasticity measures and those found by Saiz (2010) are all positive but they are stronger for our more broad-based hedonic indexes at 0.33-0.38.

To this point, we have relied on arbitrary assumptions about housing demand to aggregate the tract-level supply elasticity into the metro-level supply elasticity. This approach, while useful, cannot be used to perform an equilibrium analysis of the housing market nor a corresponding welfare analysis.

6.2 Estimating the Neighborhood Demand System

We now turn to partial structural estimation of the model, which delivers parameters that govern demand substitution patterns across neighborhoods. To understand how neighborhood-specific demand shocks feed into price growth, we must first specify the housing demand system. We take this from the model in Section 3, using its structure to determine the substitution elasticities between neighborhoods.

From (10) and (15), housing demand in each tract is given by

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

Housing demand becomes more elastic as η and $\frac{\phi}{\eta}$ 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. Evident from the housing demand equation above, the elasticity of substitution between any two neighborhoods in the same municipality is $\eta(1-\beta) + 1$, increasing in η . The elasticity of substitution between any two neighborhoods in different municipalities also depends on ϕ .

We develop an “estibration” strategy for η and ϕ using the structural equations (16) and (18) from the model for neighborhood and municipality population respectively. As inputs, we use estimates of γ_{ir} discussed above, estimates of $(\kappa\varepsilon)_r$ from gravity regressions and calibrated values for β and κ . The structural equations are estimated using GMM. (16) includes fixed effects for “municipalities”. After some experimentation with different definitions, we assign 5 municipalities to each region: one for the central city and one each for suburbs in north, south, east and westerly directions. We impose that the error term in (16) depends on all of the same control variables as the regressions in Table 2 and is orthogonal to our main instrument $d \ln \widetilde{RMA}_i$. To estimate (18), we build a municipality level data set and impose that its error term is orthogonal to $\sum_{i(m)} [S_i(1 - \frac{1-\beta}{1+\gamma_i})] d \ln \widetilde{RMA}_i$, which only includes pre-determined objects. We calibrate β to 0.2 and κ to 0.005 or 0.01.¹⁶

Table 10 presents these results. The first two columns show versions in which we assume preferences are not nested, and $\phi = \eta$. The final two columns show results assuming fully nested preferences, each for different assumptions about κ . We estimate η to be between 1.6 and 3.8 and ϕ to be between 3.8 and 8.3. Incorporation of nested choices reduces dispersion in the raises the substitutibility across neighborhoods.

¹⁶ $\kappa = 0.005$ implies that 1 minute of commuting in one direction reduces full income by 0.5%. We also tried jointly estimating κ using (17) but this yielded implied values of ε that were too low. This led us to our “estibration” strategy.

6.3 Counterfactuals

Using estimates from the model, we will consider the impacts of a spatially targeted policy that boosts amenities in one municipality. We choose the Empowerment Zone (EZ) program that targeted disadvantaged census tracts with a cocktail of government programs. This segment of the analysis is in progress.

7 Conclusion

Since Rosenthal's (1999) 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 elasticity. 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 insight and present the first set of estimates of the 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 is more elastic the further out one gets from urban centers. This pattern is in part but not entirely due to a larger fraction of land available for development as one moves away from the center. Altogether, our findings indicate that initial development density, topography and zoning regimes are among the most important determinants in local housing supply.

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. (2012). *Metropolitan land values and housing productivity*. Citeseer.
- 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*.
- Bendimerad, A. (2007). *Developing a leading indicator for the remodeling industry*. Joint Center for Housing Studies of Harvard University.
- 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. (2018). Stringency of land-use regulation: Building heights in us cities.
- 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.
- Charles, S. L. (2013). Understanding the determinants of single-family residential redevelopment in the inner-ring suburbs of chicago. *Urban Studies*, 50(8), 1505–1522.
- 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. (2017). 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.).
- 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.
- 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.
- 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.
- 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.
- Housing productivity and the social cost of land-use restrictions. (n.d.). *Journal of Urban Economics*, 107, 101.
- McMillen, D., & O'Sullivan, A. (2013). Option value and the price of teardown properties. *Journal*

of Urban Economics, 74, 71–82.

- Monte, F., Redding, S. J., & Rossi-Hansberg, E. (2018). Commuting, migration, and local employment elasticities. *American Economic Review*, 108(12), 3855-3890.
- 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.
- 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.
- Yoshida, J. (2016). Structure depreciation and the production of real estate services. *Available at SSRN 2734555*.
- Zhou, Y., & Haurin, D. (2010). On the determinants of house value volatility. *Journal of Real Estate Research*, 32(4), 377-395.
- Zillow. (2017). *Ztrax: Zillow transaction and assessor dataset, 2017-q4*. (Zillow Group, Inc. <http://www.zillow.com/ztrax/>)

Table 1: Summary Statistics

| | Mean | Growth Rate | | | | Obs |
|------------------------------------------------------------------|--------|-------------|--------|--------|-------|------------|
| | Growth | Mean | St Dev | Min | Max | |
| Tract Home Price Changes (Estimation Sample) | | | | | | |
| Repeat Sales Index, 2000-2006 | | 0.89 | 0.86 | -1.77 | 69.07 | 24,773 |
| Zillow Hedonic Index, 2000-2006 | | 0.95 | 0.77 | -2.04 | 15.28 | 24,975 |
| Repeat Sales Index, 2000-2010 | | 0.37 | 0.52 | -3.20 | 17.52 | 23,877 |
| Zillow Hedonic Index, 2000-2010 | | 0.45 | 0.56 | -10.11 | 8.57 | 24,697 |
| Census Hedonic Index, 2000-2010 | | 0.01 | 0.21 | -3.63 | 2.30 | 31,263 |
| Census Hedonic Index, 1990-2000 | | 0.00 | 0.18 | -2.38 | 2.15 | 31,263 |
| Tract Housing Quantity Changes (Estimation Sample) | | | | | | |
| Stock of Housing Units, Census, 2000-2010 | 181 | 0.07 | 0.23 | -3.18 | 5.91 | 31,250 |
| Stock of Housing Units, Census, 1990-2000 | 202 | 0.14 | 0.32 | -3.94 | 4.91 | 31,124 |
| New Units, ACS, 2000-2009 | 237 | 0.11 | 0.20 | 0.00 | 5.83 | 31,242 |
| New Units, Zillow, 2000-2006 | 193 | 0.09 | 0.19 | 0.00 | 5.68 | 30,606 |
| New Units, Zillow, 2000-2009 | 232 | 0.10 | 0.21 | 0.00 | 5.90 | 30,605 |
| Stock of Quality Adjusted Housing, 2000-2006 | | 0.33 | 1.03 | 0.00 | 9.73 | 21,198 |
| Stock of Quality Adjusted Housing, 2000-2010 | | 0.53 | 1.22 | 0.00 | 10.02 | 21,187 |
| New Units on Developed Land, ACS, 2000-2010 | 87 | 0.05 | 0.11 | 0.00 | 4.93 | 31,250 |
| Stock of Housing on the Intensive Margin, 2000-2010 | | 0.47 | 1.20 | 0.00 | 9.18 | 21,180 |
| Tract Employment and Population Variables (Full Sample) | | | | | | |
| RMA, 2000-2006 (excludes 8 regions) | | 0.04 | 0.05 | -0.88 | 0.76 | 59,752 |
| Simulated RMA, 2000-2006 | | 0.05 | 0.01 | -0.02 | 0.09 | 63,897 |
| Tract Employment, 2000-2006 (excl 8 regions) | 53 | -0.16 | 0.90 | -8.96 | 6.04 | 61,359 |
| Tract Employment, 2000-2010 | 78 | -0.17 | 0.89 | -11.25 | 6.39 | 63,616 |
| Tract Level Bartik Instrument, 2000-2006 | | 0.09 | 0.05 | -0.13 | 0.25 | 63,897 |
| Tract Population, 2000-2010 | 330 | 0.04 | 0.28 | -5.56 | 7.37 | 63,578 |
| Tract Level Supply Influencers (Full Available Sample) | | | | | | |
| Fraction of Land Area Developed, 2001 | | 0.38 | 0.25 | 0.00 | 0.90 | 63,896 |
| Fraction of Land Area Developed, 2011 | | 0.39 | 0.24 | 0.00 | 0.90 | 63,896 |
| Fraction of Land Area Flat | | 0.41 | 0.42 | 0.00 | 1.00 | 63,896 |
| Fraction of Land Area Not Flat but Slope < 8% | | 0.45 | 0.37 | 0.00 | 1.00 | 63,896 |
| Elevation Range (m) | | 64 | 147 | 0 | 4325 | 63,896 |
| Wharton Real Estate Index (municipality level variation) | | 0.13 | 0.91 | -2.02 | 4.22 | 30,526 |
| Parcel Size Weighted Residential Floor Area Ratio (8 cities) | | 2.33 | 2.02 | 0.25 | 16.00 | 8,707 |
| Metro Level Supply Influencers (Full Sample) | | | | | | |
| Fraction of Area Developed | | 0.05 | 0.04 | 0.00 | 0.38 | 306 metros |
| Fraction of Metro Area Unavailable for Development | | 0.26 | 0.20 | 0.00 | 0.87 | 306 metros |
| Fraction of Area Within 50 km of CBD Unavailable for Development | | 0.23 | 0.17 | 0.00 | 0.82 | 306 metros |
| Wharton Regulation Index (2008) | | -0.08 | 0.82 | -1.76 | 2.82 | 261 metros |

The estimation sample includes all census tracts in metropolitan regions with 1990 information on employment. Tracts are excluded from the sample if 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. This is 275 metro areas. The full sample includes all 306 metro areas with 1990 information on employment at the tract level.

Table 2: OLS Results for Housing Supply

(a) Controls for Metropolitan-Area Fixed Effects

| Quantity Measure | New Units | New Units | New Units | Δ Units | Δ Quantity Index | Δ Intensive Margin | New Units, Redev. |
|---------------------------|------------------------|-------------------------|-------------------------|-------------------------|----------------------|----------------------|-------------------------|
| Source | Zillow | Zillow | ACS | Census | Zillow | Zillow & Census | Zillow & USGS |
| Time Period | 2000-2006 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Census Hedonic Index | | -0.0908*** (0.00508) | -0.0919*** (0.00478) | -0.0921*** (0.00556) | 0.0198 (0.0368) | 0.0757** (0.0364) | -0.0306*** (0.00293) |
| Zillow Repeat Sales Index | 0.00310** (0.00131) | 0.0123*** (0.00283) | 0.0140*** (0.00266) | 0.0191*** (0.00297) | 0.0537** (0.0211) | 0.0355* (0.0208) | 0.0103*** (0.00158) |
| Zillow Hedonic Index | 0.0136*** (0.00254) | 0.0161*** (0.00320) | 0.0151*** (0.00302) | 0.0104*** (0.00337) | 0.118*** (0.0230) | 0.110*** (0.0227) | 0.00903*** (0.00178) |

(b) Controls for Metropolitan-Area-Ring Fixed Effects

| Quantity Measure | New Units | New Units | New Units | Δ Units | Δ Quantity Index | Δ Intensive Margin | New Units, Redev. |
|---------------------------|------------------------|-------------------------|-------------------------|-------------------------|---------------------|----------------------|-------------------------|
| Source | Zillow | Zillow | ACS | Census | Zillow | Zillow & Census | Zillow & USGS |
| Time Period | 2000-2006 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Census Hedonic Index | | -0.0851*** (0.00520) | -0.0874*** (0.00493) | -0.0879*** (0.00573) | -0.0518 (0.0366) | 0.00193 (0.0362) | -0.0334*** (0.00306) |
| Observations | | 30,605 | 31,242 | 31,250 | 21,187 | 21,180 | 31,250 |
| # of Regions Included | | 274 | 275 | 275 | 224 | 224 | 275 |
| Zillow Repeat Sales Index | 0.000803 (0.00132) | 0.00680** (0.00291) | 0.00948*** (0.00276) | 0.0107*** (0.00306) | 0.0115 (0.0212) | 0.000683 (0.0210) | 0.00703*** (0.00164) |
| Observations | 24,771 | 23,876 | 23,870 | 23,877 | 16,407 | 16,407 | 23,877 |
| # of Regions Included | 161 | 163 | 163 | 163 | 146 | 146 | 163 |
| Zillow Hedonic Index | 0.0108*** (0.00274) | 0.0106*** (0.00351) | 0.0109*** (0.00332) | 0.00129 (0.00371) | 0.0233 (0.0237) | 0.0254 (0.0234) | 0.00623*** (0.00197) |
| Observations | 24,969 | 24,694 | 24,689 | 24,697 | 17,457 | 17,457 | 24,697 |
| # of Regions Included | 164 | 169 | 169 | 169 | 152 | 152 | 169 |

Note: Each cell corresponds to the estimate from running an OLS regression of changes in housing quantities measured as in described in the column on changes in house price index measured as described in the row. All specifications include fixed effects indicated in panel headers, a quadratic in CBD distance, log 1990 tract employment, the 2000-2006 Bartik shock for the tract and the following tract attributes measured in 1990 and 2000: house price index, rent index, log population, log avg hh income, share black, share white, share college. Indicated sample sizes in Panel B apply to results in Panel A as well.

Table 3: Tract Level Regressions of Employment Growth on Bartik Shocks

| | $\Delta \ln \text{Emp}, 2000-2006$ | | $\Delta \ln \text{Emp}, 2000-2010$ | |
|------------------------------------------|------------------------------------|----------------------|------------------------------------|-----------------------|
| Tract Level Bartik Shock for 2000-2006 | 0.247*** (0.0940) | 0.172* (0.100) | 0.636*** (0.0920) | 0.626*** (0.0980) |
| Indicator for 0 Tract Employment in 2000 | -0.00373 (0.0209) | 0.000515 (0.0243) | 0.0496** (0.0204) | 0.0648*** (0.0236) |
| Observations | 61,359 | 61,359 | 63,616 | 63,616 |
| R-Squared | 0.058 | 0.050 | 0.055 | 0.047 |
| Region Fixed Effects | 298 | | 306 | |
| Region-Ring Fixed Effects | | | 4,392 | 4,523 |

Notes: Regressions include the indicated fixed effects, a quadratic in CBD distance, log 1990 tract employment, and the following tract attributes measured in 1990 and 2000: house price index, rent index, log population, log avg hh income, share black, share white, share college.

Table 4: Relationships Between 1990-2000 and 2000-2006/10 Tract Housing Dynamics

| | $\Delta \ln \text{House Price}$ 2000-2010 | | $\Delta \ln \text{House Quantity}$ 2000-2010 | | $\Delta \ln \text{Sim. RMA}$ 2000-2006 | |
|-----------------------------------------------|----------------------------------------------|------------------------|-------------------------------------------------|-------------------------|-------------------------------------------|---------------------------|
| $\Delta \ln \text{House Price}, 1990-2000$ | -0.277*** (0.00662) | -0.301*** (0.00706) | -0.0162** (0.00653) | -0.0260*** (0.00675) | 4.81e-05 (7.49e-05) | -4.04e-05 (6.32e-05) |
| $\Delta \ln \text{House Quantity}, 1990-2000$ | 0.0302** (0.0119) | 0.0209* (0.0124) | 0.121*** (0.0109) | 0.0779*** (0.0109) | -5.75e-05 (0.000131) | -0.000219** (0.000107) |
| Lagged Demographic Controls | Yes | Yes | Yes | Yes | Yes | Yes |
| Region All FE | Yes | No | Yes | No | Yes | No |
| Region-Ring FE | No | Yes | No | Yes | No | Yes |

Notes: Lagged demographic controls are the same as in Table 2 except for the 1990 and 2000 house price and rent indexes.

Table 5: First Stage and Reduced Form Results

| | $\Delta \ln \text{RMA, 2000-2006}$ | | | | $\Delta \text{Repeated Sales Index, 2000-2006}$ | | | |
|-----------------------------------------------|-----------------------------------------------------|----------------------|----------------------|----------------------|-----------------------------------------------------|---------------------|---------------------|---------------------|
| $\Delta \ln \text{ Simulated RMA, 2000-2006}$ | 0.694*** (0.0510) | 0.639*** (0.0507) | 0.586*** (0.0532) | 0.546*** (0.0532) | 12.43*** (2.396) | 13.68*** (2.386) | 7.966*** (3.039) | 10.60*** (3.034) |
| Observations | 28,381 | 28,381 | 28,381 | 28,381 | 24,773 | 24,773 | 24,773 | 24,773 |
| R-squared | 0.010 | 0.033 | 0.006 | 0.013 | 0.009 | 0.027 | 0.001 | 0.017 |
| | $\Delta \text{Repeated Sales Index, 2000-2010}$ | | | | $\Delta \text{Census Hedonic Index, 2000-2010}$ | | | |
| $\Delta \ln \text{ Simulated RMA, 2000-2006}$ | 7.235*** (1.382) | 6.983*** (1.273) | 2.804* (1.690) | 5.019*** (1.576) | 2.074*** (0.519) | 1.818*** (0.495) | 3.480*** (0.691) | 2.138*** (0.654) |
| Observations | 23,877 | 23,877 | 23,877 | 23,877 | 31,263 | 31,263 | 31,263 | 31,263 |
| R-squared | 0.016 | 0.174 | 0.004 | 0.144 | 0.003 | 0.099 | 0.004 | 0.1147 |
| | $\Delta \ln \text{ New Units in Zillow, 2000-2006}$ | | | | $\Delta \ln \text{ New Units in Zillow, 2000-2010}$ | | | |
| $\Delta \ln \text{ Simulated RMA, 2000-2006}$ | 2.341*** (0.449) | 1.615*** (0.415) | 2.764*** (0.562) | 2.255*** (0.533) | 2.557*** (0.499) | 1.792*** (0.461) | 3.058*** (0.621) | 2.522*** (0.587) |
| Observations | 30,606 | 30,606 | 30,606 | 30,606 | 30,605 | 30,605 | 30,605 | 30,605 |
| R-squared | 0.062 | 0.207 | 0.012 | 0.122 | 0.061 | 0.209 | 0.012 | 0.125 |
| Lagged Demographic Controls | No | Yes | No | Yes | No | Yes | No | Yes |
| Region All FE | Yes | Yes | No | No | Yes | Yes | No | No |
| Region-Ring FE | No | No | Yes | Yes | No | No | Yes | Yes |

Notes: All regressions include 1990 log employment, a quadratic in CBD distance and 2000-2006 tract Bartik shocks as controls. Lagged demographic controls are the same as in Table 2.

Table 6: IV Results for Housing Supply

(a) Controls for Metropolitan-Area Fixed Effects

| Quantity Measure Source Time Period | New Units | New Units | New Units | Δ Units | Δ Quantity Index | Δ Intensive Margin | New Units, Redev. |
|-------------------------------------------|----------------------|---------------------|----------------------|---------------------|--------------------|--------------------|----------------------|
| | Zillow | Zillow | ACS | Census | Zillow | Zillow & Census | Zillow & USGS |
| | 2000-2006 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Census Hedonic Index | | 1.014** (0.415) | 0.811** (0.339) | 1.254*** (0.456) | 15.10 (12.06) | 13.28 (10.73) | 0.332* (0.173) |
| First Stage F | | 11.63 | 13.23 | 13.29 | 1.755 | 1.752 | 13.29 |
| Zillow Repeat Sales Index | 0.174*** (0.0467) | 0.397*** (0.106) | 0.327*** (0.0940) | 0.411*** (0.110) | 4.064** (2.062) | 3.623* (1.907) | 0.136*** (0.0499) |
| First Stage F | | 32.73 | 30.09 | 30.09 | 5.467 | 5.467 | 30.09 |
| Zillow Hedonic Index | 0.217*** (0.0552) | 0.510*** (0.159) | 0.405*** (0.139) | 0.542*** (0.169) | 3.718** (1.510) | 3.375** (1.421) | 0.160** (0.0715) |
| First Stage F | | 66.29 | 19.59 | 19.48 | 19.57 | 9.692 | 19.57 |

(b) Controls for Metropolitan-Area-Ring Fixed Effects

| Quantity Measure Source Time Period | New Units | New Units | New Units | Δ Units | Δ Quantity Index | Δ Intensive Margin | New Units, Redev. |
|-------------------------------------------|----------------------|---------------------|---------------------|---------------------|--------------------|--------------------|--------------------|
| | Zillow | Zillow | ACS | Census | Zillow | Zillow & Census | Zillow & USGS |
| | 2000-2006 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 | 2000-2010 |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Census Hedonic Index | | 1.273** (0.551) | 0.872** (0.391) | 1.068** (0.467) | 12.62 (9.295) | 11.58 (8.743) | 0.232 (0.179) |
| First Stage F | | 8.516 | 10.46 | 10.33 | 2.147 | 2.075 | 10.33 |
| Zillow Repeat Sales Index | 0.252*** (0.0914) | 0.631*** (0.238) | 0.504** (0.201) | 0.562** (0.223) | 28.35 (103.2) | 25.08 (91.43) | 0.166* (0.0908) |
| First Stage F | | 12.20 | 10.15 | 10.15 | 0.0759 | 0.0759 | 10.14 |
| Zillow Hedonic Index | 0.189*** (0.0501) | 0.431*** (0.126) | 0.325*** (0.110) | 0.362*** (0.124) | 3.271** (1.411) | 2.976** (1.335) | 0.114* (0.0588) |
| First Stage F | | 81.05 | 28.56 | 28.52 | 28.52 | 9.713 | 28.52 |

Note: Each cell corresponds to the estimate from running an IV regression of changes in housing quantities measured as in described in the column on changes in house price index measured as described in the row. All specifications include the 1990 and 2000 log employment, log house value, log rent, the distance to CBD and its square, 10 and 20 years lags of tract-level demographic controls such as population, average household income, the share of African American, the share of white people and the share of college graduate students. The top panel controls for the metropolitan area fixed effects and the bottom panel controls for the within-metropolitan-area ring fixed effects.

Table 7: Heterogeneity by CBD Distance and Tract-Level Supply Conditions

(a) Over Repeated-Sales Price Growth

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) |
|------------------------|-----------|-----------|---------------------|-----------|-----------|-----------|-----------|----------|-----------------------------|----------|--------------------|
| | | | Δ Quantity (Zillow) | | | | | Δ Redev. | Δ Quality-Adjusted Quantity | | Δ Intensive Margin |
| Δ ln P | 0.209* | 0.577*** | 0.548*** | -0.0987 | -0.120 | 1.142* | 0.200*** | 9.467 | 6.742 | 5.098** | 4.075** |
| | (0.111) | (0.137) | (0.134) | (0.107) | (0.0891) | (0.642) | (0.0472) | (30.39) | (12.03) | (2.257) | (1.942) |
| Δ ln P × dis | 2.116*** | 1.912*** | 2.005*** | 0.326** | 0.353*** | 2.150** | 0.493** | 68.26 | 46.58 | | |
| | (0.632) | (0.694) | (0.698) | (0.151) | (0.135) | (0.944) | (0.246) | (230.0) | (97.66) | | |
| Δ ln P × dis × dis | -0.897*** | -0.935*** | -0.956*** | -0.550* | -0.354 | -0.899* | -0.261*** | -26.12 | -18.37 | | |
| | (0.239) | (0.265) | (0.276) | (0.281) | (0.271) | (0.511) | (0.0974) | (88.47) | (38.59) | | |
| Δ ln P × tract % Dev | | -0.698*** | -0.718*** | -0.109*** | -0.126*** | -0.669*** | -0.220*** | 3.244 | -0.808** | -0.125 | |
| | | (0.0485) | (0.0576) | (0.0278) | (0.0302) | (0.0960) | (0.0203) | (8.325) | (0.332) | (0.286) | |
| Δ ln P × tract % Flat | | | 0.0408 | 0.0470 | 0.0604** | 0.0153 | 0.00767 | -4.477 | -0.963* | -1.005** | |
| | | | (0.0645) | (0.0386) | (0.0251) | (0.0782) | (0.0227) | (9.109) | (0.503) | (0.433) | |
| Δ ln P × tract % Plain | | | 0.0747** | 0.0328 | 0.0477* | 0.0573 | 0.00710 | -1.717 | -0.656* | -0.704** | |
| | | | (0.0364) | (0.0377) | (0.0270) | (0.0959) | (0.0128) | (3.334) | (0.336) | (0.290) | |
| FARRM | | | | 0.0679*** | 0.0619*** | | | | | | |
| | | | | (0.00995) | (0.00958) | | | | | | |
| WRLURI | | | | | 0.00343 | | | | | | |
| | | | | | (0.0274) | | | | | | |
| Δ ln P × MSA % Devr50 | | | | | | -1.933 | | | | | |
| | | | | | | (4.190) | | | | | |
| Δ ln P × MSA % Unavr50 | | | | | | -2.820 | | | | | |
| | | | | | | (5.981) | | | | | |
| Observations | 23,874 | 23,874 | 23,874 | 1,512 | 1,380 | 23,874 | 23,875 | 16,398 | 16,398 | 16,398 | 16,398 |

(b) Over Hedonic-Adjusted Price Growth

| | | | | | | | | | | | |
|------------------------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|----------|---------|
| Δ ln P | 0.351** | 0.915*** | 1.286*** | -0.453*** | -0.316*** | -1.274** | 0.403*** | 2.614*** | 4.715*** | 5.068** | 4.137** |
| | (0.143) | (0.251) | (0.430) | (0.174) | (0.0987) | (0.518) | (0.143) | (0.896) | (1.807) | (2.045) | (1.812) |
| Δ ln P × dis | 1.390*** | 1.330*** | 1.811*** | 0.268 | 0.362** | -0.742* | 0.384* | 5.331* | 5.886 | | |
| | (0.327) | (0.379) | (0.636) | (0.195) | (0.156) | (0.391) | (0.211) | (3.086) | (3.983) | | |
| Δ ln P × dis × dis | -0.877*** | -1.225*** | -1.509*** | 0.248 | 0.202 | -0.0839 | -0.371** | -4.426** | -6.044* | | |
| | (0.200) | (0.278) | (0.450) | (0.357) | (0.247) | (0.237) | (0.149) | (2.203) | (3.208) | | |
| Δ ln P × tract % Dev | | -0.872*** | -1.375*** | -0.00460 | -0.0726** | -0.739*** | -0.400*** | -2.776*** | -2.850*** | -1.932** | |
| | | (0.131) | (0.282) | (0.0525) | (0.0339) | (0.0985) | (0.0937) | (0.985) | (1.037) | (0.919) | |
| Δ ln P × tract % Flat | | | 0.354*** | 0.0395** | 0.0239 | 0.249*** | 0.0984*** | 0.802** | 0.911** | 0.571 | |
| | | | (0.0702) | (0.0177) | (0.0179) | (0.0521) | (0.0232) | (0.355) | (0.403) | (0.358) | |
| Δ ln P × tract % Plain | | | 0.303*** | 0.0410** | 0.0275* | 0.151*** | 0.0762*** | 0.567** | 0.584** | 0.348 | |
| | | | (0.0706) | (0.0174) | (0.0156) | (0.0340) | (0.0233) | (0.272) | (0.257) | (0.228) | |
| FARRM | | | | 0.0858*** | 0.0762*** | | | | | | |
| | | | | (0.0226) | (0.0181) | | | | | | |
| WRLURI | | | | | -0.0340 | | | | | | |
| | | | | | (0.0236) | | | | | | |
| Δ ln P × MSA % Devr50 | | | | | | 5.781*** | | | | | |
| | | | | | | (1.438) | | | | | |
| Δ ln P × MSA % Unavr50 | | | | | | 4.387** | | | | | |
| | | | | | | (1.867) | | | | | |
| Observations | 24,690 | 24,690 | 24,690 | 2,021 | 1,880 | 24,690 | 24,693 | 17,446 | 17,446 | 17,446 | 17,446 |

Note: All specifications include the 1990 and 2000 log employment, log house value, log rent, the distance to CBD and its square, 10 and 20 years lags of tract-level demographic controls such as population, average household income, the share of African American, the share of white people, the share of college graduate students and metropolitan area fixed effects. The top panel uses the repeated sales price growth and the bottom panel uses the hedonic price growth.

Table 8: Heterogeneity by CBD Distance and Metro-level Supply Conditions

(a) Over Repeated-Sales Price Growth

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
|------------------------|---------------------|---------|-----------|----------|-----------|---------|----------|-----------|---------|
| | Δ Quantity (Zillow) | | | | | | | | |
| Δ ln P | 0.209* | 21.74 | 0.943 | 1.873 | 0.284 | 2.148 | 0.849 | -1.133*** | 1.927 |
| | (0.111) | (199.9) | (0.573) | (2.748) | (0.521) | (11.85) | (0.672) | (0.412) | (4.627) |
| Δ ln P × dis | 2.116*** | 6.633 | 2.509*** | 3.373 | 2.122*** | 4.719 | 2.508** | 2.817*** | 1.197 |
| | (0.632) | (39.27) | (0.972) | (2.458) | (0.616) | (41.92) | (1.015) | (0.435) | (6.687) |
| Δ ln P × dis × dis | -0.897*** | 0.303 | -1.083*** | -0.886** | -0.874*** | -1.196 | -1.120** | -1.228*** | -0.395 |
| | (0.239) | (12.72) | (0.377) | (0.438) | (0.296) | (8.932) | (0.529) | (0.196) | (1.268) |
| Δ ln P × MSA % Dev50 | | -119.9 | | | | 5.653 | | | -11.52 |
| | | (1,073) | | | | (166.8) | | | (17.22) |
| Δ ln P × MSA % Devr50 | | | -5.820 | | | | -5.938 | | |
| | | | (3.694) | | | | (4.292) | | |
| Δ ln P × MSA % Unav50 | | | | -13.47 | | -21.07 | | | -6.645 |
| | | | | (22.27) | | (254.2) | | | (64.88) |
| Δ ln P × MSA % Unavr50 | | | | | -0.580 | | 0.840 | | |
| | | | | | (4.098) | | (5.864) | | |
| WRLURI | | | | | | | | -0.511* | 0.879 |
| | | | | | | | | (0.266) | (3.124) |
| Observations | 23,874 | 23,874 | 23,874 | 23,874 | 23,874 | 23,874 | 23,874 | 23,874 | 23,874 |

(b) Over Hedonic-Adjusted Price Growth

| | Δ Quantity (Zillow) | | | | | | | | |
|------------------------|---------------------|-----------|-----------|----------|---------|----------|---------|-----------|---------|
| Δ ln P | 0.351** | -0.400 | -0.502* | -0.494 | 1.935 | -0.632 | 1.313 | -0.688*** | -1.024 |
| | (0.143) | (0.484) | (0.260) | (0.578) | (9.581) | (0.585) | (6.672) | (0.188) | (0.803) |
| Δ ln P × dis | 1.390*** | 2.452*** | 2.430*** | 3.116** | 6.762 | 2.908** | 6.407 | 2.551*** | -0.061 |
| | (0.327) | (0.533) | (0.443) | (1.250) | (16.20) | (1.171) | (13.77) | (0.638) | (1.045) |
| Δ ln P × dis × dis | -0.877*** | -0.821*** | -0.879*** | -1.501** | -4.055 | -1.557** | -3.670 | -1.018*** | -1.190 |
| | (0.200) | (0.149) | (0.177) | (0.728) | (10.69) | (0.607) | (8.555) | (0.354) | (0.823) |
| Δ ln P × MSA % Dev50 | | -1.254 | | | | 3.108 | | | 2.210 |
| | | (3.608) | | | | (2.864) | | | (2.142) |
| Δ ln P × MSA % Devr50 | | | -0.910 | | | | -7.127 | | |
| | | | (1.335) | | | | (25.84) | | |
| Δ ln P × MSA % Unav50 | | | | -5.670 | | -4.797 | | | 4.606 |
| | | | | (3.929) | | (3.794) | | | (6.015) |
| Δ ln P × MSA % Unavr50 | | | | | -36.24 | | -32.10 | | |
| | | | | | (107.7) | | (85.31) | | |
| WRLURI | | | | | | | | -0.492 | 0.523* |
| | | | | | | | | (0.632) | (0.281) |
| Observations | 24,690 | 24,690 | 24,690 | 24,690 | 24,690 | 24,690 | 24,690 | 24,690 | 24,690 |

Note: All specifications include the 1990 and 2000 log employment, log house value, log rent, the distance to CBD and its square, 10 and 20 years lags of tract-level demographic controls such as population, average household income, the share of African American, the share of white people, the share of college graduate students and metropolitan area fixed effects. The top panel uses the repeated sales price growth and the bottom panel uses the hedonic price growth.

Table 9: Tract Level Supply Elasticities (γ_i)

| | 5% | 25% | 50% | 75% | 95% | Mean | S.D | obs |
|-----------------------------------------|------|------|------|------|------|------|------|--------|
| Unit Supply Elasticity (RS) | 0.17 | 0.46 | 0.78 | 1.12 | 1.54 | 0.80 | 0.42 | 63,896 |
| Unit Supply Elasticity (HI) | 0.62 | 0.99 | 1.30 | 1.59 | 1.86 | 1.28 | 0.38 | 63,896 |
| Quality-Adjusted Supply Elasticity (RS) | 3.53 | 3.79 | 4.05 | 4.31 | 4.75 | 4.07 | 0.36 | 63,896 |
| Quality-Adjusted Supply Elasticity (HI) | 3.32 | 3.96 | 4.46 | 4.88 | 5.37 | 4.41 | 0.62 | 63,896 |

Table 10: Summary Statistics of Imputed Metro-Level Supply Elasticities

| | Mean | Std. Dev. | Min | Max | # Obs | Corr. ρ |
|--------------------------------|--------|-----------|--------|---------|-------|--------------|
| Less Neighborhood Substitution | | | | | | |
| Unit Supply Elasticity (RS) | .8635 | .1497 | .3541 | 1.2535 | 306 | 0.1701** |
| Unit Supply Elasticity (HI) | 1.4169 | .1395 | .8182 | 1.7681 | 306 | 0.3320** |
| Quality Supply Elasticity (RS) | 4.1959 | .2278 | 3.7521 | 4.8511 | 306 | 0.1267* |
| Quality Supply Elasticity (HI) | 4.7254 | .2693 | 3.6086 | 5.3613 | 306 | 0.3724** |
| More Neighborhood Substitution | | | | | | |
| Unit Supply Elasticity (RS) | .9428 | .1509 | .4056 | 1.3002 | 306 | 0.1837** |
| Unit Supply Elasticity (HI) | 1.4558 | .1326 | .8699 | 1.7822 | 306 | 0.3430** |
| Quality Supply Elasticity (RS) | 4.2067 | .2297 | 3.7550 | 4.8607 | 306 | 0.1168* |
| Quality Supply Elasticity (HI) | 4.7599 | .2616 | .8699 | 5.3733 | 306 | 0.3768** |
| Saiz Supply Elasticity | | | | | | |
| Saiz Supply Elasticity | 2.6093 | 1.4643 | .5953 | 12.1480 | 236 | 1 |

ρ shows the pairwise correlation between the imputed metro-level elasticities and the Saiz measure of metro-level housing supply elasticity.

** indicates 10%; * indicates 1% level significance for the null hypothesis of independency between γ and the Saiz elasticity.

Table 11: Heterogeneity in Metro-level Supply Elasticities**(a) Less Neighborhood Substitution**

| | (1) | (2) | (3) | (4) |
|---------------|----------|----------|----------|----------|
| | Unit | Unit | Quality | Quality |
| | (RS) | (HI) | (RS) | (HI) |
| MSA % Devr50 | 0.30 | -0.66*** | -0.65*** | -2.56*** |
| | (0.19) | (0.16) | (0.25) | (0.27) |
| MSA % Unavr50 | 0.02 | -0.17*** | 0.58*** | -0.48 |
| | (0.05) | (0.04) | (0.07) | (0.07) |
| WRLURI | -0.03*** | -0.05*** | -0.04** | -0.08*** |
| | (0.01) | (0.01) | (0.02) | (0.03) |
| Observations | 306 | 306 | 306 | 306 |

(b) More Neighborhood Substitution

| | (1) | (2) | (3) | (4) |
|---------------|---------|----------|----------|----------|
| | Unit | Unit | Quality | Quality |
| | (RS) | (HI) | (RS) | (HI) |
| MSA % Devr50 | 0.22 | -0.66*** | -0.65*** | -2.51*** |
| | (0.19) | (0.16) | (0.25) | (0.27) |
| MSA % Unavr50 | 0.01 | -0.17*** | 0.61*** | -0.49*** |
| | (0.05) | (0.04) | (0.07) | (0.07) |
| WRLURI | -0.03** | -0.04*** | -0.04** | -0.07*** |
| | (0.01) | (0.01) | (0.02) | (0.02) |
| Observations | 306 | 306 | 306 | 306 |

Note: The dependent variable is the imputed metro-level supply elasticity. MSA % Devr50 indicates the fraction of developed land within 50% of maximum radius from the CBD; MSA % Unavr50 indicates the fraction of area that is lost to hills, water and wetland within 50% of maximum radius from the CBD; WRLURI indicates the metropolitan-area-level WRLURI indices.

Table 12: Demand Substitution Parameters from Estimation

| | (1) One layer choice | (2) One layer choice | (3) Two layer nested choice | (4) Two layer nested choice |
|-----------------|----------------------------|----------------------------|-----------------------------------|-----------------------------------|
| κ | 0.005 | 0.01 | 0.005 | 0.01 |
| η | 2.94* (1.65) | 1.63** (0.72) | 3.83*** (1.35) | 2.08*** (0.54) |
| ϕ | | | 8.72*** (0.61) | 3.83*** (0.24) |
| ϵ min | 2.23 | 1.12 | 2.24 | 1.12 |
| ϵ max | 21.40 | 10.69 | 21.4 | 10.7 |
| ϵ mean | 5.59 | 2.80 | 8.1 | 4.05 |

Standard errors in parentheses

Figure 1: CBD Distance Patterns in Tract-Level Supply Conditions

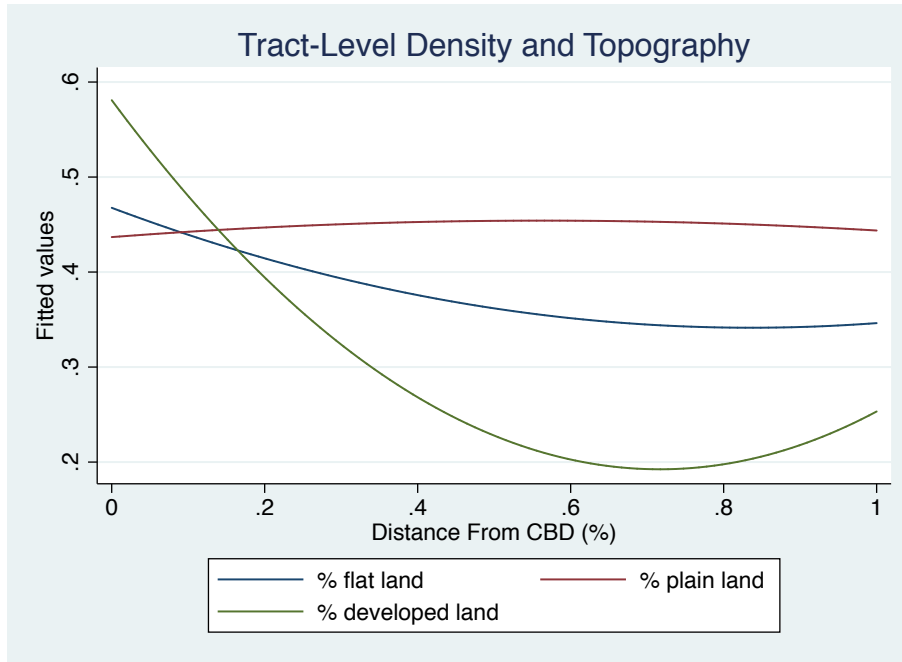


Figure 2: CBD Distance Patterns in Supply Elasticities Across All Tracts


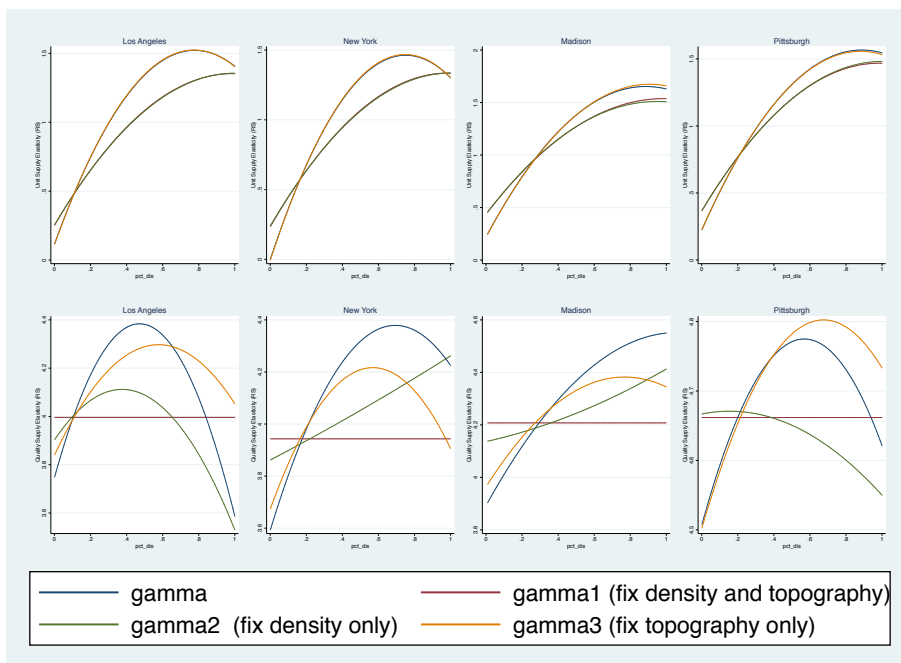


Figure 3: CBD Distance Patterns in Selected MSAs



A Model Appendix

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

$$\pi_{ij|i} = \frac{\sum_k [w_{jk} e^{-\kappa\tau_{ij}}]^\varepsilon}{\sum_k \sum_{j'} [w_{j'k} e^{-\kappa\tau_{ij'}}]^\varepsilon} \equiv \frac{\sum_k [w_{jk} e^{-\kappa\tau_{ij}}]^\varepsilon}{RMA_i} \quad (25)$$

We write this expression as a function of resident market access $RMA_i \equiv \sum_k \sum_j [w_{jk} e^{-\kappa\tau_{ij}}]^\varepsilon$, 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,

$$\bar{y}_i = \Gamma\left(1 - \frac{1}{\varepsilon}\right) (RMA_i)^{\frac{1}{\varepsilon}} \quad (26)$$

Not surprisingly, 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 m is the highest times the probability that neighborhood i is the highest utility neighborhood in municipality m . Using properties of the Frechet 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 $\frac{\sum_{i' \in m(i)} (B_{i'} P_{i'}^{\beta-1} \bar{y}_{i'})^\eta \frac{\phi}{\eta}}{\sum_m [\sum_{i' \in m(i)} (B_{i'} P_{i'}^{\beta-1} \bar{y}_{i'})^\eta \frac{\phi}{\eta}]}$. 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.

Table 13: Analysis of Pre-Treatment Trends

| | Δ Census Hedonic Index 1990-2000 | | | |
|----------------------------------|-----------------------------------------|---------|----------|----------|
| $\Delta \ln$ Sim. RMA, 2000-2006 | 0.360 | 0.277 | -0.294 | -0.361 |
| | (0.448) | (0.431) | (0.585) | (0.565) |
| Observations | 31,263 | 31,263 | 31,263 | 31,263 |
| R-Squared | 0.003 | 0.083 | 0.002 | 0.073 |
| | $\Delta \ln$ House Quantity, 1990-2000 | | | |
| $\Delta \ln$ Sim. RMA, 2000-2006 | 3.117*** | -0.109 | 2.402*** | -0.679** |
| | (0.720) | (0.248) | (0.916) | (0.334) |
| Observations | 31,124 | 31,124 | 31,124 | 31,124 |
| R-Squared | 0.093 | 0.893 | 0.037 | 0.872 |
| Region All FE | 275 | 275 | | |
| Region-Ring FE | | | 3,235 | 3,235 |
| Demographic Controls | No | Yes | No | Yes |

Notes: Each coefficient is from a separate reduced form regression of the variable listed at top on the change in \ln Simulated RMA between 2000 and 2006 and the indicated controls. All regressions include tract level Bartik shocks for 2000-2006, a quadratic in CBD distance and \ln 1990 tract employment. Demographic controls are indicated in the notes to Table 2.