The Microgeography of Housing Supply

Nathaniel Baum-Snow[*] Lu Han[†]

October 6, 2022

Abstract

We perform a comprehensive neighborhood level analysis of housing supply. Predictions of floorspace and housing unit supply elasticities using our estimates average 0.5 and 0.3 across all urban neighborhoods in the US, exhibiting greater variation within than between metro regions. New construction accounts for about 50 percent of unit supply responses, with important additional roles for teardowns and renovations. Supply responses grow with CBD distance, mostly from the increasing availability of undeveloped land, flatter land, and less regulation. Identification comes from variation in labor demand shocks to commuting destinations, as aggregated using insights from a quantitative spatial equilibrium model.

---

[*] Rotman School of Managment, University of Toronto, nate.baum.snow@rotman.utoronto.ca.

[†] Wisconsin School of Business, University of Wisconsin, lu.han@wisc.edu.

1 Introduction

Quantification of housing supply elasticities at a microgeographic scale is required to understand and analyze a wide range of phenomena that involve neighborhood level variation in housing demand within cities and regions. Targeted neighborhood investment for economic development (Busso, Gregory, & Kline, 2013; Hanson, 2009), new transportation infrastructure (Severen, 2019), changes in labor demand conditions (Fogli & Guerrieri, 2019), and changes in local amenities and public goods (Couture, Gaubert, Handbury, & Hurst, 2019; Baum-Snow & Hartley, 2020; Calabrese, Epple, & Romano, 2011) all induce shifts in housing demand that vary across neighborhoods. The extent to which these changes affect the welfare of renters versus owners depends crucially on neighborhood level estimates of housing supply elasticities. In addition, such elasticities are central for evaluating the efficacy of housing affordability policies (Favilukis, Mabille, & Van Nieuwerburgh, 2019; Davis, Gregory, & Hartley, 2019), understanding spatial variation in booms and busts within metro areas (Glaeser, Gottlieb, & Tobio, 2012; Guerrieri, Hartley, & Hurst, 2013; Genesove & Han, 2013), and the extent to which urban growth takes the form of densification or sprawl (Glaeser, Gyourko, & Saks, 2005). While the existing literature documents large differences in housing supply elasticities between cities (Saiz, 2010), little evidence exists on how housing supply elasticities differ within cities. In this paper, we provide the first comprehensive examination of housing supply at the neighborhood level, facilitating quantitative analysis of a wide range of neighborhood level phenomena and place based policies.

This paper conceptually and empirically examines housing supply and its components for all residential census tracts in US metropolitan areas. Our investigation delivers new evidence on how floorspace supply and housing unit supply decompose into teardowns, renovation of existing buildings, and new construction on newly developed and redeveloped land. Moreover, we quantify how each of the associated supply elasticities differ within cities as functions of distance to the center, land availability, building density, and zoning restrictions. As in Severen (2019), we use Bartik (1991) type variation in labor demand shocks to commuting destinations within metro areas for identification. Insights from a quantitative spatial equilibrium model help to conceptualize this identification strategy. To account for differences between metro areas, our analysis incorporates a finite mixture model (FMM) empirical setup, in which parameters that govern tract supply elasticities are allowed to flexibly differ between metro areas as functions of metro land unavailable for development, regulation, and developed land. Central to our analysis are tract level measures of housing supply components and price indices for housing services that are newly constructed using the Zillow ZTRAX assessment and transactions data sets (Zillow, 2017), married with newly organized information on land cover. Our analysis is carried out in changes for the 2000-2010 period.

Across urban census tracts in our estimation sample, we estimate an average floorspace supply elasticity of 0.42, an average housing units supply elasticity of 0.35 and an average land development elasticity of 0.09. Within floorspace and units supply, we separate out responses to demand shocks due to new construction, reductions in teardowns and renovation of existing buildings. We find that new construction accounts for 55 percent of unit supply and 69 percent of floorspace supply responses, with the remainder split roughly evenly between reduced teardowns and expansion or reconfiguration of existing structures. Only a small fraction of this new construction response comes on parcels that are already developed. The relative responsiveness across these different supply margins delivers a rich picture about shapes and densities of different neighborhoods in growing cities and sheds new light on the efficacy of policies that aim to increase supply to improve affordability.

We uncover striking within metro area differences in neighborhood-level housing supply elasticities as functions of location, available land, topography and regulation. Land development, units and floorspace supply responses all grow with distance from central business districts (CBDs), flatten out in suburban areas and then grow again at urban fringes. At CBDs, new construction accounts for a smaller share of these units and floorspace elasticities than in the average tract. Corresponding suburban elasticities are similar to the overall average supply elasticities cited above. Positive CBD distance profiles for supply elasticities are mainly driven by the fact that the fraction of land that is initially developed decreases moving away from CBDs. Tracts with more flat land and less stringent regulations also exhibit more elastic supply. Teardowns and renovation of existing units are in general not significantly related to CBD distance and developed fraction, though price growth spurs more new construction to replace teardowns on flatter land.

Using parameters estimated with data from about 50 percent of census tracts nationwide in the US, we predict supply elasticities for all 50,410 tracts in 306 metro regions. Looking across all urban census tracts nationwide and accommodating both tract and metro region variation in factors that influence supply elasticities, predicted elasticities range from 0.14 to 0.44 for for unit supply and 0.33 to 0.70 for floorspace supply in 25th and 75th percentile neighborhoods. Standard deviations are 0.20 and 0.26, respectively. Among the predicted elasticities mentioned above, within-metro area variation exceeds the corresponding variation between metro areas. We confirm evidence in Saiz (2010) that supply elasticities are decreasing in metro area land unavailable for development, existing development intensity, and regulation, even conditional on these tract level influencers.

We approach recovery of neighborhood level housing supply elasticities as the fundamentally reduced form problem of identifying coefficients in regressions of changes in tract level housing quantities on changes in a tract level home price index. The reduced form estimation is microfounded on a stylized model of neighborhood new construction supply that explicitly distinguishes between land development and intensity of construction conditional on development. The housing production literature, notably Ahlfeldt and McMillen (2014) and Combes, Duranton, and Gobillon (2021), mostly focuses on developers' choices of capital intensity conditional on land development, leaving out consideration of the parcel selection margin. Our reduced form estimates imply new construction floorspace relative to land development supply elasticities of between 2 and 4, consistent with recent land share estimates of less than 0.34 in the context of standard models of housing production. In our model, the extensive land development margin of supply also plays an important role, and depends crucially on the availability of developable sites at sufficiently low fixed costs, as shaped by topography, existing development, and regulation. In complementary work, Murphy (2018) structurally estimates a dynamic model of housing supply that predicts the timing and intensity of single-family home construction on undeveloped lots as a function of the path of expected future prices for housing services. Like in our analysis, his model accommodates cross-sectional variation in construction costs to fit intensive and extensive margin price responses. We stress, however, that at least 30 percent of new units and floorspace supply in the average tract come from endogenous teardown and renovation responses to price shocks. These are forces not captured by standard housing production or supply models, suggesting an important role for model-agnostic estimates that incorporate all margins of supply response.

The central challenge in identifying housing supply elasticities is to find an exogenous source of variation that shifts neighborhood level housing demand but not local construction costs, land use regulations, parcel size or land availability. 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, we use Bartik 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 home price growth across residential locations. These labor demand shocks are built using 1990 industry shares in commuting destinations interacted with national industry-specific employment growth rates after year 2000. We follow Tsivanidis (2022) and nest our reduced-form estimation problem into an urban quantitative 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, adjusted for labor supply competition effects from other commuting origins. 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. Construction of this simulated instrument uses predicted employment and population growth in locations beyond 2 km of tract i using 1990 tract employment shares by industry excluding construction interacted with post-2000 national employment growth by industry. Structural estimation of this demand model yields parameters of the neighborhood demand system.

We aggregate our neighborhood level housing supply elasticity estimates to the metro area level, accounting for the fact that the aggregation scheme is specific to the nature of demand shocks and to the degree of housing demand substitutability across neighborhoods. Aggregation of broad based neighborhood level demand shocks delivers average metro area level supply elasticities that are considerably smaller than those found in Saiz (2010). Allowing for a high amount of demand substitution across tracts implies an average of region level unit elasticity of 0.41 with a standard deviation of 0.11. These estimates have a rank correlation of 0.49 with Saiz' estimates. We show that our smaller elasticities come mostly because of the later and shorter time period of study than Saiz' 1970-2000 analysis. In addition, the nature of demand shocks used for identification and aggregation may be important. As neighborhoods become stronger demand substitutes, a shock affecting labor market opportunities in one location affects housing demand in a wider range of areas, as households are more willing to substitute across residential options to take advantage of lower housing prices in some places, thereby opening up more opportunities for supply elastic neighborhoods to be included. However, we note that because of the limited flexibility of parameters to differ across metros, we are somewhat less confident about levels of supply elasticities calculated using our estimates than their within-metro dispersion.

Our microgeographic evidence on all segments of housing supply facilitates improved quantitative evaluation of a wide range of neighborhood targeted policies. As an example application, we explore the welfare consequences of the Opportunity Zone (OZ) provisions of the 2017 "Tax Cuts and Jobs Act." The OZ program targets about one-quarter of low income census tracts with reduced capital gains taxes on new real estate investments. The resulting lower cost of capital associated with new construction in these neighborhoods is reflected in reduced marginal costs and outward (downward) shifts in neighborhood supply functions. The OZ program may also spur improvements in local amenities, thereby boosting local residential demand. Our analysis reveals that because OZ neighborhoods have among the lowest local supply elasticities in their metro areas, welfare gains from the program are smaller than if the program were implemented in almost any other neighborhood. In particular, we show that the potential gains in consumer surplus from implementing the same tax incentive in non-OZ neighborhoods is greater by about $5 million per tract on average. Moreover, these OZ tract gains are overstated by 34% if calculated using regional rather than tract level supply elasticity estimates. Gains in producer surplus from demand increases are also smaller in OZ tracts than they would be for the same percentage demand growth in quantity terms for other types of locations. While OZ tracts have many low income residents, they also have very inelastic supply, thereby limiting welfare gains from the program.

Our empirical work delivers magnitudes of regional supply elasticity that are in line with other recent evidence. Gorback and Keys (2020) uses variation in international capital flows to ethnic neighborhoods to identify short run local unit supply elasticities that average 0.1 across the largest 100 US metro areas. Cosman, Davidoff, and Williams (2018) confirm our evidence with the help of a calibrated dynamic theory that housing supply elasticities are increasing in the availability of buildable land at each CBD distance. Using identifying variation from foreign born residents and fertility rates across Swiss municipalities and cantons, von Ehrlich, Schöni, and Büchler (2018) shows, like us, that housing supply elasticities depend on geography and regulatory constraints. While estimated Swiss owner-occupied housing supply elasticities are similar in magnitude to our evidence for the US, rental supply elasticities are considerably higher. Finally, Orlando and Redfearn (2021) use a structural VAR to nonparameterically measure housing supply elasticities in each county nationwide in the US. They find elasticities that are quantitively similar to those discussed in this paper and declining over time. Some of this work uses housing units while other papers use housing starts as supply measures. The housing production literature uses floorspace or latent housing services in new construction instead as supply measures. Our paper helps to unify these literatures by systematically characterizing all margins of response of supply to price shocks in a unified way. In addition, we introduce an identification strategy that can be employed in many settings to structurally estimate quantatitative spatial models.

2 Data

We compile information on housing and labor markets at the census tract level for all metropolitan areas in the US. Using the the Zillow (2017) Assessor and Real Estate Database (ZTRAX) data files, supplemented with aggregate census and American Community Survey (ACS) data from 1990, 2000, 2010 and 2008-12, we construct various housing price and quantity measures. We measure local labor demand conditions using the place of work and journey to work tabulations in the 1990 and 2000 US Censuses of Population and the 2006 and 2010 LODES data plus census tract aggregate data from 1990 to 2010. Finally, we use remote sensing information on land cover in 2001 to measure baseline tract development intensity and topography and 2011 to construct changes in tract developed land. All data are keyed to 2000 definition census tracts, covering 50,410 unique tracts and 63,897 observations across 306 metro areas (with some spatial overlap across metros). The use of census tract geography allows us to conceptualize a data generating process with a uniform price per efficiency unit of housing services within each location, while also having sufficient transactions information to be able to construct a price index covering a large set of locations.[1] All statistics reported in this section apply to the estimation sample of

---

[1] For robustness analyses, we also construct a parallel data set comprising spatial units of 10-19 tracts each.

24,532 equally weighted tracts and 30,840 observations in 169 metro regions.

Below we describe in more detail how we process each data source.

2.1 Housing Prices

Our primary source for housing data is Zillow's Assessor and Real Estate Database (ZTRAX) (Zillow, 2017). This comes in the form of files for transactions, most recent assessments before 2017, and prior assessments. These data cover more of the US over time from 2000 to 2010, from coverage of at least part of 269 metro areas and about two-thirds of sample tracts to three-quarters by 2010. Because of incomplete coverage, particularly in year 2000, we supplement Zillow data with decennial Census data, as explained further below.

Transactions information is transcribed from local Recorders of Deeds and includes the sale price, location and some property attributes. To fill out property attributes, we merge in 2015 assessment data. We primarily use the resulting data set to construct price indexes at the census tract level. For the purpose of building home price indexes, we only use arm's length transactions for resale or new construction. This excludes deed transfers such as bank foreclosures and quitclaim deeds. We include all residential units, including single family houses, townhouses, and condominiums. We only consider individual buyer transactions, excluding those involving institutional buyers. We exclude homes that sell more than 9 times over our sample period and tract-year combinations with fewer than 10 sales.[2]

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

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

Here, h indexes homes in census tract i, region r, year t and month m. X_{hirtm} includes a rich set of characteristics, including quadratics in property age and log square feet plus factorization of Zillow's coding of property condition, quality and style of construction, total rooms, bathrooms,

---

[2] This second restriction drops 8% of tract-year observations in 2000 and 14% in 2010, with a minor impact on results.

bedrooms, stories, roof type, heating type, number of fireplaces, wall construction, and water pipe type. (Missing values are treated as distinct factors.) Month of sale fixed effects ρ_m^{HI} flexibly account for seasonality in market conditions.

To fill out a measure of home prices for tracts with incomplete ZTRAX coverage and to facilitate a 1990-2000 pre-trends analysis, 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:

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

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

To account for unobserved heterogeneity in quality across homes, we also use the ZTRAX data set to build a repeat sales index (RS) at the tract-year level. For this index, we exclude any sales fewer than 180 days after the prior sale. Inclusion of home fixed effects α_{hir}^{RS} in the following regression purges individual home heterogeneity that is fixed over time. Tract-year fixed effects a_{irt}^{RS} from this regression form our repeat sales index:

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

After homes are renovated, we treat them as new homes for the purpose of constructing this index. We recognize that this index may suffer from a less representative sample than the hedonic index and incorporate unwanted capitalization of unobserved home improvements.[3]

The top block in Table 1 presents summary statistics about changes in these three home price indexes for the primary estimation sample used in the empirical work. The Zillow hedonic price index growth is 0.62 on average across tracts, relative to 0.64 for repeat sales index growth during

---

[3] We also consider the tract level Federal Housing Financing Agency (FHFA) repeat sales price index. The FHFA index only covers single family house transactions involving conforming and conventional mortgages. For this reason, we focus on Zillow-based measures as they have more complete geographic coverage and include condominium sales. The correlation between the 2000-2010 growth rates of the two repeat sales indices is 0.70.

the 2000-2006 period. For 2000-2010, average growth rates are 0.25 for each, reflecting the 2007-2008 housing market crash. The correlation between the two Zillow indexes is 0.92 for the 2000-2010 period but those with the census index are only about 0.43 for both Zillow based indexes. Due to the slightly better coverage of the RS index, we use this as our primary price measure throughout the analysis.

2.2 Housing Quantities

We construct a range of housing stock and flow measures to facilitate unified decompositions of tract level residential floorspace supply responses into various components.

We begin by noting that the total residential floorspace in census tract i, S_i, is the amount of developed land L_i times the average number of housing units per developed parcel H_i/L_i times the average floorspace per housing unit S_i/H_i. Differentiating over time,

$$\Delta \ln S_i = \underbrace{\Delta \ln L_i + \Delta \ln}_{\Delta \ln H_i} \underbrace{\frac{H_i}{L_i} + \Delta \ln \frac{S_i}{H_i}}_{\Delta \ln A_i}. \tag{1}$$

The first two terms add up to changes in housing units, $\Delta \ln H_i$ and the second two terms add up to changes in floorspace per land $\Delta \ln A_i$. $\Delta \ln A$ is useful to define, as it relates to the housing production function.

We further decompose changes in units and floorspace into those from new construction, teardowns and renovations. For units, $\Delta \ln H_i = \frac{H_i^R}{H_i} + \frac{H_i^U}{H_i} + \frac{H_i^T}{H_i} + \frac{H_i^E}{H_i}$, where H_i^R refers to new units developed on already-developed land (redevelopment), H_i^U refers to new units developed on land that has not been developed before (new development), H_i^T refers to the combination of full depreciation and teardowns and is always negative, and H_i^E refers to the loss or the addition of units due to reconfiguration within existing buildings (renovation). Correspondingly, we decompose changes in floorspace $\Delta \ln S_i$ into $\frac{S_i^R}{S_i} + \frac{S_i^U}{S_i} + \frac{S_i^T}{S_i} + \frac{S_i^E}{S_i}$. While the total new construction responses include both H_i^R/H_i (S_i^R/S_i) and H_i^U/H_i (S_i^U/S_i), we separate out redevelopment (the former) from new construction on undeveloped land (the latter) as they may have different cost structures.

The primary data source used to construct all housing quantity measures is the ZTRAX current and historical assessment files. With varying coverage rates across locations, these typically have information on the stock of housing about every three years from the late 1990s until 2015. With information on year built and floorspace, we fill the data backwards when necessary to construct housing stocks as of the end of 2000 forward.[4] We also construct flows of new construction, teardowns and renovations.

As a secondary source, we construct analogous measures where possible using 2008-2012 ACS data on new construction flows (5% sample) in calendar years 2000-2009 and 100% count stocks from the 2000 and 2010 censuses, as of April 1st. We use occupied units instead of all units to be consistent with Saiz (2010) and because vacant units may be under-reported or not habitable. While it has the best neighborhood coverage, the census based measure cannot be fully decomposed.

2.2.1 Units

Summary statistics for two measures of $\Delta \ln H_i$ are presented in the second block of Table 1. The average growth rate for the 2000-2010 period, which incorporates teardowns, is 7 percent based on Zillow and 8 percent based on census data. The correlation between these two measures is 0.83. Average construction rates of new units across tracts during the same sample period ($\frac{H_i^R+H_i^U}{H_i}$) are 10 percent based on Zillow data and 12 percent based on the Census/ACS data, with 80% of this new construction occuring 2000-2006 in the average tract. The sub-section on satellite data below includes an explanation of our imputation procedure for units built through redevelopment

Changes in units due to teardowns and depreciation plus renovations are measured as a residual. They are the 2000 stock plus 2000-2010 new construction less the 2010 stock. Renovations are measured as the change in the number of units in buildings that exist in our data in both 2000 and 2010. Renovations can increase or decrease the housing stock (e.g. by converting an unfinished basement into a separate unit or by converting two duplex units into a single family home). Table 1 reports that an average 3 percent of the housing stock was lost due to teardowns and depreciation. There was no average change in units from building renovations.

---

[4] Some 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.

2.2.2 Floorspace

In much of the housing production literature, including the model we develop below, homes are viewed as differing only in the efficiency units of homogenous housing services they provide. If more recently built housing units are larger than those built in the past or some existing housing units have been renovated, housing units measures would understate the full growth in housing supply, as highlighted in equation (1). We use Zillow assessment data to construct changes in components of tract floorspace analogously to our units based measures.[5]

The average 2000-2010 growth in floorspace ($\Delta \ln S_i$) at the tract level is 14 percent. Out of this, the floorspace added through new construction is about 13 percent. These are greater than the analogous 7 and 10 percent figures for units, as new units are typically larger than existing units. Renovation of existing units expands total floorspace by 3 percent on average. The loss of floorspace due to teardowns and demolitions averages 2 percent.

2.3 Satellite Data

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

2.3.1 Topography and Land Development

We construct topographic information using the "Scientific Investigations Map 3085", derived from the US Geological Survey's National Elevation Database. This data set uses raster information on slope and elevation range for all 30X30 meter land pixels within a 0.56 km radius (1 sq. km) of each pixel to classify it into one of nine categories that describe how flat or hilly the surrounding area is. After much experimentation with various options, we focus on the fraction of land area surrounded by "flat plains" as our main topographic measure in the empirical work.

---

[5] We also explore the use of a broader housing quantity index. To build this, we regress the log of sales price on age, age-squared, square footage and square footage squared all interacted with property type (single family, townhouse, condo, co-op, prefab, small multifamily, medium/large multifamily) using 2006 data, the first year with broad coverage. Using these fixed coefficients from 2006, we predict out the log quantity of housing services provided using property attributes in the assessment data in each year. We then exponentiate, sum up within each tract, and take the log to yield the quantity index. This measure generates estimates that are very close to our floorspace based estimates (Table A3). To avoid econometric difficulties associated with a predicted dependent variable and for transparency, the main analysis uses floorspace as the efficiency units measure.

Flat plains are defined to have a slope of less than 8% in more than half of these nearby pixels and an elevation range of less than 15 meters in this 1 sq. km region. On average in our estimation sample, 42% of tract land area is surrounded by flat plains.

We construct tract developed fraction from the National Land Cover Database (NLCDB) for 2001 and 2011. For each 30X30 meter cell, the NLCDB provides one of 4 categories of development (0-19%, 20-49%, 50-79%, 80-100%). We construct the square meters of land in each tract by density of development and aggregate, assigning category midpoints, to impute average developed fraction for the land area of each census tract. $\Delta \ln L_i$ denotes the growth rate of tract land that has ever been developed. The average tract in our estimation sample experienced 8 percent 2001-2011 growth in developed land off of a base of 33% developed in 2001. (All 80-100% developed pixels in 2001 remained so in 2011.)

Figure 1 Panel A shows kernel densities of fraction flat and fraction developed. Both are bimodal. Fraction flat has modes near the extremes of 0 and 1 while fraction developed is a smoother distribution with modes near 0 and 0.4. Figure 1 Panel B shows that both decline on average with CBD distance, though fraction developed declines more rapidly.

We spatially aggregate to construct various land unavailability measures in 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 steep slopes, water or wetlands. We also construct the fraction that is developed as of 2001. We build variants of these two measures, instead aggregating to the metro area level, from the CBD to within 50%, and within 100% of the distance from the CBD to the furthest tract in each metro area. As these measures are highly correlated, our results are not sensitive to the choice of aggregation. Lutz and Sand (2019) construct similar measures of land unavailability for metro areas using some of the same data sources.

2.3.2 Redevelopment

Because the historical assessment data for 2000 is incomplete, we use satellite data and geocoded ZTRAX property data to impute whether each 2001-2010 new construction property is redevelopment. We consider a given pixel developed as of 2001 if it was coded as “high” or “medium” development intensity in the 2001 NLCDB land cover data. 5% of the 12.8 million 2001-2010 new construction properties in our data were on high intensity pixels and 14% were on medium intensity pixels. We treat these properties as redevelopment. A remaining 29% were on open space or low intensity developed pixels. The rest were on pixels coded as barren land, forest, shrubs/scrub, herbaceous, pasture, agriculture or wetlands. This imputation procedure is not perfect since small portions of 30X30 meter cells coded as medium or high intensity developed are still likely to be available for new development. However, among 2001-2010 new construction properties we note that only 0.5% of those on medium developed parcels in 2001 were coded as highly developed in 2011, with upward transition rates more than 14 times higher for other 2001 land use classifications. In our estimation sample, an average 36% of new construction units are built as redevelopment, with this distribution heavily skewed and bi-modal. The median is 0.24, the modes are near 0 and 1, and the aggregate fraction redevelopment is 0.22. [6]

2.4 Regulation

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, though only 10,016 of tracts in our estimation sample are covered. However, these tracts have wide variation in regulation by this measure.

We also incorporate separately collected information on Floor Area Ratio (FAR) restrictions on residential development from the municipalities of Atlanta, Boston, Chicago, Denver, Los Angeles, New York, San Francisco, and Washington.[7] For each residential land parcel, local zoning maps provide the residential FAR. We use the average of these within each census tract, weighted by parcel area.

Figure 1 Panel B shows that both FAR and the Wharton Index fall with CBD distance out to

---

[6] Our approach for measuring redevelopment is consistent with previous findings that in cities where building density is already high, builders can only increase housing supply through redevelopment. 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).

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

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

2.5 Population, Employment and Commutes

The Census Transportation Planning Package (CTPP) reports tabulations of 1990 and 2000 census data by residential location, work location and commuting flow. The 1990 CTPP geography determines our study regions. The 1990 CTPP assigns microgeographic units the size of census tracts 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 and 18 industry groups are reported for these 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.

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

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

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. We use these variables to measure aggregate outcomes and to control for pre-treatment trends in observables.

2.6 Estimation Sample

Our analysis requires reliable tract level information on housing quantities and prices in 2000 and 2010. To this end, we exclude all tracts from the estimation sample for which our 2000 Zillow unit stock is more than 25% below the occupied housing stock reported in the 2000 census. In addition, we must observe the RS index in 2000 and 2010. Finally, we exclude the 206 tracts which fit the above criteria but that have fewer than 500 housing units in 2000. As census tracts are typically drawn to include about 3,000 residents each, inclusion of such outliers would introduce unnecessary noise into the analysis. Because the CTPP and LODES fully cover our sample area, these data sets do not introduce any additional sample constraints. Primary estimation sample sizes for specific variables are reported in Table 1.

3 Conceptual Framework

We are interested recovering estimates of γ_{ir} in the following reduced form expression. Q^s_{ir} denotes a measure of housing quantity in tract i of metropolitan area r and P_{ir} is the observed price per unit of housing services. To accommodate taxes, we can think of the price developers receive per unit of housing services sold as $P_{ir}(1 - t_r)$. Region fixed effects θ_r thus capture potential tax wedges and other region-specific factors that influence construction costs.

$$\ln Q^s_{ir} = \theta_r + \gamma_{ir} \ln P_{ir} + u_{ir} \tag{2}$$

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

In this section, we first sketch a simple model of neighborhood housing supply for new construction. While stylized, the model delivers a natural decomposition of the residential floorspace supply elasticity into intensive (floorspace per parcel) and extensive (parcel development) margins. It delivers a theoretical basis for the floorspace and developed land supply elasticities we measure in the empirical work. It also provides rough calibrated quantification of the intensive margin component and microfounds our empirical specifications. Relative to the empirical work, the drawback of this model is that its static formulation makes it easier to understand comparisons of new development across neighborhoods in the cross-section, whereas the empirical work additionally compares changes over time for identification reasons. This requires some adjustments, developed in Section 5.4.2, to relate our empirical estimates to quantitative implications of the supply model.

We incorporate neighborhood housing supply functions into an urban spatial equilibrium model which links neighborhood housing and labor markets in an urban area. This part of the model provides theoretical support for the instruments we use to pin down consistent estimates of parameters used in the construction of γ_{ir} and helps guide our opportunity zone application in Section 7. Altogether, our conceptualization of the data generating process serves as a microfoundation that guides both our empirical strategy for estimating housing supply functions and the interpretation of empirical findings.

3.1 Housing Supply

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

The key object delivered by the model that is relevant for the empirical work is a description of relationships between housing stocks and prices across these ex-ante identical neighborhoods.

We analyze a representative developer. The developer only builds on land parcels with fixed development costs that are sufficiently low such that the variable profit minus fixed development cost is weakly positive. Each parcel's land value is the variable profit net of the fixed development cost. Conditional on development, the amount of floorspace supplied on each parcel in neighborhood i is A_i. A_i thus represents the intensity of development on any given lot in neighborhood i, incorporating both floorspace per housing unit and units per parcel. It is chosen based on tract housing productivity, parcel size and demand conditions summarized by the uniform price per unit of housing services P_i, which is the same for all parcels in neighborhood i. One motivation for using tracts as the spatial unit of analysis in the empirical work is that within-tract demand conditions across parcels are sufficiently similar so that the law of one price is a reasonable assumption.

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

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

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

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

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

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

Differentiating the developed land supply function over time yields

$$\gamma_i^{land} = \frac{f_i(\phi C_i[A_i(P_i)])}{F_i(\phi C[A_i(P_i)])} \frac{\Delta \ln A_i(P_i)}{\Delta \ln P_i} \phi P_i A_i(P_i). \tag{3}$$

Expression (3) highlights that tracts with a greater density of parcels available for development at the fixed cost that equals marginal variable profit, represented by a higher $f_i(\phi C_i[A_i(P_i)])$, exhibit more elastic land supply. Land supply is the extensive margin component of floorspace supply, which can be decomposed as $\gamma_i^{space} = \frac{\Delta \ln A_i(P_i)}{\Delta \ln P_i} + \gamma_i^{land}$, where the housing production literature has estimated the form of $A_i(P_i)$.

Though the reduced form nature of the empirical work allows us to avoid imposing functional form assumptions, to get a sense of magnitudes we parameterize with an example that delivers a convenient form for the tract *i* floorspace supply elasticity. Consistent with evidence in Combes et al. (2021), we use a Cobb-Douglas production technology with land share α and a tract-specific housing productivity, implying the parcel level housing services supply function $A_i(P_i) = \rho_i P_i^{\frac{1-\alpha}{\alpha}}$. We further assume that the fixed cost follows the Frechet distribution, where $F_i(x) = exp(-\Gamma_i x^{-\lambda})$, with the common dispersion parameter $\lambda > 1$ and the tract-specific scale parameter $\Gamma_i > 0$. With these functional form assumptions, the floorspace supply elasticity is

$$\frac{\Delta \ln S_i}{\Delta \ln P_i} = \underbrace{\frac{1-\alpha}{\alpha}}_{\frac{\Delta \ln A_i}{\Delta \ln P_i}} + \underbrace{\alpha^{-1-\lambda} \lambda \rho_i^{-\lambda} P_i^{-\frac{\lambda}{\alpha}} \Gamma_i}_{\frac{\Delta \ln L_i}{\Delta \ln P_i}}. \tag{4}$$

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

The second term relates the extensive margin positively to the scale parameter, Γ_i, and negatively to the initial price level, P_i. In the Frechet example, fixed cost distributions in tracts with a higher Γ_i have higher means and variances and hence thicker right tails. This implies a higher density of land available for development at the marginal variable profit and hence higher γ_i^{land}, which in turn boosts γ_i^{space}. We expect that tracts with lower initial development density, more flat land and less stringent regulation have fixed cost distributions with higher scale parameters, Γ_i, and hence more responsiveness along the extensive supply margin.[8] In addition, in the Frechet example a higher initial price in tract i implies a thinner right tail of the fixed cost distribution and hence less land available for development at the marginal fixed cost, which in turn causes the extensive margin to be less responsive.[9] [10]

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 (2022). 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 wages available to residents of tract i. This object can be readily calculated with counts of workers and residents in each tract. Shocks to wages in commuting destinations are reflected as shocks to RMA_i.

Second, the model makes clear the conditions required for census tract level "Bartik" shocks to represent a valid source of econometric identification. These shocks are calculated by predicting 2000 to 2006 tract level growth in employment by industry using 1990 tract level employment

---

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

[9] If P_i is too low or the minimum of the support of $F_i(x)$ is too high, the price is not sufficiently high to justify the fixed cost of developing any land and $\gamma_i^{space} = \gamma_i^{land} = 0$. There was no 2000-2010 construction in 6% of sample tracts in our data.

[10] The equilibrium split of $(A_i \equiv S_i/L_i)$ into floorspace per units (S_i/H_i) and units per parcel (H_i/L_i) depends on the composition of housing demand (families vs. single people for example), which neither our data nor identification strategy are well suited to handle.

shares by industry interacted with national industry-specific employment growth outside of the metro area in question. Originally proposed by Bartik (1991), this source of variation has been used at the metro area level in a large number of studies. However, this paper is among the first to use such identifying variation at the sub-metro level of geography.

Finally, the model delivers enough structure to estimate parameters governing neighborhood demand conditions used to perform welfare analysis of place based policies, as in Section 7. Detailed derivations and further discussion of key model equations are in Appendix B.

3.2.1 Setup

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

The model features a continuum of ex-ante identical workers indexed by ω who choose residential tract i, work tract j and industry of work k within the metro area. They receive productivity shocks $z_{ijk\omega}$ over commute origin-destination and industry triplets and preference shocks $v_{i\omega}$ over residential locations. The preference shocks are revealed first, leading agents to first choose residential locations and housing while anticipating the quality of accessible employment opportunities before productivity shocks are revealed. Productivity shocks are then revealed and agents choose work locations. In practice, the shocks primarily allow the model to generate enough variation in tract populations, home prices and employment to rationalize the data.[11]

Given Cobb-Douglas preferences, the indirect utility person ω receives from living in tract i, commuting to tract j and working in industry k is

$$v_{ijk\omega} = \frac{v_{i\omega} B_i z_{ijk\omega} w_{jk}}{P_i^{1-\beta} e^{\kappa \tau_{ij}}}, \tag{5}$$

where B_i is a local amenity, w_{jk} is the price of a unit of skill in commuting destination j and industry k, P_i is the price of one unit of housing services in i and $\kappa\tau_{ij}$ is the fraction of time spent commuting for those living in i and working in j. In the data, we observe the price P_i in year 2000 and beyond and the commuting time τ_{ij} in 1990 and 2000.

---

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

The productivity shock $z_{ij\omega}$ is drawn from the Frechet distribution with shape parameter ε. Following Tsivanidis (2022) and Couture et al. (2019), we incorporate a nested preference shock over residential locations $v_{i\omega}$. This shock is also distributed Frechet but with shape parameters η and ψ, where η is the elasticity of substitution of demand between neighborhoods in the same municipality and ψ is that between neighborhoods in different municipalities.

3.2.2 Resident Market Access

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

$$\pi_{ij|i} = \frac{\sum_k \left[w_{jk}e^{-\kappa\tau_{ij}}\right]^{\varepsilon}}{\sum_k \sum_{j'} \left[w_{j'k}e^{-\kappa\tau_{ij'}}\right]^{\varepsilon}} \equiv \frac{\sum_k \left[w_{jk}e^{-\kappa\tau_{ij}}\right]^{\varepsilon}}{RMA_i}. \tag{6}$$

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

Before the productivity shock is revealed, individuals evaluate the expected wages net of commuting costs associated with residing in each tract. Solving for the expected maximum utility tract yields the population supply function to tract i

$$\pi_i = \mu[\sum_{i' \in m(i)} \left(B_{i'}P_{i'}^{\beta-1}RMA_{i'}^{\frac{1}{\varepsilon}}\right)^{\eta}]^{\frac{\psi}{\eta}-1} \left(B_i P_i^{\beta-1}RMA_i^{\frac{1}{\varepsilon}}\right)^{\eta}. \tag{7}$$

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 tract i's municipality $m(i)$, captured by the object inside the summation. μ is an endogenous scalar that is set either to ensure everybody has a place to live (in a closed city with a fixed population) or to summarize the attractiveness of an outside option (in an open city).

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

$$\ln \pi_{ij} = a_i + b_j - (\kappa\varepsilon)\tau_{ij} \tag{8}$$

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 $\kappa\varepsilon$. We estimate $\kappa\varepsilon$ using separate flow-weighted commuting gravity regressions like (8) with origin and destination fixed effects in 2000 for each metropolitan region.[12] Because we do not observe tract level commute times after 2000, we hold commute times constant for later years.

Recognizing that the labor supply to tract j is $\sum_i \pi_i \pi_{ij|i}$, we have

$$L_j = \mu \sum_k \left[w_{jk}^{\varepsilon} \right] FMA_j, \tag{9}$$

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, in (6), we have the following system of equations.

$$FMA_j = \sum_i \frac{e^{-\kappa\varepsilon\tau_{ij}} \pi_i}{RMA_i} \tag{10}$$

$$RMA_i = \sum_j \frac{e^{-\kappa\varepsilon\tau_{ij}} L_j}{FMA_j} \tag{11}$$

Using data on employment L_j, residents π_i, the parameter cluster $\kappa\varepsilon$ and commute times τ_{ij}, we can calculate FMA_j and RMA_i by solving this system. We calculate RMA_i using (10) and (11) for 2000, 2006 and 2010.

As individuals make housing consumption decisions before productivity shocks are revealed, residents of tract i have expected housing demand $(1-\beta)\frac{\overline{y}_i}{P_i}$, where $\overline{y}_i$ is the expected income associated with living in i. Adding up, the log aggregate residential floorspace demand in tract i is thus

$$\ln S_i^d = \ln \rho_{HD} + \frac{1}{\varepsilon} \ln(RMA_i) + \ln \pi_i - \ln P_i. \tag{12}$$

This object is increasing in RMA_i conditional on population π_i because greater RMA_i is associated with greater income for tract residents. Conditional on P_i, equilibrium tract residential population π_i is also increasing in RMA_i, as seen in (7). 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

---

[12] Across the 306 regions in our broad sample, the median estimated elasticity of commuting flow with respect to one-way commuting minutes in 2000 is -0.04, the minimum is -0.11 and the maximum is -0.01. Estimates of $\varepsilon\kappa$ are about half as large in absolute value in big cities like New York and Los Angeles relative to 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.

housing demand equation formed by substituting (7) into (12). Credible identifying variation in $\ln P_i$ must come from a component of RMA_i that is cleansed of variation in housing productivities and lot sizes. Section 4 lays out how we isolate such variation using a simulated version of RMA_i based on Bartik type labor demand shocks in commuting destinations for residents of tract i.

3.2.3 Equilibrium

Combining conditions governing population supply to residential tracts (7), labor supply to work tracts (9) and imposing housing market clearing yields conditions describing equilibrium tract population and home prices. Differentiating the population condition over time yields

$$\Delta \ln \pi_i = \frac{\gamma_i + \beta}{\gamma_i + 1 + \eta(1-\beta)} \frac{\eta}{\varepsilon} \Delta \ln RMA_i + v_m^{\pi} + u_i^{\pi}. \tag{13}$$

This equation incorporates an intuitive positive relationship between growth in employment opportunities and tract population. This relationship is stronger if housing supply in tract i is more elastic and/or if there is less dispersion in idiosyncratic preferences over locations (η is larger). v_m^{π} is a municipality fixed effect that captures common population trends in all tracts in municipality m that come through their correlation in neighborhood choices delivered by the outer nest in preferences over neighborhoods. The error term u_i^{π} is a function of shocks to amenities and housing productivity in tract i. In Section 7, we use (13) as a basis for structural estimation of η, recognizing that identifying variation in $\Delta \ln RMA_i$ must be uncorrelated with tract level shocks to amenities and housing productivity for successful identification. To back out ψ, we will compare this estimate of η with that from a version of (13) without municipality fixed effects.

4 Empirical Implementation

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

$$\Delta \ln Q_{ir}^{s} = \theta_r + X_{ir}\delta + \gamma_{ir}\Delta \ln P_{ir} + \widetilde{\rho}_{ir} \tag{14}$$

Observations are for tract i in metro region r. To allow for observed heterogeneity in supply elasticities, we parameterize γ_{ir} to depend on tract-specific observables Z_{ir}.[13]

$$\gamma_{ir} = Z_{ir}\gamma \tag{15}$$

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

The first two terms in (14) 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. In tract characteristics X_{ir}, our main specification includes lagged demographic attributes, a cubic in CBD distance, 2001 tract developed fraction and share flat land, and controls for tract-specific labor demand shocks. Our controls for 1990 and 2000 tract demographic characteristics account for potential influencers of the tract regulatory environment that may be correlated with the instrument for price growth laid out below. CBD distance controls hold constant any potential spatial trends in price growth that are related to costs and are useful given the stronger 2000-2010 labor demand growth in suburban areas. 1990 and 2000 census rent and price indexes help to account for decadal mean reversion in home price growth. Controls for tract developed fraction and topography account for obvious po-

---

[13] In Section 5.3, we extend the empirical model to allow for unobserved heterogeneity. This allows us to additionally incorporate metro area level predictors of tract supply elasticities.

tential housing supply shocks and are needed as main effects in interacted specifications. Finally, 1990 log tract employment and a 2000-2006 tract-specific Bartik labor demand shock (explained below) help ensure that our IV implementation is only using variation from outside of tract *ir* for identification.

Even including this long list of control variables, observed OLS relationships between various measures of post-2000 house quantity changes and contemporaneous price growth are implausibly small. Table A1 shows that OLS estimates are at most 0.10 for housing units, 0.11 for floorspace, and 0.03 for land development and are similar when using the census based housing price index.[14] These small estimates point to several identification challenges in estimating housing supply. First, neighborhoods that experience stronger housing demand shocks may follow with unobserved changes in housing regulation in part in order to cope with these demand shocks, thereby maintaining open space and natural amenities. That is, negative supply shocks may be correlated with positive demand shocks. Similarly, the magnitude of demand shocks may be correlated with supply elasticity. Davidoff (2016) presents evidence for this, observing that more supply constrained metro areas tend to have greater productivity and housing demand growth. Second, it is possible that positive productivity shocks outside of the construction sector may simultaneously boost local housing demand through higher household earnings and reduce housing supply through higher construction costs. This would further bias the OLS relationship between quantity and price growth downward. Moreover, our price index measure, while constructed as carefully as possible, is sure to be a noisy measure of the true price of housing services. Mechanical mean reversion in decadal house price growth that could reflect classical measurement error would lead to attenuation bias.

The broad message is the possibility that tract unobservables that predict supply shocks or supply elasticity may be correlated with demand shocks, thereby generating a downward bias in OLS estimates. A valid identification strategy must hold constant such unobservables and address the classic endogeneity concern of simultaneity in demand and supply by finding variation in local housing demand shocks across neighborhoods that are uncorrelated with local construction costs.

We develop an instrument that isolates variation in tract price growth that is plausibly uncorrelated with supply factors conditional on controls. To fix ideas about the role of our instrument,

---

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

consider the tract level inverse floorsapce demand equation from the model. This equation is derived by substituting for tract population π_i (7) in (12) and solving for price.

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

$$- \frac{1}{1+\eta(1-\beta)} \ln S_i^d + \frac{\eta}{1+\eta(1-\beta)} \ln B_i \qquad (16)$$

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 home prices. However, any component of RMA_i that is correlated with tract housing productivity or land parcel size is endogenous to housing supply. Indeed, through its codetermination with FMA_i, RMA_i depends structurally on tract population, which itself depends on tract housing productivity and parcel size. As such, our identification approach is to difference over time and develop an instrument that picks out components of $\Delta \ln RMA_i$ that are likely orthogonal to levels of and shocks to productivity or other factors that influence local construction costs.

Because supply elasticity varies at the tract level, there exists structural heterogeneity in relationships between price growth and $\Delta \ln RMA_i$. For example, housing demand shocks in tracts with more land available for development may be expected to induce smaller price growth and greater quantity growth. Formally, substituting for $\ln S_i^d$ in (16) with the generic housing supply function $\ln S_i = \ln \rho_i + \gamma_i \ln P_i$ yields the equilibrium relationship $\frac{\partial \ln P_i}{\partial \ln RMA_i}$ that is positive and with a negative cross derivative in the supply elasticity (see the discussion after equation (35) in Appendix B). As a result, we also expect a heterogeneous first stage relationship between housing price growth and our instrument providing exogenous varation in $\Delta \ln RMA_i$.

As in Gorback and Keys (2020), one could alternatively use the following equivalent approach to estimate supply elasticities given an exogenous housing demand shifter. Calculate the ratio of separately estimated coefficients on the instrument in reduced form quantity and price growth regressions. Or, to estimate additional structural parameters, estimate each equation by IV instead. As these alternative approaches require reporting twice the number of estimates as our more classical IV approach, and we have many margins of supply response and sources of supply heterogeneity to study, the classical IV approach facilitates clearer exposition for our purposes.

However, we return to estimation of structural equations to recover demand parameters in Section 7.

4.1 Instrument Construction

We begin with (10) and (11) as a basis for calculating a simulated version of $\Delta \ln RMA_i$ that excludes shocks to tract housing productivity and its correlates conditional on control variables. This simulated instrument, denoted $\Delta \ln \widetilde{RMA}_i$, serves a dual purpose. First it is a reduced form housing demand shock that drives exogenous variation in tract level house price growth, as represented in (16) above. Second, it is a predictor of the structural object $\Delta \ln RMA_i$ that is unrelated to tract level shocks to local amenities or housing productivities. This use facilitates recovery of parameters governing the neighborhood demand system for use in the Opportunity Zone application in Section 6 below.

Calculation of $\Delta \ln \widetilde{RMA}_i$ is analogous to that of $\Delta \ln RMA_i$ except that it replaces actual tract employment with that predicted by 1990 tract industry compositions interacted with national industry growth rates. To solve jointly for $\widetilde{FMA}_j$, we inflate the 1990 residential population of each tract by a constant to equalize counterfactual aggregate labor supply and demand in each region R. The simulated instrument $\Delta \ln \widetilde{RMA}_i$ is conceived in the spirit of shift-share identification strategies that go back to Bartik (1991).

Beyond fixing commute times to those from 1990, we incorporate three additional elements to reduce the likelihood that the instrument is correlated with trends in construction costs conditional on controls. First, we exclude construction from the set of industries used to build the instrument. This precludes the possibility that nearby changes in productivity in the construction sector may directly affect construction costs in tract i.[15] Second, we exclude all tracts with centroids within 2 km of tract i. This exclusion mitigates the possibility that employment growth in or nearby to tract i may influence the land available for residential development in tract i or that its construction labor costs change in a way that is correlated with the instrument. Finally, we always control for predicted employment growth in tract i itself.

Putting these elements together, and following (10) and (11), we calculate the year 2000 com-

---

[15] We also construct alternative versions that either do not exclude the construction industry or additionally exclude finance, insurance and real estate. These deliver very similar estimates to our main ones, as seen by comparing results in Table 4 and Table A4.

ponent $\widetilde{RMA}_i^{2000}$ of our main instrument as:

$$\widetilde{RMA}_i^{2000} = \sum_{j\subseteq R(i)} \frac{e^{-(\widehat{\varepsilon\kappa})_{R(i)}\tau_{ij}^{90}} 1(dis_{ij} > 2km) \sum\limits_{k\neq cons} L_{jk}^{90}[E_{r'(j)k}^{2000}/E_{r'(j)k}^{1990}]}{\widetilde{FMA}_j^{2000}} \tag{17}$$

$$\widetilde{FMA}_j^{2000} = \sum_{i\subseteq R(j)} \frac{e^{-(\widehat{\varepsilon\kappa})_{R(i)}\tau_{ij}^{90}} 1(dis_{ij} > 2km) \pi_i^{90} \left[\frac{\sum\limits_{j\subseteq R(i)} \sum\limits_{k\neq cons} L_{jk}^{90}[E_{r'(j)k}^{2000}/E_{r'(j)k}^{1990}]}{\sum\limits_{j\subseteq R(i)} L_j^{90}}\right]}{\widetilde{RMA}_i^{2000}} \tag{18}$$

In these expressions, τ_{ij}^{90} is the reported or forecast commute time from i to j in the 1990 CTPP. $(\widehat{\varepsilon\kappa})_{R(i)}$ is estimated separately for each region in year 2000, as explained above in the context of equation (8). Distances from residential to work locations dis_{ij} are calculated using tract centroids. Employment in industry k in work location j, L_{jk}^{90}, is measured in the 1990 CTPP. $E_{r'(j)k}^{2000}$ and $E_{r'(j)k}^{1990}$ are 2000 and 1990 nationwide employment in industry k excluding the region of tract j. That is, $\sum\limits_{k\neq cons} L_{jk}^{90}[E_{r'(j)k}^{2000}/E_{r'(j)k}^{1990}]$ captures the predicted amount of employment that would exist in tract j if 1990 employment by industry grows at national rates (excluding region r) to year 2000. $\frac{\sum\limits_{j\subseteq R(i)} \sum\limits_{k\neq cons} L_{jk}^{90}[E_{r'(j)k}^{2000}/E_{r'(j)k}^{1990}]}{\sum\limits_{j\subseteq R(i)} L_j^{90}}$ is a constant within each region that captures the population growth rate needed to match the aggregate employment predicted by Bartik shocks in the region in year 2000. The 2006 component of the instrument is calculated analogously, with $E_{r'(j)k}^{2000}$ in (17) and (18) replaced by $E_{r'(j)k}^{2006}$.

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. Variation in 2006-2010 employment changes are not well predicted by shift-share type instruments.

4.2 Instrument Validity

The fundamental sources of identifying variation used are tract level "Bartik" (1991) type shocks in each employment location, which can be written as follows.

$$Bartik_j = \sum_{k\neq cons} \frac{L_{jk}^{90}}{\sum_k L_{jk}^{90}} [\ln E_{r'(j)k}^{06} - \ln E_{r'(j)k}^{00}] \tag{19}$$

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

Table 2 Panel A presents evidence to this effect. It presents regressions of 2000-2006 or 2000-2010 employment growth in tract j on $Bartik_j$ and controls for 1990 employment level, past demographic composition of tract residents, a cubic in CBD distance and metro region fixed effects. This tract level variation gets aggregated into the identifying variation for the first stage of our main analysis. All tracts in primary sample regions are included, as they all contribute to measures of $\widetilde{RMA}_i$ in tracts that contribute data to our main estimation exercises. We control for past employment to isolate employment growth due only to variation in industry composition. Lagged demographics and CBD distance controls account for potentially differing labor supply conditions. Results indicate that we can plausibly isolate labor demand shocks at the tract level. Each additional percentage point increase in the Bartik shock predicts 0.33 points greater 2000-2006 tract employment growth and 0.60 points greater 2000-2010 tract employment growth.[16] The difference reflects labor market frictions. With a longer period, the market has more time to respond to national productivity shocks.

One consideration we face when estimating the housing supply equation is accounting for serial correlation in home prices and quantities. The first two columns of Table 2 Panel B show that home price growth is negatively serially correlated across decades whereas unit quantity growth is positively serially correlated across decades.[17] These results may reflect supply shocks that respond to demand shocks with a lag and suggest that there could be local unobserved history that drives both relative price declines and more construction, inducing a downward bias in OLS estimation of housing supply.

A legitimate potential concern is thus that our instrument $\Delta \ln \widetilde{RMA}_i$ may be correlated with such unobserved history. However, results presented in column (3) of Table 2 Panel B help allay this concern by showing small and insignificant relationships between the instrument and pre-treatment trends in key endogenous variables. $\Delta \ln \widetilde{RMA}_i$ is not correlated with 1990-2000 housing price nor quantity growth for our main specification with region fixed effects and 1990

---

[16] Results are robust to lagging the demographic tract controls by one additional decade and/or adding 2-2.5 km CBD distance ring fixed effects interacted with metro region. Excluding demographic controls increases estimates to 0.69 and 1.04 in columns 1 and 2, respectively (Table A2, Panel A).

[17] Results in Table A2 Panel B shows that these correlations are not sensitive to conditioning on demographic controls.

and 2000 demographic controls. Table A2 Panel B column (3) shows that controls for year 2000 demographics are needed to generate the insignificant (and negative) relationship between the instrument and 1990-2000 housing quantity growth; these controls also attenuate the 1990-2000 price growth relationship, which remains statistically insignificant with and without demographic controls. Table 2 Panel C presents results of the reverse regressions with the same implication.[18] This pre-trend evidence indicates that it is unlikely that our instrument is correlated with unobserved local history in a way that biases our supply elasticity estimates.

4.3 First Stage Estimates

Table 3 presents the main first stage estimates. Results in Panel A show strong positive relationships between our primary measures of $\Delta \ln P$ and $\Delta \ln \widetilde{RMA}$. The slightly smaller first stage coefficient for 2000-2010 relative to 2000-2006 reflects the fact that the 2007-2010 period mostly saw housing market declines. Column (3) shows that we do not have strong first-stage power predicting the census hedonic index. For this reason, we only use the census index to account for pre-2000 price trends. Column (4) shows a significant estimated relationship between $\Delta \ln \widetilde{RMA}$ and $\Delta \ln RMA$ for 2000-2010 of 0.74. According to the model, this is the mechanism through which $\Delta \ln \widetilde{RMA}$ predicts $\Delta \ln P$. Tracts that appear in multiple metro regions are weighted equally to tracts that appear in just one.[19]

Standard errors are adjusted for spatial autocorrelation out to 16 km using a triangular kernel. The 16 km cutoff was selected by inspecting the spatial correlogram of errors (Figure A2) generated from a tract level IV regression of the 2000-2010 growth rate in housing units on the 2000-2010 change in the repeat sales price index using the specification reported later in Table 4 column (3). The associated first stage F statistic is 22.2, which can be calculated from estimates in column (2) of the top row of Table 3.

We undertake a number of checks to confirm that our instrument is sufficiently strong. Olea and Pflueger (2013) suggest that with dependent errors the critical F statistic for a strong first

---

[18] Goldsmith-Pinkham, Sorkin, and Swift (2020) suggest this sort of pre-trend test for evaluating the validity of Bartik instruments. Their other suggested validity tests use base year industry shares, which are not easily defined in our setting as they are nonlinearly aggregated across all potential commuting destinations into $\Delta \ln \widetilde{RMA}_i$.

[19] First stage predictions of our hedonic price index are similar to those reported in Table 3 Panel A columns (1)-(2) for the repeat sales index. We also find significant positive reduced form relationships between the instrument and 2000-2006 and 2000-2010 Zillow new construction.

stage should be adjusted upwards from the rule of thumb of 10 often used. In our context, the critical value of the F-statistic for the "worst case bias" of 10% is 23.1. Other standard statistical tests provide even more confidence that our first stage is sufficiently strong for valid inference. Lee, McCrary, Moreira, and Porter (2021) recommend that with a first stage F-statistic of 22.2, second stage standard errors should be multiplied by less than 1.3. Multiplying standard errors by this factor in Table 4 maintains statistically significant estimates in almost all cases. Conservative Anderson-Rubin test statistics associated with regressions in Table 4 using county level clustering, which increases standard errors beyond our spatially corrected estimates in all cases, also reject that our instrument is weak.

Table 3 Panel B reports extended first stage relationships between price growth and $\Delta \ln \widetilde{RMA}$ interacted with various factors that we demonstrate below influence supply elasticity. These are the first stage regressions that match the heterogeneity results reported later in Tables 5 and 6. Results in Table 3 Panel B show that there remains sufficient variation in the instrument to predict price growth over a broad range of tracts in the data with different observable characteristics, conditional on region fixed effects and our full set of controls, including those for the three main effects in the supply factors we consider (CBD distance, fraction developed and fraction of flat land). Inclusion of these three controls is crucial for identification, as doing so forces all identifying variation to come from $\Delta \ln \widetilde{RMA}$ rather than the interacted observable. For supply elasticities to differ across tracts, it must be that we can predict different amounts of price growth for the same demand shock hitting different types of locations. Therefore, we see it as important that exogenous price responses induced by the instrument differ for tracts with different observable characteristics.

Interacted first stage coefficients are of expected signs and significance. Results in column (1) indicate that tracts with a higher developed fraction experience larger price increases in response to a simulated RMA shock. The insignificant coefficient on uninteracted $\Delta \ln \widetilde{RMA}$ is as expected, as it applies to tracts with no developed land. Such locations are expected to have very elastic supply where positive demand shocks translate almost entirely into quantity rather than price responses. In other columns, estimates on the diagonal are positive and strongly significant, as expected. Anderson-Rubin test statistics assuming standard errors clustered at the county level (larger than the reported spatially corrected errors) yield evidence of strong first stages.[20]

4.4 Sources of Identifying Variation

The unified supply elasticity estimates are identified by isolating comparisons between ex-ante observationally identical tracts in the same metro area that receive different housing demand shocks because of variation in labor demand shocks in commuting destinations conditional on control variables. These labor demand shocks spur residential home price growth driven by housing demand shocks only. Unified IV estimates (in Table 4 below) are average causal effects of demand shock driven price growth on quantities, with averages taken across tracts of various underlying attributes that receive variation in treatment. This gives an idea of the average magnitudes of supply elasticities, and is particularly useful for recovery of a clean unified decomposition of supply responses into different components.

Labor demand shocks in accessible locations are spatial aggregations of 1990 industry composition interacted with 2000-2006 industry growth. As changes over time become small, within any region we can express the instrument in a way that is somewhat analogous to the tract level Bartik shocks in equation (19), as follows. [21]

$$\Delta \ln \widetilde{RMA}_i \approx \sum_{k \neq cons} [\sum_j \frac{e^{-\widehat{\varepsilon\kappa}\tau_{ij}^{90}} 1(dis_{ij} > 2km) L_{jk}^{90} [E_{r'(j)k}^{2000} / E_{r'(j)k}^{1990}]}{\sum\limits_{k \neq cons} \sum_j e^{-\widehat{\varepsilon\kappa}\tau_{ij}^{90}} 1(dis_{ij} > 2km) L_{jk}^{90} [E_{r'(j)k}^{2000} / E_{r'(j)k}^{1990}]}] \Delta \ln E_{r'(j)k}$$

If we treat the quotient as endogenous and the industry growth rates $d \ln E_{r'(j)k}$ as exogenous, the instrument does not require recentering for clean identification because the shares sum to one (Borusyak & Hull, 2020). However, it may alternatively (or additionally) be reasonable to treat the shares as exogenous, as they are predetermined. Our set of controls is oriented toward making these initial industry shares in commuting destinations uncorrelated with unobserved supply factors in tract i conditional on controls. The lack of pre-trends seen in Table 2 Panel C allays such potential concerns.

Of further interest is how average supply elasticities are built from aggregating impacts in

---

[20] Evaluation of first stage strength in an environment with multiple endogenous regressors and non-spherical errors is an unsettled and active area of research in econometrics (Andrews, Stock, & Sun, 2019).

[21] To build intuition, we leave out the multilateral resistence term that depends on $d \ln \widetilde{FMA}_j$. $\Delta \ln \widetilde{RMA}_i$ has a correlation of 0.98 with the reduced form analog that leaves $\widetilde{FMA}_j^Y$ out of the denominator.

tracts with different baseline attributes. Consider comparisons between tracts that receive the same demand shock and are ex-ante identical in all factors mentioned above except one supply factor, like initial land developed fraction. This type of comparison only cleanly applies for the interacted specifications (in Table 5 below). The interacted first stage results in Table 3 Panel B show how the same shock affects their prices differently. The tract with a low developed fraction sees prices rise by less than that with the higher developed fraction. Inspection of the result in Table 5 shows that for a given demand shock driven price rise in those tracts with a greater developed fraction, quantity rises by less.

Identifiying variation is not primarily from comparisons between neighboring or nearby tracts because such tracts typically have very similar shocks and attributes. Moreover, since we allow for spatial autocorrelation in errors out to 16 km, adjacent tracts will be treated almost as identical observations. Instead, comparisons for identification purposes are made between tracts that have the same observables and are in the same metro region but are not adjacent. For example, a counterfactual for a tract in the Hyde Park neighborhood of Chicago could be a tract in the Uptown neighborhood. These two neighborhoods have similar demographics, developed fraction, fraction flat and CBD distance.

5 Main Results

In this section, we present our main tract supply elasticity estimates and explore the factors that drive their heterogeneity across locations.

5.1 Unified Supply Elasticity Estimates

Table 4 presents unified supply elasticity estimates and their decomposition. While evidence below confirms that supply elasticities are heterogeneous across neighborhoods, we begin with unified estimates to provide a sense of magnitudes and of the relative importance of new construction versus other margins of response to positive demand shocks. Table 4 reports coefficients in regressions of various measures of housing quantity growth constructed using Zillow data on home price growth. The top row uses the repeat sales index and the bottom row uses the hedonic index as the price measure. In all regressions, we instrument for $\Delta \ln P$ with $\Delta \ln \widetilde{RMA}$ and control for region fixed effects and the same tract level factors as in Table 3 Panel A.

Columns (1)-(2) cover 2000-2006 for unit supply and floorspace supply respectively. The remaining columns cover 2000-2010, with columns (3)-(7) on unit supply, columns (8)-(11) on floorspace supply, and column (12) on land supply. Following our discussion in Section 2.2, we decompose changes in housing units (column 3) into that from new construction (column 4) and from changes to existing buildings (column 6). For the former, we further distinguish between new construction on developed land (column 5) and that on undeveloped land (column 4 minus column 5). We decompose the latter into expansion and renovation of existing buildings (column 7) and teardowns (column 6 minus column 7). Similarly, we decompose changes in floorspace of existing buildings (column 8) into additions through new construction (column 9) and changes through teardowns and renovations (column 10). We separate out renovations in existing buildings (column 11) from fewer teardowns (column 10 minus column 11).

Before discussing the estimates, we highlight a few observations that inform our subsequent analysis. First, unlike the corresponding OLS estimates reported in Table A1, estimated coefficients in Table 4 are of reasonable magnitudes and follow a coherent pattern. This is consistent with our narrative that OLS relationships between quantities and prices in part reflect movement along demand rather than supply curves. Second, 2000-2010 supply responses are considerably larger than those during 2000-2006, even though both are identified using 2000-2006 labor demand shocks. Given the minimal number of 2007-2010 housing starts, these gaps appear to mostly reflect construction lags after 2000-2006 price shocks. Third, estimated housing supply elasticities are remarkably robust to different price growth measures in all specifications. Table A3 provides additional checks. Columns (1)-(3) show that estimated unit supply elasticities are robust to different housing quantity measures constructed from the Census and ACS data. Columns (4)-(5) show that estimated floorspace supply elasticities are robust to the quality-adjusted housing quantity measure described in footnote 5.

Table 4 Column (3) shows that the estimated total unit supply elasticity during 2000-2010 is 0.35. That is, between two ex-ante identical census tracts, if Tract A experiences home price growth that is 10 percentage points higher than Tract B, changes in total housing units in Tract A would exceed that in Tract B by 3.5 percentage points. Estimates in columns (4) and (6) show that, out of these 3.5 points, 1.9 points comes from new construction and 1.6 points comes from fewer teardowns and reconfiguration of existing buildings into multi-unit buildings.

Out of the portion due to new construction, only a small and statistically insignificant portion is estimated to be redevelopment, as reported in column (5). This comes despite the fact that on average 36% of new construction in estimation sample tracts with some 2000-2010 new construction was through redevelopment. Our evidence points to factors other than price signals as the primary drivers of redevelopment, with unobserved land use and zoning restrictions likely to be centrally important.

Out of the portion due to teardowns and renovations, about half is through expansion net of consolidation of existing housing units, as indicated in column (7). An example is converting an unfinished basement or an attic in an existing house into a separate apartment. The remaining half is through fewer teardowns. These estimates are consistent with evidence that neighborhood renewal is largely driven by the deterioration and subsequent renovation of the existing housing stock (Rosenthal, 2018; Brueckner & Rosenthal, 2009).

Table 4 Column (8) shows the estimated total floorspace elasticity of 0.42. Of this, over two-thirds comes from floorspace in newly constructed units (Column 9); the remaining part comes from a combination of fewer teardowns and expansions of existing buildings (Column 10). Over two-thirds of this last component is through building renovation and expansion, accounting for a noisily estimated 21 percent of the total floorspace elasticity (Column 11 divided by Column 8).

Column (12) reports the estimated elasticity of observed developed land is a statistically insignificant 0.09. Recall that $\gamma_i^S = \gamma_i^{S/L} + \gamma_i^L$, where $\gamma_i^{S/L}$, estimated at 0.33, is the elasticity of floorspace per ever developed parcel with respect to price. In Section 5.4.3, we square this evidence with our supply model's predictions about the magnitude of land use intensification with price increases.

5.2 Tract Level Heterogeneity

The local average treatment effect tract-level supply estimates in Table 4 mask substantial variation across neighborhoods. This sub-section shows how tract-level housing supply elasticities vary as a function of distance to CBD, land availability, topographical features, and land use regulations.

Table 5 repeats the IV regressions in columns (3)-(12) of Table 4 with the addition of a set of interactions between price growth and tract-level factors that may influence supply. Price growth is constructed using the repeat sales index and housing supply responses are constructed using 2000-2010 Zillow data. All specifications include interactions with CBD distance, which is measured as the fraction of the way from the CBD to the furthest census tract in the region from the CBD. Table A5 presents specifications that are analogous to Table 5 but quadratic in tract developed fraction.

Results in column (1) summarize the average relationship between unit supply elasticity and CBD distance. Unit supply elasticity increases with CBD distance at a marginally decreasing rate. At the CBD, the implied average supply elasticity is estimated to be only 0.11, rising to 0.59 near halfway to the region edge. As only 18 percent of observations fall beyond the halfway point, the quadratic coefficient is mostly identified from variation near the CBD and predicted elasticities using this specification are thus most accurate in that region. These microgeographic estimates provide a supply-side explanation for the recent finding of more price growth in the center of metropolitan areas in the post-2012 boom (Glaeser et al., 2012; Yoshida, 2016). They are also consistent with the observation that a given common increase in demand throughout an urban area leads to relatively smaller price responses and relatively greater quantity responses further from CBDs (Genesove & Han, 2013).

This positive CBD distance profile can be largely explained by neighborhood-level factors that affect development costs. The model predicts that the extensive margin of supply is more responsive in tracts where fixed cost distributions have a fatter right tail. Such easier development conditions are associated with sparser initial development, flatter land and looser regulation. Panel B of Figure 1 shows that the average tract in our data is almost 60% developed at the CBD but less than 10% developed at the region edge. However, flat land declines from 45% to 38% from CBDs to region edges and land use is more regulated at 30% of the way to region edges than at CBDs.

To see how much these tract-level factors matter for the CBD distance effect, column (2) of Table 5 expands the specification in column (1) by adding interactions between $\Delta \ln P$ and the 2001 fraction of land developed in each tract and the 2001 fraction of flat land in each tract.[22] Consistent with the model, we find that supply elasticity declines with developed fraction and increases with the fraction of flat land. Moreover, CBD distance coefficients turn negative. This change in

---

[22] As the CBD distance squared interaction is no longer significant, we drop this variable. Neither fraction flat squared nor fraction flat X fraction developed price interactions are significant, so we exclude these variables as well. Appendix Table A5 presents analogous parameter estimates for specifications that are quadratic in tract developed fraction.

coefficient sign reflects the significant negative correlation between the fraction of developed land and CBD distance seen in Figure 1.

As in Table 4, we decompose total unit supply responses in column (2) into components. Results in Columns (3)-(4) show the influences of tract-level factors on new construction elasticities, both for all new development and redevelopment, are attenuated versions of those for total unit supply in column (2). At CBDs, predicted supply elasticities using estimates in Columns (2)-(3) and quantities in Figure 1 Panel B show that the entire unit supply response to price growth is through fewer teardowns and renovations. Conditional on developed fraction and topography, the new construction elasticity declines with CBD distance, likely reflecting increasing regulation. However, the declining developed fraction in CBD distance outweighs this residual influence of CBD distance such that at 50 percent of the way to metro region edges, the predicted overall unit supply elasticity is 0.40 while that for new construction rises to 0.23. Developed fraction and topography affect the redevelopment supply elasticity in the same direction as the new construction elasticity, consistent with model predictions. Interestingly, topography matters for teardowns. Comparisons of results in columns (4), (5) and (6) indicate that positive demand shocks precipitate more teardowns and redevelopment in neighborhoods with flatter land, possibly due to the fact that land assembly for demolition and rebuilding is easier in flatter areas (Dye & McMillen, 2007).

Estimates in columns (7)-(10) for floorspace supply mostly mirror those for unit supply in columns (2)-(6). In particular, developed land and CBD distance have negative effects on supply elasticities based on total changes and new construction but have no effects on teardowns and renovation of existing floorspace. Similar to units, we observe a larger loss of existing floorspace supply due to teardowns and full depreciation in flatter areas. Finally, column (11) of Table 5 focuses on land development. These patterns are remarkably consistent with those for the unit supply (column 2) and for the floorspace supply (column 7), though with attenuated coefficients.[23]

Conditional on topography and developed fraction, the negative CBD distance effects on unit supply, floorspace supply and land development in Tables 5 and A5 likely reflect the impact of local regulations, as regulations increase on average with CBD distance within metro regions over the range well covered by our data (Figure 1 Panel B). Columns (1)-(8) of Table A6 expand the

---

[23] Estimates analogous to those in Table 5 using the Zillow hedonic price index instead are very similar (unreported).

main specification in Table A5 to include interactions of price changes with the Wharton Regulation Index (WRLURI), measured at the municipality level, or tract level residential floor area ratio (FAR) building restrictions. As expected, estimated impacts of regulation are if anything negative. Moreover, incorporating regulation moves coefficients on CBD distance from significantly negative to insignificantly negative or positive, depending on the outcome. Our evidence is consistent with border discontinuity evidence in Shanks (2021) and Chiumenti, Kulka, and Sood (2022), finding that greater municipal regulation increases lot sizes and prices of single-family homes. As the WRLURI is only observed for 40 percent of our primary sample observations and FAR is only observed for 6 cities, Table A6 estimates are not useful for predicting supply elasticities for most tracts.[24]

5.3 Introducing Unobserved Heterogeneity

So far we have controlled for metropolitan area level heterogeneity by including region fixed effects. However, we have not allowed parameters governing tract level heterogeneity in supply elasticities to vary across metropolitan areas. The same fraction developed land in two neighborhoods may be associated with different local supply elasticities in different cities, as in topographically and regulatory constrained San Francisco versus relatively unconstrained Wichita. To more flexibly incorporate metro-level heterogeneity in tract supply elasticities, we extend the environment to a finite mixture model (FMM) with two latent classes for coefficients on main and price-interaction supply factors in the "second stage" equation of the IV model described above. That is, we recover class-specific coefficients on price growth, developed fraction, fraction flat, CBD distance, and these three supply factors interacted with price growth. Coefficients on all other controls and region fixed effects are constrained to be the same across classes. To maintain statistical power, we also retain a single "first stage" equation. Because estimation mechanics uses both generated and residualized regressors, we must bootstrap standard errors. We use a spatial block bootstrap in which blocks are constructed as 8X8 km grid cells, with one centered on each CBD. [25]

---

[24] We explore a number of plausible supply factors as additional price growth interactions. One of note is that supply elasticity is larger in tracts that are bordering or crossed by a highway (Table A6). Better highway access may lower the cost of accessing construction materials and workers or lead to less land use regulation.

[25] We also experimented with specifications which index more parameters by class and with three class models. Adding the flexibility of indexing control variables and fixed effects by class increases standard errors and makes model convergence more difficult. Adding a third class results in key interaction parameter estimates that are very

Beyond adding additional parameter flexibility, the finite mixture formulation allows us to more closely connect our analysis to existing metropolitan area level evidence on supply elasticities. To achieve this, we allow the probability of being assigned to each latent class to depend on three metro area level attributes, following Saiz (2010). These are the fraction of land that is developed within 50 km of the CBD, the fraction of area that is lost to hills, water and wetland within the same radius, and the metropolitan area level Wharton Regulation Index. For robustness, we also construct the first two variables for areas within 10 km and 20 km of the CBD, 10%, 50% and 100% of the maximum distance from the CBD to the metro edge, and for the entire metropolitan area. Conceptually, these metro area level predictors may influence neighborhood supply elasticities through their impacts on initial prices and attributes of fixed development cost distributions that are common across tracts in a given metro area. We emphasize, however, that this is primarily a predictive exercise. There may well be correlates of these three metro level factors that are the true drivers of cross-metro variation in housing supply elasticities.[26] [27]

Table 6 presents the finite mixture model estimates. Panel A presents logit coefficients predicting membership in latent class 2, which is the less elastic class. Panel B presents class-specific coefficients. We show results for four key outcomes of interest.

Results in Panel A echo evidence from Saiz (2010) that metro area level factors matter for supply elasticities. For all supply components investigated, greater metro developed fraction increases the probability of being in the less elastic latent class. Conditional on initial development density, metros with a higher fraction lost to hills, water and wetland also have a higher probability of belonging to the less elastic class in total units and total floorspace supply elasticities, with insignificant positive estimates for new construction. Metro-level regulation is associated with less elastic new construction supply for both units and floorspace. However, total floorspace supply elasticities are positively related to regulation. This is evidence that building owners in more regulated areas are more likely to expand floorspace through renovation. Overall, results in Panel A show that natural and policy constraints at the metro level are key determinants of

---

similar across two of the three classes, thereby providing little additional information about supply elasticities.

[26] In the parameteric example from the model, the land development component of the supply elasticity is $\alpha^{-1-\lambda}\lambda\rho_i^{-\lambda}P_i^{-\frac{\lambda}{\alpha}}\Gamma_i$. We imagine, for example, that the Frechet shape parameter on the fixed development cost distribution λ may differ by metro area.

[27] An alternative (or complementary) approach would be to more flexibly allow the intercept to depend on metro area level factors, identifying off of metro area level variation in average price growth using Bartik type instruments. Unfortunately in our sample of metros, first stage power is not sufficient for this procedure to deliver precise estimates.

neighborhood supply elasticities.

Results in Table 6 Panel B are fully consistent with their IV counterparts reported in Table 5. Class-specific coefficients replicate the patterns in Table 5, though class 2 coefficients are muted. Across all supply measures, we find that the estimated effects of tract-level factors on supply elasticities conditional on being assigned to the more elastic class are 10-20 times larger than those conditional on being assigned to class two. Metros with environments that are not conducive to development have much more similar supply elasticities across tracts than those with lots of developable land and low regulation.

The bottom three rows of Table 6 provide summary statistics about the in-sample class-specific elasticities calculated with the finite mixture model estimates. The average unit supply elasticity is 0.56 conditional on a tract being in the more elastic class and 0.21 conditional on a tract being in the less elastic class. Standard deviations of predicted supply elasticities are much greater in class 1 than class 2 for all outcomes. Table A7 presents analogous parameter estimates for the specifications that are quadratic in tract developed fraction.

5.4 Predicted Supply Elasticities

We now characterize tract-level supply elasticities predicted using our empirical model. Because the FMM-IV estimates in Table 6 accommodate richer heterogeneity, we focus on the elasticities produced using these estimates. While the estimation sample is limited, we use coefficient estimates to predict supply elasticities for all 54,310 census tracts in the 306 metro regions in our data. We focus on supply responses for "units", new construction units ("new units"), floorspace, and floorspace in new construction units ("new floorspace"). Prediction standard errors are calculated using a parametric bootstrap of 100 draws from the joint normal distribution of FMM-IV estimated parameters. Recall that the new units and new floorspace supply responses we calculate are not technically elasticities, but instead are calculated as $\frac{d[Q_{ir}^N/Q_{ir}^{2000}]}{d\Delta \ln P_{ir}}$, where Q_{ir}^N is the number of new units or amount of floorspace in new units constructed in calendar years 2001-2010. d denotes cross-sectional differences and Δ denotes differences over time. Figure 2 shows kernel density graphs of these four measures and their lower and upper 95% confidence bands. Average predicted units, new units, floorspace and new floorspace supply elasticities are 0.29, 0.15, 0.51 and 0.19, respectively.

For comparison, we also make available predicted supply elasticities based on the FMM-IV model that is quadratic in developed fraction and the simple IV linear and quadratic in developed fraction models. Figure A3 depicts associated kernel densities. Compared with the FMM-IV estimates, distributions of tract IV estimates typically have longer left tails. IV estimates disproportionately weight large metros, which tend to have lower predicted elasticities given their more intensely developed land. As such, the FMM-IV estimates are likely to apply more accurately in smaller metros. The quadratic estimates are somewhat wilder, as seen in their greater dispersion across tracts. As only 18 percent of observations fall beyond halfway to the city edge, the quadratic coefficient is mostly identified from variation near the CBD and is thus more noisily estimated. The linear estimates are better powered but do not do as well at capturing nuanced differences between tracts. Figure A3 also depicts predicted land development elasticities. As discussed above, we suspect these are under-estimates given imperfect measurement.

Because our empirical setup is oriented toward exploiting within rather than between metro variation, we emphasize that our predicted supply elasticities may not fully capture variation in average tract elasticities between metros. For example, we identify the constant (uninteracted) impact of price changes on quantities by comparing tracts in the same metro at the same CBD distance, developed fraction and topography with different exogenous price growth. The FMM-IV estimates show very different constants for the two latent classes and indicate that supply parameters are likely to be different across different types of metro areas. Because of the necessarily limited flexibility of parameters to differ across metros, variation in levels of supply elasticities calculated using our estimates are less well identified than their within-metro dispersion. Nevertheless, in Section 6 we demonstrate that the FMM-IV estimates are considerably more informative than the simple IV estimates about aggregate metro supply elasticities.

While we use empirical estimates to predict tract-specific supply elasticities for each tract in our data, these elasticities should be viewed as averages across tracts with similar observables. Tracts with the same CBD distance, developed fraction and topography in the same metro are assigned the same predicted elasticity even if they may be subject to different zoning restrictions. As a result, rather than being used to evaluate policies affecting specific tracts, our predicted supply elasticity estimates are more suitable in policy evaluation for a broad set of tracts, across which idiosyncratic differences in local regulation can average out.

5.4.1 Spatial Patterns

Figure 3 shows how and why supply elasticities differ by CBD distance. Each panel shows local first-degree polynomial smoothed plots of actual predicted supply elasticities and three counterfactual elasticities holding certain tract attributes constant. Predicted supply elasticities (solid lines) are almost indistinguishable from supply elasticities that hold the fraction of flat land constant with CBD distance at metro means (long dash-dot lines). Each supply elasticity increases with CBD distance, flattens in the suburbs, then increases again toward the urban fringe. To understand this pattern, note our evidence from Figure 1 that developed fraction declines monotonically in CBD distance, yet land use regulation is higher in suburbs than in central cities. These two forces offset to keep average supply elasticities constant in the suburbs from approximately 30 to 90 percent of the way from CBDs to metro area edges, before the developed fraction effect again dominates, pushing supply elasticities up at metro edges.[28]

The remaining two lines in each panel represent predicted supply elasticities when holding the initial developed fraction at each metro region's mean and when holding both developed fraction and topography at the metro means. The two lines coincide, indicating that, on average, topography alone does not play a big role in explaining the CBD distance pattern in local supply elasticities. Under both counterfactuals, we see the supply elasticity falling over the full range of CBD distance for all quantity measures. Mechanically, this is because of the negative CBD distance coefficients in Table 6, which likely capture increasing regulation with CBD distance. Figure A4 shows predicted unit and floorspace supply elasticities by CBD distance in six select metros. It indicates marked divergence of these two objects with CBD distance in each city, reflecting the more floorspace intensive construction in suburban areas. Patterns of elasticities comparing across cities are as expected given differences in initial development intensity and regulation.

5.4.2 Relating Estimates to Model Predictions

Our parameterized supply model in Section 3.1 predicts that the elasticity of floorspace supply per parcel for new construction is $\frac{1-\alpha}{\alpha}$, where α is the land share in a Cobb-Douglas housing production function. The literature estimates the land share to be at most 0.35 (Combes et al.,

---

[28] Plots using predictions from the specification that is quadratic in developed fraction instead look similar except for sharper increases near the edge.

2021), which implies the elasticity of floorspace per parcel in new construction should be at least 1.86.[29] Predicted floorspace and unit elasticities reported in Figure 2 do not map directly to the model, as they apply to changes in the total housing stock over time, not just changes due to new construction. The "new units" and "new floorspace" plots in Figure 2 are not elasticities but are built as components of total changes. As our focus is on calculating supply elasticities from all sources, we do not have the most appropriate empirical setting to estimate elasticities that apply to new construction flows only. These must come entirely from cross-sectional comparisons, whereas our empirical strategy exploits relative changes across tracts for identification. As such, our goal here is only to show that our estimates are roughly commensurate with evidence in the housing production literature.

To make such connections, we calibrate new construction elasticities using the new floorspace, new units and land development predictions reported in Figures 2 and A3. We calculate $\frac{d\ln S_i^n}{d\ln P_i}$ as $\frac{S_i}{S_i^n}\frac{d\widehat{[S_i^n/S_i]}}{d\Delta\ln P_i}$ and $\frac{d\ln H_i^n}{d\ln P_i}$ as $\frac{H_i}{H_i^n}\frac{d\widehat{[H_i^n/H_i]}}{d\Delta\ln P_i}$, where $n = N$ denotes new construction of all types and $n = U$ denotes that on undeveloped land only. Given evidence in Tables 4 and 5 that redevelopment supply elasticities are near 0, we assume $\frac{d\widehat{[S_i^U/S_i]}}{d\Delta\ln P_i} = \frac{d\widehat{[S_i^N/S_i]}}{d\Delta\ln P_i}$ and $\frac{d\widehat{[H_i^U/H_i]}}{d\Delta\ln P_i} = \frac{d\widehat{[H_i^N/H_i]}}{d\Delta\ln P_i}$ for the purpose of these calculations. For land, we can only measure $\frac{d\ln L_i^U}{d\ln P_i} = \frac{L_i}{L_i^U}\frac{d\widehat{[L_i^U/L_i]}}{d\Delta\ln P_i}$. Adjustments use 2000-2010 flows and 2000 stocks. Identification from comparison of ex-ante observationally identical tracts prior to experiencing demand shock variation justifies equating $d\ln P_i$ with $d\Delta\ln P_i$ and allows 2000 stocks to cancel.

As there is wide dispersion across tracts in the ratio of stocks to new construction or land development flows, implied new construction elasticities vary quite sensitively across tracts. Nevertheless, our estimates for the median tract are sensible and in line with housing production function estimates. Calculated using the linear FMM-IV specification, medians of implied tract distributions of the new construction floorspace and units elasticities $\frac{d\ln S_i^N}{d\ln P_i}$ and $\frac{d\ln H_i^N}{d\ln P_i}$ are both 1.9. For new construction on undeveloped land ($\frac{d\ln S_i^U}{d\ln P_i}$ and $\frac{d\ln H_i^U}{d\ln P_i}$), these estimates are 3.2 and 2.8, respectively, with a corresponding new land development elasticity $\frac{d\ln L_i^U}{d\ln P_i}$ of 0.2. Medians of the tract distributions of new construction floorspace and units on undeveloped land relative to the new land development elasticity ($\frac{d\ln S_i^U}{d\ln P_i} - \frac{d\ln L_i^U}{d\ln P_i}$ and $\frac{d\ln H_i^U}{d\ln P_i} - \frac{d\ln L_i^U}{d\ln P_i}$) are 3.0 and 2.8. Given Cobb-

---

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

Douglas production, the floorspace estimate thus justifies a land share in housing production of 0.26.

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 metro level estimates, in this section we explore the aggregation of tract based supply elasticities to larger spatial units. Aggregation brings up a number of conceptual and practical challenges. The typical approach, for example in Saiz (2010), has been to use metro level labor demand and/or population supply shocks to deliver housing demand shocks. However, metro level studies may find supply elasticities that weigh certain types of neighborhoods above their share of metro populations. As neighborhoods are linked in the residential demand system, metro level demand shocks of the same size but aggregated from different combinations of changes in neighborhood fundamentals can imply different aggregate housing supply elasticities. Because of this sensitivity to setting, here we provide a few examples of the macro supply elasticities implied from some simple broad-based shocks. Context matters and neighborhood level supply elasticities must be aggregated as appropriate to the application at hand. While our narrative in this section is about aggregation from census tract to metro region levels, the same logic can be applied to recover elasticities for any tract aggregate spatial units.

The tract level supply elasticity for supply measure Q, γ_{ir}^{Q}, generically aggregates to a metro region level elasticity, γ_{r}^{Q}, as follows. Aggregating tract level supply growth to the metro level means taking a sum weighted by initial neighborhood shares of the housing stock:

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

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

$$\gamma_r^{Q} \equiv [\sum_i \frac{Q_{ir}}{Q_r} \gamma_{ir}^{Q} \Delta \ln P_{ir}] / [\sum_i \frac{Q_{ir}}{Q_r} \Delta \ln P_{ir}]. \tag{20}$$

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

We first consider the case in which all neighborhoods simultaneously experience identical housing demand shocks. Because of differing housing supply elasticities, these shocks manifest themselves as different combinations of housing price and quantity changes, depending on the neighborhood. In the context of the model from Section 3.2, the tract level expenditure share on housing is $\frac{S_{ir}P_{ir}}{\overline{y}_{ir}\pi_{ir}} = 1 - \beta$. Therefore, if the demand shock changes aggregate expected income by the same percentage in every neighborhood, $\Delta \ln S_{ir} + \Delta \ln P_{ir}$ is a constant; call this quantity x. This would happen if the outside option $\ln \mu$ in (7) changes, thereby leaving no scope for households to substitute across neighborhoods in response to this shock.

The resulting metro housing supply elasticity is a weighted average of tract-level elasticities, where the weight is the initial housing share adjusted for the neighborhood supply elasticity.[30] Using $\Delta \ln P_{ir} = x/(1 + \gamma_{ir})$ and (20) to aggregate over tracts, we have

$$\frac{\Delta \ln Q_r}{\Delta \ln P_r} = \gamma_r^1 = \sum_i \frac{\frac{H_{ir}}{1+\gamma_{ir}}}{\sum\limits_i \frac{H_{ir}}{1+\gamma_{ir}}} \gamma_{ir}. \tag{21}$$

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

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

---

[30] To have complete coverage, we use housing units in the 2000 census for all weighting.

are perfect substitutes. The resulting metro-level supply elasticity is

$$\frac{\Delta \ln Q_r}{\Delta \ln P_r} = \gamma_r^2 = \sum_i \left[\frac{H_{ir}}{H_r} \gamma_{ir} \right]. \tag{22}$$

Figure 4 presents the distributions of our two aggregate supply elasticity measures γ_r^1 and γ_r^2, focusing on the unit and floorspace elasticities $\frac{\Delta \ln H_r}{\Delta \ln P_r}$ and $\frac{\Delta \ln S_r}{\Delta \ln P_r}$. We report two versions of each, derived from linear FMM-IV and linear IV empirical specifications. Vertical lines in Figure 4 show that the linear FMM-IV specification yields a mean unit supply elasticity across metro regions of 0.38 for γ_r^1 and 0.41 for γ_r^2. Analogous mean floorspace elasticities are 0.61 and 0.63, respectively. As it accommodates more demand substition across neighborhoods, γ_r^2 first order stochastically dominates γ_r^1 in all cases. When neighborhoods become less substitutable, tracts with smaller γ_{ir} receive more weight in aggregation, as households in these tracts have less room to be relocated to more elastic tracts to accommodate positive demand shocks. Thus $\gamma_r^1 < \gamma_r^2$.

Patterns in Figure 4 highlight the advantages of using FMM-IV elasticities to capture cross-regional differences in supply elasticities. Because IV elasticities are constrained to be the same for tracts with the same observable supply factors in different regions, any cross-region differences must be driven by variation in the distribution of tract level supply factors. Moreover, as larger regions have more observations and more opportunity to provide the within-region identifying variation needed for estimation, IV estimated parameters are more heavily influenced by variation within larger regions and thus may not apply well to smaller regions. Gaps between IV and FMM-IV distributions in Figure 4 reflect these two forces. The result is more reasonable FMM-IV region supply elasticity estimates. For example, only 1 region (Newark at -0.04 for γ_r^1 and -0.02 for γ_r^2) has negative FMM-IV unit elasticity estimates relative to 5 regions with more strongly negative IV based estimates. In addition, the fact that floorspace FMM-IV elasticities are typically greater than those for units is in line with conceptual predictions, in contrast to the almost identical IV based counterparts.

Comparing the distributions of the metro-level elasticities in Figure 4 with those of the tract-level elasticities in Figure 2 reveals comparable or larger dispersion in housing supply elasticities within than between regions. Both distributions of FMM-IV unit elasticities in Figure 4 have a standard deviation of 0.11 while that for both distributions of floorspace elasticity is 0.13. Analogous numbers for the full tract distributions in Figure 2 are 0.20 and 0.26. Mechanical variance decompositions of the distributions in Figure 2 reveals that 30% of the variation in tract unit elasticities and 24% of tract floorspace elasticities are from cross-region variation. This evidence highlights that the metro-level variation in supply elasticities alone is inadequate for evaluating economic consequences of neighbohood-level changes in market conditions.

Saiz (2010) reports estimates of unit supply elasticities in the period 1970-2000 calculated for 237 metro regions that coincide with those in our data.[31] Saiz' estimates are greater in magnitude and exhibit more dispersion (with a mean of 2.54 and a s.d. of 1.43) than do our metro region estimates (with means of 0.38-0.41 and a s.d. of 0.11). This difference can be explained by the shorter later time period, aggregation, and differences in price and quantity measures. Table A8 replicates Saiz' estimation for the 2000-2010 period using the same set of metros, specifications and variables. The average implied metro supply elasticity declines to 1.25 for 2000-2010, indicating that many metropolitan areas have become increasingly supply constrained over time. Moreover, as seen in (21) and (22), our aggregation weights tracts with initial housing stocks (H_{ir}) more heavily. Since these tracts are typically associated with smaller subsequent housing supply responses, the resulting γ_r^1 and γ_r^2 should be less than γ_r^{Saiz}, which implicitly uses quantity variation from neighborhoods that were more supply responsive to the metro level Bartik shocks. A third difference between two approaches is the underlying measure of home price growth.[32] Figure A5 presents a comparison of the FMM-IV metro elasticities (2000-2010) with the original Saiz elasticities (1970-2000) and their updated counterparts (2000-2010). Despite the differences in magnitudes, the Spearman rank correlation of FMM-IV metro elasticities to Saiz' 1970-2000 elasticities is 0.49. The FMM-IV versions have much higher correlations with Saiz' than do the simpler IV versions.

7 Opportunity Zone Application

Tract level supply elasticities are essential to carry out welfare analysis of neighborhood targeted place based policies. In this section, we work through an example focusing on the census tracts

---

[31] Saiz (2010) regresses the 1970-2000 difference in home price on the corresponding difference in the log number of households at the MSA level, where home price is measured as the metro area median self-reported home value in the census. Change in households is instrumented with a Bartik labor demand shock, log average hours of sun in January and the share of new immigrants in the metro area.

[32] For 2000-2010, the mean (median) MSA-level price growth is 45% (40%) based on census median self-reported home values but 24% (24%) based on the FHFA repeat-sales price index.

designated for economic development as Opportunity Zones (OZ) under the auspices of the US federal "Tax Cuts and Jobs Act" of 2017.

7.1 Recovering the Neighborhood Demand System

In anticipation of our evaluation of the OZ program, we use the structure of our demand model in Section 3.2 to recover estimates of parameters that govern demand substitution patterns across neighborhoods. From (7) and (12), aggregate housing demand in each tract is given by

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

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

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

$$\frac{\Delta \ln S_i^d}{\Delta \ln P_i} = [\eta(\beta - 1) - 1] + s_i[\frac{\psi}{\eta} - 1], \tag{23}$$

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

We develop an "estibration" strategy for η using the structural equation (13) from the model. Using FMM-IV estimates of γ_{ir}^{space} discussed above, estimates of $(\kappa\varepsilon)_r$ from gravity regressions and calibrated values of 0.8 for β and 0.01 for κ as inputs,[33] we estimate various specifications of

---

[33] $\kappa = 0.01$ implies that 2 minutes of commuting reduces full income by 1%. We also tried estimating κ jointly with neighborhood demand parameters, but this yielded implied values of ε that were too low, leading us to our

equation (13) with GMM. This includes a metro-region fixed effects specification and a municipality fixed effects specification where each metro area is divided into 5 municipalities: one for the central city and one each for suburbs in north, south, east and westerly directions. We impose that the error term in (13) is orthogonal to our main instrument $\Delta \ln \widetilde{RMA}_i$. In theory, we can adopt the same strategy to structurally estimate ψ. In practice, we do not have sufficient statistical power to precisely estimate ψ using a municipality level estimation equation. Instead, we infer the importance of municipalities by estimating (13) with and without municipality fixed effects. We use the same estimation sample and tract weights as in Tables 4-6.

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

With municipality fixed effects, the estimate of η is 8.5 ($se = 1.2$). Without municipality fixed effects, we estimate η^R to be 3.9 ($se = 0.2$). This smaller estimate reflects less substitutability between neighborhoods in different municipalities than between neighborhoods in the same municipality. The resulting implied average elasticity of demand for floorspace in each neighborhood is $-0.2 * 3.9 - 1 = -1.8$ while that for units is -0.8. These estimates are similar to estimates in Hanushek and Quigley (1980) and Couture et al. (2019).

7.2 Opportunity Zones

The OZ program was created to incentivize investment in economically distressed communities. Among other incentives, the program provides preferential tax treatment of capital gains for new real estate investments within the census tracts designated by state governors to be in an OZ. Governors could designate 25% of eligible census tracts in their states for OZ status. Eligible

---

"estibration" strategy.

tracts are those in "low income communities" (LICs), which have an individual poverty rate of at least 20 percent and a median family income that is at most 80 percent of the area median, plus adjacent tracts that are sufficiently low income.[34]

The OZ program may boost both the supply and demand for housing in OZ tracts. The reduction in the capital gains tax liability for investors reduces the financing costs of real estate development in these areas, which we model as a reduction in the marginal cost of building housing by Δs log points relative to other tracts. OZ status may additionally spur local governments to invest in tract amenities. We treat this as an outward shift in housing demand by Δd log points in terms of quantities relative to observationally equivalent tracts that do not attain OZ designation. As the OZ program applies to 3,957 urban census tracts in our sample area, we see this treatment as broad based enough such that our calculated impacts of the program can reasonably apply to an average OZ tract.

Our analysis concludes that because of relatively inelastic housing supply in OZ tracts, capital gains tax reductions impart lower average welfare gains than if the same policy were implemented in other neighborhoods. In addition, using metro level rather than tract supply elasticities would deliver misleading conclusions about the welfare consequences of the program. Other evidence on the efficacy of the OZ program is mixed. Comparing OZ to similar looking non-OZ tracts, Chen, Glaeser, and Wessel (2022) finds that the OZ program had little effect on home prices. Arefeva, Davis, Ghent, and Park (2020) finds that the OZ program promoted job creation in OZ relative to non-OZ tracts. However, Freedman, Khanna, and Neumark (2021) finds that OZ status had little effect on resident outcomes. Whatever the true benefits of the OZ program, the primary goal of our analysis is to use the OZ program as an example to demonstrate that variation in housing supply elasticities within metro regions can matter for evaluation of place based policies.

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

$$\ln Q_i^d = d_i + \varepsilon_{Di} \ln P_i$$

$$\ln Q_i^s = s_i \varepsilon_{Si} + \varepsilon_{Si} \ln P_i$$

---

[34] In the US, at least 95 billion dollars are spent annually on place-based policies (Kline & Moretti, 2014). In Europe, 2007-2013 expenditure on "Structural Funds", much of which is targeted at neighborhoods, was 278 billion euros.

Shocks of Δd to demand (in terms of $d \ln Q^d$) and Δs to supply (in terms of $d \ln P^s$) yield the equilibrium price change $\Delta \ln P_i = \frac{\Delta d_i - \varepsilon_{Si} \Delta s_i}{\varepsilon_{Si} - \varepsilon_{Di}}$. To avoid relying strongly on functional form assumptions, this setup is most sensible for analyzing welfare impacts of small supply shocks on consumers and small demand shocks on producers. So that the likely irreversibility of the existing housing stock does not affect conclusions, we only consider shocks that increase housing quantities. Dollar changes in CS and PS for small supply and demand shocks respectively are:

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

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

Associated percentage changes in CS and PS are $\Delta d_i + (1 + \varepsilon_{Di}) \Delta \ln P_i$ and $\varepsilon_{Si} \Delta s_i + (1 + \varepsilon_{Si}) \Delta \ln P_i$.

We measure base year prices using the 2016 repeat sales index and base year quantities of units and floorspace in each tract using Zillow data from 2016. All values are in 2010 dollars. We impose tract unit and floorspace demand elasticities of -0.8 and -1.8 respectively, as calculated in Section 7.1. We use predicted linear FMM-IV tract supply elasticities based on results in Table 6 and tract developed fraction from 2011. To get a sense of the importance of local heterogeneity, we compare results using predicted tract supply elasticities to those using region supply elasticities γ_r^2. Under these assumptions, we calculate changes in CS and PS for all census tracts for which we have repeat sales price index information in 2016, including $3,957$ OZ designated tracts, given either $\Delta s = 0.05$ or $\Delta d = 0.05$. Assuming a capital gain of 25% on an average property and savings of the 20 percent capital gains tax via the OZ program, $\Delta s = 0.05$ seems reasonable.[35]

Table 7 presents the results, in which Panel A examines the market for housing units and Panel B examines the market for floorspace. OZs are primarily located in more urban and developed areas. As a result, supply elasticities in these areas are lower than in the average neighborhood and the average OZ tract has a smaller scope for gains through the program. Because they are greater and uniform across types of neighborhoods, using metro region level supply elasticities overstates gains from the program and understates the dispersion in these gains.

Results in Panel A show that in the market for housing units our assumed supply shock would increase CS by an estimated \$3.3 million on average in OZ tracts, \$3.7 million in other low income

---

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

tracts, \$6.7 million in adjacent tracts and of \$8.3 million in fully ineligible tracts. Imposing metro level supply elasticities instead would imply CS changes of \$4.4 million in OZ tracts and \$7.0 million in fully ineligible tracts, with smaller standard deviations for all types of locations.

Increases in PS for 5% demand shocks are greater in magnitude than their CS counterparts, as housing demand is more elastic than housing supply. But again, because of the relatively inelastic supply in OZ tracts, PS increases the least on average in these locations. The average PS increase of \$18.7 million in OZ tracts is far below the \$28.5 million average increase in fully ineligible tracts, with other LICs and tracts adjacent to LICs in-between. For PS, implications using region rather than tract supply elasticities are more similar than for CS.

Results for the floorspace market in Panel B exhibit similar patterns, though with smaller magnitudes due to the larger demand and supply elasticities for floorspace than units. In particular, the average OZ tract would experience an increase in CS of \$3.2 million relative to \$7.1 million for a typical non-OZ tract. These numbers are slightly more uniform using region supply elasticities. For PS, again our evidence indicates that developers in the average fully ineligible tract would benefit almost twice as much from a demand shock of this type as those in OZ tracts.

8 Conclusions

Since DiPasquale (1999)'s lament about the limited amount of research on housing supply, large and distinct literatures on housing production and housing supply have developed. Evidence from the former finds support for a near Cobb-Douglas form of the production function in land and capital for newly constructed housing conditional on land development. Evidence from the latter shows that land use regulation and topographic conditions influence housing supply elasticities at the metro area level. This supply literature emphasizes how selection of land parcels into development influences overall supply elasticities.

This paper performs a comprehensive analysis of housing supply for the US, moving this literature forward in two main ways while also more closely connecting it to the housing production literature. First, we decompose total floorspace supply responses to price shocks into eight margins, painting a rich picture of the components of housing supply. This aspect of the analysis includes separate consideration of land development, new construction, redevelopment, teardowns, renovation and floorspace per unit. While the new construction component of our analysis conforms quantitatively with evidence in the housing production literature, we demonstrate that new construction accounts for at most two-thirds of the supply response to price shocks in the average neighborhood. Because of housing's durability and variation in land development costs across locations, a complete understanding of the production function for housing is not sufficient to have a full picture of housing supply. As has been emphasized elsewhere, land availability matters. We show that it even matters at the neighborhood level. We also show that the maintenance and renovation margins matter as well, as they can reduce teardowns, thereby keeping vintage properties in the housing stock longer.

Second, we recover variation in each component of neighborhood level supply elasticities as functions of both neighborhood and metro area topographical conditions, development intensity, and regulation. Using our FMM-IV estimates, average predicted elasticities of floorspace and housing unit supply with respect to price are 0.51 and 0.29 across 50,410 census tracts in 306 US metropolitan areas for the 2000-2010 period, with just over half of the unit supply response in the typical tract due to new construction. We find that each of these components of housing supply becomes more elastic moving out from urban centers and that there is more variation within than between metro areas in housing supply elasticities. This pattern is in part but not entirely due to the increasing fraction of land available for development with CBD distance. Initial development intensity, availability of flat land and zoning regimes are all important determinants of local housing supply. Identification comes from variation in labor demand shocks to commuting destinations, as aggregated using insights from a quantitative spatial equilibrium model.

We hope this new evidence on housing supply informs housing affordability policies and aids in assessing the welfare consequences of place based policies more generally. On affordability, our results indicate that intensive margin supply responses, measured in various ways, are important components of supply. Building renovations that add units and reduced teardown rates together account for about 40% of unit supply responses to price growth. As these segments of supply are more likely to serve lower income households and are less sensitive to land availability constraints, policies that make them easier are likely to contribute to improved affordability, particularly in neighborhoods with low new construction supply elasticities because of limited land availability. On place based policies, the Opportunity Zone example shows the implicit costs of targeting inelastic supply neighborhoods with subsidies for housing construction. Our evidence of high within metro variation in housing supply elasticities indicates the importance of considering neighborhood supply conditions in evaluations of the efficacy of neighborhood targeted policies.

Acknowledgements: We thank Camilo Acosta, Rolando Campusano, Ramin Forouzandeh, Ian Herzog and James Macek for excellent research assistance. We thank the editor and three referees for their insightful comments that have greatly improved the paper. We also thank Morris Davis, Gilles Duranton, Chris Foote, Stuart Gabriel, Laurent Gobillon, Dan McMillen and Lindsay Relihan for helpful comments. We gratefully acknowledge financial support from the Social Sciences and Humanities Research Council of Canada. All errors are our own.

References

Ahlfeldt, G. M., & McMillen, D. P. (2014). New estimates of the elasticity of substitution of land for capital.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., & Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, *83*(6), 2127–2189.

Albouy, D., Ehrlich, G., & Shin, M. (2018). Metropolitan land values. *Review of Economics and Statistics*, *100*(3), 454–466.

Andrews, I., Stock, J. H., & Sun, L. (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, *11*, 727–753.

Arefeva, A., Davis, M. A., Ghent, A. C., & Park, M. (2020). Who benefits from place-based policies? job growth from opportunity zones. *Job Growth from Opportunity Zones (July 7, 2020)*.

Bartik, T. J. (1991). Who benefits from state and local economic development policies?

Baum-Snow, N., & Hartley, D. A. (2020). Accounting for central neighborhood change, 1980-2010. *Journal of Urban Economics*.

Been, V., Ellen, I., & Gedal, M. (2009). Teardowns and land values in new york city. In *Lincoln institute working paper.*

Borusyak, K., & Hull, P. (2020). *Non-random exposure to exogenous shocks: Theory and applications* (Tech. Rep.). National Bureau of Economic Research.

Brooks, L., & Lutz, B. (2016). From today's city to tomorrow's city: An empirical investigation of urban land assembly. *American Economic Journal: Economic Policy*, *8*(3), 69–105.

Brueckner, J. K., & Rosenthal, S. S. (2009). Gentrification and neighborhood housing cycles: Will america's future downtowns be rich? *The Review of Economics and Statistics*, *91*(4), 725–743.

Brueckner, J. K., & Singh, R. (2020). Stringency of land-use regulation: Building heights in us cities. *Journal of Urban Economics*, *116*, 103239.

Busso, M., Gregory, J., & Kline, P. (2013). Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review*, *103*(2), 897–947.

Calabrese, S. M., Epple, D. N., & Romano, R. E. (2011). Inefficiencies from metropolitan political and fiscal decentralization: Failures of tiebout competition. *The Review of Economic Studies*, *79*(3), 1081–1111.

Chen, J., Glaeser, E., & Wessel, D. (2022). Jue insight: The (non-) effect of opportunity zones on housing prices. *Journal of Urban Economics*, 103451.

Chiumenti, N., Kulka, A., & Sood, A. (2022). How to increase housing affordability? understanding local deterrents to building multi-family housing.

Combes, P.-P., Duranton, G., & Gobillon, L. (2021). The production function for housing: Evidence from france. *Journal of Political Economy*, *129*(10), 2766–2816.

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.

Davis, M., Gregory, J., & Hartley, D. (2019). The long-run effects of low-income housing on neighborhood composition.

DiPasquale, D. (1999). Why don't we know more about housing supply? *The Journal of Real Estate Finance and Economics*, *18*(1), 9–23.

Dye, R. F., & McMillen, D. P. (2007). Teardowns and land values in the chicago metropolitan area. *Journal of urban economics*, *61*(1), 45–63.

Epple, D., Gordon, B., & Sieg, H. (2010). A new approach to estimating the production function for housing. *American Economic Review*, *100*(3), 905–24.

Favilukis, J., Mabille, P., & Van Nieuwerburgh, S. (2019). Affordable housing and city welfare.

Fogli, A., & Guerrieri, V. (2019). *The end of the american dream? inequality and segregation in us cities* (Tech. Rep.). National Bureau of Economic Research.

Freedman, M., Khanna, S., & Neumark, D. (2021). Jue insight: The impacts of opportunity zones on zone residents. *Journal of Urban Economics*, 103407.

Genesove, D., & Han, L. (2013). A spatial look at housing boom and bust cycles. In *Housing and the financial crisis* (pp. 105–141). University of Chicago Press.

Glaeser, E. L., Gottlieb, J. D., & Tobio, K. (2012). Housing booms and city centers. *American Economic Review*, *102*(3), 127–33.

Glaeser, E. L., & Gyourko, J. (2005). Urban decline and durable housing. *Journal of political economy*, *113*(2), 345–375.

Glaeser, E. L., Gyourko, J., & Saks, R. E. (2005). Urban growth and housing supply. *Journal of economic geography*, *6*(1), 71–89.

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, *110*(8), 2586–2624.

Goodman, A. C. (2005). Central cities and housing supply: Growth and decline in us cities. *Journal of Housing Economics*, *14*(4), 315–335.

Gorback, C. S., & Keys, B. J. (2020). *Global capital and local assets: House prices, quantities, and elasticities* (Tech. Rep.). National Bureau of Economic Research.

Guerrieri, V., Hartley, D., & Hurst, E. (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics*, *100*, 45–60.

Hanson, A. (2009). Local employment, poverty, and property value effects of geographically-targeted tax incentives: An instrumental variables approach. *Regional Science and Urban Economics*, *39*(6), 721–731.

Hanushek, E. A., & Quigley, J. M. (1980). What is the price elasticity of housing demand? *The Review of Economics and Statistics*, 449–454.

Kline, P., & Moretti, E. (2014). Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valley authority. *The Quarterly Journal of Economics*, *129*(1), 275–331.

Lee, D. S., McCrary, J., Moreira, M. J., & Porter, J. R. (2021). *Valid t-ratio inference for iv* (Tech. Rep.). National Bureau of Economic Research.

Lutz, C., & Sand, B. (2019). Highly disaggregated land unavailability. *Available at SSRN 3478900*.

McMillen, D., & O'Sullivan, A. (2013). Option value and the price of teardown properties. *Journal* *of Urban Economics*, *74*, 71–82.

Murphy, A. (2018). A dynamic model of housing supply. *American Economic Journal: Economic Policy*, *10*(4), 243–67.

Olea, J. L. M., & Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics*, *31*(3), 358–369.

Orlando, A., & Redfearn, C. (2021). Housing supply elasticities: A structural vector autoregression approach. *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.

Roback, J. (1982). Wages, rents, and the quality of life. *Journal of political Economy*, *90*(6), 1257–1278.

Rosenthal, S. S. (2018). Owned now rented later? housing stock transitions and market dynamics.

Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics*, *125*(3), 1253–1296.

Severen, C. (2019). Commuting, labor, and housing market effects of mass transportation: Welfare and identification. *The Review of Economics and Statistics*, 1–99.

Shanks, B. (2021). Land use regulations and housing development: Evidence from tax parcels and zoning bylaws in massachusetts.

Tsivanidis, N. (2022). The aggregate and distributional effects of urban transit infrastructure: Evidence from bogotas transmilenio. *Unpublished manuscript*.

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

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

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

Table 1: Summary Statistics

	Mean	St Dev	Obs	Tracts
Tract Home Price Changes, Estimation Sample				
Repeat Sales Index, 2000-2006	0.64	0.35	30,502	24,233
Hedonic Index, 2000-2006	0.62	0.34	29,274	23,239
Repeat Sales Index, 2000-2010	0.25	0.38	30,840	24,532
Hedonic Index, 2000-2010	0.25	0.35	29,424	23,378
Census Index, 2000-2010	0.54	0.28	30,694	24,417
Tract Housing Quantity Changes, Estimation Sample				
Stock of Housing Units, Census, 2000-2010	0.08	0.22	30,840	24,532
New Units, Census/ACS, 2000-2009	0.12	0.20	30,829	24,521
Stock of Housing Units, Zillow, 2000-2010	0.07	0.19	30,840	24,532
New Units, Zillow, 2000-2006	0.08	0.15	30,840	24,532
New Units, Zillow, 2000-2010	0.10	0.18	30,840	24,532
New Units on Developed Land, 2000-2010	0.03	0.06	30,840	24,532
Teardown+Renov. Units, Zillow, 2000-2010	-0.03	0.10	30,836	24,528
Renov. Units, Zillow, 2000-2010	-0.00	0.06	30,380	24,095
Floorspace, Zillow, 2000-2010	0.14	0.26	30,384	24,099
New Floorspace, Zillow, 2000-2010	0.13	0.22	30,395	24,110
Teardown/Renovation Flrspc, Zillow, 2000-2010	0.01	0.16	30,379	24,094
Renovation Flrspc, Zillow, 2000-2010	0.03	0.13	30,380	24,095
Developed Land, 2001-2011	0.08	0.12	30,840	24,532
Tract Level Supply Influencers, Estimation Sample				
Fraction of Land Area Developed, 2001	0.33	0.21	30,840	24,532
Fraction of Land Area Flat	0.41	0.43	30,840	24,532
Wharton Real Estate Index (municipality level variation)	0.28	1.02	12,367	10,016
Residential Floor Area Ratio (8 cities)	1.76	1.50	2,128	1,714
Fraction of Way from CBD to Metro Edge	0.27	0.21	30,840	24,532
Tract Employment and Population Variables				
Tract Employment, 2000-2010, Regions in Est. Sample	-0.19	0.87	56,043	42,755
Tract Level Bartik Instrument, 2000-2006, Regions in Est. Sample	0.08	0.05	56,274	42,902
RMA, 2000-2010, Estimation Sample	0.04	0.05	30,840	24,532
Simulated RMA, 2000-2006, Estimation Sample	0.04	0.01	30,840	24,532

All changes are in percentage terms. The full study region includes the 50,410 unique census tracts in the 306 partially overlapping metro regions with 1990 information on tract employment. The estimation sample includes 24,532 equally weighted unique tracts in 169 regions with at least 10 housing market transactions in 2000 and 2010 in the ZTRAX data. It excludes tracts for which the 2000 ZTRAX housing unit counts are more than 25% below the 2000 census count or tracts with fewer than 500 ZTRAX housing units in 2000. 2.6% of estimation sample tracts experienced zero growth in developed land and 6.1% experienced zero new construction.

Table 2: Tract Level Employment and Housing Market Dynamics

Panel A: Tract Level Regressions of Employment Growth on Bartik Shocks

	(1)	(2)
	Change in Log Employment	
	2000-2006	2000-2010
Tract Bartik Shock, 2000-2006	0.33***	0.60***
	(0.10)	(0.10)
No 1990 Emp. Info.	0.13	0.27
	(0.24)	(0.23)
Observations	51,900	56,043
R-squared	0.06	0.05
Number of Regions	158	169

Panel B: Tract-Level Housing Market Dynamics

	(1)	(2)	(3)
	Δln House Price, 00-10	Δln House Quantity, 00-10	Δln Simulated RMA, 2000-2006
Δln House Price, 1990-2000	-0.25***	0.01	0.0000
	(0.01)	(0.01)	(0.0001)
Δln House Quantity, 1990-2000	-0.02	0.24***	-0.0002
	(0.03)	(0.04)	(0.0002)

Panel C: Analysis of Pre-Treatment Trends, 1990-2000

	(1)	(2)
	Δln House Price, 90-00	Δln House Quantity, 90-00
Δ ln Simulated RMA, 2000-2006	0.26	-0.31
	(0.52)	(0.23)

Panel A: Regressions also include metro region fixed effects, fraction developed in 2001, fraction of tract land that is flat, a cubic in fraction of the way to region edge, log 1990 tract employment and the following tract attributes from 1990 and 2000: census home price index, rent index, log population, log average household income, share black, share white, share college. Sample includes all tracts in metro regions that are in the primary sample. Each tract receives equal weight. Robust standard errors.

Panel B: Each entry is from a separate regression of the variable at top on the variable at left with the indicated fixed effects. Regressions include metro region fixed effects, a cubic in fraction of the way from the CBD to the metro edge, fraction of tract land that is flat, log 1990 tract employment, the 2000-2006 Bartik shock for the tract and the following tract attributes measured in 1990 and 2000: log population, log average household income, share black, share white and share college.

Panel C: 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 region fixed effects. Controls are the same as in Panel A with the addition of a tract 2000-2006 Bartik shock. Standard errors are corrected for spatial autocorrelation up to 16 km.

Table 3: First Stage Results for Baseline Specifications

	(1)	(2)	(3)	(4)
Panel A: Unified Specification				
	Repeat Sales Index		Census Index	RMA
	2000-2006	2000-2010	2000-2010	2000-2010
Δ ln Simulated RMA, 2000-2006	5.94***	5.18***	0.29	0.74***
	(1.26)	(1.10)	(0.96)	(0.17)
Observations	30,500	30,838	30,692	30,838
R-squared	0.14	0.23	0.08	0.04
Number of Fixed Effects	166	167	167	167
Panel B: Interacted Specification				
	ΔlnP	ΔlnP x (CBD Dis)	ΔlnP x %Dev	ΔlnP x %Flat
Δln Simulated RMA	-2.04	-3.56***	-4.82***	-6.42***
	(1.76)	(0.90)	(1.28)	(1.78)
Δln Simulated RMA x (CBD Dis)	2.92	11.12***	0.55	5.12**
	(2.34)	(1.73)	(1.23)	(2.17)
Δln Simulated RMA x %Dev	19.44***	2.93**	20.01***	8.67***
	(4.13)	(1.34)	(3.44)	(2.81)
Δln Simulated RMA x %Flat	-0.06	0.48	0.64	13.01***
	(1.58)	(0.60)	(1.04)	(3.32)
Observations	30,838	30,838	30,838	30,838
R-squared	0.24	0.38	0.28	0.25
Number of Fixed Effects	167	167	167	167

Regressions include metro region fixed effects and the same controls as in Table 2 Panel C. Tracts are equally weighted, even if they appear in multiple metro regions. Standard errors are corrected for spatial autocorrelation up to 16 km using a Bartlett kernel. $\Delta \ln P$ in Panel B refers to the 2000-2010 repeat sales index. Analogous estimates using the hedonic price index instead are very similar.

Table 4: Unified IV Results for Housing Supply

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Total	Total	Total					Total				Total
Subset of Total				New		Remain.			New	Remain.		
Subset of New/Remain.					Redev.		Expan.				Expan.	
	Units	Floorsp.	Units	Units	Units	Units	Units	Floorspace	Floorspace	Floorspace	Floorspace	Land
Time Period	2000-2006	2000-2006	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2001-2011
Repeat Sales Index	0.24***	0.26**	0.35***	0.19**	0.03	0.16**	0.08	0.42***	0.29**	0.13	0.09	0.09
	(0.09)	(0.10)	(0.12)	(0.08)	(0.03)	(0.08)	(0.06)	(0.16)	(0.12)	(0.11)	(0.09)	(0.06)
Obs	30,500	30,048	30,838	30,838	30,838	30,834	30,377	30,381	30,392	30,376	30,377	30,838
Hedonic Index	0.24***	0.27***	0.37***	0.18**	0.03	0.19**	0.09	0.44***	0.32***	0.12	0.08	0.09
	(0.09)	(0.10)	(0.12)	(0.09)	(0.03)	(0.08)	(0.06)	(0.16)	(0.11)	(0.11)	(0.09)	(0.07)
Obs	29,272	29,272	29,422	29,422	29,422	29,418	29,422	29,422	29,422	29,421	29,422	29,422

Regressions include the same controls as those in Table 3 Panel A. The estimation sample for the repeat sales index uses data from 167 metro regions while that for the hedonic index uses data from 165 regions. These samples are reduced 1 region for the repeat sales index floorspace outcomes due to missing floorspace information for some tracts. Entries in Columns 1 and 2 use 164-166 regions. The hedonic index sample excludes tracts in the repeat sales sample containing homes for which age and/or floorspace are not observed in 2000 or 2010. All outcomes are measured using the ZTRAX data except developed land, which uses USGS land cover information. Standard errors are corrected for spatial autocorrelation to 16 km. First stage F-statistics can be determined from results in Table 3.

Table 5: Linear IV Model: Heterogeneity in Supply Elasticities by CBD Distance and Tract Condition

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Total	Total					Total				
Subset of Total			New		Remain.			New	Remain.		
Subset of New/Remain.				Redev.		Expan.				Expan.	
	Units	Units	Units	Units	Units	Units	Floorspace	Floorspace	Floorspace	Floorspace	Land
ΔlnP	0.11	0.94***	0.62***	0.07	0.33**	0.16	0.91***	0.70***	0.15	0.13	0.30**
	(0.12)	(0.25)	(0.20)	(0.05)	(0.14)	(0.10)	(0.31)	(0.26)	(0.20)	(0.17)	(0.13)
ΔlnP x (CBD Dis)	1.89**	-0.58**	-0.43**	-0.13**	-0.17*	-0.08	-0.46*	-0.50**	0.12	0.11	-0.29**
	(0.73)	(0.24)	(0.21)	(0.05)	(0.10)	(0.08)	(0.27)	(0.25)	(0.15)	(0.12)	(0.12)
ΔlnP x %Dev		-1.54***	-1.33***	-0.13	-0.23	-0.15	-1.39***	-1.32***	0.08	-0.06	-0.61***
		(0.40)	(0.36)	(0.09)	(0.15)	(0.11)	(0.45)	(0.40)	(0.23)	(0.18)	(0.21)
ΔlnP x %Flat		0.30**	0.37***	0.08**	-0.09**	-0.01	0.35**	0.46***	-0.15**	-0.07*	0.19***
		(0.12)	(0.12)	(0.03)	(0.05)	(0.03)	(0.14)	(0.15)	(0.07)	(0.04)	(0.07)
ΔlnP x (CBD Dis)^2	-1.84**										
	(0.78)										
Observations	30,838	30,838	30,838	30,838	30,834	30,377	30,381	30,392	30,376	30,377	30,838
Kleib-Paap F-Stat	17.49	11.01	11.01	11.01	11.01	10.18	10.22	10.13	10.17	10.18	11.01

Regressions are the same specification as in Table 4 with the addition of indicated interaction terms. The repeat sales price index measure is used throughout. SE adjusted for spatial autocorr. to 16 km. If included where omitted, coefficients on ΔlnP x (CBD Dis)^2 would be insignificant.

Table 6: Linear Finite Mixture IV Model Housing Supply Interacted Regressions

Outcome	Units	New Units	Floorspace	New Floorspace
Panel A: Logit Parameters for Membership in Class 2 (More Inelastic Supply)				
	(1)	(2)	(3)	(4)
Fraction Developed within 50km of the CBD	5.34***	4.36***	4.37***	3.86***
	(0.61)	(0.64)	(0.60)	(0.59)
Fraction of land within 50km of the CBD Unavail	0.52*	0.29	1.11***	0.40
	(0.29)	(0.26)	(0.31)	(0.27)
Metro Wharton Index	0.03	0.12**	-0.68***	0.10**
	(0.06)	(0.06)	(0.09)	(0.05)
Constant	0.60***	1.10***	0.79***	1.20***
	(0.12)	(0.11)	(0.14)	(0.11)

Outcome	Units		New Units		Floorspace		New Floorspace	
Panel B: Second Stage Estimates								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2
ΔlnP	1.79**	0.29**	1.88**	0.08	2.12**	0.57***	1.43	0.13
	(0.71)	(0.14)	(0.81)	(0.07)	(1.07)	(0.22)	(1.21)	(0.09)
ΔlnP x (CBD Dis)	-1.14**	0.03	-1.14**	0.12**	-1.70**	0.15	-0.90	0.10
	(0.52)	(0.11)	(0.54)	(0.06)	(0.69)	(0.21)	(0.66)	(0.08)
ΔlnP x %Dev	-3.35***	-0.30*	-4.14***	-0.19**	-3.09***	-0.71***	-3.09**	-0.21*
	(0.95)	(0.17)	(1.14)	(0.09)	(1.19)	(0.26)	(1.46)	(0.12)
ΔlnP x %Flat	0.47*	0.03	0.78**	0.06	0.75**	0.09	1.03***	0.08
	(0.27)	(0.05)	(0.32)	(0.04)	(0.34)	(0.08)	(0.39)	(0.05)
Mean Class Probability	0.22	0.78	0.16	0.84	0.24	0.76	0.15	0.85
Mean implied γ	0.56	0.21	0.52	0.08	0.94	0.41	0.59	0.11
SD implied	0.59	0.06	0.74	0.06	0.56	0.16	0.59	0.06

Notes: Sample sizes are the same as in Table 5. Standard errors are bootstrapped with spatially clustered sampling with replacement. Clusters are defined as 8 by 8 km grid squares with one grid square centered at each metro region's CBD.

Table 7: Welfare Consequences of the OZ Program (millions of 2016)

Opportunity Zone	Yes	No, LIC	No, Adjacent	No, Other
Sample Size	2,580	8,776	3,704	15,939
Panel A: Market for Housing Units				
Tract Supply Elasticity	0.20	0.22	0.31	0.32
	(0.17)	(0.17)	(0.19)	(0.17)
Based on Tract Supply Elasticities				
CS ($\Delta s = 0.5, \Delta d = 0$)	3.28	3.74	6.69	8.28
	(4.63)	(4.66)	(7.41)	(8.68)
PS ($\Delta s = 0, \Delta d = 0.5$)	18.65	19.46	23.71	28.47
	(16.76)	(16.31)	(19.46)	(25.76)
Region Supply Elasticity	0.28	0.28	0.29	0.25
	(0.12)	(0.11)	(0.12)	(0.11)
Based on Region Supply Elasticities				
CS ($\Delta s = 0.5, \Delta d = 0$)	4.39	4.67	6.61	7.00
	(3.95)	(4.14)	(6.26)	(6.71)
PS ($\Delta s = 0, \Delta d = 0.5$)	17.30	18.33	23.81	30.03
	(15.82)	(15.36)	(19.37)	(26.86)
Panel B: Market for Floorspace				
Tract Supply Elasticity	0.40	0.43	0.53	0.56
	(0.25)	(0.24)	(0.24)	(0.23)
Based on Tract Supply Elasticities				
CS ($\Delta s = 0.5, \Delta d = 0$)	3.19	3.54	5.76	7.14
	(3.77)	(3.80)	(5.88)	(6.99)
PS ($\Delta s = 0, \Delta d = 0.5$)	8.35	8.77	11.06	13.30
	(7.43)	(7.24)	(8.97)	(11.74)
Region Supply Elasticity	0.51	0.51	0.52	0.47
	(0.15)	(0.15)	(0.15)	(0.14)
Based on Region Supply Elasticities				
CS ($\Delta s = 0.5, \Delta d = 0$)	3.90	4.14	5.70	6.28
	(3.37)	(3.55)	(5.20)	(5.70)
PS ($\Delta s = 0, \Delta d = 0.5$)	7.96	8.45	11.09	13.76
	(7.12)	(6.92)	(8.93)	(12.06)

Tract means with standard deviations in parentheses. Entries are calculated using assumptions about demand and supply shocks indicated in the leftmost column. Estimates use a floorspace demand elasticity of -1.8 and a units demand elasticity of -0.8. In the floorspace market, indicated increases amount to 0.7 percent for OZs and LICs and 0.9 percent in other types of locations. PS increases by 3.2 percent in OZs and LICs and 3.3 percent in other locations. These percentages are undefined in the market for housing units because inelastic demand makes CS infinite.

Figure 1: Development, Topography and Regulation

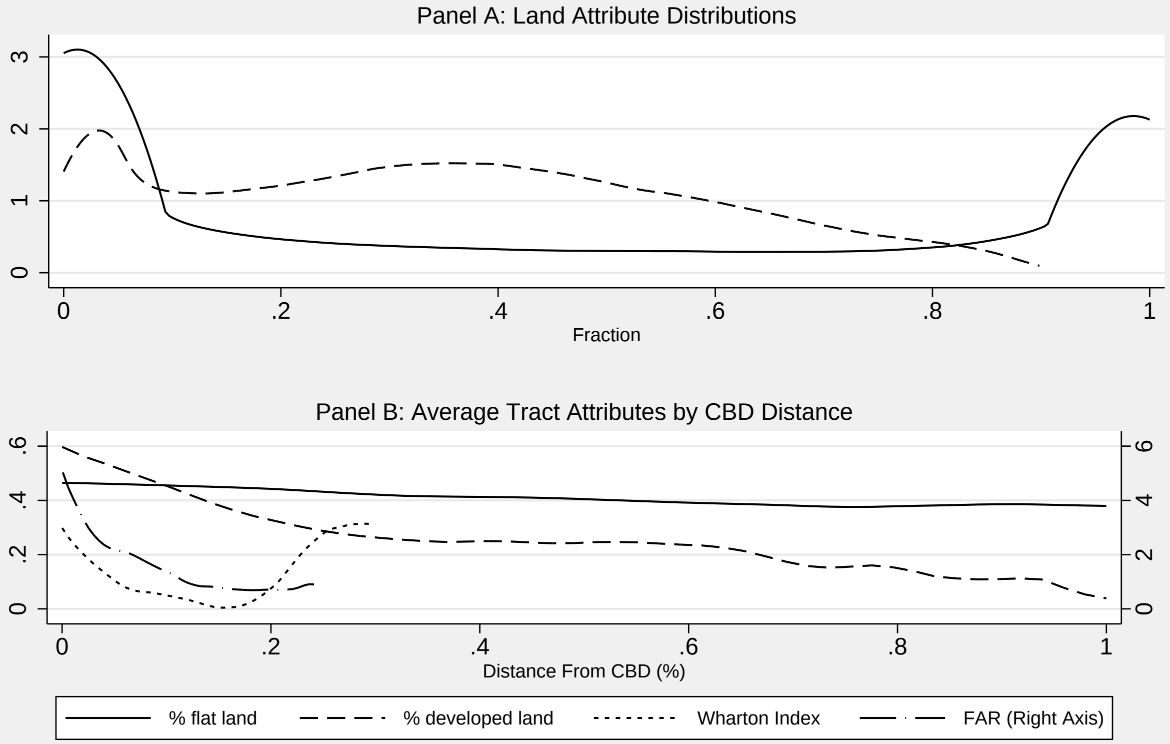

Panel A: Land Attribute Distributions

Panel B: Average Tract Attributes by CBD Distance

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

Figure 2: Kernel Densities of Predicted Elasticities

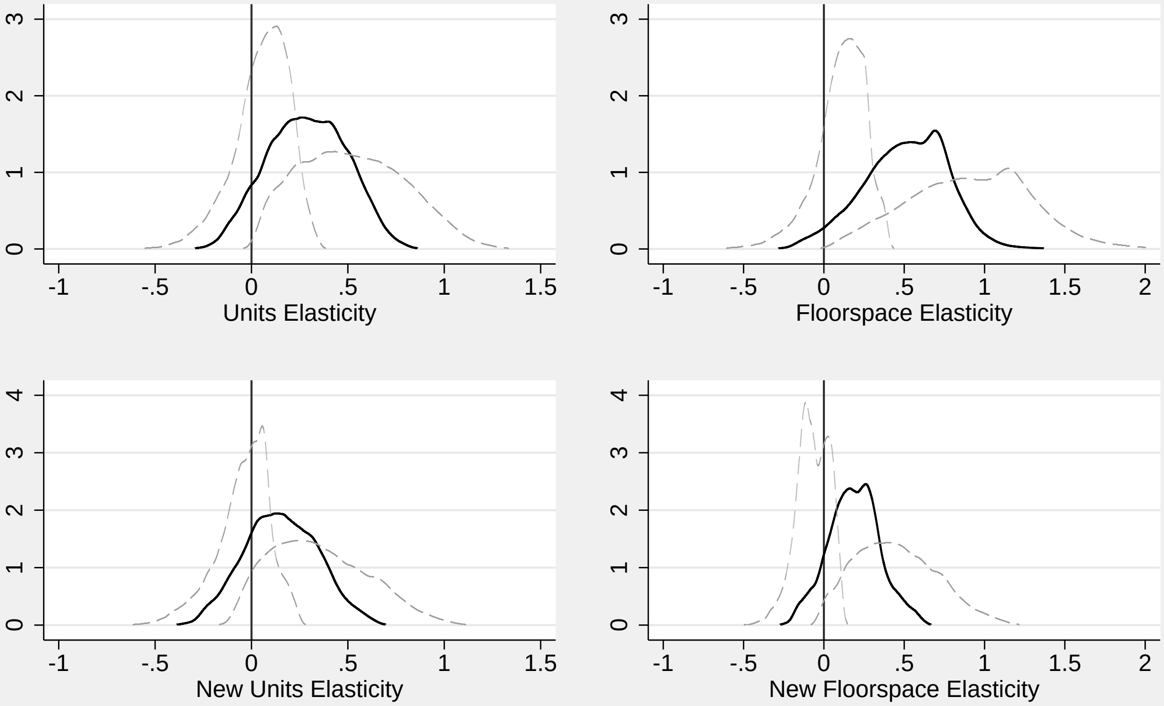

Elasticities are predicted for all metro tracts nationwide using estimates in Table 6. Dashed lines indicate 95% confidence intervals. Comparisons of estimates across all specifications are in Figure A3.

Figure 3: Supply Elasticities by CBD Distance

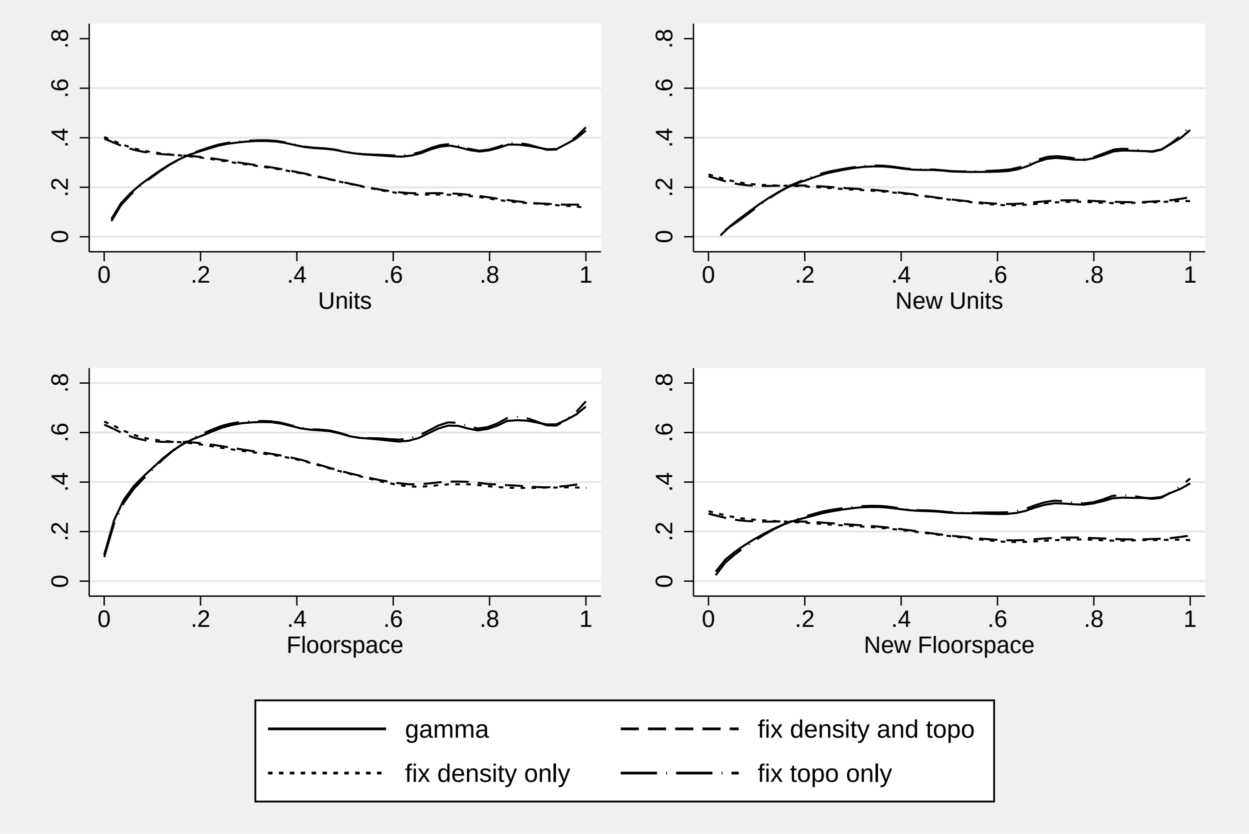

Figure 4: Kernel Densities of Region Supply Elasticities

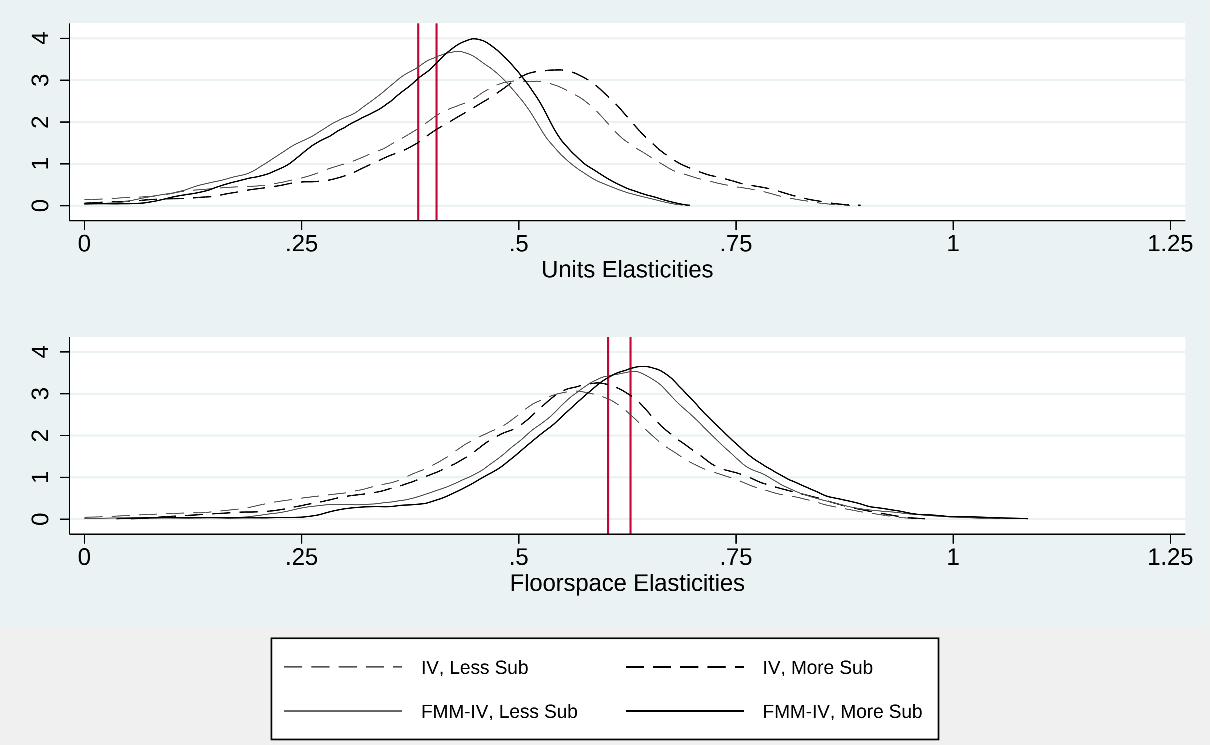

Predicted elasticities in 306 metro regions contribute to all plots. 5 regions have negative units IV elasticities and 1 have negative space IV elasticities. Means of FMM-IV elasticities are indicated with vertical lines.

A Housing Supply Model

This appendix fills in many of the details from the model presented in Section 3.1.

A.1 Setup

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

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

The developer is subject to marginal cost pricing $P_i = \frac{dC_i(A_i)}{dA_i}$ and chooses the optimal amount of floorspace $A_i^* = A_i(P_i)$ accordingly. Fixed lot size and Cobb-Douglas housing production in land and capital would imply $\frac{d \ln C_i(A_i)}{d \ln A} > 1$. The 0 profit condition requires that the bid for each plot of undeveloped land is, after substituting $\frac{dC_i(A_i)}{dA_i}$ for P_i:

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

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

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

A.2 Tract Housing Services Supply

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

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

Differentiating, the supply elasticity decomposes as

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

This expression reflects intensive and extensive margin responses respectively.

A.3 Parameterization with Cobb-Douglas Production

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

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

Floorspace per parcel is

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

and variable cost is

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

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

Plugging in, the resulting elasticity expression is

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

A.4 Parameterization with a Frechet fixed Cost Distribution

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

Plugging into the expression above, the overall supply elasticity is

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

Under Cobb-Douglas production,

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

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

A.5 Land Redevelopment

To see how the model can incorporate redevelopment, consider an environment in which $f_i^L(x)$ is the density of fixed development costs across all parcels in tract i as if they had no prior development and r is an additional fixed redevelopment cost.

The fixed cost distribution for redevelopment, $f_i^R(x)$, is the left tail of the $f_i^L(x)$ distribution up to ϕC_i but shifted by r to reflect the additional redevelopment cost rescaled to integrate to 1. With fraction $\frac{M_i^R}{M_i}$ of parcels up to fixed cost $\overline{g_i}$ previously developed, the new density of redevelopment fixed costs can thus be written as $f_i^R(z) = \frac{f_i^L(z-r))}{M_i^R/M_i}1(z \leq \overline{g_i} + r)$, where $M_i^R/M_i = F_i^L(\phi C_i)$. The fixed cost distribution for developing previously undeveloped land $f_i^U(x)$ is the right tail of $f_i^L(x)$ at fixed costs above ϕC_i only rescaled to integrate to 1. The new density of fixed costs of developing undeveloped land is $f_i^U(z) = \frac{f_i^L(z)}{M_i^U/M}1(z \geq \overline{g_i})$. This is the right tail of the $f_i^L(x)$ distribution. Supports of both distributions include the range between fixed costs ϕC_i and $\phi C_i + r$. Using these ideas, we decompose the fixed cost distribution $F_i(z)$ from Section A.1 into $F_i(z) = \frac{M_i^R}{M_i}F_i^R(z) + \frac{M_i^U}{M_i}F_i^U(z) = F_i^L(z-r)1(z \leq \overline{g_i} + r) + F_i^L(z)1(z \geq \overline{g_i})$.

A corresponding version of equation (3) under redevelopment is then:

$$\gamma_i^{land} \equiv \frac{d\ln L_i(P_i)}{d\ln P_i} = (\frac{M_i^R f_i^R(\phi C_i)}{M_i^{R'}} + \frac{M_i^U f_i^U(\phi C_i)}{M_i^{R'}})\frac{d\ln A_i(P_i)}{d\ln P_i}\phi P_i A_i(P_i), \tag{24}$$

where $M_i^{R'} = M_i^R F_i^R(\phi C_i) + M_i^U F_i^U(\phi C_i)$ is the amount of land that is newly developed. Equation (24) decomposes the land supply response into land redevelopment and new construction on undeveloped land respectively.

Figure A1 shows plots of $f^L(x)$, $f^R(x)$ and $f^U(x)$ for an example tract. A few implications follow. First, as P_i rises, marginal land parcels are developed left to right in the region of overlapping support of the $f^R(x)$ and $f^U(x)$ distributions. Developers in neighborhoods with greater price growth carry out both additional redevelopment and additional development of previously undeveloped land, relative to developers in neighborhoods with smaller price growth, to supply new housing. Second, the relative magnitudes of the land redevelopment versus new land development elasticities depend on the density of the fixed cost distribution $f_i^L(x)$ at ϕC_i and $\phi C_i + r$ and the relative amounts of previously undeveloped versus developed land in the tract. Third, given that $f_i^R(\phi C_i)$ and $f_i^U(\phi C_i)$ both depend on parameters that govern the $F^L(x)$ distribution, tract characteristics affect the redevelopment supply elasticity in the same way as the extensive margin unit supply elasticity in equation (4).

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

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

Relative to an initial baseline, prices rise such that $\overline{g_i} < \phi C_i < \overline{g_i} + r$. As a result, developers draw from both the previously developed and undeveloped land for their new developments.

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

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

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

Differentiating,

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

Therefore,

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

Analogously,

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

A.6 FAR Restriction

One particular tract characteristic considered in the empirical work is a floor-area-ratio (FAR) restriction that constrains developers from building beyond some maximum intensity $\overline{A}_i$. A binding FAR constrains the intensive margin supply response to 0.

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

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

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

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

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

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

$$\frac{d\ln S_i^{FAR}}{d\ln P_i} = \frac{d\ln L_i^{FAR}}{d\ln P_i} = \frac{f_i(P_i\overline{A}_i - C_i(\overline{A}_i))}{F_i(P_i\overline{A}_i - C_i(\overline{A}_i))}\overline{A}_i. \tag{25}$$

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

Equation (25) shows two forces through which a less binding FAR (increase in $\overline{A}_i$) affects the supply elasticity. First, the mechanical effect of being allowed to build more increases supply elasticity by making more parcels viable for development with a marginal price increase. Second, a higher $\overline{A}_i$ attracts more development on available parcels, thereby changing the supply of developable parcels at the marginal fixed cost $P_i\overline{A}_i - C_i(\overline{A}_i)$. This could dampen or increase supply elasticity depending on the form of the fixed cost distribution $f_i(x)$. The net effect of a relaxation of a FAR on supply elasticity is thus an empirical question, which is examined in Section 5.2.

B Housing Demand Model

This appendix fills in many of the details from the model presented in Section 3.2.

B.1 Setup

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

$$v_{ijk\omega} = \frac{v_{i\omega} B_i z_{ijk\omega} w_{jk}}{P_i^{1-\beta} e^{\kappa \tau_{ij}}}, \tag{26}$$

where B_i is a local amenity, w_{jk} is the price of a unit of skill in commuting destination j and industry k, P_i is the price of one unit of housing services in i and $\kappa\tau_{ij}$ is the fraction of time spent commuting for those living in i and working in j. In the data, we observe the price P_i in year 2000 and beyond and the commuting time τ_{ij} in 1990 and 2000. The preference shocks $v_{i\omega}$ are revealed first, leading agents to first choose residential locations anticipating the quality of employment opportunities nearby but before productivity shocks are revealed. Productivity shocks $z_{ijk\omega}$ are then revealed and agents choose work locations second.

The productivity shock $z_{ij\omega}$ is drawn from the Frechet distribution with shape parameter ε.

$$F_z(z_{ijk\omega}) = e^{-z_{ijk\omega}^{-\varepsilon}}, \varepsilon > 1 \tag{27}$$

Productivity dispersion is decreasing in ε.

Following Tsivanidis (2022) and Couture et al. (2019), we specify a nested preference shock over residential locations $v_{i\omega}$. This shock is also distributed Frechet but with shape parameters η and ψ. This nested structure allows individuals to have different elasticities of substitution in demand between neighborhoods within versus between municipalities, where municipalities are indexed by m and $i(m)$ refers to neighborhood i in municipality m.

$$F_v(v_{i\omega}) = exp[-\sum_m [\sum_{i(m)} v_{i\omega}^{-\eta}]^{-\frac{\psi}{\eta}}], \psi > 1, \eta > 1 \tag{28}$$

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 $\frac{B_i z_{ijk\omega} w_{jk}}{P_i^{1-\beta} e^{\kappa\tau_{ij}}}$ alone. As a practical matter, it also delivers a convenient expression for mean income net of commuting cost in each tract, as is derived below. If the distribution functions for the two shocks are identical and $\eta = \psi = \varepsilon$, the utility shock becomes redundant and this model reduces to one similar to that in Ahlfeldt et al. (2015).

Commensurate with the housing supply model in Appendix A, there is a constant elasticity housing supply function with tract-specific elasticity γ_i expressed as $H_i^s = \rho_i P_i^{\gamma_i}$.

We solve the model backwards, first determining commute flows conditional on residential locations and then determining population supply to residential locations anticipating wages net of commuting costs in each residential location.

B.2 Work Location Choice

The fraction of residents of tract i that work in tract j and industry k can be expressed as $\pi_{ijk|i} = Pr(\frac{v_{i\omega} B_i z_{ijk\omega} w_{jk}}{P_i^{1-\beta} e^{\kappa\tau_{ij}}} >= \max_{j',k'} \frac{v_{i\omega} B_i z_{ij'k'\omega} w_{j'k'}}{P_i^{1-\beta} e^{\kappa\tau_{ij'}}})$. Using the properties of the Frechet draws $z_{ijk\omega}$, we have:

$$\pi_{ijk|i} = \frac{\left[w_{jk} e^{-\kappa\tau_{ij}}\right]^{\varepsilon}}{\sum_k \sum_{j'} \left[w_{j'k} e^{-\kappa\tau_{ij'}}\right]^{\varepsilon}}$$

Summing over industries k, we have the commuting probabilities from i to j conditional on living in i.

$$\pi_{ij|i} = \frac{\sum_k \left[w_{jk} e^{-\kappa\tau_{ij}}\right]^{\varepsilon}}{\sum_k \sum_{j'} \left[w_{j'k} e^{-\kappa\tau_{ij'}}\right]^{\varepsilon}} \equiv \frac{\sum_k \left[w_{jk} e^{-\kappa\tau_{ij}}\right]^{\varepsilon}}{RMA_i} \tag{29}$$

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

B.3 Residential Location Choice

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

$$\overline{y}_i = \Gamma(1 - \frac{1}{\varepsilon})(RMA)_i^{\frac{1}{\varepsilon}} \tag{30}$$

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

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 $\frac{\left(B_i P_i^{\beta-1} \overline{y}_i\right)^{\eta}}{\sum_{i' \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} \overline{y}_{i'}\right)^{\eta}}$. The first object is $\frac{\sum_{i' \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} \overline{y}_{i'}\right)^{\eta}]^{\frac{\psi}{\eta}}}{\sum_m [\sum_{i'' \in m} \left(B_{i''} P_{i''}^{\beta-1} \overline{y}_{i''}\right)^{\eta}]^{\frac{\psi}{\eta}}]}$. Plugging $\overline{y}_i$ into both objects yields the population supply function to tract i

$$\pi_i = \mu [\sum_{i\'{} \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} RMA_{i'}^{\frac{1}{\varepsilon}}\right)^{\eta}]^{\frac{\psi}{\eta}-1} \left(B_i P_i^{\beta-1} RMA_i^{\frac{1}{\varepsilon}}\right)^{\eta}, \tag{31}$$

where $\mu = 1/\sum_m [\sum_{i'' \in m} \left(B_{i''} P_{i''}^{\beta-1} RMA_{i''}^{\frac{1}{\varepsilon}}\right)^{\eta}]^{\frac{\psi}{\eta}}$. 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.

B.4 Labor Supply and Market Access

Labor supply, or the number of people working in tract j, can be calculated as $L_j = \sum_m \sum_{i(m)} \pi_i \pi_{ij|i}$. This calculation yields

$$L_j = \mu \sum_k \left[w_{jk}^{\varepsilon} \right] \sum_m \sum_{i(m)} e^{-\kappa\varepsilon\tau_{ij}} \left[[\sum_{i' \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} RMA_{i'}^{\frac{1}{\varepsilon}} \right)^{\eta}]^{\frac{\psi}{\eta}-1} \left(B_i P_i^{\beta-1} \right)^{\eta} RMA_i^{\frac{\eta}{\varepsilon}-1} \right]. \tag{32}$$

Define "Firm Market Access" $FMA_j = \mu \sum_m \sum_{i(m)} e^{-\kappa\varepsilon\tau_{ij}} \left[[\sum_{i' \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} RMA_{i'}^{\frac{1}{\varepsilon}} \right)^{\eta}]^{\frac{\psi}{\eta}-1} \left(B_i P_i^{\beta-1} \right)^{\eta} RMA_i^{\frac{\eta}{\varepsilon}-1} \right]$, a measure of the access to workers enjoyed by firms in tract j. Plugging $\sum_k \left[w_{jk}^{\varepsilon} \right] = L_j / FMA_j$ into the definition $RMA_i = \sum_j e^{-\varepsilon\kappa\tau_{ij}} \left[\sum_k w_{jk}^{\varepsilon} \right]$ yields the relationship

$$RMA_i = \sum_j \frac{e^{-\kappa\varepsilon\tau_{ij}} L_j}{FMA_j}.$$

Going back to its definition, note that $FMA_j = \sum_m \sum_{i(m)} \frac{e^{-\kappa\varepsilon\tau_{ij}} \pi_i}{RMA_i}$. For notational simplicity, we write this as

$$FMA_j = \sum_i \frac{e^{-\kappa\varepsilon\tau_{ij}} \pi_i}{RMA_i}.$$

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 of equations above.

B.5 Housing Demand

An individual who lives in i and works in j in industry k has housing demand of $(1-\beta)\frac{\overline{y}_i}{P_i}$ 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. After multiplying individual housing demand by tract population, the log aggregate residential floorspace demand in tract i can be expressed as

$$\ln S_i^d = \ln \rho_{HD} + \frac{1}{\varepsilon} \ln(RMA_i) + \ln \pi_i - \ln P_i. \tag{33}$$

Here, $\rho_{HD} = (1-\beta)\Gamma(1-\frac{1}{\varepsilon})$. $\ln S_i^d$ is increasing in RMA_i conditional on population π_i because greater RMA_i is associated with greater income for tract residents. Conditional on P_i, equilibrium tract residential population π_i is also increasing in RMA_i. Thus, shocks to RMA_i result in housing demand shocks. This is the key insight used for identification in the empirical work.

Equalizing with the constant-elasticity housing supply function, the equilibrium price of floorspace is given by

$$\ln P_i = \frac{1}{\varepsilon(\gamma_i+1)} \ln RMA_i + \frac{1}{\gamma_i+1} \ln \pi_i + \frac{1}{\gamma_i+1} \ln[\frac{(1-\beta)}{\rho_i}\Gamma(1-\frac{1}{\varepsilon})] \tag{34}$$

This expression has the intuitive features that more elastic supply is associated with lower prices and better labor market opportunities and higher population are associated with higher prices.

B.6 Equilibrium Neighborhood Change

Substituting for (31) into (34), solving for price and time differencing yields the growth rate in tract house price, expressed as follows.

$$\Delta \ln P_i = \frac{1+\eta}{\varepsilon[\gamma_i+1+\eta(1-\beta)]} \Delta \ln RMA_i + \frac{\psi/\eta - 1}{\gamma_i+1+\eta(1-\beta)} \Delta \ln \sum_{i\acute{} \in m(i)} \left(B_{i'} P_{i'}^{\beta-1} RMA_{i'}^{\frac{1}{\varepsilon}} \right)^{\eta}] + \frac{\Delta \ln[B_i^{\eta}/\rho_i]}{\gamma_i+1+\eta(1-\beta)} \tag{35}$$

The second term of this equation can be viewed as a municipality fixed effect if γ_i does not differ across tracts within municipalities. The final term can be viewed as an error term, while recognizing that RMA_i includes tract amenity B_i, tract housing productivity ρ_i and tract supply elasticity γ_i. The first term shows that tracts within a municipality that have larger positive shocks to employment opportunities have more rapid home price growth, as is intuitive. Implicit differentiation of (35) yields that $\frac{\partial \ln P_i}{\partial \ln RMA_i} > 0$ unconditional on municipality fixed effects as well provided that tracts within the same municipality are stronger demand substitutes than tracts in different municipalities. The magnitude of capitalization of improved labor market opportunities is decreasing in the floorspace supply elasticity γ_i. This capitalization is also decreasing in the dispersion of amenity draws, which is negatively related to η. This happens because for higher η, population inflow frictions to tract i are smaller in response to increases in RMA_i.

Time differencing (31), the growth rate in tract population can be written in general terms as

$$\Delta \ln \pi_i = \Delta \ln \mu + G_i + (\frac{\psi}{\eta} - 1) \sum_{i' \in m(i)} S^m_{i'} G_{i'},$$

where $G_i = \eta(\beta - 1)\Delta \ln P_i + \frac{\eta}{\varepsilon}\Delta \ln RMA_i + \eta \Delta \ln B_i$. This expression captures the housing price, labor market opportunity and amenity impacts in tract i and tract i's municipality respectively. $S^m_{i'}$ is the initial share of municipality m's population in tract i'. Summing over all tracts in municipality m, the resulting municipality level population growth rate is

$$\Delta \ln \pi_m = \Delta \ln \mu + \frac{\psi}{\eta} \sum_{i \in m} S^m_i G_i.$$

This expression reflects a common growth rate and a population share weighted tract growth rate that depends on tract housing prices, labor market opportunities and amenities.

Substituting for $d \ln P_i$ yields the following equilibrium relationship between $\Delta \ln RMA_i$ and the population growth rate in tract i.

$$\begin{aligned}
\Delta \ln \pi_i &= \frac{\gamma_i + \beta}{\gamma_i + 1 + \eta(1-\beta)} \frac{\eta}{\varepsilon} \Delta \ln RMA_i + \frac{\gamma_i + \beta}{\gamma_i + 1 + \eta(1-\beta)} (\frac{\psi}{\eta} - 1) \sum_{i' \in m(i)} S^m_{i'} \widetilde{G}_{i'} + u^\pi_i, \text{ where} \\
\widetilde{G}_{i'} &= \frac{\eta}{\varepsilon} \frac{\gamma_{i'} + \beta}{\gamma_{i'} + 1} \Delta \ln RMA_{i'} - \frac{\eta(1-\beta)}{\gamma_{i'} + 1} \Delta \ln \pi_{i'} + \frac{\eta(1-\beta)}{\gamma_{i'} + 1} \Delta \ln \rho_{i'} + \eta \Delta \ln B_{i'} \text{ and} \\
u^\pi_i &= \frac{\gamma_i + \beta}{\gamma_i + 1 + \eta(1-\beta)} [\Delta \ln \mu + \eta \Delta \ln B_i - \frac{1-\beta}{\gamma_{i'} + 1} \eta \Delta \ln \rho_i]
\end{aligned}$$

As with (35) above, we have a tract-specific component, a municipality-specific component and an idiosyncratic component. Summing up over tracts within each municipality m, we derive the following implicit equation which describes the relationship between $\Delta \ln RMA_i$ and municipality level aggregates of tract population growth $\Delta \ln \pi_i$.

$$\sum_{i(m)} S^m_i [1 + \frac{\psi(1-\beta)}{1+\gamma_i}] \Delta \ln \pi_i - \Lambda - \frac{\psi}{\varepsilon} \sum_{i(m)} [S^m_i (1 - \frac{1-\beta}{1+\gamma_i})] \Delta \ln RMA_i = u^m \tag{36}$$

As ψ rises, dispersion in preferences across municipalities falls. As a result, positive shocks to RMA in any neighborhoods within m result in more rapid population growth in this municipality.

Appendix Tables

Table A1: OLS Housing Supply Elasticity Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Total	Total	Total					Total				Total
Subset of Total				New		Remain.			New	Remain.		
Subset of New/Remain.					Redev.		Expan.				Expan.	
	Units	Floorsp.	Units	Units	Units	Units	Units	Floorspace	Flooorspace	Floorspace	Floorspace	Land
Time Period	2000-2006	2000-2006	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2001-2011
Panel A: Housing Quantities Measured Using Zillow Data Only												
Repeat Sales Index	0.06***	0.07***	0.07***	0.06***	0.01***	0.00	0.01***	0.09***	0.09***	-0.01	-0.00	0.02***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)
Obs	30,500	30,048	30,838	30,838	30,838	30,834	30,377	30,381	30,392	30,376	30,377	30,838
Hedonic Index	0.08***	0.09***	0.10***	0.09***	0.03***	0.01*	0.01***	0.11***	0.12***	-0.02**	-0.01	0.03***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)
Obs	29,272	29,272	29,422	29,422	29,422	29,418	29,422	29,422	29,422	29,421	29,422	29,422
Panel B: Housing Quantities Measured Using Census and ACS Data Only												
Repeat Sales Index			0.09***	0.06***		0.03***						
			(0.01)	(0.01)		(0.01)						
Obs			30,838	30,827		30,827						
Hedonic Index			0.12***	0.08***		0.04***						
			(0.01)	(0.01)		(0.01)						
Obs			29,422	29,411		29,411						

Regressions include metro region fixed effects, a cubic in fraction of the way from the CBD to the metro edge, fraction tract developed in 2001, fraction of tract land that is flat, log 1990 tract employment, the 2000-2006 Bartik shock for the tract and the following tract attributes measured in 1990 and 2000: census home price index, rent index, log population, log avg hh income, share black, share white and share college. The estimation sample for the repeat sales index uses data from 167 metro regions while that for the hedonic index uses data from 165 regions. These samples are reduced 1 region for the repeat sales index floorspace outcomes due to missing floorspace information for some tracts. Entries in Columns 1 and 2 use 164-166 regions. The hedonic index sample excludes tracts in the repeat sales sample containing homes for which age and/or floorspace are not observed in 2000 or 2010. All outcomes are measured using the ZTRAX data except developed land, which uses USGS land cover information. Analogous regressions using census and/or ACS based measures for changes in units or new units are statistically identical.

Table A2: Robustness Checks for Table 2 on Employment and Housing Dynamics

Panel A: Tract Level Regressions of Employment Growth on Bartik Shocks

	(1)	(2)
	Change in Log Employment	
	2000-2006	2000-2010
Tract Bartik Shock, 2000-2006	0.69***	1.04***
	(0.10)	(0.10)
No 1990 Emp. Info.	0.18**	0.32***
	(0.09)	(0.09)
Demographic controls	No	No
Observations	51,900	56,043
R-squared	0.01	0.01
Number of Regions	158	169

Panel B: Tract-Level Housing Market Dynamics

	(1)	(2)	(3)
	Δln House Price, 00-10	Δln House Quantity, 00-10	Δln Simulated RMA, 2000-2006
Δln House Price, 1990-2000	-0.25***	0.01	0.0000
	(0.01)	(0.01)	(0.0001)
Δln House Quantity, 1990-2000	-0.00	0.25***	-0.0002
	(0.03)	(0.04)	(0.0002)
Demographics	90+00	90+00	90+00
Δln House Price, 1990-2000	-0.22***	0.01	0.0001
	(0.01)	(0.01)	(0.0001)
Δln House Quantity, 1990-2000	-0.01	0.20***	0.0003***
	(0.01)	(0.01)	(0.0001)
Demographics	90	90	90
Δln House Price, 1990-2000	-0.22***	0.00	0.0001
	(0.01)	(0.01)	(0.0001)
Δln House Quantity, 1990-2000	0.00	0.25***	0.0002***
	(0.01)	(0.01)	(0.00)
Demographics	No	No	No

Panel C: First-Stage in 1990-2000

	(1)	(2)	(3)
		Δln House Quantity, 1990-2000	
Δln Simulated RMA, 1990-2000	4.32***	4.00***	-0.29
	(0.64)	(0.49)	(0.23)
Demographic controls	No	90	90+00
Number of Regions	167	167	167

Panel A: Regressions also include region fixed effects, fraction developed in 2001, fraction flat plains, a cubic in fraction of the way to region edge, log 1990 tract employment and the following tract attributes from 1990 and 2000: census home price index, rent index, log pop, log avg hh income, fraction black, fraction white, fraction college. Sample includes all tracts in metro regions that are in the primary sample. Each tract receives equal weight. Robust standard errors.

Panel B: Each entry is from a separate regression of the variable at top on the variable at left with the indicated fixed effects. Controls are the same as in Table 2, excluding 2001 developed fraction and the 1990 and 2000 census house price and rent indexes.

Panel C: Regressions include metro region fixed effects and the same controls as in Table 2. Tracts are equally weighted, even if they appear in multiple metro regions. Standard errors are corrected for spatial autocorrelation up to 16 km using a Bartlett kernel.

Table A3: Robustness Checks for Unified IV Results

	(1) Total Units 2000-2010 Census and ACS data	(2) New Units 2000-2010 Census and ACS data	(3) Remain Units 2000-2010 Census and ACS data	(4) Total Floorspace 2000-2010 Zillow data	(5) Total Quality-adjusted floorspace 2000-2010 Zillow data
Repeat Sales Index	0.31***	0.23**	0.09	0.42***	0.40***
	(0.11)	(0.09)	(0.06)	(0.16)	(0.14)
Obs	30,838	30,827	30,827	30,381	30,838
Hedonic Index	0.27**	0.21**	0.08	0.44***	0.44***
	(0.11)	(0.09)	(0.06)	(0.16)	(0.14)
Obs	29,422	29,411	29,411	29,422	29,422

The specification is the same as Table 4.

Table A4: Unified IV Results with Alternative IVs

	(1) Total	(2) Total	(3) Total	(4) New	(5) New	(6) Total
	Units	Units	Floorspace	Units	Floorspace	Land
Time Period	2000-2010 Census data	2000-2010	2000-2010	2000-2010	2000-2010	2001-2011
			Panel A: Including all industries			
Repeat Sales Index	0.30***	0.32***	0.42***	0.19**	0.29**	0.08
	(0.11)	(0.11)	(0.15)	(0.08)	(0.12)	(0.06)
First stage F	24.17	24.17	22.68	24.17	22.51	24.17
Obs	30,838	30,838	30,381	30,838	30,392	30,838
			Panel B: Excluding construction industries			
Repeat Sales Index	0.30***	0.35***	0.42***	0.19**	0.29**	0.09
	(0.11)	(0.12)	(0.16)	(0.08)	(0.12)	(0.06)
First stage F	22.01	22.01	20.35	22.01	20.24	22.01
Obs	30,838	30,838	30,381	30,838	30,392	30,838
			Panel C: Excluding construction and FIRE industries			
Repeat Sales Index	0.30***	0.35***	0.39**	0.19**	0.29**	0.09
	(0.11)	(0.12)	(0.16)	(0.08)	(0.12)	(0.06)
First stage F	21.71	21.71	20.06	21.71	19.97	21.71
Obs	30,838	30,838	30,381	30,838	30,392	30,838
			Panel D: Excluding construction and real estate industries			
Repeat Sales Index	0.30***	0.35***	0.40**	0.18**	0.28**	0.08
	(0.11)	(0.12)	(0.16)	(0.08)	(0.12)	(0.06)
First stage F	21.68	21.68	20.05	21.68	19.96	21.68
Obs	30,838	30,838	30,381	30,838	30,392	30,838

Estimation sample and specifications are the same as in Table 4.

Table A5: Quadratic IV Model: Heterogeneity in Supply Elasticities by CBD Distance and Tract Condition

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Total	Total					Total				
Subset of Total			New		Remain.			New	Remain.		
Subset of New/Remain.				Redev.		Expan.				Expan.	
	Units	Units	Units	Units	Units	Units	Floorspace	Floorspace	Floorspace	Floorspace	Land
ΔlnP	0.11	1.40***	0.96***	0.13*	0.47**	0.21	1.42***	0.96***	0.35	0.17	0.38**
	(0.12)	(0.37)	(0.29)	(0.07)	(0.20)	(0.13)	(0.45)	(0.36)	(0.27)	(0.22)	(0.19)
ΔlnP x (CBD Dis)	1.89**	-0.79***	-0.65***	-0.17***	-0.15	-0.08	-0.68**	-0.72***	0.13	0.13	-0.37***
	(0.73)	(0.25)	(0.22)	(0.05)	(0.11)	(0.08)	(0.28)	(0.25)	(0.16)	(0.12)	(0.13)
ΔlnP x %Dev		-4.52***	-4.02***	-0.59***	-0.51	-0.24	-4.66***	-3.89***	-0.31	0.09	-1.49***
		(0.91)	(0.77)	(0.19)	(0.38)	(0.20)	(1.01)	(0.87)	(0.55)	(0.36)	(0.43)
ΔlnP x %Dev^2		3.84***	3.54***	0.60***	0.30	0.09	4.20***	3.41***	0.39	-0.24	1.17***
		(0.81)	(0.69)	(0.22)	(0.37)	(0.18)	(0.91)	(0.78)	(0.59)	(0.34)	(0.37)
ΔlnP x %Flat		0.28**	0.38***	0.08**	-0.12**	-0.02	0.33**	0.49***	-0.19**	-0.09*	0.20***
		(0.12)	(0.13)	(0.03)	(0.05)	(0.03)	(0.14)	(0.16)	(0.08)	(0.05)	(0.07)
ΔlnP x (CBD Dis)^2	-1.84**										
	(0.78)										
Observations	30,838	30,838	30,838	30,838	30,834	30,377	30,381	30,392	30,376	30,377	30,838
Kleib-Paap F-Stat		5.330	5.330	5.330	5.330	5.060	5.083	5.013	5.049	5.060	5.330

Regressions are the same specification as in Table 5 with the addition of the quadratic term. The repeat sales price index measure is used throughout. SE adjusted for spatial autocorr. to 16 km.

Table A6: Impacts of Regulation and Highway on Supply Elasticities

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Units	New Units	Redev.Units	Floorspace		New Floorspace		Dev. Land	Units	Units
ΔlnP	4.24	3.66	0.63	5.93	2.27	4.05	0.04	2.06	0.87***	0.89***
	(2.92)	(2.43)	(0.53)	(4.65)	(7.82)	(3.17)	(2.95)	(1.44)	(0.25)	(0.25)
ΔlnP x (CBD Dis)	0.17	-0.06	-0.04	0.62	3.01	-0.01	-1.02	-0.14	-0.57**	-0.57**
	(0.69)	(0.55)	(0.12)	(1.03)	(7.95)	(0.65)	(3.12)	(0.33)	(0.24)	(0.24)
ΔlnP x %Dev	-12.91	-11.96*	-1.85	-17.18	-2.76	-13.01	-2.18	-6.86*	-1.46***	-1.52***
	(8.07)	(6.79)	(1.47)	(12.01)	(12.67)	(8.35)	(5.82)	(3.90)	(0.40)	(0.40)
ΔlnP x %Dev^2	10.33	9.74*	1.36	14.21	0.68	10.78*	2.38	5.55*		
	(6.41)	(5.40)	(1.18)	(9.29)	(9.52)	(6.49)	(5.23)	(3.06)		
ΔlnP x %Flat	0.51	0.46	0.14	0.39	-1.35	0.49	0.80	0.29	0.29**	0.31**
	(0.42)	(0.36)	(0.09)	(0.55)	(5.78)	(0.40)	(2.37)	(0.22)	(0.12)	(0.12)
ΔlnP x WRLURI	-0.03	-0.04	-0.01	-0.02		-0.04		-0.02		
	(0.04)	(0.03)	(0.01)	(0.06)		(0.04)		(0.02)		
Robust SE	[0.03]	[0.02]**	[0.01]	[0.04]		[0.03]		[0.01]**		
ΔlnP x (Res. FAR)					0.07*		0.01			
					(0.04)		(0.05)			
Robust SE					[0.04]		[0.02]			
ΔlnP x Highway									0.09***	
									(0.02)	
ΔlnP x Interstate Highway										0.11***
										(0.02)
Observations	12,361	12,361	12,361	12,149	2,128	12,155	2,128	12,361	30,838	30,838
Number of Regions	144	144	144	144	13	144	13	144	167	167

Regressions are the same specification as in Table 5 with the addition of indicated interaction terms. For WRLURI interactions, the main effect is excluded to maintain statistical power. Its inclusion increases the magnitude of the related interacted coefficient. FAR main effects are included in those interacted regressions. The repeat sales price index measure is used throughout. SE in parenthesis. adjusted for spatial autocorr. to 16 km.

Table A7: Quadratic Finite Mixture Model Housing Supply Interacted Regressions

Outcome	Units		New Units		Floorspace		New Floorspace		Developed Land	
Panel A: Logit Coefficients for Membership in Class 2 (More Inelastic Supply)										
Fraction Developed within 50 km of the CBD	5.38***		4.31***		4.40***		3.81***		5.30***	
	(0.61)		(0.64)		(0.61)		(0.60)		(0.64)	
Fraction Developed within 50 km of the CBD Unavail.	0.51*		0.28		1.08***		0.38		0.34	
	(0.29)		(0.26)		(0.32)		(0.27)		(0.29)	
Metro Wharton Index.	0.03		0.13**		-0.67***		0.12**		0.11*	
	(0.06)		(0.06)		(0.09)		(0.05)		(0.06)	
Constant	0.58***		1.09***		0.77***		1.19***		1.04***	
	(0.12)		(0.11)		(0.14)		(0.11)		(0.12)	
Panel B: Second Stage Estimates										
	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2
ΔlnP	2.93**	0.37*	2.83**	0.15	3.15*	0.85**	1.88	0.17	-0.05	0.14
	(1.15)	(0.21)	(1.26)	(0.11)	(1.74)	(0.35)	(1.78)	(0.14)	(0.70)	(0.13)
ΔlnP x (CBD Dis)	-1.39**	-0.07	-1.52***	0.04	-1.80**	0.02	-1.12	-0.01	-0.98**	-0.03
	(0.61)	(0.11)	(0.58)	(0.07)	(0.75)	(0.23)	(0.69)	(0.08)	(0.41)	(0.06)
ΔlnP x %Dev	-9.13***	-1.05***	-10.82***	-0.88***	-9.45***	-2.16***	-7.58**	-0.96***	-0.83	-0.43
	(2.24)	(0.40)	(2.85)	(0.21)	(3.12)	(0.62)	(3.59)	(0.25)	(1.49)	(0.26)
ΔlnP x %Dev^2	7.55***	0.96***	8.96***	0.89***	8.40***	1.81***	6.51**	0.98***	-0.12	0.36*
	(1.93)	(0.31)	(2.63)	(0.18)	(2.72)	(0.52)	(2.97)	(0.23)	(1.49)	(0.21)
ΔlnP x %Flat	0.41	0.04	0.78**	0.08**	0.76**	0.06	1.04***	0.11**	0.47**	0.03
	(0.29)	(0.04)	(0.33)	(0.04)	(0.36)	(0.08)	(0.40)	(0.05)	(0.20)	(0.02)
Mean Class Probability	0.22	0.78	0.16	0.84	0.24	0.76	0.15	0.85	0.16	0.84
Mean implied γ	0.85	0.16	0.53	0.04	1.13	0.45	0.50	0.04	-0.42	0.06
SD implied	0.73	0.08	0.86	0.07	0.70	0.20	0.63	0.07	0.25	0.03

Notes: Models use the same specification as in Table 6, with the addition of the quadratic interaction term in developed fraction and a control for developed fraction squared. Sample sizes are the same as in Table 6.

Table A8: Metro Level Supply Regressions

Quantity Measure	Households	Households	Households	Households	Households	Households	Households	Households	Households	Households	Households	Units	Units	Units	Units
Price Measure	Median	Median	Median	Median	Median	Median	Median	Median	Median	Median	Median	Median	Median	FHFA Index	FHFA Index
Time Period	1970-2000 (from Saiz, 2010)				2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010	2000-2010
ΔlnQ	0.650***	0.336***	0.06	1.246***	0.675***	1.035***	-0.195	-0.142		-0.146	-0.0257	-0.0676	0.0381	0.298	
	(-0.107)	(-0.116)	(-0.215)		(0.222)	(0.250)	(0.284)	(0.363)	(0.578)		(0.583)	(0.330)	(0.505)	(0.476)	(0.677)
$\Delta lnQXUnavailLand$		0.560***	0.511***	-5.260**		2.182***	0.910	1.931***	0.335	-8.624	-0.506	1.478***	0.249	0.142	-1.009
		(-0.118)	(-0.214)	(-1.396)		(0.546)	(0.736)	(0.514)	(0.704)	(6.468)	(6.678)	(0.417)	(0.679)	(0.313)	(0.626)
$\Delta lnQXUnavail$				0.475**						0.818	0.0536				
X ln Base Yr Pop				(-0.119)					(0.507)	(0.508)					
$\Delta lnQXln(WRLURI + 3)$			0.237*	0.280**				0.715**	1.225	0.525**	1.209*	0.555*	1.081	-0.0972	0.0580
			(-0.13)	(-0.077)				(0.339)	(0.764)	(0.216)	(0.730)	(0.285)	(0.666)	(0.340)	(0.631)
Unavail Land							0.203*		0.257**		-2.291**		0.217*		0.215**
							(0.123)		(0.111)		(1.168)		(0.116)		(0.0989)
ln(WRLURI+3)									-0.104		-0.0937		-0.113		-0.0336
									(0.104)		(0.0979)		(0.103)		(0.0750)
First Stage F					44.41	31.37	18.37	12.33	12.48	25.00	8.311	15.66	10.19	20.79	9.053
Mean of Implied elasticity	1.54	2.19	2.56	2.54	0.803	0.910	0.803	1.270	0.949	1.253	0.855	1.300	1.011	-58.29	2.990
SD of Implied elasticity	0	0.48	0.88	1.43	0	0.299	0.111	0.780	0.485	0.683	0.341	0.603	0.444	664.8	42.22

Robust standard errors in parentheses. The table shows the coefficient of 2SLS estimation of a metropolitan housing supply equation with Census region fixed effects. The left block reproduces Tables III and V in Saiz (2010). Remaining columns replicate similar specifications using 2000-2010 data. Instruments used for demand shocks are a shift-share of the 1974 metropolitan industrial composition, the magnitude of immigration shocks, and the log of January average hours of sun. While 2000-2010 data constraints reduce the estimation sample from 269 to 193 metros (234 in the final two columns), predicted distributions with means and SD reported at the bottom use predicted values for the full Saiz sample of 269 metros.

Figure A1: Example Fixed Cost Distributions

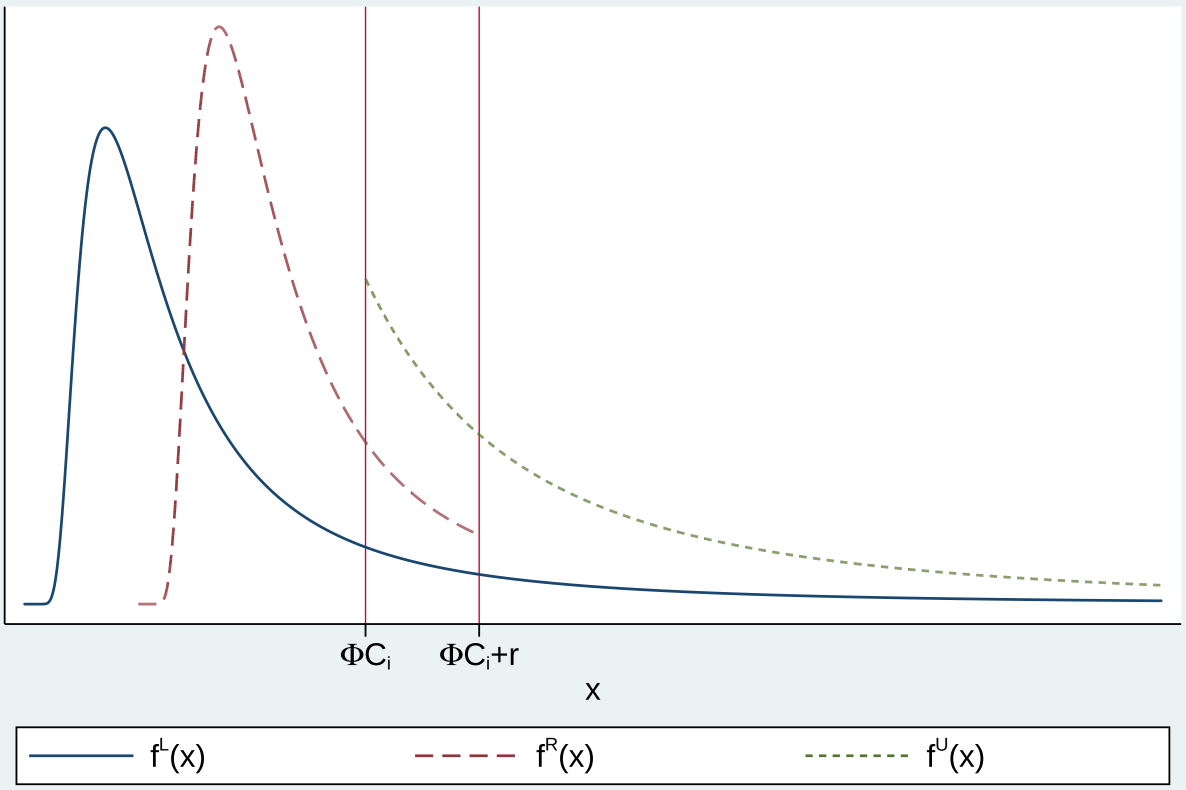

Figure A2: Spatial Covariance Function

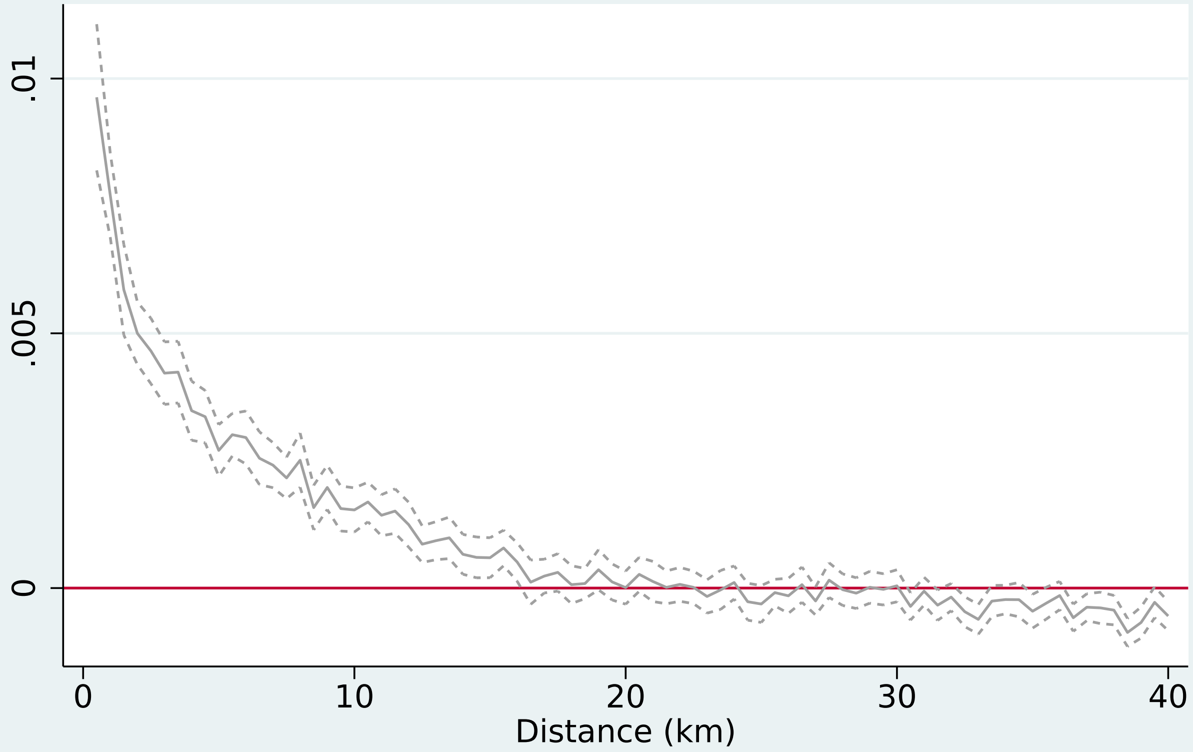

Plot uses errors associated with the regression in Table 4 column (3).

Figure A3: Kernel Densities of Predicted Elasticities

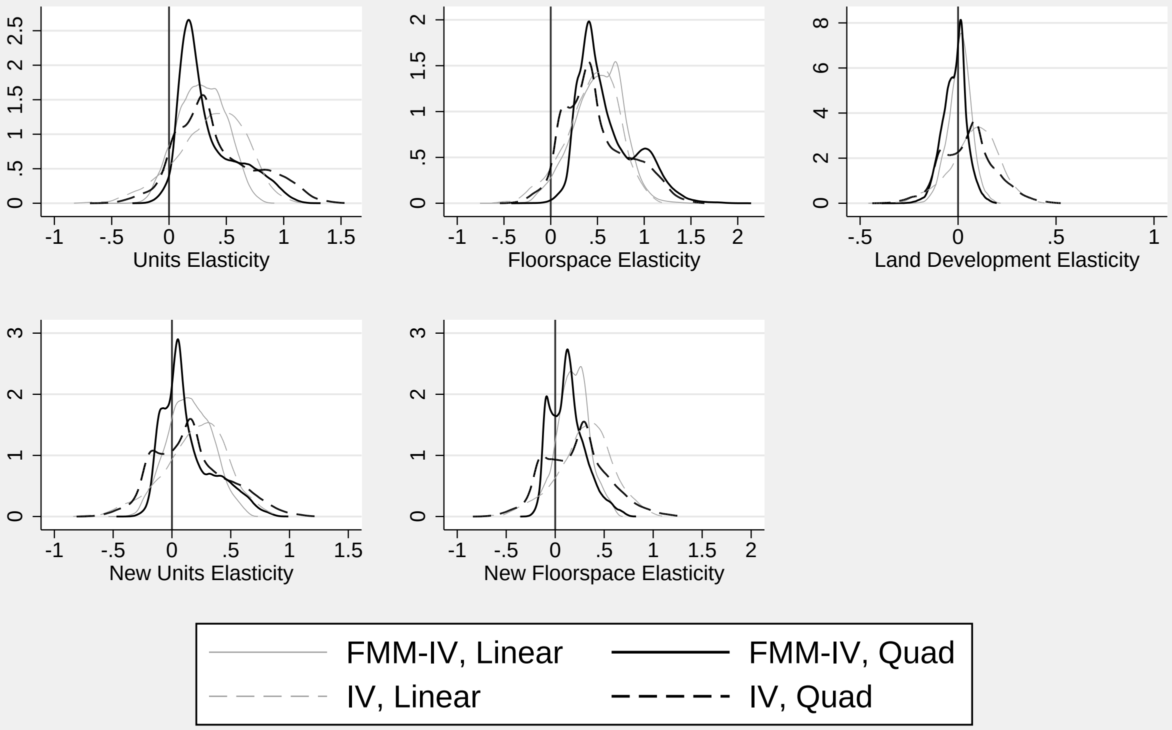

Elasticities are predicted for all metro tracts nationwide. Linear specifications use estimates in Table 5 and 6. Quadratic specifications use estimates in Tables A5 and A6.

Figure A4: Supply Elasticities by CBD Distance

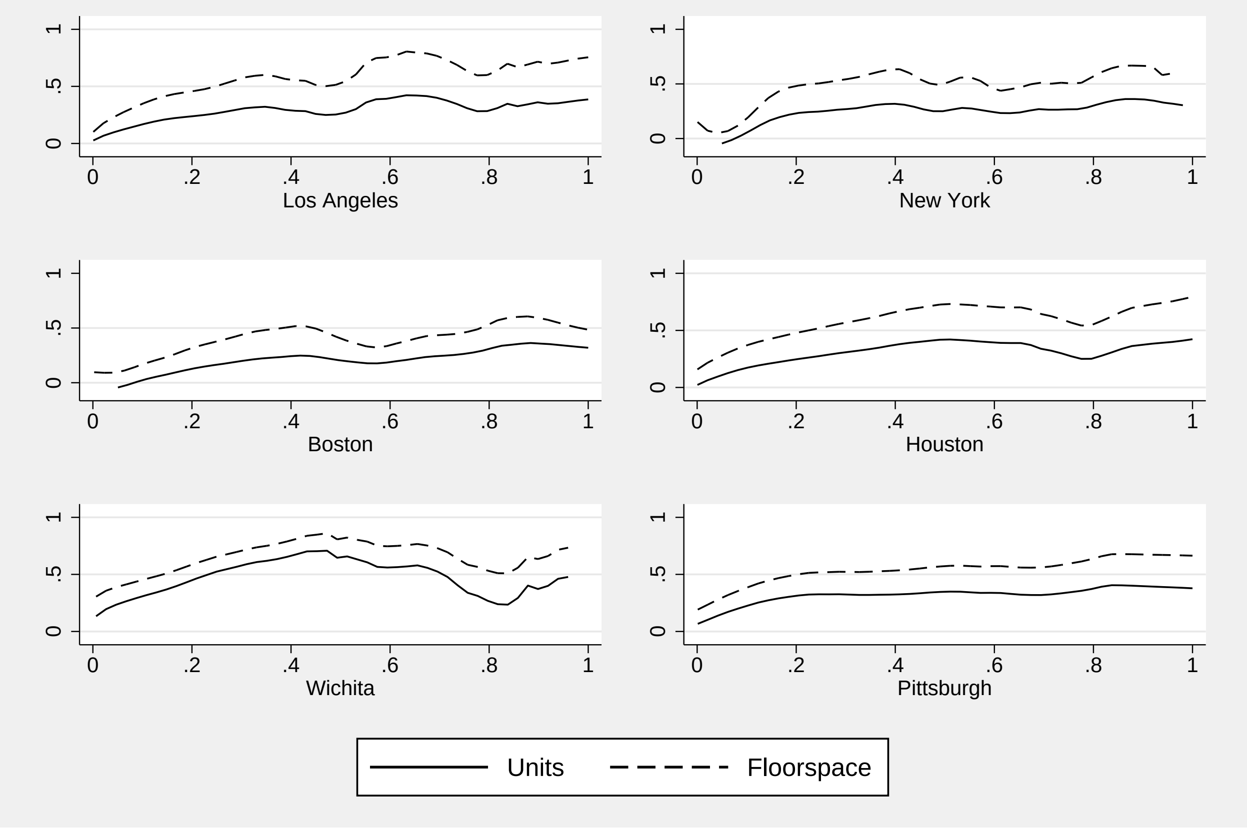

Figure A5: Comparisons of Elasticities Calculated Using Saiz (2010) Methods to FMM-IV Metro Units Elasticities

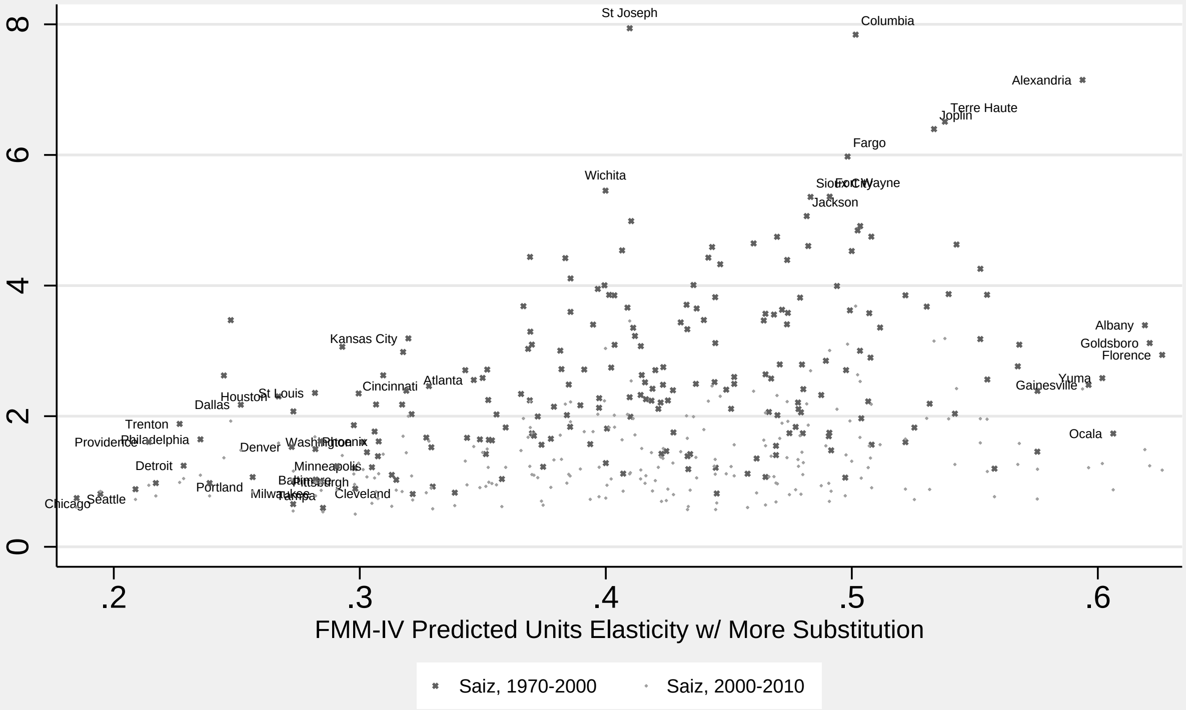

Select cities indicated next to 1970-2000 elasticities with corresponding 2000-2010 elasticities indicated underneath.
All 237 Saiz (2010) elasticities are shown except those for Pine Bluff and Newark, which are off the chart.