Moving Mountains: Geography, Neighborhood Sorting, and Spatial Income Segregation

Victor Yifan Ye*†‡
Charles M. Becker§

March 20, 2021

Abstract

Using a novel geospatial panel combined with data from the 2015 American Community Survey (ACS), we investigate the effect of topography – altitude and terrain unevenness – on income segregation at the neighborhood level. Specifically, we perform large-scale counterfactual simulations by estimating household preferences for topography, altering the topographical profile of each city, and observing the resulting neighborhood sorting outcome. We find that unevenness contributes to the segmentation of markets: in the absence of hilliness, rich and poor households experience greater mixing. Hillier cities are more income-segregated because of their unevenness; the opposite is true for flatter cities.

JEL: C63, R20, R32

Keywords: Computation, Geography, Counterfactual, Household Income, Neighborhood Sorting, Spatial Modelling

*Department of Economics, Boston University, 270 Bay State Road, Boston, MA, 02215, United States `victor@yifanye.com`

†Corresponding Author

‡We are thankful to Patrick Bayer, Seth Benzell, Peter Blair, Laurence Kotlikoff, session attendees at the 12th Meeting of the Urban Economics Association, and the 8th and 9th European Meeting of the Urban Economics Association for advice and comments. All errors and misinterpretations remain our own.

§Department of Economics, Duke University, Campus Box 90097, 213 Social Sciences, Durham, NC, 27708, United States `cbecker@duke.edu`

1 Introduction

It is documented that high-income households in US urban areas concentrate in high-altitude, topographically uneven neighborhoods (Ye and Becker 2018; Lee and Lin 2018). Elevation and unevenness – more specifically, better scenery, lack of crime, and microclimate – likely are amenities enjoyed mainly by the wealthy, with an income elasticity well above unity. Moreover, elevation variance also may impose a cost to poor households by constraining walkability and access to public transit.

This paper analyzes the economic consequences of topography – altitude and unevenness of terrain – in a neighborhood sorting framework. Using data from 25 major Metropolitan Statistical Areas (MSAs) containing more than a quarter of all US urban census tracts, we demonstrate the extent to which topography affects within-city income sorting equilibria. Specifically, we estimate not only demand-side topography-induced neighborhood amenity effects by household income, but also supply-side effects of topography as a constraint to the quality and composition of local housing stock.

This is achieved by combining the 2015 American Community Survey (ACS) with novel geospatial panel data to estimate household preferences for topographical characteristics and other neighborhood amenities. Using these preferences, we specify two large-scale counterfactual simulation models: one which "flattens" MSAs to a baseline, city-average altitude and zero elevation variance, and a second "reverse" counterfactual which doubles both within and across-tract unevenness. In both cases we observe the resulting neighborhood sorting outcome and, in so doing, separate topography-based income sorting from that of preferences for other neighborhood amenities.

Removing topographical features decreases MSA-level income segregation. Own-neighborhood income falls for the rich and rises for the poor, and households outside of the top income decile are more likely to live alongside those in the adjacent, richer decile. The opposite is true when we double unevenness: rich households converge to hilly areas, enjoy higher neighborhood income levels, and experience higher prices. Poor and middle-income households live in flatter, cheaper, and poorer locations and display greater mixing among themselves.

Counterfactual results are largely consistent when we remove neighborhood-specific housing quality effects from the model to simulate a scenario where utility derived from neighborhood amenities is fully determined by a tract's location and the income level of its households. However, preferences of the richest households become further differentiated from those of the middle class when fixed housing quality effects are eliminated, leading to the centralizing of upper-middle class households which, consequently, forces a fraction of the poorest out to the far suburbs.

Additionally, to simulate the housing supply effect of topography, we decorrelate topography and compositions of housing units by spatially smoothing households' utility derived from qualities and types of local housing stock and randomly re-assigning the residuals within each MSA. The outcome of the "housing supply" counterfactual is comparable to that of flattening cities: the rich experience less unevenness, lower altitudes, and poorer neighborhoods, and the reverse applies to the poor. This similarity suggests that the relationship between neighborhood income and topography is best explained by a combination of demand and supply-side effects.

A key benefit of our approach is the ability to explicitly account for sorting equilibrium effects and the role of income-specific tastes for residential location qualities in affecting the counterfactual outcome. For example, eliminating unevenness leads households to sort more strongly on other aspects of housing quality. In this case, gains to very poor households may be partially canceled out as their preferences for non-topographic amenities becomes further differentiated from those of poor and middle-income groups.

The salient implication of our findings is that topography is neither only reflective of a historical amenity for the rich nor merely a locational determinant of where rich households locate: income sorting equilibria of current cities are less segregated when they are flattened. Residential sorting outcomes of uneven cities (San Francisco, Portland) are qualitatively different from those of flat cities (Chicago, Miami) *because of* their unevenness.

Section 2 presents a review of prior literature related to topographical effects, neighborhood sorting and spatial modelling. Section 3 discusses data sources and methodologies for calibrating and simulating the counterfactuals. Results are outlined and discussed in Section 4, and Section 5 concludes.

2 Literature

We believe this is the first paper to investigate the relationship between elevation gradient effects and neighborhood sorting outcomes. This relationship is absent even in modern treatments of urban spatial equilibrium models (for example, Lucas and Rossi–Hansberg 2002). A limited yet growing body of literature addresses the role of elevation gradients in the formation and development of cities: examples include the desirability of coastal living (Rappaport and Sachs 2003), flood risk associated with low-lying areas (Scawthorn, Iemura, and Yamada 1982; Shilling, Sirmans, and Benjamin 1989; Bin et al. 2011), and geographical features as a cause of initial locational choice of current European cities (Bosker and Buringh 2015). Although only tangentially related to elevation gradient effects as discussed in this paper, these analyses nonetheless suggest the possibility of elevation affecting distributional outcomes of neighborhood income at the city-level.

The classic literature on the economic consequences of elevation gradients primarily focuses on elevation and, more broadly, geographical features as a constraint to land supply. Rose (1989) studies land supply effects caused by large bodies of water such as lakes and oceans. Kok, Monkkonen, and Quigley (2014) investigate determinants of land value in San Francisco and find evidence of elevation effects, though their primary focus is on land use regulations.

Utilizing satellite data and a broad, 73-MSA dataset, Saiz (2010) presents evidence that undevelopable land on the city periphery is a strong predictor of low housing supply elasticity. While this is a seminal paper, it emphasizes differences at the MSA-level and not intra-urban locational choice or supply elasticities. The approach of using hilliness as an instrument for housing supply has also been applied in a number of recent papers such as (Baum-Snow and Han 2019).

Bleakley and Lin (2012) emphasize path dependence and persistence for urban agglomerations at portages but focus on counties or cities rather than at the neighborhood level. Ananat (2011) uses variation in historic railroad track location to explain patterns of variation in racial segregation across cities. Berger and Enflo (2017) provide a similar study for Sweden that emphasizes the importance of initial railroad lines for city growth. Duranton and Turner (2012) explore the impact on urban growth patterns of the expansion of the US interstate highway system. Pierce and Kolden (2015) provide a variety of hilliness measures for 100 US cities, but do not explore the associated economic implications.

Beyond the role of elevation as a land supply constraint, Lee and Lin (2018) build on prior literature (Bleakley and Lin 2012; Lin 2015) on the economic consequences of geographic features to present evidence that "natural" amenities influence spatial income distributions within urban areas. They model natural geographical features as immutable points of attraction for rich households and demonstrate that proximity to hills and a range of other geographical features such as coastal proximity and lakes is correlated with higher income levels. Consistent with the immutable role of hilliness, Lee and Lin (2018) find that flatter cities experience substantially more change in the social composition of neighborhoods and in intra-neighborhood income distribution over time than do their hillier counterparts.

This discussion is developed by Ye and Becker (2017a). They present evidence using transaction-level housing price data from Hong Kong that the undesirability of walking up or downhill to public transit stations is robustly factored into sales prices: other factors constant, a 1-decimal-degree increase in the slope between a middle to middle-low income class apartment and the closest metro stop decreases its selling price by up to 1.9%. While Hong Kong may be an outlier among major cities worldwide because of its extremely uneven topography, the findings nonetheless suggest that elevation may not only be an attraction to the rich through natural amenity effects but also a deterrent to the poor by increasing the difficulty of accessing public transit.

Ye and Becker (2018b) further develop the discussion in two ways. First, they show that the economic influence of topography is not limited to cities that are conventionally considered to be rich in natural amenities. In other words, cities do not need to be as uneven as San Francisco or Portland for elevation effects to play a nontrivial role in income and population distributions. Second, that preferences for terrain unevenness and higher altitudes by the rich and distastes by the poor can be broken down into preferences for intermediary amenities such as lower crime (further evidenced by Kelsay and Haberman 2020a; 2020b) superior microclimate, lack of traffic congestion, and difficulty of accessing public transit. Through such effects, both middle and low-income households display preferences or distaste for unevenness, even though they are not all wealthy enough to value scenery *per se*.

In a paper perhaps most closely related to ours, Allen and Arkiolakis (2014) develop an equilibrium trade model to explain variation in spatial inequality of income due to physical location at the US county (but not within-MSA) level. Their many findings include the implication that extreme amenities are an important source of inequality for a limited set of prosperous counties; overall, geographic variation appears to be responsible for as much as one-fifth of spatial variation in US income. However, this result is driven by location rather than topography. In another related work, Andreoli and Peluso (2017) focus on variation in neighborhood inequality for 50 MSAs, but do not include a topographic component.

Our contribution to this literature is threefold. First, we identify a relationship between unevenness and the degree of neighborhood income stratification through a counterfactual framework. While the possibility that elevation gradients merely determine where stratification occurs cannot be ruled out by a regression-based analysis, our counterfactual simulation explicitly introduces a quasi-experiment comparing the same set of MSAs with and without topographical features. This approach allows us to show that holding all other factors constant, more uneven cities are indeed more stratified income-wise. Additionally, once preferences lead to stratification, they could be reinforced by local spatial externalities across residents, as documented by Guerrieri, Hartley and Hurst (2013) and Rossi-Hansberg, Sarte and Owens (2010).

Second, we distinguish between demand and supply-side effects of elevation gradients by employing a general equilibrium approach with an explicit treatment of housing supply quality. In our baseline simulation, unevenness enters only through household demand for location and, because housing quality effects are fixed, is solely responsible for the counterfactual outcome. Conversely, the only source of variation in the "housing supply" counterfactual is the elimination of direct correlation between unevenness and specific aspects of the quality of a tract's housing stock. Hence, we achieve relatively clean treatments of unevenness both as a natural amenity for the rich and as a constraint for the quality of local housing stock.

Finally, our approach allows for multiple counterfactual simulations where amenities and preferences can be selectively altered to reflect different scenarios. By both simulating the flat-city case and a scenario where all MSAs are twice as uneven, we present a much stronger case for our inferences on income-decile-specific effects. We also selectively restrict the functionality of parts of the simulation, such as neighborhood housing quality effects, to reflect assumptions about the extent to which such amenities can be considered exogenous or endogenous.

3 Data and Methodology

3.1 Data

We use census tract-level data from the 2015 American Community Survey (ACS) for our multi-MSA panel and select 25 major MSAs for the simulation model.[1] The sample of MSAs is selected to maximize diversity in terms of geographical location, with MSAs being approximately equally distributed among Census Bureau statistical regions, and with each of the ten Standard Federal Regions being represented by at least one MSA in the data.

All dataset MSAs are selected such that they contain a minimum of 100 census tracts. While we do not select for MSAs with substantial terrain unevenness, extremely flat cities (Chicago, Miami) are excluded from the dataset. In addition, New York MSA is omitted because it has by far the largest number of census tracts among all US MSAs (3,288), more than twice as large as the largest dataset MSA (Washington DC, 1,281), and is likely an outlier in terms of sorting dynamics, housing market conditions as well the distribution of quality across tracts. MSA-specific summary statistics are presented in Table 1.

[1] Albuquerque , Atlanta , Austin, Baltimore, Boston, Charlotte, Cincinnati, Colorado Springs, Denver, Kansas City, Los Angeles, Louisville, Memphis, Nashville, Omaha, Phoenix, Pittsburgh, Portland (OR), Salt Lake City, San Diego, San Francisco, Seattle, St. Louis, Tucson, and Washington DC

Table 1: Summary statistics of dataset MSAs

	Albuquerque	Atlanta	Austin	Baltimore	Boston
No. of Tracts	203	938	345	662	945
No. of households	343,434	1,945,508	682,841	1,014,043	1,690,304
Avg housing price ($)	138,745	139,507	156,898	215,012	270,338
Avg tract income ($)	65,427	78,980	86,314	92,277	102,488

	Charlotte	Cincinnati	CO Springs	Denver	Kansas City
No. of Tracts	424	496	135	608	533
No. of households	690,992	818,462	253,756	1,026,271	807,165
Avg housing price ($)	146,731	124,153	164,179	197,808	124,787
Avg tract income ($)	77,685	73,981	76,808	87,713	76,495

	Los Angeles	Louisville	Memphis	Nashville	Omaha
No. of Tracts	789	315	310	361	254
No. of households	1,275,802	508,694	486,008	629,204	343,656
Avg housing price ($)	176,274	125,563	103,485	156,761	118,556
Avg tract income ($)	72,431	68,626	66,633	75,540	75,502

	Phoenix	Pittsburgh	Portland	Salt Lake City	San Diego
No. of Tracts	981	706	488	234	622
No. of households	1,556,535	988,201	882,439	381,795	1,087,836
Avg housing price ($)	145,333	114,790	187,710	189,700	271,771
Avg tract income ($)	72,968	71,769	79,372	81,570	86,906

	San Fran.	Seattle	St. Louis	Tucson	Washington DC[2]
No. of Tracts	953	709	616	233	1,281
No. of households	1,621,912	1,380,387	1,105,213	377,987	2,014,619
Avg housing price ($)	401,876	231,156	131,049	125,176	289,247
Avg tract income ($)	116,182	92,840	74,468	63,729	119,433

This dataset is merged with high-resolution Digital Elevation Models (DEMs) collected using the Microsoft Representational State Transfer (REST) Application Programming Interface (API). We use REST to sample altitude values over a 1,000-by-1,000 grid covering the entire respective MSA areas. Sample points are joined to census tracts using tract boundary data from the Census Bureau's Topologically Integrated Geographic Encoding and Referencing (TIGER) database, with a sampling density of 1,114.4 observations per tract.[3] Elevation variance of a given tract is estimated as the variance of all internal altitude samples. We present an example of the elevation sampling and the spatial distribution of elevation variance with data from Boston in Figure 1.[4]

[2]The DC MSA is the largest in the dataset because it spans both Maryland and Virginia, in addition to the District of Columbia area.

[3]Since variance across sample points is a per-area metric, there is no inherent bias toward higher elevation variance for larger tracts. Larger tracts nonetheless tend to be on the periphery of MSAs, which coincides with areas of higher variance. This also means that given the high sampling density, our measure of elevation variance is comparable across tracts in different MSAs, despite the fixed total resolution of samples.

[4]Two more examples (San Francisco, Washington DC) are provided in Figures A3 and A2.

Figure 1: Elevation contour and elevation variance spatial distribution, Boston

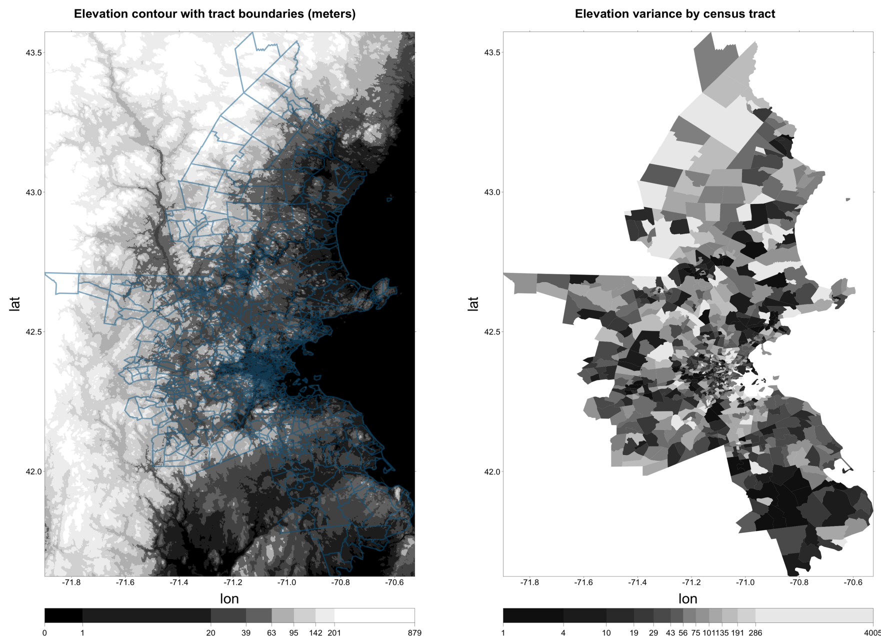

Elevation sample points at the local water level are omitted.[5] We also exclude census tracts with less than 20 total elevation samples from the dataset, resulting in a total of 14,141 tracts in the merged dataset: this is approximately 19.2% of all 2010-boundary US census tracts or 30.1% of all urban tracts. To provide a consistent estimation of altitude across different MSAs, we transform altitude data by standardizing at the MSA mean (meters above or below the MSA average). Figure 2 provides boxplots of elevation variance distribution across dataset MSAs, and Figure A1 presents boxplots for tract (nominal) altitude.

[5]The local water level is approximated by the minimum altitude value for the MSA. For coastline cities, we omit all elevation samples that round to a sea-level altitude of 0.

Figure 2: Boxplot of log elevation variance, all dataset MSAs

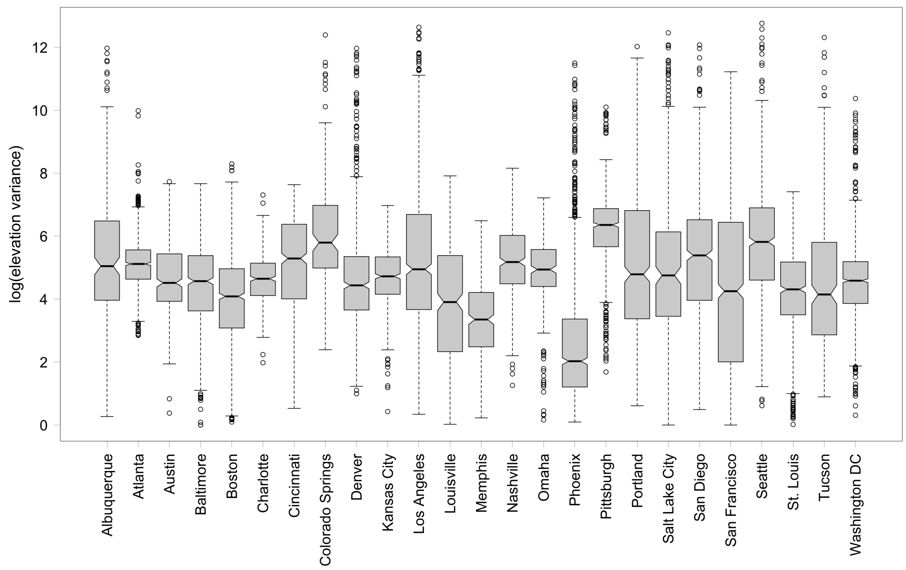

REST is also utilized to generate distance and time estimates by driving from each census tract to their respective CBDs.[6] While not all residents in all tracts regularly commute to the center, proximity to the CBD is highly correlated with a rich variety of spatial amenities. Distance to center is also a proxy for valuing marginal housing consumption versus marginal costs of commuting of local households: downtown residents are more likely to trade housing for low commuting costs while the opposite holds for suburban residents. By controlling for monocentricity, we account for these preferences in the simulation.

Finally, we construct tract coastline distance variables for coastal dataset MSAs, defined as MSAs where the tract nearest to the coastline is less than 2km away, using National Oceanic and Atmospheric Administration (NOAA) coastline profiles and calculate tract distances to the coastline.[7] Distances are estimated from respective tract centers to the nearest point on the coastline.

3.2 Model and Simulation

Our structural specifications and simulation model are based on the neighborhood sorting framework proposed by Bayer, McMillan and Rueben (2004) and developed in Bayer and McMillan (2005). Households are categorized by 10 income bins matching respective income deciles of the entire dataset. It is

[6]The default choice among route alternatives is shortest driving time. The shorter route is chosen if driving time is identical to the minute. For consistency, all driving times assume no local traffic.
[7]See NOAA definition for "coastline". http://shoreline.noaa.gov/glossary.html

crucial that bins are matched to the household income distribution of the entire dataset and not individual MSAs, since distributions of household income vary substantially across cities, and matching income bins to individual MSAs results in bins that are not comparable across MSAs.[8] Since ACS does not report cumulative densities by income decile, we linearly interpolate between ACS reported income brackets to estimate the CDF of households by income in each tract.[9]

Each simulated household's income is assigned by fitting the MSA-level cumulative income density function to a Burr distribution (Burr 1942). For each MSA, income values are drawn from this distribution and, according to income bin boundaries, allocated to simulated households. We further assume that households are of a single ethnicity and not specifically renters or owners. Additionally, all households in a particular income bin have identical preference functions over the quality of neighborhood amenities.[10]

For each income bin i, a household living in census tract t is assumed to derive utility over residential location choices from a composite tract housing stock quality variable $H_t = \alpha_1 h_{1t} + \alpha_2 h_{2t} + \ldots$, housing expenditure P_t, and tract level of terrain unevenness T_t.[11] The household also derives utility from "unobserved" amenities not reflected in the quality of housing stock (e.g. schooling, parks, quality of restaurants), reflected in the tract median income level $\bar{I}_t$, and a trade-off between tract-specific per-unit-cost housing consumption and expected commuting cost, approximated by the driving distance to a central location, which we term downtown, D_t. Correspondingly, we specify the following log utility function:

$$U_{it} = \beta_{1_i} \log(P_t) + \beta_{2_i} \log(\bar{I}_t) + \beta_{3_i} \log(T_t) + \beta_{4_i} \log(H_t) + \beta_{5_i} \log(D_t) \qquad (1)$$

Since the residence location decision of any individual household is not expected to significantly influence either the market-clearing price or median tract income, the marginal household for a specific income bin take as given the endogenously determined price P_t and $\bar{I}_t$. The household chooses to locate in

[8] Our matching approach does result in unbalanced bins at the MSA level. However, this does not adversely affect the simulation as groups optimize taking tract-level prices and amenities as given, assuming that no individual bin is too small to significantly influence demand in at least some tracts.

[9] The number of households in each bin is rounded to whole households. When rounding leads to the total tract household count exceeding or becoming less than the original value, the bin that is rounded upwards or downwards the most is rounded in the opposite direction.

[10] Since both race and ownership are highly correlated with household income, it may be useful to think of households within each income bin as being a similar composite of races and renter-versus-owner status. However, we do not explicitly allow for such sorting mechanisms in the simulation model to reduce computational complexity. Concurrent racial sorting, along with housing stock aging and gradual expansion of the MSA, are topics for subsequent analysis.

[11] h_t variables represent specific aspects of housing stock quality such as fraction of single owner units and age composition of units.

the tract where the endogenous amenities, in combination with fixed amenities such as terrain, housing stock quality and distance to downtown, provides the greatest amount of utility among all tracts within the MSA.

As is often the case in Dynamic Discrete Choice models, as pioneered by Pakes (1986), we estimate the parameters of equation (1) by assuming that the deviation of log shares represents utility derived from each unique choice of tract. Specifically, we assume that the deviation of income bin i's log share in each tract t, $\log(S_{it})$, from the log share of tract t's household count of the MSA, $\log(S_t^{msa})$, is a representation of i utility derived from tract t, U_{it}, plus a stochastic, EVT1 component ε_{it}. Intuitively, the ratio S_{it}/S_t^{msa} represents how desirable (or undesirable) a tract is for a particular income bin as a deviation from the "average" level of utility that the income bin derives from the MSA. Hence, the tract-bin specific estimating equation is:

$$\log(S_{it}/S_t^{msa}) = \beta_{1_i}\log(P_t) + \beta_{2_i}\log(\bar{I}_t) + \beta_{3_i}\log(T_t) + \beta_{4_i}\log(H_t) + \beta_{5_i}\log(D_t) + \varepsilon_{it} \quad (2)$$

Solving this equation for each MSA and income bin within an MSA yields MSA-bin-specific preferences for each specific amenity. To construct a single measure of housing stock quality H_t, we use a number of control variables including fractions of housing units by bedroom count, fraction of single-household detached homes, fraction of owner occupied units, fraction of mobile homes as well as vacancy and the age distribution of tract housing stock. D_t is characterized by driving time to the CBD, and T_t is characterized by tract elevation variance, tract relative altitude, and the interaction between the two variables.

Instead of allowing discrete choices over tracts, we assign household locational choices fractionally by a multinomial distribution over all tracts with weights determined by relative log utility derived from each tract.Specifically, the fraction of a household h_i in tract t is given by:

$$Frac_{h_i t} = \frac{\log(S_{it}/S_t^{msa})}{\sum_t \log(S_{it}/S_t^{msa})} \quad (3)$$

where the denominator sums over log share ratios of i across all tracts within the MSA. A significant issue with updating entire households is that the response to small changes in a tract's bin-specific utility may result in large changes to the tract's composition of households, greatly restricting the speed at which the model can be solved. Using fractional updates, entries and exits into tracts always happens at the margin, allowing for faster, smoother updating and less need for scaling up the model to real-world sizes.

Prices clear markets. In our model, the amount of housing supply for each tract is fixed. Hence, prices are driven solely by fluctuations in demand. We adjust price incrementally upwards in oversubscribed tracts and downwards in undersubscribed tracts until each tract recovers, precisely, the original number of residing households. Since the total number of households in the simulation is always preserved by the aforementioned updating process, at the market-clearing state each tract should have exactly as many households as in the original data. We note that this implies that the price sensitivity term, β_{1i}, must be strictly negative for all income bins. Otherwise, prices adjust upwards perpetually and some tracts remain oversubscribed.

Tract median income is estimated as the income of the household whose fraction straddles the 50th percentile of the income CDF.[12] Intuitively, preferences for this value are likely to be strictly positive for households with income levels substantially above the MSA-level average, as rich households should not dis-prefer their own marginal move-in effect on a neighborhood. It is possible that poor households see such amenities as either desirable or undesirable, depending on the strength of the amenity effect relative to the desire of households to match their budget with neighborhood amenities.

We perform the counterfactual simulations by first running each MSA's model to its steady state, defined as the point when the maximum deviation between current and previous-iteration prices and tract median incomes among tracts is less than 0.5%, and markets clear to within 0.5% of each tract's expected number of simulated households. After performing changes specific to a particular counterfactual, we then continue to iteratively update the model until a new steady state is reached.

3.3 Computation

The primary challenge in estimating preferences for locational choice is the strong correlation between prices and the quality of a neighborhood, including both observable and unobservable amenities. In a regression model, it is uncertain that explanatory power will be correctly distributed between different preferences, potentially resulting in price preferences being biased upward or downward and leading to unrealistic simulation steady states. A related concern is that ACS does not report prices faced by each income bin. Consequently, our price sensitivity measures are in effect not sensitivities to actual prices experienced by each income bin but bin-specific sensitivities to prices which are reflected by movements of the tract median price.

This distinction has significant implications for the model. Low-income households experience prices significantly below the tract median, facing less than unit change in the actual price for housing when the median changes.[13] Conversely, rich households face greater than unit change when the median changes by one unit.[14] The regression model does not perfectly capture this

[12] In the case that fractions add up to exactly 0.5, the lower income of the two households closest to 0.5 in the CDF is used.

[13] This is because if we consider prices of units in each tract as distributed roughly log-normally, the density center of the truncated area below the median does not respond linearly to changes of the median.

[14] This also means that it is ambiguous as to whether rich or poor households actually display greater price sensitivity to changes in tract median prices, even though rich households are less sensitive to changes in the nominal price.

distinction, and is biased in its estimations of price sensitivities. Performing the simulation using preferences extracted from OLS results in market-clearing prices that are unrealistically high: many tracts do not clear until prices are 200-1,000 times higher than the maximum price in the original data.

Our solution to this issue is threefold. First, we use an iterative calibration approach to estimate preferences for price, β_1, by adjusting preferences until the simulation steady state aligns with distributions in the census data over a set of targets. To avoid the possibility of the starting guess of price preferences influencing the calibrated preferences, we initialize the model with an extremely small price sensitivity of -0.05 for the richest income bin. To prevent divergent sequences of guesses, we adjust this value downward slightly for other bins, increasing the sensitivity according to relative tract average prices experienced by each bin.[15]

Calibration targets are tract median prices and income levels by income bin: each value is estimated as the average for all households in the bin or 20 targets in total per MSA for the 10 income bins.[16] We increment β_{1_i} by each bin according to the amount of deviation between average prices experienced by each bin in the simulation steady state and that of the data. An income bin i facing prices that are overall too high implies that β_{1_i} is biased upwards, i.e. not sensitive enough and need to be decreased, and vice versa.

When calculating prices, we omit the top 0.5% of tracts by price to prevent tracts that are persistently expensive from driving the calibration process. Tracts can become resistant to changes in price sensitivities for a number of reasons not easily addressed in the simulation: scaled-down small tracts being fully subscribed by large fractions of a few simulated households, data errors, or large local amenity bundles (e.g. cultural landmarks, museums) that cannot be accounted for through tract income alone. This procedure yields substantially smoother updating at the cost of only a small fraction of tracts: the maximum number of MSAs not accounted for is 6 tracts for Washington, DC.

New β_1 guesses are proposed using Newton-Raphson with cross-partials assumed to be zero. After obtaining new guesses for β_1, the remaining preferences – tastes for unobservables, elevation, housing stock, and distance to center – are estimated with a regression model with β_1 treated as a fixed coefficient. From this set of new preferences, we re-run the simulation, and the process iterates until the steady state is reached. Hence, the model is iteratively calibrated until all targets fall within 10% of the original data, which we consistently achieve across all dataset MSAs.

Instead of calibrating preferences for unobservables, β_{2_i}, by iterative incre-

[15] These values are estimated by taking shares of each bin in each tract and obtaining share-weighted housing prices for all bins. Assuming that the richest decile faces an average price of P^{10}, decile i's initial price sensitivity is set to $-0.05 \cdot \frac{P^{10}}{P^i}$.

[16] For example, a household with the same fraction assigned to n tracts experiences price and tract income that are the simple averages of those of all n tracts. Households in each bin are then averaged to derived the average price and tract income experienced by each bin.

mentation, we estimate β_{2_i}'s directly from the regression. However, we compare tract median incomes experienced by each bin in the steady state against that of the data, and use the deviations as clearing conditions. The calibration is not stopped until both targets for prices and tract median incomes clear at the 10% level. If, in the calibrated steady state, bins experience prices that align with the data but live, on average, in neighborhoods that are too rich or too poor, then it is likely that the β_1 values are systematically mis-calibrated. On the other hand, income and price targets being simultaneously met when we do not specifically calibrate against the former suggests that model is being correctly parameterized. A flowchart summary of the entire calibration procedure is presented in Figure 3.

Figure 3: Outline of calibration and simulation procedure

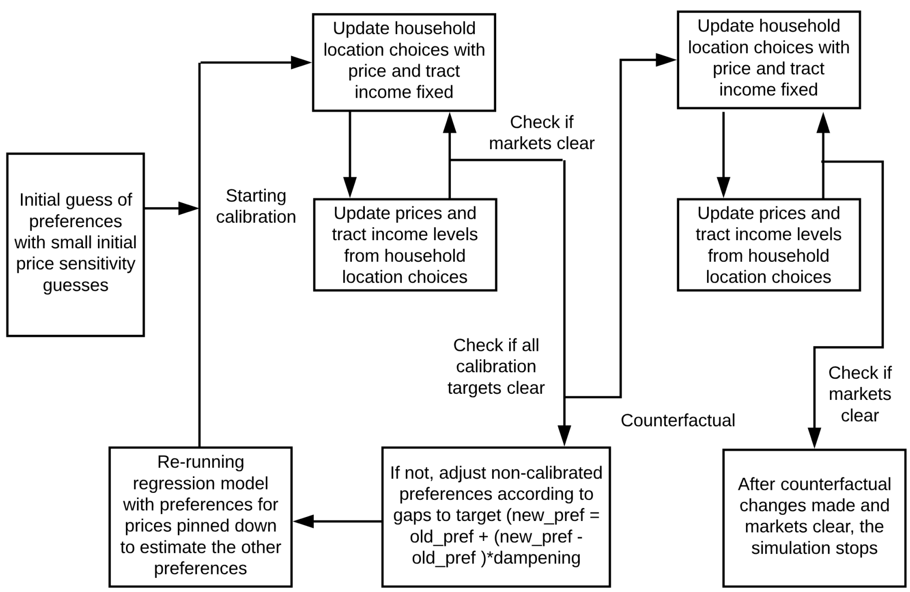

Our second solution component is that we exploit higher moments of the MSA-level price distribution to discipline the convergence process. Intuitively, preferences that are well-calibrated should not result in prices in the steady state that are far above the maximum or below the minimum price observed in the data. To this end, we increment price sensitivities for all income bins upward by 2% whenever the price of the most expensive tract in the simulation steady state exceeds 1 standard deviation above that of the most expensive tract in the data.[17]

If no tracts are more expensive than the most expensive tract and the cheapest simulated tract is less than half as expensive as the cheapest dataset tract,

[17] This adjustment, if applicable, occurs concurrently with new guesses for β_1, and is applied to the new set of guesses before the next calibration iteration.

we decrease sensitivity by 2% for all bins. Correspondingly, we check maximum prices when determining whether the model has been successfully calibrated. The clearing threshold is set as the price of the 99.5th percentile most expensive simulated tract not exceeding 1 standard deviation of that of the 99.5th percentile dataset tract.[18]

The final component of the solution is that we substitute OLS with a LASSO regression model (Tibshirani 1996) with a λ parameter obtained by 20-fold cross validation. The LASSO or L2 penalization regression optimizes for predictive performance by penalizing all coefficients toward zero, and a subset of coefficients to exactly zero. In effect, the process assumes that some parameters of the model are too small to be meaningfully distinguished from zero, and setting them to exactly zero optimizes predictive power. Hence, LASSO prevents the model from overfitting and extracts more generalizable preferences, which is highly desirable from the perspective of a structural model.

While somewhat complex, this three-fold approach results in rapid convergence of the steady state with respect to calibration targets: in the baseline model, all 25 MSAs hit price and tract income level targets within 10% tolerance and satisfy maximum price conditions within 300 calibration iterations.[19] Additionally, the elevation gradient-income relationship is well behaved despite not being explicitly calibrated, with the average absolute deviation between steady state and original data expected elevation variance being only 3.0% across all bins in all MSAs in the steady state.[20]

Relative altitude is also well-recovered in the calibrated steady state, albeit slightly less so than variance. Among MSA-income bins that on average live more than 10 meters away from the MSA average altitude, the average absolute deviation is 7.5% and the maximum deviation 25.5%.[21] These results suggest that our calibration procedure is successful in allowing the simulation to reach a steady state that is well-behaved with regard to real-world distributions of income, prices, and locational choice by elevation gradients, and by extension extracting relatively useful estimations of income bin-specific preferences.

Specifically, we set a scaling level for number of simulated households that allocates at least 1000 simulated households to the smallest income bin in a given MSA. This yields a total of 369,111 simulated households at an average scaling ratio of 133.1 real households per simulated household. To ensure convergence of the tract median income targets, we scale all simulated household incomes by a fixed factor so that the average of simulation steady state tract median incomes equals exactly the average of the data. This adjustment does not introduce any restrictions on sorting choice but simply guarantees that if one bin lives in tracts

[18] Standard Deviations are calculated from dataset price distributions.

[19] Specifications of different counterfactuals are discussed in detail in Section 4.

[20] SD = 3.0%. Both mean and standard deviation weighted by MSA-income bin sizes. The maximum deviation is 14.4% and 11 out of the 250 MSA-bins have a deviation of greater than 10%.

[21] Maximum absolute deviation for MSA-income bins that live in tracts less than 10 meters away from the MSA average altitude is 8.4 meters, and average abs. deviation is 0.73 meters.

that are richer than they do in the original data, at least one other bin must live in tracts that are poorer.

We also standardize prices to the average of all tracts' median prices after every new guess of prices. Conceptually, since prices adjust to clear markets, uniformly increasing or decreasing prices everywhere should not affect the sorting equilibrium or relative preference of each income-bin. Similar to the standardizing of incomes, this process guarantees that if one bin lives, on average, in more expensive tracts, one or more other bins must on average live in cheaper tracts.

To prevent the simulation from potentially being stuck in local optima, we update substantially more aggressively for the first 20 iterations in each simulation run. For each new guess of preferences or "calibration iteration", we also do not allow the simulation run to stop until reaching 150 iterations, even if the market clearing conditions have been reached. We use a maximum of 800 iterations for each calibration iteration (except the final one) and move to new preference guesses if the model has not converged by iteration 800. All MSAs converge with the final set of preferences before this limit is reached.

4 Results

4.1 Flattening Cities

Our first and primary counterfactual scenario estimates the effect of flattening elevation gradients assuming fixed tract-specific housing stock. As described in Section 3, we assume that spatial amenities - the presence of water areas, distance to the coast and costs of commuting to the CBD - as well as tract-specific housing stock quality variables such as the structure composition and age of units, are fixed for all tracts. However, unobserved amenities adjust with local income levels. In other words, it is assumed that rich households are able to bring a certain amount of non-spatial local amenities with them as they re-sort but cannot alter the quality of the local stock of housing itself.

Preferences extracted from the target-cleared steady state under these assumptions are presented in Table 2.[22] Tastes for tract median income are generally more positive for higher income bins, and unobservables are strictly an amenity for almost all bins across MSAs. Similarly, richer households are less sensitive to prices. The richest households are approximately twice as sensitive to changes in unobserved amenities, as reflected in log tract median income, and half as sensitive to changes to housing prices as the poorest households.

[22] Price and income preferences by MSA and income bin are provided in Tables A1 and A2.

Table 2: Select preferences averaged across MSAs, baseline model

Decile	1	2	3	4	5	6	7	8	9	10
median price[23]	-0.425	-0.412	-0.407	-0.390	-0.368	-0.351	-0.342	-0.306	-0.274	-0.220
median income	0.171	0.186	0.250	0.297	0.306	0.329	0.343	0.317	0.332	0.396
elev variance	0.028	0.016	0.020	0.016	0.009	0.008	0.008	0.014	0.026	0.037
relative altitude	-0.012	-0.005	-0.004	-0.002	-0.001	0.001	0.007	0.015	0.023	0.028
elv-ral interaction	0.001	0.000	-0.000	-0.000	-0.000	0.000	-0.001	-0.002	-0.002	-0.002
CBD Drive dist	0.114	0.142	0.118	0.094	0.075	0.040	0.022	-0.020	-0.133	-0.338
Income cutoff ($)	14,556	25,948	37,387	49,486	62,984	79,013	98,617	124,974	179,410	-

Elevation preferences are weakly upward sloping with respect to own income. Preferences for altitude is positive for rich households and negative for poor households, consistent with findings of Ye and Becker (2018) that flatness and low altitude may be an amenity to the poor because of walkability and access to public transit. While the bottom income deciles prefer elevation variance more strongly than middle class households, this effect is dominated by preferences for altitude for tracts that are above the MSA average altitude. The altitude-variance interaction is small overall and negative for rich households. We speculate that the interaction reflects diminishing returns to living in particularly high-altitude and low-lying areas for the rich and poor, respectively, though this effect does not appear to be significant for most MSAs.

Finally, a downwards-sloping preference to driving time to the CBD suggests that the rich have higher costs of commuting, holding marginal consumption of housing fixed. This is consistent with the general monecentric city model pattern of the rich either suburbanizing, where per-unit-area cost of housing drops sharply and households can trade commuting costs for large quantities of housing, or residing at the very center, where commuting costs are the lowest.

We perform the counterfactual by gradually reducing the log elevation variances and relative altitudes for all tracts to zero over 50 iterations. The post and pre-counterfactual steady-state differences between expected tract median prices, income levels, driving distance to CBD's and elevation variance levels are presented in Figure 4.[24] Importantly, expected elevation variance in this figure is estimated at *the original data's* elevation variance levels, or otherwise expected variance would be exactly zero in the counterfactual.

[23] Tract median prices, incomes and elevation variance logged. We add 1 to all values so that setting the log value to zero is equivalent to zeroing the value. Driving distance to CBD in km is also logged. Relative altitude is per 10-meter increment. Cutoff of own-group income in 2015 dollars.

[24] Values are percentage changes to each bin caused by the counterfactual. Each household's expected price and tract income is estimated as the mean across tracts weighted by location choice fractions. Values for households in each bin are averaged again to calculate differences.

Figure 4: Counterfactual outcomes from flattening cities by income decile.

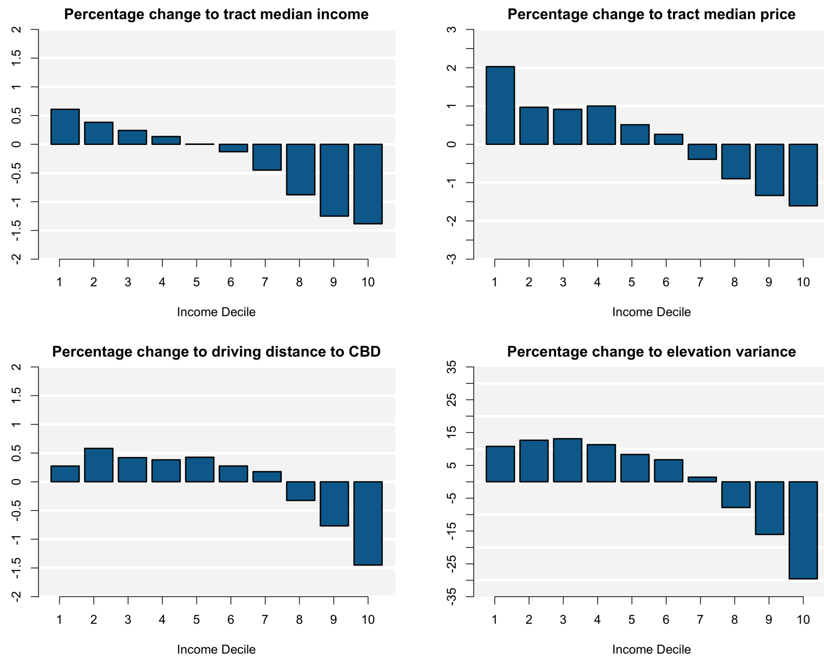

In the counterfactual, the top income bin's expected tract median income falls by approximately 1.4%, and the bottom bin's expected tract median income rises by approximately 0.6%. All bins above the dataset's median income level live in poorer tracts in the counterfactual. Prices fall for rich income bins and rise for poor income bins : the top bin lives in tracts that are approximately 1.6 % cheaper, and the bottom bin lives in tracts that are 2% more expensive. Additionally, the rich move away from the suburbs when cities are flattened: the top decile lives approximately 1.4% or 0.5 km closer to the center in the counterfactual.[25]

Changes to the pattern of income distribution by elevation variance and altitude are much larger. In the counterfactual, the top decile lives, on average, in tracts that were, pre-flattening, 30% flatter and deciles 1 through 4 all live in tracts that were more than 10% hillier. Households in the top decile live in tracts that were 7.5 meters or 69.5% lower in altitude, and households in the bottom decile live in tracts that were 1.4 meters higher in altitude.

Strikingly, the gradient of own group income against elevation variance is almost completely reversed in the counterfactual. The top deciles live in the flattest areas in the counterfactual while in the initial steady state and the original data they live in the most uneven (measured at the pre-counterfactual

[25]We choose to not report welfare effects here because of the difficulty in contrasting an estimate of welfare across MSAs with both different scaling and unbalanced income bins. However, this would be an interesting topic for subsequent analysis.

level). The gap of 19.7 meters between the expected relative altitude of decile 10 and decile 1 shrinks to only 8.5 meters, and in the counterfactual decile 10 no long lives at higher altitudes than deciles 7-8. We present plots of relative altitude and elevation variance gradient by income level in Figure 5.[26]

Figure 5: Relative altitude and elevation variance gradient by income level, flattening cities

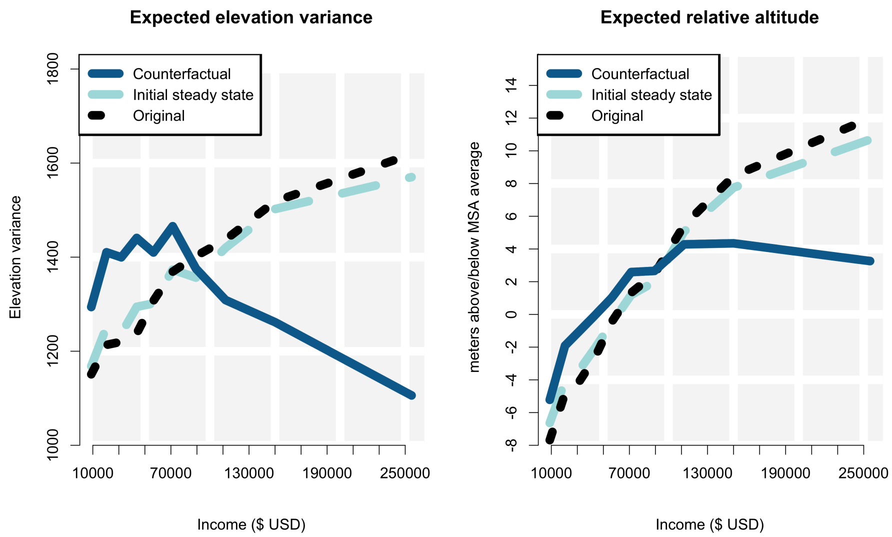

The story is best illustrated by observing the differences in cross-decile exposure between the counterfactual and the initial steady state, where the exposure of decile i to to decile j is the expected fraction of households in decile j, in the average tract, for the average decile i household. We present the percentage difference between cross-decile exposure of the counterfactual and the initial steady state in Table 3. Here we observe that exposure to the top and 9th decile go up for deciles 1 to 6. The poorest decile expects to live with 3.9% more households in the top decile, and 2.6% more households in the 9th decile.

Correspondingly, self-exposure of the top income bins drop. The top 10% of households by income live with 3.6% fewer households in their own group and 2.8% fewer households in the 9th income decile. We note that, consistent with Figure 4, deciles 1-5 gain by being exposed more to richer bins and less to poorer bins. In decile 1's case, exposure to deciles 1-5 drops while exposure to deciles 6-10 rises. As rich households cease to concentrate in high-altitude, high-variance areas, they sort into locations with larger fractions of low-income households.

[26] Expected relative altitude is not perfectly centered at zero because we do not weight tracts when we calculate MSA average altitude, but weight tracts by number of households when calculating expected altitude levels.

[27] Differences are estimated as (counterfactual - initial)/initial. Note that cross-decile exposure matrix is symmetric.

Table 3: Average percentage difference in cross-decile exposure between counterfactual and baseline, flattening cities[27]

Decile	1	2	3	4	5	6	7	8	9	10
1	-1.04	-	-	-	-	-	-	-	-	-
2	-1.02	-0.78	-	-	-	-	-	-	-	-
3	-0.74	-0.70	-0.52	-	-	-	-	-	-	-
4	-0.52	-0.52	-0.49	-0.38	-	-	-	-	-	-
5	-0.25	-0.25	-0.27	-0.24	-0.10	-	-	-	-	-
6	0.07	0.01	-0.09	-0.13	-0.10	-0.11	-	-	-	-
7	0.45	0.31	0.17	0.07	-0.03	-0.13	-0.18	-	-	-
8	1.27	0.98	0.72	0.47	0.14	-0.11	-0.33	-0.59	-	-
9	2.57	1.94	1.50	1.05	0.46	0.00	-0.42	-1.17	-1.73	-
10	3.85	3.13	2.44	1.98	0.98	0.36	-0.21	-1.42	-2.81	-3.60

Prices fall for the rich as they sort among a larger group of tracts and no longer compete for hilly tracts. As the rich enter poorer neighborhoods, competition increases in these neighborhoods and prices are bid up. Additionally, the rich centralize as cities are generally more hilly in the suburbs than at the center.

The outcome of rich households living in even flatter areas than the poor in the counterfactual is most likely caused by elevation variance being negatively associated with high quality housing stock and neighborhood amenities. Correlation is positive and significant between elevation variance and median age of structure (0.049), fraction of mobile homes (0.104), as well as the fraction of homes with five or more bedrooms (0.029).[28] While we do not have data for schools and recreational facilities, one would also expect them to locate in flatter areas because of lower construction costs, all else being equal. This effect also draws the poor to uneven areas in the counterfactual, as the same housing stock qualities are potentially an amenity to the poor.

4.2 More Mountains

Our second, "reverse" counterfactual simulation uses the same initial steady state and procedure as the first, but with one major distinction: instead of gradually setting elevation variance and relative altitude to zero across 50 iterations, we gradually adjust both values to twice that of the original data. Conceptually, this means that for a given MSA, we not only make every tract twice as uneven in terms of variance but also the entire distribution of tract altitudes twice as varied: low-lying areas are twice as low-lying and vice versa.

Post- and pre-counterfactual steady state differences are presented in Fig-

[28] All correlations report $p<0.001$. We note that while more bedrooms are typically an amenity, homes with more than 5 bedrooms are either of extremely high quality or have been converted to multi-family use, the latter of which is associated with low neighbor income.

ure 6, corresponding to Figure 4 in the previous counterfactual. Plots of relative altitude and elevation variance gradient by income level are presented in Figure A4. Similar to Figure 4 we calculate expected elevation variance and altitude at original dataset levels instead of doubled levels.

Figure 6: Counterfactual outcomes from doubling elevation variance and relative altitude by income decile.

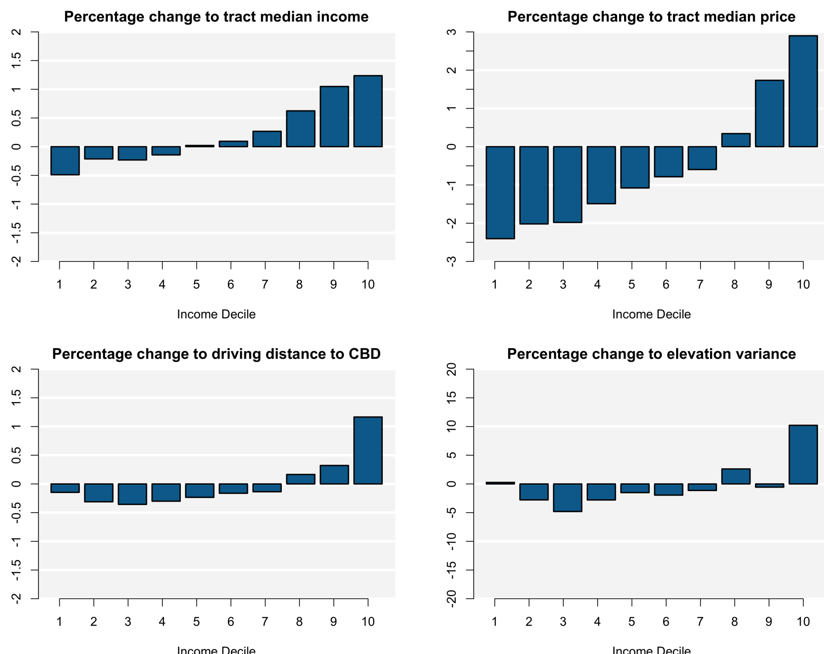

The richest decile see an increase of approximately 1.2% in tract median income and the poorest decile sees a decrease of 0.5%. Deciles 1 to 4 live in poorer neighborhoods and deciles 5 to 10 in richer ones. Top earners must pay more for living in richer areas: prices are bid up 3% for decile 10. The rich suburbanize, with the top decile moving 1.2% further from the center. Change to expected elevation variance and altitude is concentrated in the top income decile: they live in tracts that are 10.2% more uneven and 4.6 meters higher in altitude. Decile 9 also lives in higher altitude locations (0.8 meters) but experiences very little change in expected variance (-0.6%).

Cross-decile exposure is presented in Table 4. The outcome is the opposite of the flat city scenario: own-group exposure of the 10th decile increases by 3.1% while exposure to decile 10 of decile 1 decreases by 2.2%. The richest households select themselves into uneven tracts, sorting away from all other groups and increasing income stratification at the MSA-level. While the effect is most significant for the top 20% of households against the remaining 80%, lower-income households also sort among themselves, with exposure to adjacent deciles rising for all groups.

Table 4: Average percentage difference in cross-decile exposure between counterfactual and initial steady state, doubling elevation variance and relative altitude

Decile	1	2	3	4	5	6	7	8	9	10
1	0.90	-	-	-	-	-	-	-	-	-
2	0.59	0.50	-	-	-	-	-	-	-	-
3	0.52	0.46	0.51	-	-	-	-	-	-	-
4	0.35	0.36	0.39	0.34	-	-	-	-	-	-
5	0.12	0.14	0.20	0.19	0.17	-	-	-	-	-
6	-0.07	0.05	0.11	0.14	0.14	0.17	-	-	-	-
7	-0.37	-0.17	-0.11	-0.01	0.08	0.17	0.22	-	-	-
8	-0.84	-0.58	-0.52	-0.30	-0.04	0.12	0.28	0.58	-	-
9	-1.55	-1.03	-1.07	-0.74	-0.36	-0.16	0.18	0.66	1.39	-
10	-2.19	-1.80	-1.80	-1.43	-0.88	-0.74	-0.25	0.52	1.86	3.12

These results strongly contrast those of Section 4.1. When cities are flattened, rich households in the top 10%-20% of earners suburbanize, mix with poorer households, move to relatively flatter areas and, by not competing for unevenness and high altitude areas, enjoy lower equilibrium prices. On the other hand, when we introducing greater unevenness, rich households concentrate and compete for hilly locations, causing other groups to both live in poorer areas and face lower prices.

4.3 Endogenous Housing Stock

Thus far, we have assumed that households cannot influence the composition of neighborhood housing stock in the sorting process. As an extension, we relax this assumption to estimate the size of elevation gradient effects when neighborhood amenities are completely reflected in neighborhood income level. In other words, in addition to exerting influence on non-housing attributes such as school quality, restaurants and crime rates, households determine all non-spatial amenities associated with the neighborhood and can modify the type and quality of local structures to maximize utility.

We conduct this exercise for two reasons. First, a realistic static simulation of neighborhood sorting would incorporate some flexibility of housing stock quality: age of structure is not a perfect reflection of unit quality, and units can be meaningfully improved even in the very short run. By locking down housing stock quality, we disallow such changes and constrain how strongly households can respond to movements in neighborhood income. Hence, our counterfactual outcome in sections 4.1 and 4.2 underestimates the general equilibrium effects associated with altering elevation profiles. Fully flexible housing stock as we assume in this section, on the other hand, overestimates but provides an upper bound to the problem.

Second, we do not explicitly calibrate preferences for housing stock quality in the simulation and take the regression-estimated housing stock quality preferences as given. Running the simulation without housing stock variables provides an external check of the sensitivity of our previous results to the method of deriving these preferences. Results that are substantially inconsistent with the previous sections would suggest that such preferences are mis-estimated.

Specifically, we calibrate and simulate dataset MSAs using the same setup as 4.1-4.2, but without any housing stock quality variables in the set of household preferences. We fully calibrate the model from the same initial guesses as described in Section 3.3. This is because households preferences are relative not only to those of other deciles but also preferences for other goods, and hence loading the quality of housing stock onto tract median income necessarily requires new preferences and sensitivities for all groups. We use the same clearing conditions and, similar to the original setup, all MSAs clear expected prices and tract income level targets within 200 calibration iterations.

Results for the "flat cities" and "more mountains" counterfactuals without constraining housing stock are summarized in Figures A5 and A6, respectively. Figure A7 and A8 contrasts relative altitude and elevation variance gradient by income level with and without fixed housing stock quality. In the new flatness counterfactual, the richest decile experiences a greater drop in expected tract income (2.1% instead of 1.4%), pays more (3.2 %) instead of less for housing, centralizes more strongly by living a further 0.5% closer to the CBD, and still lives in much flatter and lower altitude areas compared to the original data. Notably, the richest decile centralizes - they now live in locations that have the lowest average elevation variance *and* lower than MSA-average altitude.

Higher prices for the rich are consistent with the rich being a greater amenity to the rich when we endogenize housing stock quality. When cities are flattened, the rich face two price effects: that of demand decreasing among themselves in originally highly uneven and high-altitude tracts and that of demand from the poor increasing in such tracts. When housing stock quality is endogenous, both the rich and poor face fewer inherently desirable or undesirable locations and experience larger demand changes in response to flattening the city. If this response to demand is proportionately much larger for the rich, the effect of rich households moving out of previously uneven areas will dominate that of poor households sorting into such areas. It follows that prices adjust upwards in tracts to which rich households relocate.

Income changes, price changes and location choice relative to the CBD in the reverse counterfactual with endogenous housing stock are consistent with those without: when cities are twice as uneven, rich households suburbanize, live in richer neighborhoods and pay higher prices. However, we observe that rich households live in flatter neighborhoods on average while the poorest decile live in much more uneven tracts (as shown in Figures A6 and A8). Both the richest and poorest decile live in areas that are higher altitude, although the direction of the altitude-income relationship is largely preserved.

The reason for the bottom, and only the bottom decile experiencing much more uneveneness in this counterfactual is that a fraction of the poorest households are crowded out, from the center, into tracts that are far away from the center, cheap and highly uneven. This is reflected in the poorest decile living both about 1% further away from the CBD on average and approximately 5 meters higher in altitude. It is important to note that the top decile is not the group that pushes the poor away from the CBD, but the 3rd to 9th deciles, as illustrated in Figure A6. As the top income bin sort much more strongly among themselves and occupy tracts that have the most desirable combination of elevation and spatial amenities, the upper-middle and middle class move into the center, and a fraction of the poorest are forced out into the far suburbs.

This result is in contrast to that of Section 4.2 as shown in Figure 6, where the bottom decile centralizes and experiences virtually no change in expected elevation variance. When housing stock is fully endogenized, preferences for local amenities - housing stock quality included - is represented entirely by price sensitivity and tastes for tract median income. Now, the richest decile, effectively being able to adjust local housing quality at will, is further differentiated from other income groups. Hence, they sort along with households in the 7th-9th decile when housing stock quality is exogenous, but sort *away* from them when such preferences are endogenized. The difference in outcomes of the poorest income bin follow as they occupy locations that are less desirable to richer households.

4.4 Supply Side Effects

To provide contrast to the demand-side effects of Sections 4.1-4.3, our final counterfactual exercise studies the supply-side consequences of elevation gradients. Specifically, we simulate the mechanism discussed by Saiz (2010) and, more recently, Baum-Snow and Han (2019), where unevenness leads to higher costs of construction and particularly so for multi-unit structures, larger plots and lower utilization of land, all of which are amenities to the wealthy and disamenities to the poor. Conceptually, the most straightforward approach to addressing this question would be to adjust housing quality in each tract toward an MSA-level average, while keeping elevation gradients unchanged.

However, we cannot simply simulate a "flattening" of housing stock qualities as a direct parallel to Section 4.1. This is for three reasons. First, our calibrated utility function assigns far greater importance - as it should - to the combined effects of housing stock variables compared to elevation gradient effects. This is fairly intuitive, since households are likely to heavily value the local types and quality of structures as a direct proxy to that of their own residence. Yet it also implies that large changes to the relative levels of amenities in each tract introduce drastic shifts to the MSA-level sorting equilibrium.

Second, the level of homogeneity among tracts is massively increased by setting a uniform level of housing stock quality throughout each MSA. Hence, the counterfactual simulations are prone to experiencing multiple equilibria issues that are not observed when flattening elevation gradients. Third, as rich households tend to reside in the suburbs and CBDs are (usually) located in the flattest parts of an MSA, one would expect to see *some* correlation between housing stock quality and elevation gradients even if there is no direct effect of elevation on housing stock. This implies that any counterfactual that reduces the correlation coefficient between elevation and housing stock quality to exactly zero overstates the importance of elevation gradients.

Our solution to these issues appeals to the fact that the utility that each income bin derives from housing stock is influenced by the monocentricity of the MSA. For example and very generally, downtown areas have more units per structure, older housing stock, and units with fewer bedrooms. However, adjusting for distance, the relative amount of utility provided by each location differs. This variation introduces variance among income levels of tracts at any given distance to the CBD. Hence, we break down the utility that each income bin derives from the single measure of housing stock quality:

$$U_{H_{it}} = \beta_i \log(H_t) = \beta_i \log(\alpha_1 h_{1t} + \alpha_2 h_{2t} + \dots)$$

Into two components:

$$U_{H_{it}} = \hat{U}_{H_{it}} + \epsilon_{H_{it}}$$

Where $\hat{U}_{H_{it}}$ represented the predicted utility that a household in i should derive from tract t given its distance to the CBD, and ϵ_{H_it} represents the t-specific amount of utility beyond or below $\hat{U}_{H_{it}}$. We extract these components respectively by the predicted values and residuals of an OLS regression between $U_{H_{it}}$ and $\log(D_t)$, log driving time to the CBD. The residuals are hence interpreted as the level of utility (or disutility) that t provides to i controlling for D_t. We then perform the counterfactual simulation by randomizing $\epsilon_{H_{it}}$ over all tracts within each MSA.

This approach has three advantages. First, the presence of $\hat{U}_{H_{it}}$ preserves the monocentricity of the MSA by allowing rich households to derive more utility from bundles commonly found in the MSA's periphery and vice versa, and hence limits the extent to which the new sorting equilibrium differs from the baseline. Second, the overall amount of variance among - in other words, shape of distribution of - utility derived from housing bundles within the MSA is also preserved, ruling out the possibility of one or a few tracts incurring market clearing problems because of an extreme excess or lack of demand.

Third, the counterfactual remains a direct analogy to the altering of elevation gradients discussed earlier in this section. In the case of sections 4.1 and 4.2, housing stock quality is held constant while elevation gradients are altered with regard to a given overall target level (zero or twice as much). Here, elevation gradients are unchanged but we permute the utility derived from each tract's housing stock with regard to a distance gradient that is smooth across space.

We note that the counterfactual outcome is still likely overstated. Other natural amenities (lakes, rivers, coastal views) correlate with both elevation gradients and high quality housing stock, and the permutation approach removes the relationship between such amenities and housing stock quality. In addition, our simplified regression approach may not capture distributions of utility derived from housing stock with multiple local optima(the very rich preferring to either be extremely close to the center or quite far away), which would potentially result in the residual component being over-estimated.

The income sorting effect on inequality of this counterfactual, as presented in Figure A9, is roughly comparable to that of flattening the city with endogenous housing stock, and slightly above that of flattening the city with fixed housing stock. The top income bin lives in tracts that are 2.2% poorer and the bottom bin lives in tracts that are 1.5% richer. However, the price effect is quite large - approximately 50% higher for households in the top income decile - as a consequence of the rich bidding intensely for a few select locations where elevation variance and good housing stock bundles still exist simultaneously post-randomization.[29] The top decile live in tracts that are 20% flatter, and the bottom decile in tracts that are 15% hillier.

Table A7 summarizes cross-decile exposure of households by permuting housing stock amenities. Notably, while the top and 9th decile both increase exposure to themselves, cross-exposure among individuals above the MSA median income decrease broadly, as does cross-exposure for those earning below the MSA median. By de-linking elevation gradients with housing stock preferable to the rich, the top decile becomes split among those who can afford the premium of the remaining tracts with an "ideal" combinations of housing stock quality and elevation, and those who are forced to mingle with poorer households.

The final distinction between this counterfactual and that of 4.1 is that the richest decile suburbanizes as opposed to centralizing. This is because here, the desirable elevation bundles of the suburbs remain intact, but the particularly desirable housing bundles of the CBD are permuted. While we cannot conclude precisely whether the demand or housing supply effect is stronger in determining the correlation between income and hills, the broad similarity of the income and expected elevation gradient outcomes of 4.1 and 4.4 suggests that both effects are not unimportant. The relationship between neighborhood income and topography is most likely best explained by a combination of demand and supply-side effects.

[29] As households in our model respond to log prices with linear sensitivity, in reality the top decile is likely far more sensitive to large increases in housing prices. However, higher marginal sensitivity would not fundamentally change the sorting equilibrium, only the price at which the equilibrium is reached by equalizing demand and supply.

5 Conclusion

In this paper, we demonstrate with a simulation-based approach that topographical structure - the distribution of elevation within and among locations - matters in terms of neighborhood sorting outcomes. By simulating household location choices for 25 major MSAs containing 29% of all US urban tracts, we show that rich households sort to the hilliest locations when cities become hillier, and sort away from such locations when cities are perfectly flattened. When cities are flattened, rich households centralize, living in relatively poorer neighborhoods at lower prices. When cities become more uneven, rich households suburbanize, live in richer neighborhoods at higher prices.

We consider two main assumptions regarding the interaction between quality of local housing stock and neighborhood income. In the base case, we assume that housing stock quality is perfectly immutable and cannot be adjusted as households move in and out of a neighborhood. We consider a second scenario where housing stock quality is perfectly adjustable and fully reflected in the median income level of a given tract. In the second case, both the rich and poor respond more strongly in terms of changes to income, price, and elevation gradient outcomes in the two counterfactuals.

Additionally, we simulate a scenario where hills remain, but we break the link between hills and housing bundles desirable to rich households. This simulation yields income sorting effects and changes to the amount of hilliness experienced by the rich and poor that are broadly similar to those of perfectly flattening cities. This suggests that the relationship between rich households and hills is likely best explained by a combination of the demand-side effects of hills as a natural amenity, and supply-side effects of housing bundles desirable to such households being concentrated in hilly areas.

In so doing, we present a strong case that elevation is neither merely a historical determinant of high-income locations nor only a natural amenity that only applies to substantially uneven cities. Topography in itself plays a non-trivial role in determining the spatial distribution of rich and poor neighborhoods as well as housing prices at the equilibrium; this holds true for a broad range of cities spanning all major US geographical areas, many with only moderate amounts of elevation variance.

We conclude that topography plays a significant, ongoing, and nuanced role in shaping income and housing price patterns in cities. Not only is income segregation in uneven cities qualitatively different from that of extremely flat cities, certain locations in such cities are also fundamentally attractive or unattractive to high-income households and, other locations, to low-income ones. Redistributive economic policies will be less effective in more uneven cities because of their unevenness, and they must also struggle with an immutable dimension to neighborhood inequality.

References

Allen, Treb and Costas Arkolakis. (2014). "Trade and the Topography of the Spatial Economy". In: *Quarterly Journal of Economics* 129 (3), pp. 1085–1140.

Ananat, Elizabeth Oltmans (2011). "The wrong side(s) of the tracks: The causal effects of racial segregation on urban poverty and inequality". In: *American Economic Journal: Applied Economics* 3 (2), pp. 34–66.

Andreoli, Francesco and Eugenio Peluso (2017). "So close yet so unequal: Spatial inequality in American cities". In: *LISER*.

Baum-Snow, Nathaniel and Lu Han (2019). "The Microgeography of Housing Supply". In: *Working Paper*.

Bayer, Patrick and Robert McMillian (2005). "Racial Sorting and Neighborhood Quality". In: *NBER Working paper* 11813.

Bayer, Patrick, Robert McMillian, and Kim Rueben (2004). "An Equilibrium Model of Sorting in an Urban Housing Market". In: *NBER Working paper* 10865.

Berger, Thor and Kerstin Enflo (2017). "Locomotives of local growth: The short- and long-term impact of railroads in Sweden". In: *Journal of Urban Economics* 98, pp. 124–138.

Bin, Okmyung et al. (2011). "Measuring the Impact of Sea-level Rise on Coastal Real Estate: A Hedonic Property Model Approach". In: *Journal of Regional Science* 51.4, pp. 751–767.

Bleakley, Hoyt and Jeffrey Lin (2012). "Portage and Path Dependence". In: *The Quarterly Journal of Economics*. DOI: 10.1093/qje/qjs011. URL: http://qje.oxfordjournals.org/content/early/2012/04/19/qje.qjs011.abstract.

Bosker, Maarten and Eltjo Buringh (2015). "City seeds: Geography and the origins of the European city system". Unpublished Work. Journal of Urban Economics.

Burr, Irving W. (1942). "Cumulative Frequency Functions". In: *The Annals of Mathematical Statistics* 13.2, pp. 215–232.

Duranton, Gilles and Matthew A. Turner (2012). "Urban growth and transportation". In: *Review of Economic Studies* 79 (4), pp. 1407–1440.

Guerrieri, Veronica, Daniel Hartley, and Erik Hurst (2013). "Endogenous gentrification and housing price dynamics". In: *Journal of Public Economics* 100, pp. 45–60.

Kelsay, James D. and Cory P. Haberman (2020a). "The Influence of Street Network Features on Robberies Around Public Housing Communities". In: *Crime and Delinquency*. ISSN: 0011128720928915.

— (2020b). "The Topography of Robbery: Does Slope Matter". In: *Quantitative Criminology* 21 (1).

Kok, Nils, Paavo Monkkonen, and John M. Quigley (2014). "Land use regulations and the value of land and housing: An intra-metropolitan analysis". In: *Journal of Urban Economics* 81.1, pp. 136–148.

Lee, Sanghoon and Jeffrey Lin (2018). “Natural Amenities, Neighbourhood Dynamics, and Persistence in the Spatial Distribution of Income”. In: *The Review of Economic Studies* 85.1, pp. 663–694.

Lin, Jeffrey (2015). “The Puzzling Persistence of Space”. In: *Federal Reserve Bank of Philadelphia Business Review* (Q2), pp. 1–8.

Lucas, Robert E. and Esteban Rossi–Hansberg (2002). “On the internal structure of cities”. In: *Econometrica* 70 (4), pp. 1445–1476.

Pakes, Ariel (1986). “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks”. In: *Econometrica* 54.4, pp. 755–784. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1912835.

Pierce, Joseph and Crystal A. Kolden (2015). “The Hilliness of US Cities”. In: *Geographical Review* 105 (4), pp. 581–600.

Rappaport, Jordan and Jeffery D. Sachs (2003). “The United States as a Coastal Nation”. In: *Journal of Economic Growth* 8, pp. 5–46.

Rose, Louis A. (1989). “Topographical Constraints and Urban Land Supply Indexes”. In: *Journal of Urban Economics* 26, pp. 335–347.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Raymond Owens III (2010). “Housing externalities”. In: *Journal of Political Economy* 118 (3), pp. 485–535.

Saiz, Albert (2010). “The Geographic Determinants of Housing Supply”. In: *Quarterly Journal of Economics* 125.3, pp. 1253–1296.

Scawthorn, Charles, Hirokazu Iemura, and Yoshikazu Yamada (1982). “The influence of natural hazards on urban housing location”. In: *Journal of Urban Economics* 11.2, pp. 242–251.

Shilling, James D., C. F. Sirmans, and John D. Benjamin (1989). “Flood insurance, wealth redistribution, and urban property values”. In: *Journal of Urban Economics* 26.1, pp. 43–53.

Tibshirani, Robert (1996). “Regression Shinkage and Selection via the Lasso”. In: *Journal of the Royal Statistical Society, Series B (Methodological)* 58.1, pp. 267–288.

Ye, Yifan and Charles M. Becker (2017). “The (Literally) Steepest Slope: Spatial, Temporal, and Elevation Variance Gradients in Urban Spatial Modelling”. In: *Journal of Economic Geography*.

— (2018). “The Z-axis: Elevation Gradient Effects in Urban America”. In: *Regional Science and Urban Economics*.

Appendix

Figure A1: Boxplot of log tract average altitude (meters), all dataset MSAs

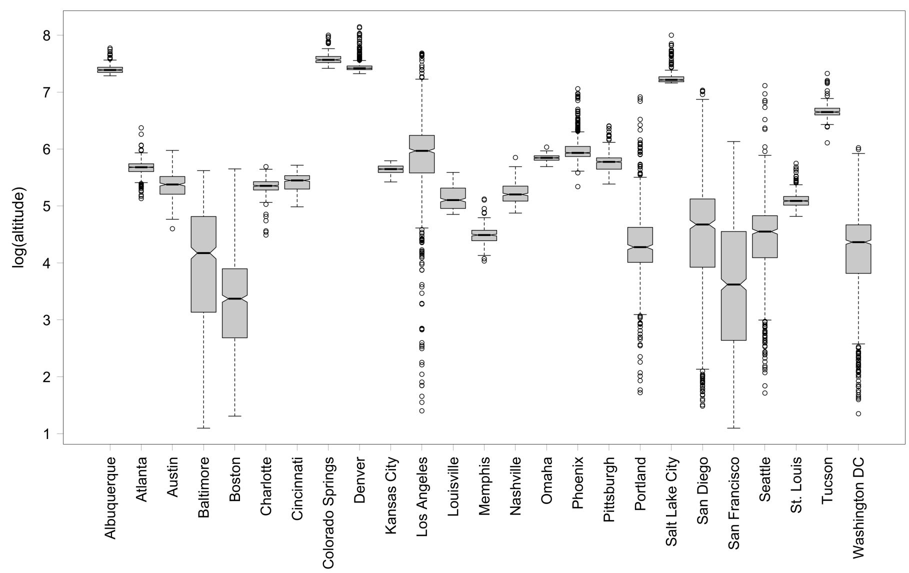

Figure A2: Elevation contour and elevation variance spatial distribution, Washington DC

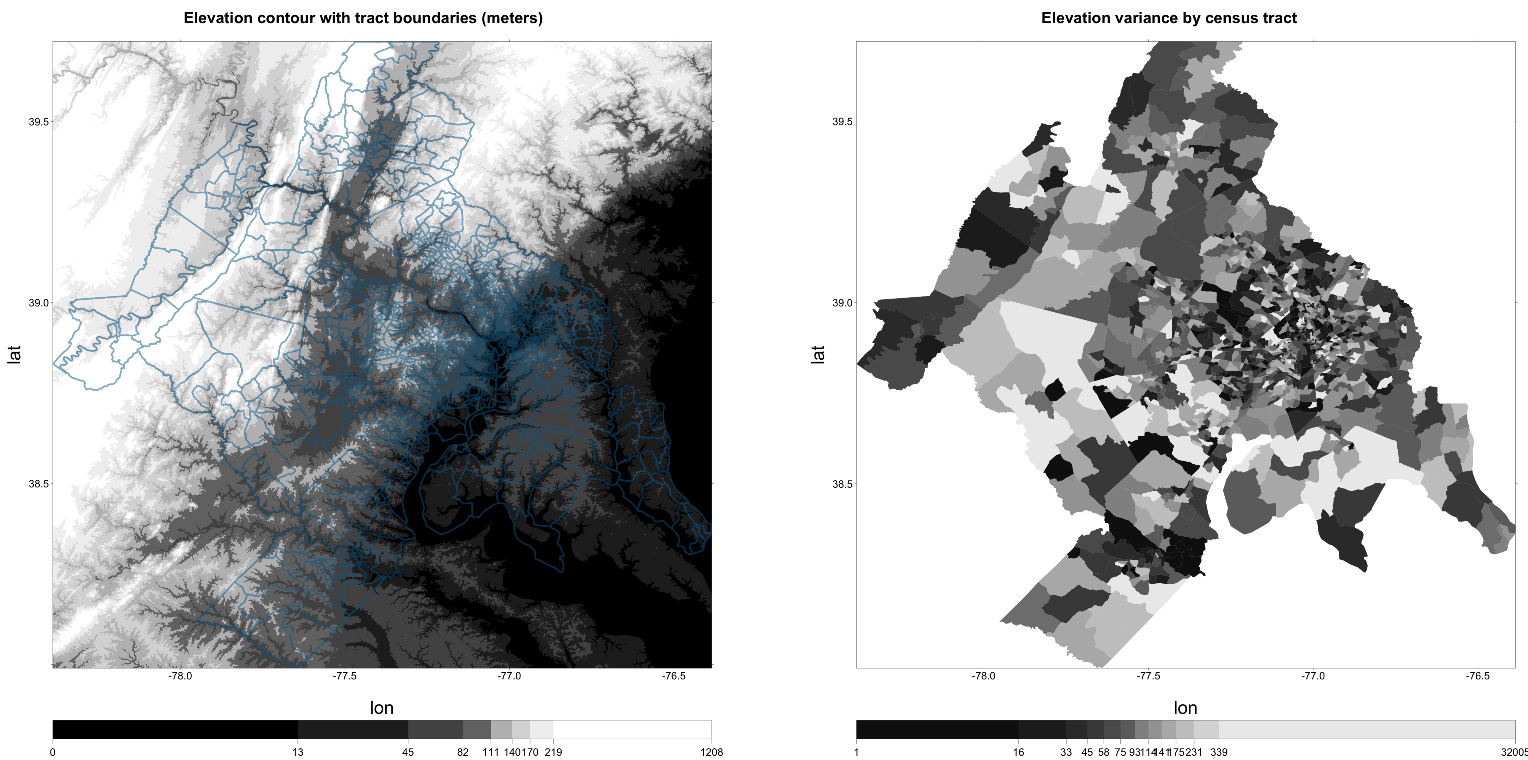

Figure A3: Elevation contour and elevation variance spatial distribution, San Francisco

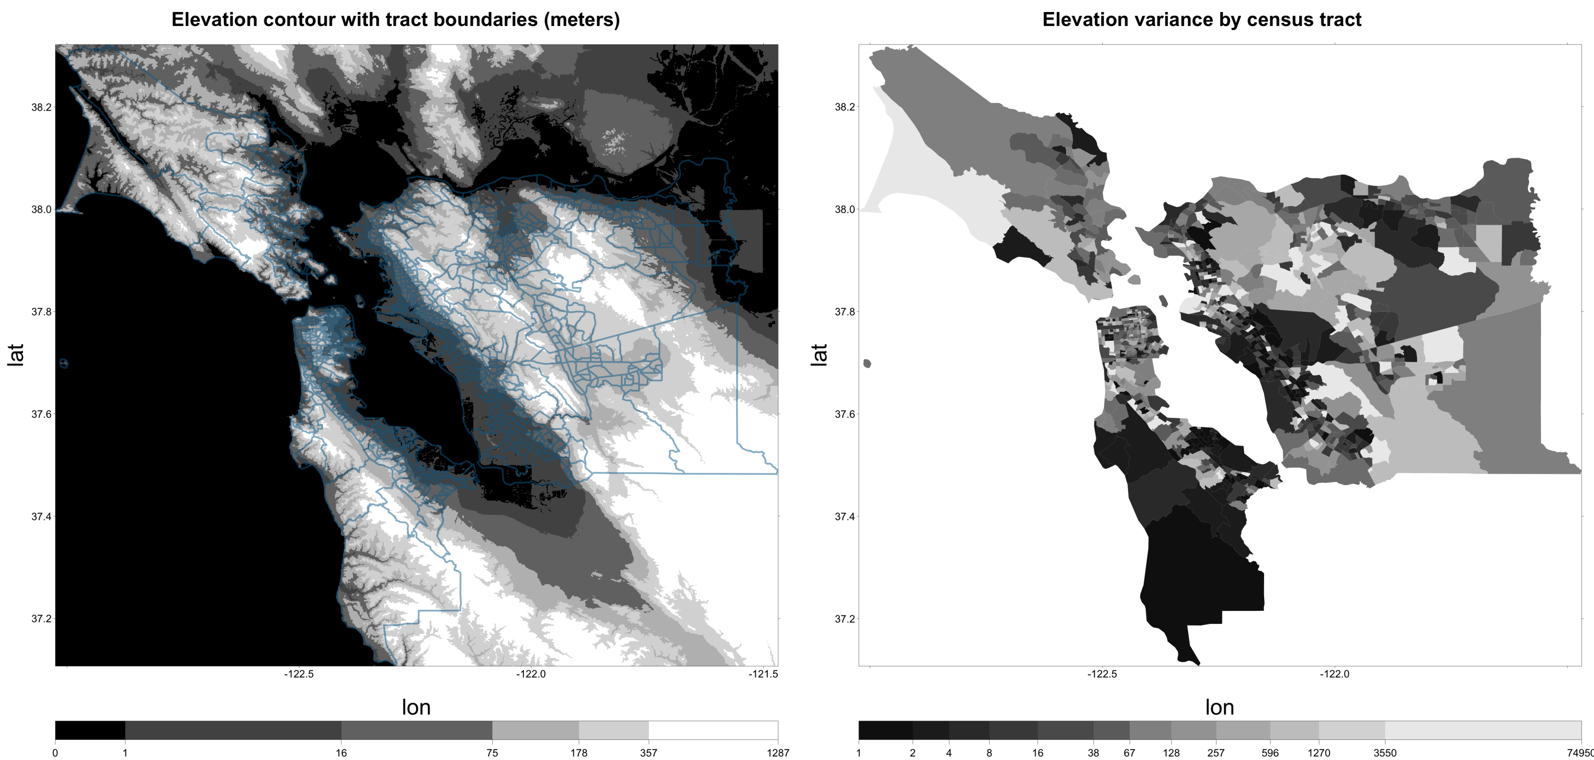

Table A1: Price sensitivity by MSA, baseline model

Decile	1	2	3	4	5	6	7	8	9	10
Albuquerque	-0.263	-0.255	-0.248	-0.254	-0.236	-0.234	-0.222	-0.198	-0.182	-0.138
Atlanta	-0.444	-0.413	-0.436	-0.405	-0.386	-0.364	-0.356	-0.308	-0.288	-0.242
Austin	-0.350	-0.368	-0.345	-0.338	-0.314	-0.322	-0.314	-0.261	-0.217	-0.157
Baltimore	-0.449	-0.420	-0.411	-0.388	-0.360	-0.348	-0.315	-0.294	-0.252	-0.221
Boston	-0.524	-0.491	-0.506	-0.541	-0.486	-0.470	-0.498	-0.459	-0.425	-0.335
Charlotte	-0.394	-0.378	-0.356	-0.338	-0.329	-0.293	-0.280	-0.264	-0.210	-0.150
Cincinnati	-0.554	-0.499	-0.465	-0.424	-0.393	-0.360	-0.352	-0.327	-0.268	-0.210
CO Springs	-0.498	-0.469	-0.470	-0.423	-0.359	-0.354	-0.340	-0.241	-0.195	-0.160
Denver	-0.555	-0.570	-0.574	-0.565	-0.539	-0.576	-0.533	-0.498	-0.404	-0.327
Kansas City	-0.431	-0.418	-0.381	-0.378	-0.346	-0.314	-0.316	-0.263	-0.249	-0.205
Los Angeles	-0.445	-0.437	-0.436	-0.405	-0.373	-0.338	-0.306	-0.256	-0.228	-0.207
Louisville	-0.276	-0.281	-0.270	-0.257	-0.242	-0.230	-0.219	-0.196	-0.153	-0.106
Memphis	-0.714	-0.516	-0.491	-0.366	-0.315	-0.281	-0.244	-0.159	-0.125	-0.081
Nashville	-0.342	-0.349	-0.364	-0.317	-0.323	-0.292	-0.286	-0.235	-0.220	-0.167
Omaha	-0.342	-0.333	-0.317	-0.292	-0.271	-0.247	-0.236	-0.203	-0.181	-0.172
Phoenix	-0.527	-0.484	-0.468	-0.444	-0.409	-0.381	-0.387	-0.315	-0.305	-0.239
Pittsburgh	-0.484	-0.429	-0.415	-0.387	-0.333	-0.319	-0.321	-0.271	-0.228	-0.186
Portland	-0.329	-0.351	-0.348	-0.358	-0.348	-0.330	-0.320	-0.291	-0.298	-0.229
Salt Lake City	-0.323	-0.338	-0.337	-0.320	-0.326	-0.282	-0.288	-0.269	-0.232	-0.192
San Diego	-0.360	-0.408	-0.406	-0.379	-0.383	-0.368	-0.353	-0.299	-0.291	-0.229
San Francisco	-0.376	-0.372	-0.371	-0.374	-0.342	-0.340	-0.344	-0.320	-0.285	-0.246
Seattle	-0.321	-0.340	-0.338	-0.341	-0.344	-0.319	-0.317	-0.301	-0.280	-0.226
St. Louis	-0.362	-0.330	-0.312	-0.277	-0.271	-0.258	-0.248	-0.232	-0.196	-0.133
Tucson	-0.260	-0.266	-0.241	-0.234	-0.224	-0.197	-0.186	-0.168	-0.149	-0.115
Washington DC	-0.404	-0.422	-0.425	-0.406	-0.403	-0.393	-0.365	-0.368	-0.326	-0.261

Table A2: Preference for tract median income by MSA, baseline model

Decile	1	2	3	4	5	6	7	8	9	10
Albuquerque	0.046	0.085	0.160	0.171	0.213	0.166	0.190	0.254	0.243	0.335
Atlanta	0.165	0.211	0.326	0.360	0.345	0.355	0.378	0.374	0.416	0.542
Austin	0.063	0.144	0.154	0.236	0.290	0.304	0.326	0.268	0.232	0.269
Baltimore	0.191	0.212	0.271	0.314	0.285	0.311	0.279	0.271	0.201	0.200
Boston	0.278	0.281	0.319	0.447	0.452	0.462	0.562	0.536	0.557	0.586
Charlotte	0.144	0.178	0.203	0.248	0.313	0.254	0.233	0.193	0.220	0.170
Cincinnati	0.273	0.382	0.404	0.400	0.361	0.341	0.365	0.335	0.302	0.221
CO Springs	0.317	0.285	0.335	0.372	0.299	0.343	0.286	0.147	0.003	0.000
Denver	0.454	0.414	0.403	0.537	0.461	0.595	0.548	0.516	0.443	0.681
Kansas City	0.181	0.259	0.318	0.357	0.319	0.243	0.260	0.222	0.228	0.180
Los Angeles	0.267	0.257	0.388	0.413	0.408	0.385	0.347	0.259	0.243	0.358
Louisville	0.000	0.042	0.181	0.195	0.229	0.219	0.193	0.189	0.167	0.046
Memphis	0.428	0.318	0.444	0.344	0.264	0.195	0.205	0.090	0.063	0.227
Nashville	0.000	0.028	0.299	0.274	0.322	0.283	0.306	0.226	0.260	0.206
Omaha	0.000	0.183	0.131	0.142	0.224	0.227	0.187	0.120	0.156	0.209
Phoenix	0.340	0.280	0.336	0.371	0.396	0.389	0.407	0.391	0.605	0.725
Pittsburgh	0.116	0.205	0.343	0.328	0.285	0.307	0.303	0.308	0.219	0.251
Portland	0.117	0.114	0.158	0.171	0.263	0.323	0.325	0.240	0.229	0.249
Salt Lake City	0.189	0.073	0.115	0.247	0.315	0.158	0.274	0.263	0.172	0.074
San Diego	0.129	0.154	0.240	0.262	0.320	0.406	0.421	0.352	0.508	0.648
San Francisco	0.109	0.113	0.136	0.212	0.164	0.234	0.321	0.330	0.308	0.339
Seattle	0.000	0.000	0.022	0.163	0.231	0.352	0.349	0.319	0.352	0.458
St. Louis	0.019	0.082	0.157	0.173	0.197	0.234	0.214	0.170	0.163	0.247
Tucson	0.000	0.049	0.135	0.182	0.183	0.145	0.173	0.120	0.267	0.592
Washington DC	0.163	0.134	0.159	0.200	0.240	0.321	0.321	0.383	0.410	0.483

Table A3: Price sensitivity by MSA, endogenous housing stock

Decile	1	2	3	4	5	6	7	8	9	10
Albuquerque	-0.399	-0.398	-0.353	-0.348	-0.321	-0.319	-0.305	-0.280	-0.251	-0.217
Atlanta	-0.491	-0.450	-0.445	-0.423	-0.423	-0.397	-0.371	-0.327	-0.308	-0.254
Austin	-0.336	-0.350	-0.349	-0.350	-0.316	-0.323	-0.309	-0.273	-0.230	-0.162
Baltimore	-0.451	-0.427	-0.388	-0.379	-0.348	-0.338	-0.322	-0.291	-0.274	-0.221
Boston	-0.656	-0.633	-0.633	-0.630	-0.611	-0.605	-0.589	-0.560	-0.550	-0.447
Charlotte	-0.409	-0.372	-0.334	-0.328	-0.296	-0.277	-0.281	-0.240	-0.200	-0.148
Cincinnati	-0.623	-0.621	-0.550	-0.475	-0.397	-0.444	-0.369	-0.347	-0.293	-0.246
CO Springs	-0.528	-0.493	-0.484	-0.472	-0.341	-0.343	-0.329	-0.239	-0.183	-0.141
Denver	-0.438	-0.447	-0.430	-0.442	-0.418	-0.409	-0.405	-0.409	-0.329	-0.265
Kansas City	-0.512	-0.469	-0.445	-0.418	-0.397	-0.375	-0.353	-0.310	-0.281	-0.250
Los Angeles	-0.423	-0.357	-0.394	-0.371	-0.334	-0.309	-0.282	-0.237	-0.227	-0.220
Louisville	-0.280	-0.285	-0.269	-0.256	-0.237	-0.240	-0.227	-0.190	-0.153	-0.105
Memphis	-0.632	-0.500	-0.462	-0.351	-0.306	-0.279	-0.245	-0.167	-0.126	-0.082
Nashville	-0.374	-0.353	-0.369	-0.336	-0.309	-0.290	-0.275	-0.237	-0.211	-0.162
Omaha	-0.340	-0.323	-0.296	-0.279	-0.246	-0.235	-0.227	-0.186	-0.183	-0.181
Phoenix	-0.377	-0.329	-0.336	-0.305	-0.288	-0.279	-0.270	-0.224	-0.212	-0.169
Pittsburgh	-0.529	-0.487	-0.461	-0.436	-0.363	-0.370	-0.372	-0.294	-0.265	-0.221
Portland	-0.408	-0.449	-0.443	-0.444	-0.423	-0.421	-0.423	-0.386	-0.352	-0.307
Salt Lake City	-0.334	-0.339	-0.348	-0.338	-0.335	-0.300	-0.292	-0.249	-0.239	-0.192
San Diego	-0.335	-0.357	-0.355	-0.359	-0.334	-0.333	-0.305	-0.284	-0.262	-0.205
San Francisco	-0.336	-0.343	-0.347	-0.352	-0.335	-0.320	-0.320	-0.303	-0.289	-0.230
Seattle	-0.398	-0.435	-0.426	-0.423	-0.423	-0.398	-0.389	-0.371	-0.348	-0.270
St. Louis	-0.348	-0.335	-0.300	-0.273	-0.246	-0.250	-0.233	-0.238	-0.201	-0.143
Tucson	-0.390	-0.378	-0.358	-0.341	-0.309	-0.283	-0.276	-0.263	-0.232	-0.171
Washington DC	-0.383	-0.393	-0.409	-0.403	-0.405	-0.385	-0.369	-0.361	-0.330	-0.260

Table A4: Preference for tract median income by MSA, endogenous housing stock

Decile	1	2	3	4	5	6	7	8	9	10
Albuquerque	0.312	0.375	0.354	0.361	0.339	0.327	0.316	0.279	0.226	0.167
Atlanta	0.051	0.056	0.224	0.284	0.377	0.409	0.370	0.388	0.556	0.806
Austin	0.001	0.071	0.157	0.287	0.297	0.350	0.331	0.307	0.290	0.362
Baltimore	0.321	0.224	0.283	0.323	0.277	0.329	0.281	0.195	0.173	0.162
Boston	0.551	0.501	0.490	0.593	0.635	0.611	0.662	0.639	0.680	0.721
Charlotte	0.074	0.060	0.129	0.249	0.302	0.284	0.273	0.233	0.331	0.433
Cincinnati	0.149	0.380	0.437	0.437	0.405	0.450	0.362	0.322	0.360	0.456
CO Springs	0.581	0.261	0.327	0.465	0.275	0.366	0.341	0.248	0.079	0.000
Denver	0.250	0.156	0.208	0.353	0.303	0.421	0.430	0.502	0.451	0.673
Kansas City	0.338	0.337	0.363	0.401	0.411	0.299	0.365	0.263	0.431	0.540
Los Angeles	0.043	0.000	0.226	0.301	0.346	0.341	0.293	0.222	0.323	0.535
Louisville	0.000	0.084	0.182	0.259	0.122	0.220	0.208	0.157	0.123	0.148
Memphis	0.316	0.279	0.451	0.356	0.219	0.265	0.218	0.169	0.231	0.394
Nashville	0.000	0.000	0.266	0.199	0.303	0.298	0.280	0.231	0.228	0.396
Omaha	-0.390	-0.087	0.000	0.000	0.145	0.235	0.300	0.492	0.813	1.296
Phoenix	0.036	0.000	0.109	0.167	0.253	0.270	0.316	0.367	0.615	0.797
Pittsburgh	0.355	0.363	0.452	0.441	0.291	0.341	0.364	0.275	0.342	0.437
Portland	0.070	0.169	0.266	0.261	0.357	0.441	0.470	0.405	0.369	0.436
Salt Lake City	0.115	0.000	0.154	0.256	0.334	0.142	0.279	0.220	0.189	0.291
San Diego	0.000	0.034	0.091	0.189	0.258	0.366	0.321	0.251	0.364	0.483
San Francisco	0.001	0.026	0.093	0.173	0.153	0.219	0.309	0.322	0.313	0.390
Seattle	-0.052	0.000	0.000	0.058	0.201	0.363	0.420	0.457	0.709	0.924
St. Louis	0.000	0.000	0.021	0.079	0.156	0.210	0.216	0.348	0.487	0.754
Tucson	0.000	0.000	0.178	0.302	0.324	0.268	0.282	0.491	0.657	0.826
Washington DC	0.018	0.000	0.016	0.118	0.127	0.257	0.294	0.390	0.362	0.520

Figure A4: Relative altitude and elevation variance gradient by income decile, doubling elevation variance and relative altitude

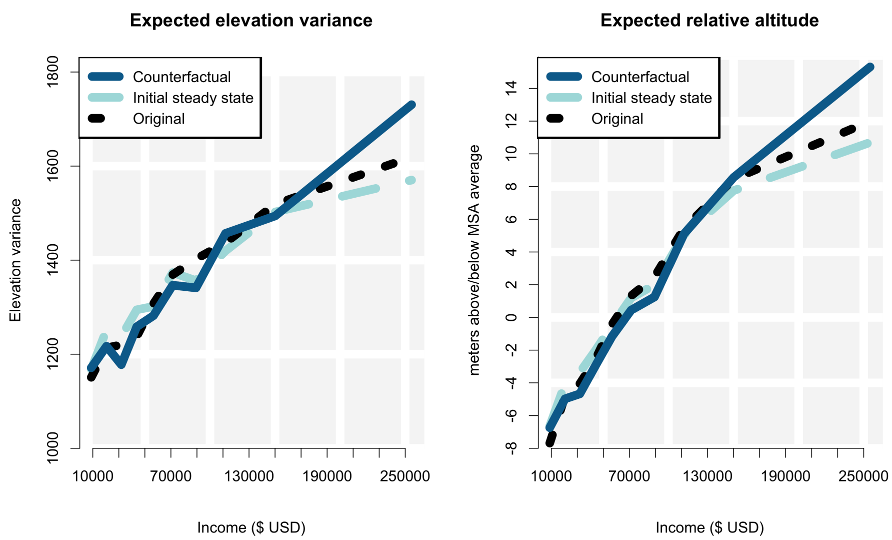

Figure A5: Counterfactual outcomes from flattening cities by income decile, endogenous housing stock

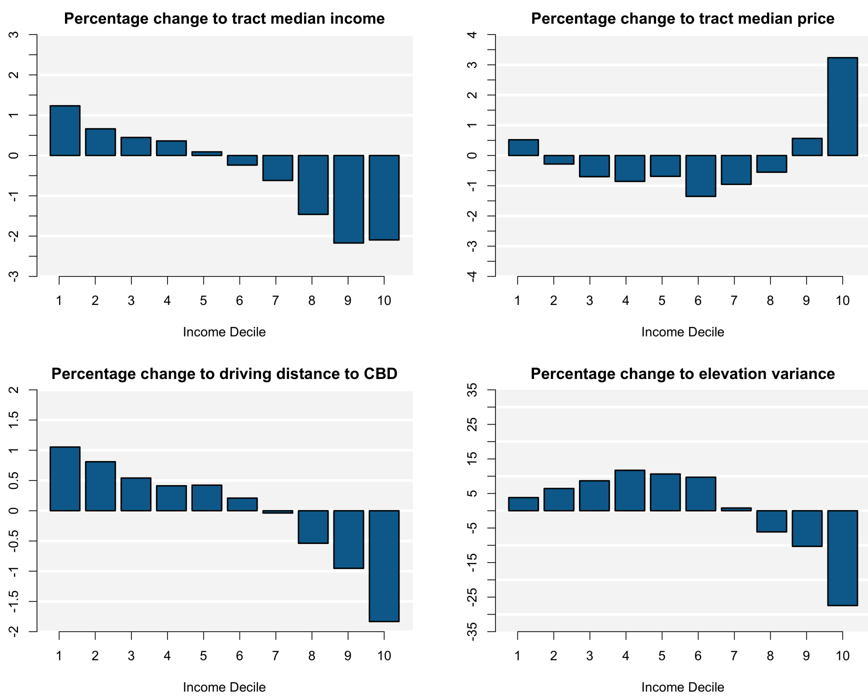

Figure A6: Counterfactual outcomes from doubling elevation variance and relative altitude by income decile, endogenous housing stock.

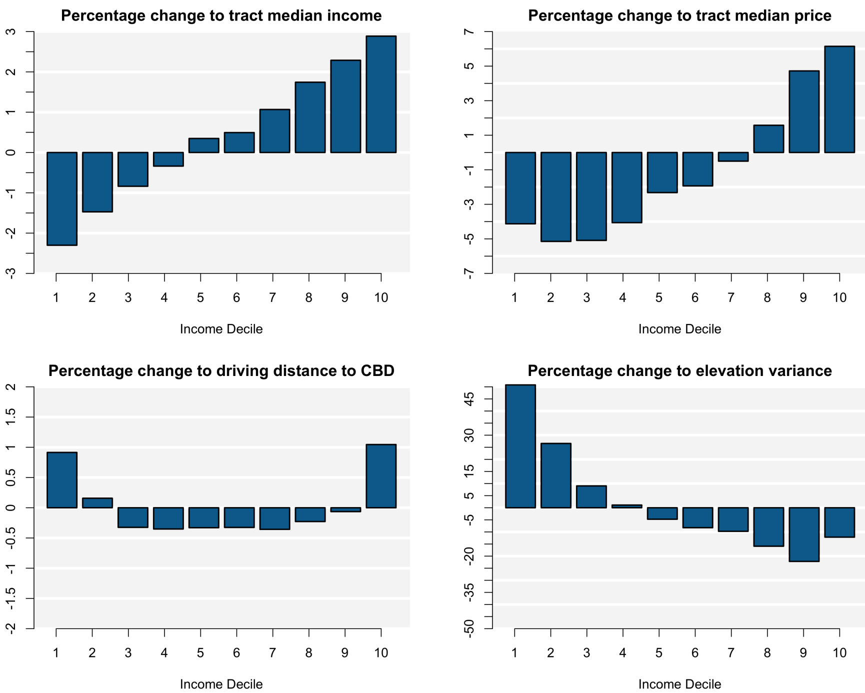

Figure A7: Relative altitude and elevation variance gradient by income decile, flattening cities with endogenous housing stock.

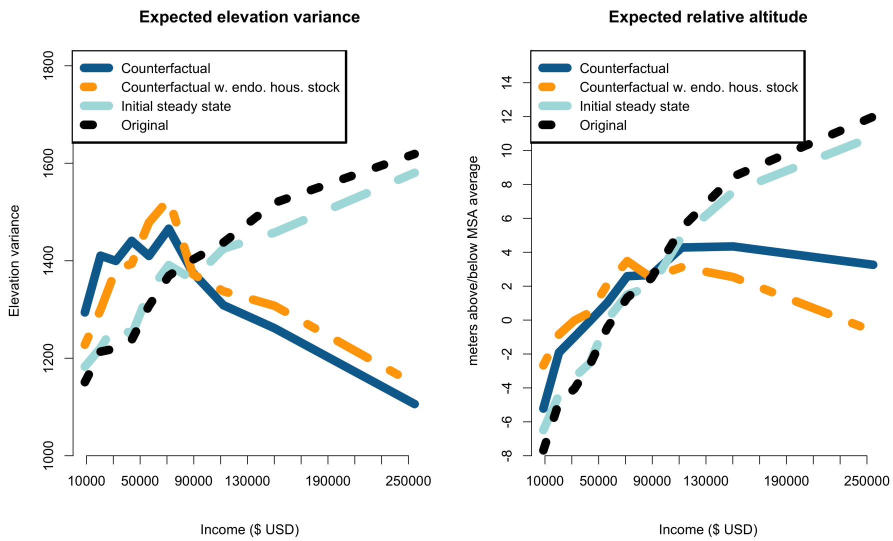

Figure A8: Relative altitude and elevation variance gradient by income decile, doubling elevation variance and relative altitude with endogenous housing stock.

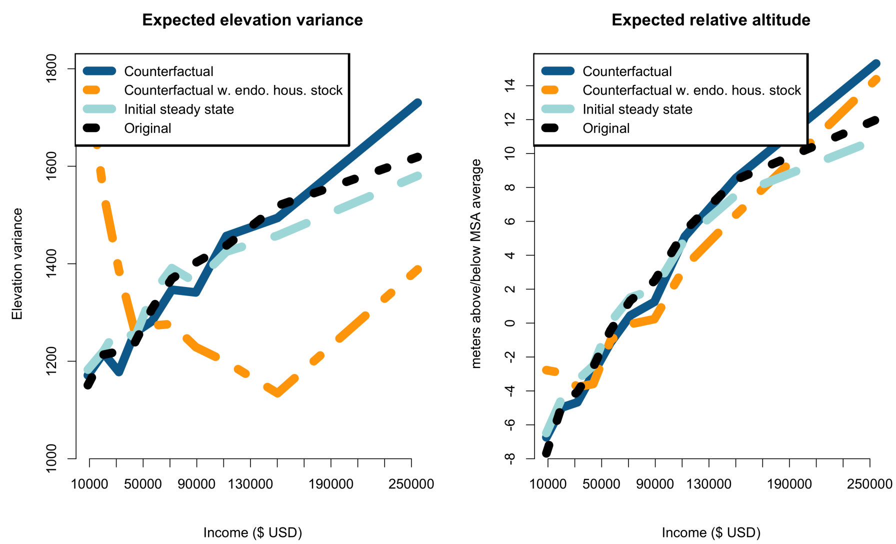

Table A5: Average percentage difference in cross-decile exposure between counterfactual and baseline, doubling elevation variance and relative altitude with endogenous housing stock

Decile	1	2	3	4	5	6	7	8	9	10
1	-2.77	-	-	-	-	-	-	-	-	-
2	-2.39	-1.35	-	-	-	-	-	-	-	-
3	-1.47	-1.25	-0.81	-	-	-	-	-	-	-
4	-0.90	-0.92	-0.80	-0.50	-	-	-	-	-	-
5	0.07	-0.23	-0.39	-0.32	-0.14	-	-	-	-	-
6	0.68	0.23	-0.02	-0.10	-0.16	-0.14	-	-	-	-
7	1.47	0.85	0.46	0.21	-0.18	-0.30	-0.34	-	-	-
8	2.76	1.97	1.35	0.81	0.03	-0.35	-0.68	-0.76	-	-
9	4.58	3.28	2.42	1.57	0.34	-0.37	-1.05	-1.99	-1.98	-
10	6.57	5.00	3.73	2.72	1.11	0.02	-0.87	-2.68	-4.08	-4.46

Table A6: Average percentage difference in cross-decile exposure between counterfactual and baseline, flattening cities with endogenous housing stock

Decile	1	2	3	4	5	6	7	8	9	10
1	4.22	-	-	-	-	-	-	-	-	-
2	3.05	2.37	-	-	-	-	-	-	-	-
3	1.64	1.65	1.41	-	-	-	-	-	-	-
4	0.7	0.94	1.06	0.89	-	-	-	-	-	-
5	-0.31	0.04	0.31	0.44	0.39	-	-	-	-	-
6	-0.93	-0.44	0	0.26	0.42	0.49	-	-	-	-
7	-1.83	-1.11	-0.56	-0.09	0.34	0.55	0.71	-	-	-
8	-3.76	-2.64	-1.71	-0.9	0.00	0.47	0.95	2.01	-	-
9	-5.68	-4.22	-2.97	-1.78	-0.46	0.09	0.92	2.51	3.46	-
10	-7.31	-5.85	-4.49	-3.09	-1.36	-0.75	0.34	2.38	4.22	5.65

Figure A9: Counterfactual outcomes from permuting tract-level amenities by income decile

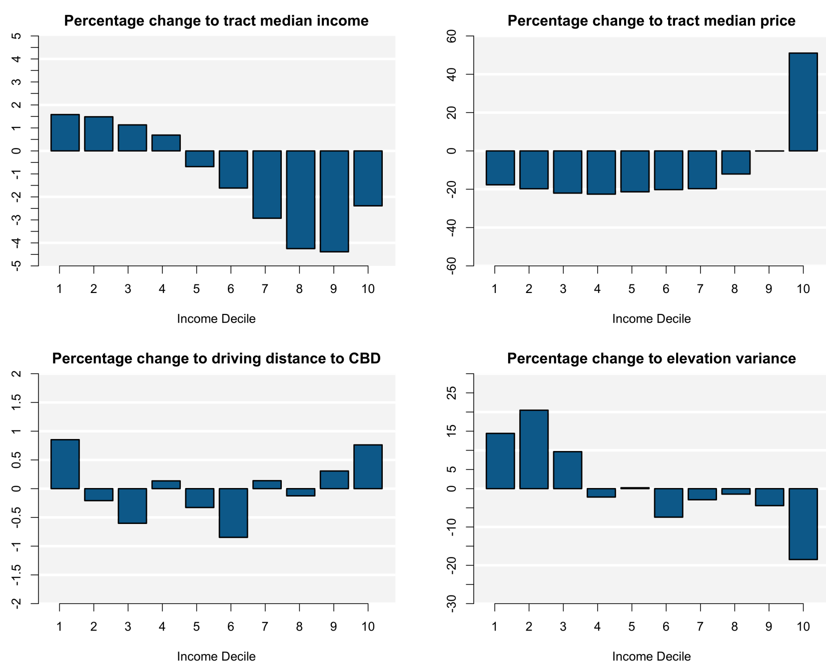

Table A7: Average percent difference in cross-decile exposure between counterfactual and baseline, permuted tract amenities

Decile	1	2	3	4	5	6	7	8	9	10
1	6.90	-	-	-	-	-	-	-	-	-
2	-10.05	8.01	-	-	-	-	-	-	-	-
3	-6.56	-5.61	4.18	-	-	-	-	-	-	-
4	-4.70	-3.90	-2.58	2.46	-	-	-	-	-	-
5	-2.02	-2.08	-1.08	-0.63	1.72	-	-	-	-	-
6	0.54	-0.26	-0.08	-0.09	-0.28	1.41	-	-	-	-
7	2.42	1.79	1.65	1.05	0.34	-0.50	1.71	-	-	-
8	5.33	4.18	3.38	2.22	0.50	-0.72	-2.36	3.82	-	-
9	8.70	6.83	4.90	3.36	1.14	-0.90	-3.26	-7.02	7.61	-
10	9.42	9.47	6.73	5.00	1.45	-1.29	-4.96	-9.53	-15.61	8.63