# WORKING PAPERS

RESEARCH DEPARTMENT

**WORKING PAPER NO. 13-48**

# NATURAL AMENITIES, NEIGHBORHOOD DYNAMICS, AND PERSISTENCE IN THE SPATIAL DISTRIBUTION OF INCOME


Sanghoon Lee
University of British Columbia

Jeffrey Lin
Federal Reserve Bank of Philadelphia


December 6, 2013


RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/research-and-data/

# Natural Amenities, Neighborhood Dynamics, and Persistence in the Spatial Distribution of Income*

Sanghoon Lee† Jeffrey Lin‡

December 6, 2013

**Abstract**

We present theory and evidence highlighting the role of natural amenities in neighborhood dynamics, suburbanization, and variation across cities in the persistence of the spatial distribution of income. Our model generates three predictions that we confirm using a novel database of consistent-boundary neighborhoods in U.S. metropolitan areas, 1880–2010, and spatial data for natural features such as coastlines and hills. First, persistent natural amenities anchor neighborhoods to high incomes over time. Second, downtown neighborhoods in coastal cities were less susceptible to the suburbanization of income in the mid-20th century. Third, naturally heterogeneous cities exhibit spatial distributions of income that are dynamically persistent.

Keywords: Neighborhood change, suburbanization, locational fundamentals, multiple equilibria
JEL classification: R23, N90, O18

*The views expressed here are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. We thank Nick Reynolds for excellent research assistance, Dan Hartley and Stuart Rosenthal for generously sharing data, Hoyt Bleakley, Jeff Brinkman, Jan Brueckner, Satyajit Chatterjee, Tom Holmes, Stuart Rosenthal, Ralph Winter, and participants at seminars and conferences at AMES, AREUEA, FRBP, IEB, NARSC/UEA, UBC, and WEAI for helpful comments. This paper is available free of charge at http://www.philadelphiafed.org/research-and-data/publications/working-papers/.

†Assistant Professor, Sauder School of Business, University of British Columbia. Post: 2053 Main Mall, Vancouver, BC, Canada, V6T 1Z2. Electronic mail: sanghoon.lee@sauder.ubc.ca.

‡Senior Economist, Research Department, Federal Reserve Bank of Philadelphia. Post: Ten Independence Mall, Philadelphia, PA, 19106, USA. Electronic mail: jeff.lin@phil.frb.org.

# 1 Introduction

Neighborhood change is common and contentious. Of neighborhoods in 35 U.S. metropolitan areas studied by Rosenthal (2008), more than two-thirds transitioned to a different income quartile between 1950 and 2000. In declining areas, homeowners fear deteriorating values even as entrants enjoy new opportunities; in gentrifying areas, rising prices cause anxiety for longtime renters. And in response to shifting neighborhood needs, policymakers often act to preserve neighborhood quality or quicken the pace of change.

Though changes in neighborhood quality are widespread, it is less well known that neighborhood change varies across cities. While in some cities neighborhoods seem immune from change—leading to overall persistence in the internal structure of the city—other cities feature quickly-changing neighborhoods and spatial patterns of income. For example, Los Angeles has long featured a stable arrangement of high incomes and prices along its beaches and in its foothills; between 1970 and 1980, the average neighborhood in the Los Angeles metropolitan area moved just 9 percentile points across the city's income distribution. In contrast, over the same period, the average neighborhood in the Dallas metropolitan area moved 21 percentile points.

In this paper, we examine why the geographic distribution of income is persistent for some neighborhoods and cities but turns over frequently elsewhere. Our explanation highlights the role of natural geographic features that have persistent amenity value—for example, oceans, mountains, and lakes. We begin with the idea that persistent natural amenities can "anchor" neighborhoods to high incomes, even as they experience various shocks over time. A key implication is that for cities as a whole, the presence of an especially valuable natural amenity can hold back neighborhood tipping and suburbanization. Thus, in naturally heterogeneous Los Angeles, the spatial distribution of income is persistent, but in flat Dallas, the spatial distribution of income churns quickly.[1]

We present a dynamic model of household neighborhood choice to formalize our thinking. Neighborhoods derive amenity value from both natural features and endogenous characteristics such as safety, school quality, or shopping. High-income households outbid low-income households for neighborhoods with greater overall amenity value. Neighborhoods are also subject to idiosyncratic shocks to amenity

[1] In Appendix Figure A1, maps illustrate differences in neighborhood change in Dallas versus Los Angeles, 1970–1980.

value over time. We characterize conditions when these shocks can potentially reverse the historical spatial pattern of income.

We test and confirm several implications of our theory using a new database of consistent-boundary neighborhoods in many U.S. metropolitan areas, spanning census years from 1880 to 2010. We match these data to spatial information on the location of many persistent natural features, including shorelines, mountains, lakes, rivers, temperate climates, and floodplains.

Our first main result is that persistent natural amenities anchor neighborhoods to high incomes over time. For example, conditional on current income, neighborhoods with superior natural amenities are more likely to remain high-income neighborhoods. Intuitively, this result confirms a simplified version of the folk wisdom among realtors that a beachfront home will better retain its value versus a similar home with a mundane view.

A second main result is that cities with downtowns adjacent to strong natural amenities, such as a coastline, experienced a less-severe suburbanization of income in the early and mid-20th century and a more pronounced gentrification of inner-city areas in the late 20th century. To our knowledge, our paper is the first in the extensive literature on suburbanization to document these differences across cities in suburbanization and gentrification patterns and relate them to differences in natural heterogeneity.

Our third main result is that cities with dominant natural features (e.g., a coastline or mountain range) exhibit internal spatial distributions of income that are dynamically stable. In other words, neighborhood incomes tend to fluctuate less over time in a city like Los Angeles, with its beaches, hills, and valleys, than in a city like Dallas, which more closely resembles a flat, featureless plain. Intuitively, a shock to a neighborhood's amenity value—due to idiosyncratic migration, effects of policy, depreciation of fixed assets, natural disasters, etc.—has the potential to reverse the historical distribution of income across neighborhoods. But in cities where some neighborhoods have overwhelming natural advantages, small shocks or interventions are unlikely to undo history.

Thus, our results relate to a central debate in economics about the roles of natural fundamentals versus endogenous amenities in the spatial distribution of income. In many models featuring endogenous amenities (e.g., benefits from agglomeration economies), multiple equilibria are possible.[2] Intuitively, if households and firms care

[2]For example, see Krugman (1991), Rauch (1993), and Arthur (1994).

only about being near other households and firms, then they might crowd together in any location. We characterize conditions when multiple equilibria exist—namely, when spatial variation in natural amenity value is muted. Our results are therefore consistent with the view that when there is sizable natural variation, fundamentals uniquely determine the spatial distribution of income.

Our model also delivers several implications that are consistent with well-known patterns. First, neighborhoods that have superior natural amenity value are more likely to be, but are not necessarily, high-income neighborhoods. Second, neighborhood quality is usually persistent over time. Third, there is mean reversion in neighborhood quality.

We address several identification issues in evaluating our evidence. One important empirical challenge is that we do not directly observe the value of natural features. For example, a natural feature can be either an amenity or a disamenity: A river used for industrial purposes can detract from surrounding neighborhoods. First, we focus on high-value natural amenities such as proximity to the ocean, where we believe benefits are obvious and large relative to the value of other amenities. A second strategy, developed from our model, is to focus on natural features that were surrounded by high-income neighborhoods in an initial year. Intuitively, if a stretch of river was surrounded by high incomes, we can be more confident that the river is a positive versus negative amenity. Third, we develop a hedonic weighting method to aggregate the amenity value from many natural features at once.

A related challenge is that the value of natural features changes over time. An alternative view of our results is that neighborhoods with superior natural amenities tend to increase in income because of an increasing taste for natural amenities over time. However, this view is inconsistent with our third main result, as it predicts that these neighborhoods experience greater changes over time in neighborhood income. In addition, our main analysis focuses mostly on 10-year changes in the location of income: While today's tastes and technologies are different from the 19th century's, it is less likely that tastes for natural amenities have evolved significantly over a single decade. Finally, by examining the relative likelihood that high-income neighborhoods are proximate to natural amenities, we find no evidence that tastes of high-income versus low-income households for natural features have fluctuated significantly over time.

Neighborhoods in growing cities tend to experience greater fluctuations in income. Our view, consistent with theories featuring the "filtering" of aging structures (e.g., Brueckner and Rosenthal, 2009), is that the addition of new houses to a city acts as a negative shock to the relative value of existing homes and neighborhoods. An alternative view is that in flat cities, the supply of housing is more elastic, and therefore the causal link between geography and neighborhood stability is mediated by city growth, not the value of natural amenities (see Saiz, 2010). In our empirical work, we adjust for changes in city size by examining only changes in *relative* rankings within a fixed group of neighborhoods in each ten-year period. In addition, we show that our theoretical and empirical results are robust to controlling for city growth and the age of housing.

## 1.1 Related Work

To our knowledge, our paper is the first in an extensive literature on neighborhood change to document and relate variation *across* cities in neighborhood dynamics to differences in natural heterogeneity. Of course, past work has recognized that cities are not flat, featureless plains. However, geographic heterogeneity has typically been treated as a background or control variable as researchers have emphasized other factors in neighborhood change, such as aging homes filtering from high- to low-income households (e.g., Brueckner and Rosenthal, 2009); spillovers among neighborhoods (Aaronson, 2001; Guerreri, Hartley, and Hurst, 2010); changes in transportation technology or infrastructure (LeRoy and Sonstelie, 1983; Baum-Snow, 2007); African-American migration to cities (Boustan, 2010); or a combination of factors (Kolko, 2007). In addition, despite a deep theoretical literature and empirical literature explicitly examining suburbanization and gentrification (Jackson, 1985; Mieszkowski and Mills, 1993), ours is the first to test for systematic differences across cities in these patterns. Moreover, our results also suggest that evidence for alternative theoretical channels is often stronger when considering cities where churning is more salient: that is, those cities that closely resemble flat, featureless plains.

Our paper is also related to the large literature on place-based policies, such as enterprise zones (e.g., Neumark and Kolko, 2010; Ham et al., 2011), targeted tax credits (Freedman, 2012), or the Tennessee Valley Authority (Kline and Moretti, 2013). The results of our paper suggest that the effects of place-based policies may be sensitive to natural heterogeneity in the areas studied. In regions that are naturally undifferentiated, households and firms may migrate easily in response to relatively small place-specific interventions. However, in regions where there is great natural heterogeneity, even large place-based interventions may not result in lasting changes in the spatial distribution of activity.

A broad literature examines the geographic sorting of different types of households (e.g., Tiebout, 1956; Epple and Sieg, 1999). Of course, the cross-sectional implications of variation in natural value are well known. Brueckner, Thisse, and Zenou (1999) propose a static theory that links the spatial distribution of exogenous amenities to the distribution of household income.[3] The authors note that in their monocentric city model, multiple equilbria can exist (with the rich living in either the city center or the suburbs) if the exogenous amenity advantage of the center is small. We extend this intuition to a dynamic setting. In addition to focusing on the dynamic versus static implications of our model, our work departs from theirs in testing these implications empirically.

Finally, our work is related to the literature in development and geography concerned with persistence in the spatial distribution of income and population. In theory, locational persistence might be caused by fundamental geographic features, sunk and durable factors, or amenities that are endogenous to location decisions (e.g., Davis and Weinstein, 2002; Rappaport and Sachs, 2003; Redding, Sturm, and Wolf, 2011; Bleakley and Lin, 2012). Our work departs from this literature in focusing on persistence in the within-city distribution of income, versus the distribution of income and population across cities or other subnational regions. Moreover, our results suggest that variation in natural features may be an important explanation for differences in locational persistence at other spatial scales.

## 2 Theory

We present a stylized model of neighborhood change that highlights the role of natural amenities in neighborhood dynamics. Our theory predicts that natural amenities anchor neighborhood incomes over time. Moreover, our theory also implies that natural heterogeneity can explain differences *across cities* in neighborhood dynamics and persistence in the spatial distribution of income. In order to clearly illustrate our key economic mechanism and its implications, we discuss here a simple two-

[3] Other theoretical papers share a similar intuition about transitions between (static) equilibria, with Fujita and Ogawa (1982) being an early example. Krumm (1980) provides and tests a static model with endogenous amenities and location choices. Bond and Coulson (1989) note that neighborhoods are more likely to "tip" from high- to low-income (or vice versa) when housing quality is more homogenous across neighborhoods.

neighborhood city with uncorrelated amenity shocks. This simple model abstracts from other important theoretical channels emphasized elsewhere in the literature, but we do control for these omitted channels in our empirical analysis. In Appendix A.1, we relax these assumptions and show that our theoretical predictions are robust to settings with more than two neighborhoods and correlated amenity shocks over time.

### 2.1 Model

Consider a closed city with two neighborhoods, indexed by $j = b, d$, a beach and a desert. Each neighborhood has one unit measure of land, owned by absentee landlords. There are two unit measure of workers, heterogeneous in income $\theta$.

In each period $t$, workers choose neighborhoods by comparing the *aggregate* amenity levels $A_{j,t}$ (i.e., overall quality of life) and land rents $R_{j,t}$ of the two neighborhoods. Each worker consumes one unit of land in her chosen neighborhood and receives utility $A_{j,t} \cdot c_{j,t}$ from the neighborhood's amenities and numeraire consumption $c_{j,t}$. Note that this utility specification implies that aggregate amenities and numeraire consumption are complements; this feature will be important for generating sorting by income across neighborhoods. Thus, a type-$\theta$ worker solves the following problem in each period:[4]

$$\begin{aligned}\max_j V_{j,t} &\equiv A_{j,t} \cdot c_{j,t} \text{ subject to } c_{j,t} + R_{j,t} = \theta \\ &= A_{j,t} \cdot (\theta - R_{j,t}).\end{aligned}$$

The aggregate amenity level $A_{j,t}$ is composed of both natural and man-made amenities.

$$\begin{aligned}A_{j,t} &\equiv \{a_j + \varepsilon^a_{j,t}\} + \{m_t + E(w|j,t) + \varepsilon^m_{j,t}\} \\ &= a_j + m_t + E(w|j,t) + \varepsilon_{j,t}.\end{aligned} \tag{1}$$

Two components make up a neighborhood's natural amenity value. A natural amenity $a_j$ is persistent, as in proximity to a beach. (We assume that the beach has a persistent natural advantage versus the desert, $a_b > a_d$.) Or there may be idiosyncratic shocks to natural amenities $\varepsilon^a_{j,t}$, as in natural disasters such as fires or

[4]There are no moving costs or savings; thus, maximizing the discounted sum of utilities is equivalent to maximizing current utility in each period.

sinkholes.

Man-made amenities include city-level trends common to all neighborhoods $m_t$, such as those that may result from citywide improvements in transportation infrastructure.[5] Man-made amenities may also include endogenous amenities that depend positively on the average income of residents that choose to live in that neighborhood—for example, safety, school quality, or shopping. $E(w|j,t)$ is the average income of neighborhood $j$ residents in period $t$ and captures these endogenous amenities; note that we normalize aggregate amenity levels so that the endogenous amenity level $E(w|j,t)$ has a unit coefficient. Finally, there may be idiosyncratic shocks to man-made amenities $\varepsilon^m_{j,t}$, such as unexpected changes to the quality of local governance.

The aggregate amenity shock $\varepsilon_{j,t} \equiv \varepsilon^a_{j,t} + \varepsilon^m_{j,t}$ is independent and identically distributed with a cumulative distribution function $G(-\infty, \infty)$, a symmetric distribution with mean 0. In the appendix, we show that our theoretical results are robust to correlated amenity shocks. The essential structure is that there are occasionally changes in neighborhood amenity levels that are not perfectly correlated with existing amenity levels, and shocks to neighborhoods are not always permanent. Thus, in its basic setup, our theory resembles core static models in urban economics and economic geography in that endogenous location choices depend on both the fundamental and endogenous amenity values of neighborhoods. Our departure is to include dynamic shocks to neighborhood amenity values that may occasionally reverse the historical pattern of amenities.

## 2.2 Equilibrium within a Period

We characterize equilibria of the model within a period. First, we solve for workers' neighborhood choices given neighborhoods' aggregate amenity levels.

Recall that in utility, aggregate amenities $A_{j,t}$ and numeraire consumption $c_{j,t}$ are complements. This complementarity implies that higher-income workers are willing to pay more than lower-income workers to live in superior-amenity neighborhoods. Thus, higher-income workers sort into better aggregate amenity neighborhoods by

[5] Note that any citywide trends $m_t$ cancel out when workers make neighborhood choices within the city, and thus do not affect our theoretical results. We include $m_t$ here only to account for the components of aggregate amenities that affect (or do not affect) our theoretical results.

outbidding lower-income workers.[6]

Since each neighborhood has one unit of land and each worker consumes one unit of land, each neighborhood accommodates one unit measure of workers. Thus, the superior *aggregate* amenity neighborhood (i.e., the neighborhood featuring the best combination of natural and man-made amenities) will have the top 50 percent of workers in terms of their income and the other neighborhood will have the bottom 50 percent. We define $\Theta_H$ to be the set of $\theta$ in the top 50 percent and $\Theta_L$ to the set of $\theta$ in the bottom 50 percent.

**Lemma 1** *(Sorting) In each period, high-income $\Theta_H$ workers live in the superior aggregate amenity neighborhood and low-income $\Theta_L$ workers live in the inferior aggregate amenity neighborhood.*

This perfect sorting between income and aggregate amenities implies that there are only two candidate equilibrium states in each period. In State 1, high-income workers ($\Theta_H$) live in the beach and low-income workers ($\Theta_L$) live in the desert. In State 2, high income workers ($\Theta_H$) live in the desert and low-income workers ($\Theta_L$) live in the beach.

So far, we have characterized workers' neighborhood choices as a function of the distribution of aggregate amenities across neighborhoods. In turn, workers' neighborhood choices affect aggregate amenity levels through endogenous amenities such as school quality. Thus, in equilibrium, the neighborhood ordering by aggregate amenity levels must be consistent with the aggregate amenity levels generated by workers' choices. This implies that the following condition must hold for State 1 to be an equilibrium.

$$S_1 : A_{b,t} = \alpha_b + \bar{w}_H + \epsilon_{b,t} \geq A_{d,t} = \alpha_d + \bar{w}_L + \epsilon_{d,t}, \tag{2}$$

where $\bar{w}_H \equiv E(\theta|\theta \in \Theta_H)$ and $\bar{w}_L \equiv E(\theta|\theta \in \Theta_L)$. Similarly, the following condition must hold for State 2 to be an equilibrium:

$$S_2 : A_{b,t} = \alpha_b + \bar{w}_L + \epsilon_{b,t} \leq A_{d,t} = \alpha_d + \bar{w}_H + \epsilon_{d,t}. \tag{3}$$

[6]More formally, the single crossing property holds between aggregate amenities and rents:

$$\partial(-\frac{\partial V/\partial A}{\partial V/\partial R})/\partial\theta > 0.$$

Together, conditions 2 and 3 show an equilibrium always exists. If $\alpha_b + \epsilon_{b,t} \geq \alpha_d + \epsilon_{d,t}$, condition 2 is satisfied and thus $S_1$ is an equilibrium. If $\alpha_b + \epsilon_{b,t} < \alpha_d + \epsilon_{d,t}$, condition 3 is satisfied, and thus $S_2$ is an equilibrium.

Conditions 2 and 3 also show that there can be multiple equilibria. As a simple example, suppose that $\alpha_b + \epsilon_{b,t} = \alpha_d + \epsilon_{d,t}$. In this case, conditions 2 and 3 are both satisfied, and thus both $S_1$ and $S_2$ are equilibria. Note that the multiple equilibria exist because of the endogeneity of amenity values. A high income neighborhood leads to a superior endogenous amenity level, which in turn attracts high income workers. The existence of multiple equilibria here is analogous to a familiar result in the context of agglomeration economies; see, e.g., Krugman (1991). Thus, we obtain the following Proposition.

**Proposition 2** *(i) There exists an equilibrium in each period.*
*(ii) There can be multiple equilibria in each period.*

Finally, rents are determined in each period so that the marginal worker (i.e., the median-income worker at the border between $\Theta_H$ and $\Theta_L$) is indifferent between the beach and the desert. We normalize the rent for the inferior aggregate amenity neighborhood to be 0. An immediate result is that the ordering of neighborhoods by equilibrium rents follows the same ordering by average incomes, since both increase with the aggregate amenity level of a neighborhood. We use this result in our empirical work, since for some historical years data on neighborhood incomes are unavailable.

## 2.3 Equilibrium Selection and History Dependence

Over time, depending on realizations of the amenity shocks, the equilibrium state switches back and forth between states $S_1$ and $S_2$. When both $S_1$ and $S_2$ are possible equilibria, we allow history to determine the equilibrium outcome—the selected equilibrium is simply the outcome from the previous period. This means that a selected equilibrium state will remain the chosen outcome until amenity shocks *rule out* the state as no longer a possible equilibrium.[7]

Since the current period's selected equilibrium depends on the previous period's selected equilibrium, the selected equilibrium path follows a Markov chain. From

[7]The idea that equilibrium selection might be determined by history is an old one; see, e.g., Krugman (1991). Redding, Sturm, and Wolf (1991) and Bleakley and Lin (2012) provide evidence consistent with this idea.

conditions (2) and (3), we obtain the probability of transitioning from $S_1$ to $S_2$, $\Pr(S_2|S_1)$, and vice versa.

$$\Pr(S_2|S_1) = \Pr(\varepsilon_{d,t+1} - \varepsilon_{b,t+1} > a_b - a_d + \bar{w}_H - \bar{w}_L) \tag{4}$$

$$\Pr(S_1|S_2) = \Pr(\varepsilon_{b,t+1} - \varepsilon_{d,t+1} > a_d - a_b + \bar{w}_H - \bar{w}_L) \tag{5}$$

In other words, $\Pr(S_2|S_1)$ and $\Pr(S_1|S_2)$ are the probabilities that states $S_1$ and $S_2$ are no longer possible equilibria in period $t+1$, respectively. Further, $\Pr(S_1|S_2)$ is greater than $\Pr(S_2|S_1)$, because $a_b - a_d > 0 > a_d - a_b$ and both $\varepsilon_{d,t+1} - \varepsilon_{b,t+1}$ and $\varepsilon_{b,t+1} - \varepsilon_{d,t+1}$ follow the same probability distribution. Intuitively, because the beach has a persistent natural advantage, similarly sized shocks are less likely to reverse the historical equilibrium when high-income workers live at the beach.

**Lemma 3** $\Pr(S_1|S_2) > \Pr(S_2|S_1)$.

We write the Markov transition matrix as

$$M \equiv \left\{ \begin{array}{cc} \Pr(S_1|S_1) = 1 - \Pr(S_2|S_1) & \Pr(S_1|S_2) \\ \Pr(S_2|S_1) & \Pr(S_2|S_2) = 1 - \Pr(S_1|S_2) \end{array} \right\} \tag{6}$$

and define the steady state vector $\pi$ as

$$\pi = M\pi \tag{7}$$

where the elements of $\pi$ are positive and sum to 1. The steady state vector $\pi$ is a time-invariant probability distribution over the two states, which we can also interpret as the long-run probability distribution over the states in the city.

Any Markov chain with a regular transition matrix (defined as a matrix whose elements are all positive for some power of the matrix) is known to converge to a steady state. Since $M$ is a regular Markov matrix, the probability distribution over states converges to the steady state $\pi$. By solving equation (7), we obtain $\pi$:

$$\pi \equiv \left\{ \begin{array}{c} p^*_{S_1} \\ p^*_{S_2} \end{array} \right\} = \frac{1}{\Pr(S_2|S_1) + \Pr(S_1|S_2)} \left\{ \begin{array}{c} \Pr(S_1|S_2) \\ \Pr(S_2|S_1) \end{array} \right\}. \tag{8}$$

## 2.4 Theoretical Implications

In our empirical analysis, we use within-city percentile rankings of neighborhoods by average income, versus raw nominal incomes. In this way, we control for wage level differences across cities and over time. Moreover, neighborhood percentile rankings also better capture changes in the *relative* spatial distribution of income within cities over time.[8] Therefore, we cast our theoretical implications using the within-city percentile ranking $r_j$ of each neighborhood by average income.

The income percentile ranking of a neighborhood is defined as the fraction of neighborhoods that have the same or lower income than the neighborhood. With two neighborhoods, this means that the income percentile ranking of the high-income neighborhood is $r^H = 1$ while that of the low-income neighborhood is $r^L = 0.5$.

Our first theoretical implication is that natural amenities anchor neighborhood income over time. For example, the beach is more likely to remain a high-income neighborhood than the desert. If the beach becomes a low-income neighborhood, it is more likely to return to a high-income neighborhood than the desert.

To illustrate the anchoring effect, we calculate the expected change (in steady state) in the percentile rank of neighborhood $j = b, d$ conditioned on its current ranking $r = r^H, r^L$. For example, suppose that the beach is initially inhabited by high-income workers. This happens when the city is in state 1. If the city remains in state 1 in the next period, the income percentile ranking of the neighborhood does not change. If the city changes into state 2, its percentile rank declines from $r^H$ to $r^L$. Thus, we have

$$\begin{aligned} E(\Delta r | j = b, r = r_H) &= -(r^H - r^H)P(S_1|S_1) + (r^L - r^H)P(S_2|S_1) \\ &= -(r^H - r^L)P(S_2|S_1). \end{aligned}$$

Table 1 summarizes steady-state expected changes in neighborhood percentile ranks, conditioned on possible combinations of neighborhood and initial income. Because $Pr(S_2|S_1)$ is smaller than $Pr(S_1|S_2)$ (Lemma 3), Table 1 implies that the high-income beach is more likely to remain a high-income neighborhood than the high-income desert, and the low-income beach is more likely to become a high-income neighborhood than the low-income desert. Note, too, that there is mean reversion in expected neighborhood change: the expected change in neighborhood income is negative for high-income neighborhoods and positive for low-income neighborhoods. Thus, we obtain the following Proposition:

[8]Our approach also allows us to abstract from cross-city variation in growth, described more fully in the next section. In short, we rank each neighborhood among the fixed group of extant neighborhoods in the city, in an initial year.

**Proposition 4** *Natural amenities anchor neighborhood income.*
*(i) Conditioned on initial income percentile ranks, a superior natural amenity neighborhood tends to outperform an inferior natural amenity neighborhood.*
*(ii) A high-income neighborhood tends to go down in income percentile rank while a low-income neighborhood tends to go up in income percentile rank.*

Our next theoretical implications are derived from comparative statics analyses, where we vary the heterogeneity in natural amenity levels across neighborhoods $|a_b - a_d|$ within a city. In our empirical analysis, we rely on these comparative statics to understand differences across cities in patterns of neighborhood change.

The second implication is that the relative performance of the beach versus the desert improves as $|a_b - a_d|$ increases (i.e., the greater the natural advantage of the beach). Equations (4) and (5) imply that, as $|a_b - a_d|$ increases, $Pr(S_2|S_1)$ decreases and $Pr(S_1|S_2)$ increases. Thus, it follows from Table 1 that, as $|a_b - a_d|$ increases, the expected conditional change in income of the beach rises while that of the desert falls. In other words, we expect the beach to conditionally outperform the desert *even more* in cities where the beach is especially naturally advantageous. (Empirically, we investigate this implication by comparing downtown neighborhoods in cities where downtowns are near oceans versus downtown neighborhoods in interior cities whose downtowns are less naturally advantageous.)

**Proposition 5** *The relative performance of a superior natural amenity neighborhood versus that of an inferior natural amenity neighborhood improves further as the natural amenity heterogeneity across neighborhoods increases.*

Finally, our third implication is that a city's spatial income distribution becomes more dynamically stable (i.e., more persistent) as natural heterogeneity within the city increases. To see this clearly, denote by $Var(r_{j,t})$ the over-time variance in percentile income rank of a neighborhood $j$. In steady state, the over-time variance of income ranking in the beach can be written as

$$\begin{aligned} Var(r_{j,t}|j=b) &= \{p^*_{S_1}(r^H)^2 + (1-p^*_{S_1})(r^L)^2\} - \{p^*_{S_1}r^H + (1-p^*_{S_1})r^L\}^2 \\ &= (1-p^*_{S_1}) \cdot p^*_{S_1} \cdot (r^H - r^L)^2. \end{aligned}$$

Since the average income of the desert takes exactly the opposite value to that of the desert, the over-time variance of income ranking in the desert is equal to that of the beach. Thus, average over-time income ranking variance for the city can be written as

$$E(Var(r_{j,t}|j)) = (1 - p^*_{S_1})p^*_{S_1} \cdot (r^H - r^L)^2.$$

This average variance is maximized when $p^*_{S_1} = 0.5$ and decreases monotonically as $p^*_{S_1}$ moves away from 0.5. Equations (4), (5), and (8) imply that, conditional on $\bar{w}_H - \bar{w}_L$, $p^*_{S_1}$ increases from 0.5 with $|a_b - a_d|$. Therefore, $E(Var(\bar{w}_{j,t}|j))$ decreases with $|a_b - a_d|$. Intuitively, the city's spatial distribution of income experiences the least persistence over time when there is no natural heterogeneity within the city (i.e., the city is in a flat, featureless plain). As the beach's natural advantage increases, the likelihood of churning between states declines, leading to stability and persistence in the spatial distribution of income.

**Proposition 6** *Conditioned on income dispersion $\bar{w}_H - \bar{w}_L$, the expected over-time variance of neighborhood income for the city $E(Var(r_{j,t}|j))$ decreases with the differences in natural advantages $|a_b - a_d|$.*

### 2.5 Full Model

In Appendix A.1, we show that our theoretical results are robust in a full model relaxing several assumptions of the simple model presented here. Instead of two neighborhoods, the city has $J \in \mathbb{N}$ neighborhoods, and the aggregate amenity shock $\epsilon_{j,t}$ follows an AR(1) process such that $\epsilon_{j,t+1} = \rho\epsilon_{j,t} + \nu_t$, where $\nu_t$ is independent and identically distributed. We also extend the equilibrium selection rule in Section 2.3: when multiple equilibria are possible, we choose the one that is closest to the selected equilibrium in the previous period, in terms of Euclidean distance in the vector of average incomes across neighborhoods. We analytically prove Lemma 1 and Proposition 2. We use numerical methods to demonstrate that Propositions 4, 5, and 6 hold widely when the aggregate amenity shock follows a stationary process (i.e., $\rho < 1$). Note that stationarity is not restrictive, because an overall trend in aggregate amenity level can be captured by $m_t$ in equation 1.

## 3 Empirics

We confirm the testable implications of our theory using a novel database of consistent-boundary neighborhoods in U.S. metropolitan areas, 1880–2010. Section 3.1 and Appendix A.2 discuss the construction of our database and its key features. Section 3.2

tests Proposition 4 that natural amenities anchor neighborhoods to high incomes, and it also describes our responses to several identification challenges, including the unknown and changing amenity value of natural features. Section 3.3 shows that, consistent with Proposition 5, downtown neighborhoods in coastal cities were less susceptible than interior cities to the suburbanization of income in the middle 20th century. Finally, Section 3.4 tests Proposition 6 that naturally heterogeneous cities feature more persistent spatial distributions of income.

## 3.1 Data

### 3.1.1 Census Data and Geographic Normalizations

We construct a panel database of consistent-boundary neighborhoods in many U.S. metropolitan areas from 1880 to 2010. We use census tracts as neighborhoods, as tracts are relatively fine geographic units and data are available at the tract level over our sample period, even in historical census years. Since census tract boundaries change from one decade to the next, we normalize historical data to 2010 census tract boundaries.[9] For each census tract, we collect information about household income, population, and housing from decennial censuses between 1880 and 2000 and the American Community Surveys between 2006 and 2010.[10]

An important limitation is that our panel is unbalanced. Cities expand and add neighborhoods over time, and new cities and neighborhoods emerge as the census adds tract coverage of more cities in later years. In addition, our ability to match census households to neighborhoods is limited by the availability of maps showing the spatial location of historical census tracts or enumeration districts. Table 2 shows the number of metropolitan areas and neighborhoods (i.e., census tracts) available in each year. Overall, we have observations of 60,757 neighborhoods across 308 metropolitan areas and 12 census years from 1880 to 2010. However, the number of observations used in our empirical analysis varies across tests with data availability.[11] The data are most complete for later census years, especially after 1960, and we do not have any data for census years 1890 and 1900.

[9] In 2010, our average sample neighborhood is 28 square kilometers and contains 4,300 people.

[10] Because of small annual sample sizes and privacy concerns, the ACS data represent five-year averages of residents and houses located in each tract. For convenience, we refer to these data as coming from the year 2010, though they really represent an average over 2006–2010.

[11] Small boundary normalization errors account for the small number of tracts in 2000 that do not appear in 2010, but these tract fragments are ultimately dropped in our regressions.

We assign each neighborhood to a single metropolitan area, using the Office of Management and Budget's definitions of core-based statistical areas (CBSAs) from December 2009. We refer to each metropolitan area as a "city." (We address changes in metropolitan area boundaries over time by dropping the nonurbanized areas in each period within present-day boundaries, as described in the Appendix. Thus, neighborhoods are unlikely to appear in our panel until they are urbanized and part of the metropolitan economy.) When relevant, we aggregate CBSAs to consolidated statistical areas. For example, we combine the Los Angeles-Long Beach-Santa Ana CBSA with the Oxnard-Thousand Oaks-Ventura and Riverside-San Bernardino-Ontario CBSAs.

Finally, we spatially match neighborhoods to a variety of persistent natural geographic features. While there are innumerable natural features that might have amenity value, our work features a long list of highly visible and important physical attributes. For each neighborhood, we separately calculate the (i) distance from the tract centroid to the nearest coastline (i.e., the Atlantic or Pacific Ocean, the Gulf of Mexico, or a Great Lake), (ii) the nearest (non-Great) lake, and (iii) the nearest major river. We also calculate (iv) the average slope, (v) the flood-hazard risk, and (vi) the 1971–2000 average annual precipitation, (vii) July maximum temperature, and (viii) January minimum temperature. The Appendix describes sources for these data.

### 3.1.2 Tract Percentile Ranks and Other Variables

Because we are interested in the performance of neighborhood income relative to other neighborhoods within the same city, we rank tracts within each metropolitan area and census year. We use neighborhoods' percentile rank $r_{i,t}$, a variable bounded by 0 and 1. Thus, in 2010, both Malibu (within the Los Angeles metropolitan area) and the Upper East Side (within the New York metropolitan area) have values of $r_{i,2010}$ near 1. In this way, we also control for differences in wage levels across cities and years.

We use average household income to rank tracts within each metropolitan area, except in historical census years 1880–1940 when income data are not available. For 1910–1940, we use median housing rents to rank tracts. In 1880, lacking both data on income or prices, we use an imputed occupational income score. Recall that our theory predicts that neighborhood rents follow the same ordering as neighborhood incomes and that the percentile rank of a neighborhood in terms of rent is the same as that in terms of income. In unreported robustness checks, we have verified that our results are robust to using housing rents or household incomes, when both are available, or using the literacy rate in 1880.

We also calculate a number of metropolitan area statistics based on these tract data, including log levels and changes in (i) population, (ii) land area, (iii) total housing units and the housing age distribution, and various other characteristics. Appendix Table A1 reports summary statistics.

### 3.2 The Anchoring Effect

We test Proposition 4 that superior natural amenities anchor neighborhoods to high incomes over time. Recall that the Proposition predicts that (i) conditional on current income, a superior natural amenity neighborhood tends to outperform an inferior natural amenity neighborhood and (ii) high-income neighborhoods tend to go down in average income while low-income neighborhoods tend to go up.

To test this proposition, consider the following neighborhood-level regression:

$$\Delta r_{i,t} = \beta_0 + \beta_1 \mathbf{1}(natural\_amenity_i) + \beta_2 r_{i,t} + \epsilon_{i,t} \tag{9}$$

where $\Delta r_{i,t}$ is the change in neighborhood $i$'s percentile rank from $t$ to $t+1$, $\mathbf{1}(natural\_amenity_i)$ is an indicator for superior natural amenity neighborhood, and $r_{i,t}$ is initial percentile rank in $t$.[12] Below, we explain how we identify superior natural amenity neighborhoods. We include the initial rank as a regressor since our theory suggests that in steady state, it is *conditional on initial rank* that neighborhoods with superior natural amenities tend to increase in income versus other neighborhoods. This term also allows us to test for mean reversion—the second part of Proposition 4. Thus, Proposition 4 predicts that $\beta_1 > 0$ and $\beta_2 < 0$.

One concern with this specification, derived from the model, is that the model abstracts away from many features that may affect neighborhood dynamics. Some of

[12] We calculate the change in percentile rank $\Delta r_{i,t}$ for each tract by subtracting its current rank $r_{i,t}$ from next period's rank $r_{i,t+1}$. Obviously, this change can only be calculated for tracts that exist in both the current period and the next period. Therefore, tracts that are added to the metropolitan area are *not* included in this calculation. *Rank changes are based only on the metropolitan area footprint in the initial year, and they represent ranks calculated only among extant tracts in the initial year.* (See the Appendix for further discussion.) For example, Levittown, New York, was a neighborhood that appeared for the first time in 1960. In computing rank changes among New York metropolitan area neighborhoods between 1950 and 1960, Levittown is *excluded from both groups of 1950 and 1960 neighborhoods.* This is one way in which our empirical analysis abstracts from differences in city growth rates. (Later, we also condition on city growth in our regressions.)

these features may be correlated with neighborhoods' natural value. For example, proximity to the ocean may be correlated with proximity to the central business district (CBD), since many cities and their downtowns were founded near significant natural features such as harbors. Since proximity to a central business district affects changes in income (c.f. Brueckner and Rosenthal, 2009), estimation of equation (9) will be contaminated by omitted-variables bias.

In order to control for these factors, we control for a number of additional neighborhood and city characteristics:

$$\Delta r_{i,t} = \beta_0 + \beta_1 \mathbf{1}(natural\_amenity_i) + \beta_2 \ln r_{i,t} + \mathbf{X}'_{i,t}\beta_3 + \epsilon_{i,t}. \tag{10}$$

Here, $\mathbf{X}'_{i,t}$ is a vector that could contain many tract characteristics. In order to control for the effect of downtown proximity and the internal structure of the city, we control for initial tract log population density, or, alternatively, distance to the CBD.[13] We also include average housing age to control for the housing filtering process and metropolitan area level changes in population and land area to control for city size growth. We cluster the standard errors at the metropolitan area level.

Next, we identify superior natural amenity neighborhoods with three different strategies. Our first strategy is to assign $\mathbf{1}(natural\_amenity_i) = 1$ if neighborhood $i$ is within 1 kilometer from a particular natural feature. Our natural features include (separately) oceans or Great Lakes, other lakes, and major rivers. (We vary the 1-kilometer distance threshold as robustness tests.) We also define thresholds for hilly neighborhoods (average slope greater than 15 degrees), flood hazard (an average annual probability of flooding less than 1 percent), and a moderate climate (not too hot or cold, and not too rainy; see the Appendix).

A concern with this strategy is that some observable natural features can be noxious rather than pleasant: For example, we cannot observe whether a river is polluted and thus a disamenity. In order to alleviate this concern, our second strategy is to look at natural features near top-income neighborhoods in some initial year. Because workers can observe whether a particular natural feature is an amenity, natural features near top-ranked neighborhoods in some initial year are more likely to be positive amenities than natural features initially surrounded by slums. Thus, if a section of river is surrounded by neighborhoods whose average income is in the

[13] Data on tract distances to metropolitan area CBDs are from Fee and Hartley (2012). Because they are missing for a handful of metropolitan areas, we do not present them in our baseline results.

90th percentile among all of the neighborhoods in the city, the river is *more likely* to be a positive amenity versus a disamenity. So in our second strategy, we assign $\mathbf{1}(natural\_amenity_i) = 1$ if, and only if, the neighborhood is proximate to the natural feature and the neighborhood was initially in the top decile of neighborhoods by average income. While we are still uncertain about the amenity value of natural features, this strategy relies on the idea that conditioning features on the initial location of incomes increases the likelihood that such features are in fact amenities.

So far, we have evaluated each feature individually. Our third approach is to combine these features together into an index of aggregate natural value by predicting rent from our various observed natural features. We regress the logarithm of neighborhood median housing rent, reported in censuses from 1930 to 2010, against a vector of variables indicating proximity to all of our natural features, log population density, log distance to the CBD, log number of housing units, average housing age, and city–year effects. Then, we predict values for housing rents based on the estimated coefficients from this regression. Of course, this hedonic regression raises endogeneity concerns, particularly about omitted endogenous factors such as school quality. However, the resulting predicted values may be unbiased estimates of the aggregate *natural* value of neighborhoods if omitted factors are related to the observed factors in the same way. For example, if school quality is related to coastal proximity but not hills, then the estimated coefficients on coastal proximity and hilliness will be biased, relative to each other. However, if school quality is related to the overall natural advantage of neighborhoods, then the estimated coefficients on coastal proximity and hilliness will be biased in the same way, but the relative weights will be unbiased. In this case, predicted rents may be a good indicator for the aggregate natural value of neighborhoods.

A final and related concern is that the value of natural features may change over time. An alternative view is that neighborhoods with superior natural amenities may increase in income because of increasing taste for natural amenities over time. However, this view is inconsistent with our later result, as it predicts that these neighborhoods experience greater fluctuations over time in neighborhood income. In addition, our regressions here focus mostly on 10-year changes in the location of income: while today's tastes and technologies are different from the 19th century's, it is less likely that tastes for natural amenities have evolved significantly over a single decade. We also find no evidence that the tastes of high-income versus low-income households for natural features have fluctuated significantly over time. If that were the case, high-income neighborhoods to be increasingly concentrated near superior natural amenities like the coast. In Figure 1, we show that the relative likelihood that a high-income neighborhood is within 1 kilometer of an ocean or Great Lake coastline has remained roughly constant over the 130-year span of our sample.

Table 3 shows results of conditional neighborhood change for neighborhoods near and far from an ocean or Great Lake. Recall that the first part of Proposition 4 predicts that the coefficient on $\mathbf{1}(natural\_amenity_i)$ is positive. The first column shows the result from estimating the baseline equation (9): the coefficient on $\mathbf{1}(natural\_amenity_i)$ turns out to be negative, counter to our theoretical prediction. However, once we control for the internal structure of the city using the logarithm of neighborhood population density in column (2), the conditional effect of coastal proximity is estimated to be positive and statistically significant. In column (3), we also control for the logarithm of city population growth, the logarithm of land area growth, and neighborhoods' average housing age. These additional controls do not change the estimated coefficients of interest. In column (4), we additionally condition coastal areas on their proximity to top-decile neighborhoods in the initial year. Recall that this exercise relies on the idea that this condition increases the likelihood that such coastal areas are positive rather than negative amenities. As column (4) shows, the coefficient estimate on proximity to the coast is even stronger, consistent with the increased likelihood that these coastlines are positive amenities.

The estimated effect of initial income on the change in neighborhood rank is negative, consistent with the mean reversion predicted by the second part of Proposition 4.

In Table 4, we show that estimates of the conditional effect of superior natural amenities on neighborhood change are robust to various natural features, including our hedonic rent index that combines all of our observed natural features. In all cases except rivers, natural features are associated with a conditional increase in neighborhood income. (This result is consistent with the view that, on average, U.S. rivers are disamenities.) When we condition these natural features on initial proximity to top-decile neighborhoods, the coefficient estimate increases, consistent with the greater likelihood that these natural features are actually amenities.

In the Appendix, we show that our results are robust to varying the time horizon and base year. We also show that our results are robust to varying the definition of our indicator variable thresholds.

### 3.3 The Suburbanization and Gentrification of Income

Next, we test Proposition 5 that the relative performance of superior natural amenity neighborhoods improves further as natural heterogeneity among neighborhoods increases. We test this implication by comparing how downtown neighborhoods in coastal cities performed relative to those of interior cities.

Coastal cities are cities that developed on the shores of the Atlantic or Pacific oceans, one of the Great Lakes, or the Gulf of Mexico. Thus, their downtowns continue to be located near these bodies of water. Interior cities tend not to be devoid of surface water, but are typically river cities.

We maintain two assumptions throughout this test. First, coastal cities tend to have more heterogeneity in natural amenity level across neighborhoods than interior cities. A justification for this assumption is that oceans tend to be stronger natural amenities than rivers; the results reported in the previous section are consistent with this assumption. Second, the central business districts of coastal cities tend to have developed along coastlines (i.e., areas of high natural value).

Proposition 5 suggests that neighborhoods near the central business districts of coastal cities tended to suffer relatively less versus those of interior cities during the suburbanization process common in many U.S. cities. Further, Proposition 5 also suggests that the CBDs of coastal cities tended to experience faster gentrification in the late 20th and early 21st centuries.

We have neighborhood-level data from 1880 for 12 coastal and 17 interior cities, listed in Appendix Table =refclassify.

In Figure 2, we show a series of snapshots, one for each census decade where data are available, of the spatial pattern of income versus distance to the city center.[14] The horizontal axis measures distance from the city center, in meters, and the vertical axis measures average household income, on a percentile rank scale. The lines represent the results of lowess regressions, fitted separately for coastal versus interior cities.

In 1880, both coastal and interior cities display similar declining income gradients with distance from the city center. These patterns are consistent with the fact that many of these cities were still recently founded as of 1880, and the best-developed

[14]Data on distances to city centers were graciously provided to us by Dan Hartley. Fee and Hartley (2012) identify the latitude and longitude of city centers by taking the spatial centroid of the group of census tracts listed in the 1982 Census of Retail Trade for the central city of the metropolitan area. For metropolitan areas not in the 1982 Census of Retail Trade, they use the latitude and longitude for central cities using ArcGIS's 10.0 North American Geocoding Service.

areas would have been clustered near downtown.

However, as early as 1930, we see a divergence in the fortunes of downtowns in coastal versus interior cities. Downtowns in interior cities tend to decline faster than downtowns in coastal cities, at least until 1960. Then, from 1970 onward, coastal city downtowns tend to improve faster than interior city downtowns. Thus, our evidence on suburbanization and gentrification patterns in coastal versus interior cities is consistent with the theoretical prediction that greater within-city heterogeneity improves the performance of superior natural amenity neighborhoods.

### 3.4 Persistence in the Spatial Distribution of Income

Finally, we test Proposition 6 that cities with greater variation among neighborhoods in natural value tend to have more dynamically stable (i.e., persistent) spatial distributions of income.

We begin with the following hierarchical linear model.

$$Var(r_i) = \delta_m + \varepsilon_i \tag{11}$$

$$\hat{\delta}_m = \gamma_0 + \gamma_1 \mathbf{\Gamma}_m + \gamma_2 Var(w_j | j \in m) + \mathbf{Z}'_m \gamma_3 + \mu_m \tag{12}$$

Here, $Var(r_i)$ is the *over-time* variance of neighborhood $i$'s percentile ranking within city $m$, $\delta_m$ is a city-level effect estimated in the first level and used as the dependent variable in the second level. The estimated robust errors $\hat{\varepsilon}_i$ are clustered at the city level. Following Wooldridge (2003), the minimum distance estimator is equivalent to estimating the second step using weighted least squares, where the weights are $1/\widehat{Avar}(\hat{\delta}_m)$. $\mathbf{\Gamma}_m$ is a city-level measure of variation in natural value among neighborhoods within city $m$. Finally, we control for differences across cities in the distribution of nominal income using $Var(w_j|j \in m)$, a city-level measure of the variance in income across neighborhoods $j$ within city $m$. Thus, the source of identification of $\gamma_1$ is cross-sectional variation across cities in terms of within-city natural heterogeneity. Proposition 6 predicts that $\gamma_1 < 0$.

We calculate $Var(r_i)$, the over-time variance of neighborhood $i$'s percentile rank, using the 6 observations from 1960 to 2010. We begin using 1960 as our base year because the number of neighborhoods in the data set decreases as we go back in time as shown in Table 2. For example, our data has 38,669 neighborhoods in 1960 but only 17,681 in 1950. In later years, the gain in the number of observed neighborhoods comes at the cost of lost precision in calculating the over-time variance in neighborhood rank. We show later that our results are robust to the choice of base year. For all choices of base year, we fully balance our panel—thus, for a base year of 1960, our regressions and computations of ranks over time exclude any neighborhoods or cities that do not appear in our sample in 1960. In historical census years, we also drop any cities that are missing from subsequent years. Thus, while we have observations of 10 cities in 1930, only nine cities appear in all of the same nine census years between 1930 and 2010.

To measure $\mathbf{\Gamma}_m$, we take three approaches. First, we use an indicator variable for a coastal city. The ocean is arguably a dominant natural feature affecting natural amenity levels. Thus, we expect coastal cities to have higher internal variance in natural amenities than noncoastal cities.

Second, we use the within-city standard deviation in log neighborhood distance to the coast. In the case of distance to an ocean or Great Lake, we take the logarithm in order to deemphasize variation in interior cities where all neighborhoods are far from the coast. In using logarithms, all neighborhoods in Denver have equally poor access to the coast, while there is much greater variation among neighborhoods in Miami.

Third, we use the within-city standard deviation in the predicted rent index, discussed earlier. This measure has the advantage of summarizing all of our observed natural features into a measure of natural heterogeneity.

Finally, to control for omitted factors in the model, we add other city-level covariates related to over-time volatility in neighborhood income in $\mathbf{Z}_m$. For example, we control for differences across cities in growth rates by adding the logarithms of city growth in population and land area.

Table 5 shows the results of estimating equation (12). Each column shows a separate regression. We multiply the dependent variable by 100 for presentation purposes, so the units are percentile points. Column 1 shows that, on average, neighborhoods in coastal metropolitan areas experience smaller fluctuations in income over time. The coefficient on a metropolitan indicator for proximity to the ocean is negative and precisely estimated. The magnitude of the effect is small but not negligible—the reduction in over-time variance is approximately 19% of one standard deviation across neighborhoods. The coefficient on within-city income inequality is negative, too, but it is imprecisely estimated.

An alternative explanation of our result in column 1 is that heterogeneous cities are also land-supply constrained. That alternative view suggests that in flat cities, the supply of housing is more elastic, and therefore the causal link between geography and neighborhood stability is mediated by city growth, not the value of natural amenities (c.f. Saiz, 2010). To address this concern, we control for metropolitan area growth in population and area. Column 2 shows the result. The estimates do suggest that growing cities are less stable, consistent with the alternative view. However, even conditioned on city growth, coastal cities are more stable. Thus, we do not view our results as being spuriously caused by differences in land-supply elasticity across cities.

In columns 3 and 4, we repeat the first two regressions but use the standard deviation within each city in the logarithm of neighborhood distance to an ocean or Great Lake. The coefficients on this regressor are negative and precisely estimated. Again, neighborhoods in naturally heterogeneous cities tend to experience smaller fluctuations in income over time.

Table 6 shows that these results are robust when we vary starting year to calculate $Var(r_i)$ and use different natural features other than the coast. Each cell reports the estimated coefficient on the within-metropolitan area standard deviation natural value from a separate regression, with a specification identical to Table 5, column 4. Thus, that estimate is repeated in the first row, column 5 of Table 6.

Each row shows results where the explanatory variable of interest is noted by the row heading. Each column displays results for regressions using the base year indicated. For example, in column 1, we rely on cross-sectional variation in 1880 across 29 cities. Note that in column 1, we use seven census years to calculate the over-time variance in neighborhood income—these are the census years for which we can create a balanced panel of city observations. Thus, we drop neighborhood and city observations in years where we have incomplete coverage of cities.

Neighborhoods in rugged cities (second row) or naturally heterogeneous cities (third row) tended to experience smaller fluctuations in income over 1960–2010 (column 5), consistent with our earlier result for coastal cities. These estimates are negative and precisely estimated as well. Overall, all of the precisely estimated effects are negative, and the results are especially strong when considering periods between 1950 and 1980. The lack of precision in historical census years is consistent with the small number of metropolitan areas observed in those years. In 1990, the small number of over-time observations used to calculate the dependent variable may introduce noise and bias the estimates towards zero, contributing to the lack of precision.

Finally, Figure 3 illustrates our main result that neighborhoods in naturally heterogeneous cities tended to experience smaller over-time fluctuations in income over 1960–2010. Each point represents a metropolitan area. The vertical axis measures the metropolitan-level residual from a regression of mean variance in percentile rank over time on controls as in Table 5, column 4. The horizontal axis measures the within-city standard deviation in our predicted rent index; Los Angeles and the San Francisco Bay Area (labeled San Jose) are the two most naturally heterogeneous metropolitan areas by this index. The slope of the fitted line is the same as the estimate reported in Table 6, column 5. Thus, naturally heterogeneous cities exhibit more persistent spatial distributions of income over time.

# 4 Conclusions

We combine new theory and a novel database of consistent-boundary neighborhoods to study both neighborhood dynamics and differences across cities in patterns of neighborhood change, suburbanization, and persistence. Our theory and results highlight the role of natural amenities in neighborhood dynamics. Persistent natural amenities anchor neighborhoods to high incomes over time, and they can affect neighborhood dynamics citywide. Downtown neighborhoods in coastal cities were both less susceptible to suburbanization and more responsive to gentrification versus interior cities. Finally, cities with greater internal natural heterogeneity tend to exhibit more persistent spatial distributions of income.

Our results are also broadly related to the literature on place-based policies. The effects of such policies may be sensitive to underlying geographic heterogeneity in the targeted areas; in regions with great natural heterogeneity, even large interventions may not result in churning in the spatial distribution of activity, while small interventions may cause factors to migrate in flat cities. However, we have intentionally said little about the welfare implications of our results, although understanding optimal neighborhood sorting and the structure of cities is of paramount interest. Recent studies have suggested that the degree of neighborhood sorting may be related to outcomes such as intergenerational income mobility (e.g., Chetty et al., 2013). In future work, it would be useful to better understand the importance of natural heterogeneity for place-based policies and outcomes.

## 5 Works Cited


Aaronson, Daniel. "Neighborhood Dynamics," *Journal of Urban Economics* 49 (2001), 1–31.

Arthur, W. Brian. *Increasing Returns and Path Dependence in the Economy.* Ann Arbor: University of Michigan Press, 1994.

Baum-Snow, Nathaniel. "Did Highways Cause Suburbanization?" *The Quarterly Journal of Economics* 122:2 (2007), 775–805.

Bleakley, Hoyt and Jeffrey Lin. "Portage and Path Dependence," *The Quarterly Journal of Economics* 127:2 (2012), 587–644.

Bogue, Donald. *Census Tract Data, 1940: Elizabeth Mullen Bogue File* [Computer file]. ICPSR version. University of Chicago, Community and Family Study Center [producer], 1975. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor] (2000a).

Bogue, Donald. *Census Tract Data, 1950: Elizabeth Mullen Bogue File* [Computer file]. ICPSR version. University of Chicago, Community and Family Study Center [producer], 1975. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor] (2000b).

Bogue, Donald. *Census Tract Data, 1960: Elizabeth Mullen Bogue File* [Computer file]. ICPSR version. University of Chicago, Community and Family Study Center [producer], 1975. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor] (2000c).

Bond, Eric W., and N. Edward Coulson. "Externalities, filtering, and neighborhood change," *Journal of Urban Economics* 26 (1989), 231–249.

Boustan, Leah Platt. "Was postwar suburbanization 'White flight'? Evidence from the Black migration," *Quarterly Journal of Economics* 125:1 (2010), 417–443.

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

Brueckner, Jan K., Jacques-François Thisse, and Yves Zenou. "Why is central Paris rich and downtown Detroit poor? An amenity-based theory," *European Economic Review* 43 (1999), 91–107.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez (2013). Equality of Opportunity project, http://www.equality-of-opportunity.org/.

Davis, Donald, and David Weinstein. "Bones, Bombs, and Break Points: The Geography of Economic Activity," *American Economic Review* 92:5 (2002), 1269–1289.

Epple, Dennis, and Holger Sieg. "Estimating Equilibrium Models of Local Jurisdictions," *Journal of Political Economy* 107:4 (1999), 645–681.

Federal Emergency Management Agency. *National Flood Hazard Layer* [Computer file] (2012). Data accessed July 17, 2012.

Fee, Kyle and Daniel Hartley. "The Relationship between City Center Density and Urban Growth or Decline," Working paper, 12–13 (2012), Federal Reserve Bank of Cleveland.

Freedman, Matthew. "Teaching new markets old tricks: The effects of subsidized investment on low-income neighborhoods," *Journal of Public Economics* 96:11–12 (2012), 1000–1014.

Fujita, Masahisa, and Hideaki Ogawa. "Multiple equilibria and structural transition of non-monocentric urban configurations," *Regional Science and Urban Economics* 12 (1982), 161–196.

Guerrieri, Veronica, Daniel Hartley, and Erik Hurst. "Endogenous gentrification and housing-price dynamics," Working paper, 10–08 (2010), Federal Reserve Bank of Cleveland.

Ham, John, Charles Swensom, Ayse Imrohoroglu, and Heonjae Song. "Government programs can improve local labor markets: Evidence from State Enterprise Zones, Federal Empowerment Zones, and Federal Enterprise Community." *Journal of Public Economics* 95:7–8 (2011), 779–797.

Jackson, Kenneth T. *Crabgrass Frontier: The Suburbanization of the United States.* Oxford University Press, 1985.

Kline, Patrick, and Enrico Moretti. "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority," *The Quarterly Journal of Economics*, forthcoming.

Kolko, Jed. "The Determinants of Gentrification," Working paper (2007), Public Policy Institute of California.

Krugman, Paul. "History Versus Expectations," *The Quarterly Journal of Economics* 106:2 (1991), 651–667.

Krumm, Ronald J. "Neighborhood amenities: An economic analysis," *Journal of Urban Economics* 7 (1980), 208–224.

LeRoy, Stephen F., and Jon Sonstelie. "Paradise lost and regained: Transportation innovation, income, and residential location," *Journal of Urban Economics* 13 (1983), 67–89.

Logan, John R., Jason Jindrich, Hyoungjin Shin, and Weiwei Zhang. "Mapping America in 1880: The Urban Transitional Historical GIS Project," *Historical Methods* 44:1 (2011), 49–60.

Logan, John R., Zengwang Xu, and Brian Stults. "Interpolating US Decennial Census Tract Data from as Early as 1970 to 2010: A Longitudinal Tract Database," *Professional Geographer*, forthcoming.

Mieszkowski, Peter, and Edwin S. Mills. "The Causes of Metropolitan Suburbanization," *Journal of Economic Perspectives,* 7:3 (1993), 135–147.

Minnesota Population Center. *National Historical Geographic Information System: Version 2.0.* Minneapolis: University of Minnesota, 2011.

National Oceanic and Atmospheric Administration. *Coastal Geospatial Data Project* [Computer file, http://coastalgeospatial.noaa.gov/data_gis.html] (2012). Data accessed September 26, 2012.

Neumark, David, and Jed Kolko. "Do Enterprise Zones Create Jobs? Evidence from Californias Enterprise Zone Program." *Journal of Urban Economics* 68:1 (2010), 1–19.

Office of Management and Budget. *Metropolitan and Micropolitan Statistical Areas and Components, December 2009, with Codes* [Computer file, http://www.census.gov/population/metro/files/lists/2009/List1.txt] (2009).

PRISM Climate Group, Oregon State University [Computer file, http://prism.oregonstate.edu]. Data accessed January 15, 2013.

Rappaport, Jordan, and Jeffrey D. Sachs. "The United States as a Coastal Nation," *Journal of Economic Growth* 8 (2003), 5–46.

Rauch, James E. "Does History Matter Only When It Matters Little? The Case of City-Industry Location," *The Quarterly Journal of Economics* 108:3 (1993), 843–867.

Redding, Stephen J., Daniel M. Sturm, and Nikolaus Wolf. “History and Industry Location: Evidence from German Airports,” *The Review of Economics and Statistics* 93:3 (2011), 814–831.

Rosenthal, Stuart S. “Old homes, externalities, and poor neighborhoods: A model of urban decline and renewal,” *Journal of Urban Economics* 63 (2008), 816–840.

Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek. *Integrated Public Use Microdata Series: Version 5.0* [Computer file]. Minneapolis: University of Minnesota, 2010..

Saiz, Albert. “The Geographic Determinants of Housing Supply,” *The Quarterly Journal of Economics*, 125:3 (2010), 1253–1296.

Tatian, P. A. *Neighborhood Change Database (NCDB) 1970-2000 Tract Data: Data Users Guide*, Washington, DC: Urban Institute and Geolytics, Inc, 2003.

Tiebout, Charles M. “A Pure Theory of Local Expenditures,” *Journal of Political Economy* 64 (1956), 416–424.

Wooldridge, Jeffrey M. Cluster-Sample Methods in Applied Econometrics, *American Economic Review, Papers and Proceedings* 93:2 (2003), 133–138.

A. Share of neighborhoods within 1km of coast

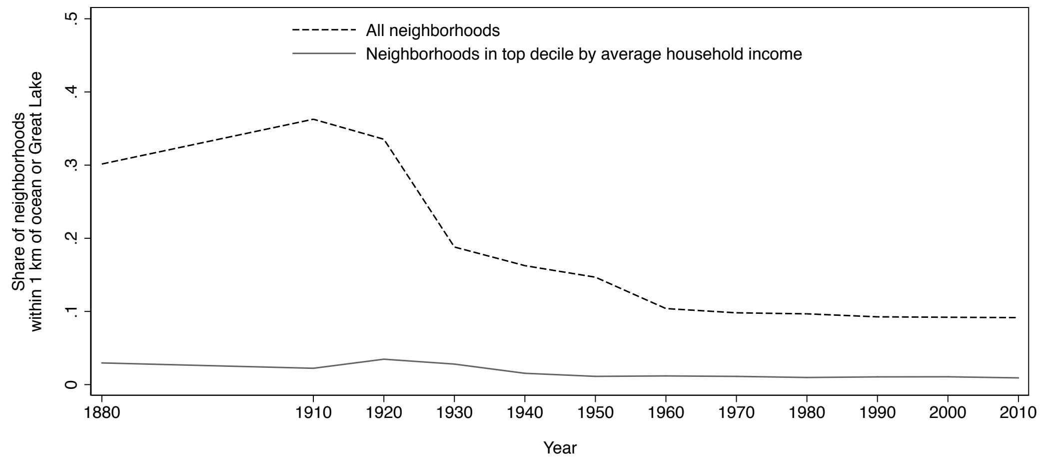


B. Relative likelihood that high-income neighborhood is within 1km of coast

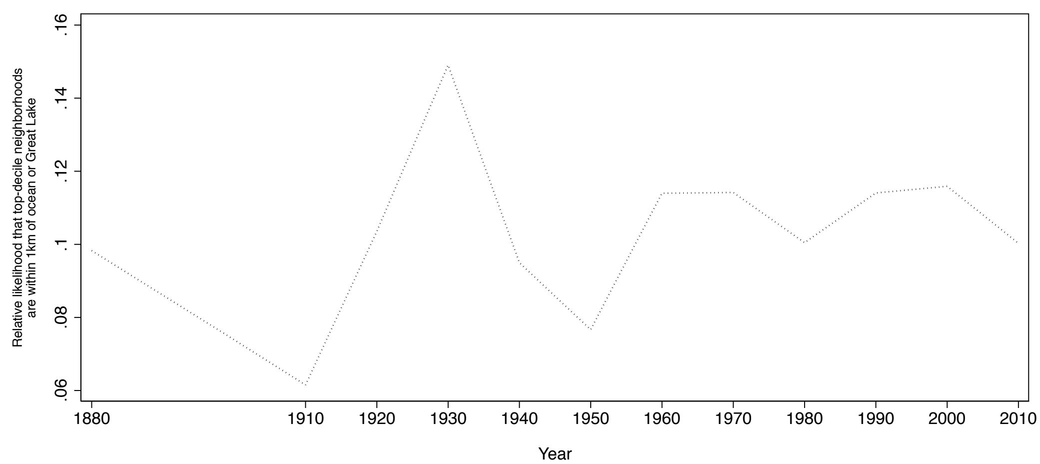


Figure 1: Stability in the likelihood that a high-income neighborhood is near coast

Panel A shows, by census year, the share of all neighborhoods within 1km of an ocean or Great Lake (dashed) and the share of high-income neighborhoods (top decile by average household income) within 1km of an ocean or Great Lake. Panel B shows, by census year, the relative likelihood that a high-income neighborhood is within 1km of an ocean or Great Lake (i.e., the ratio of the two series shown in Panel A).

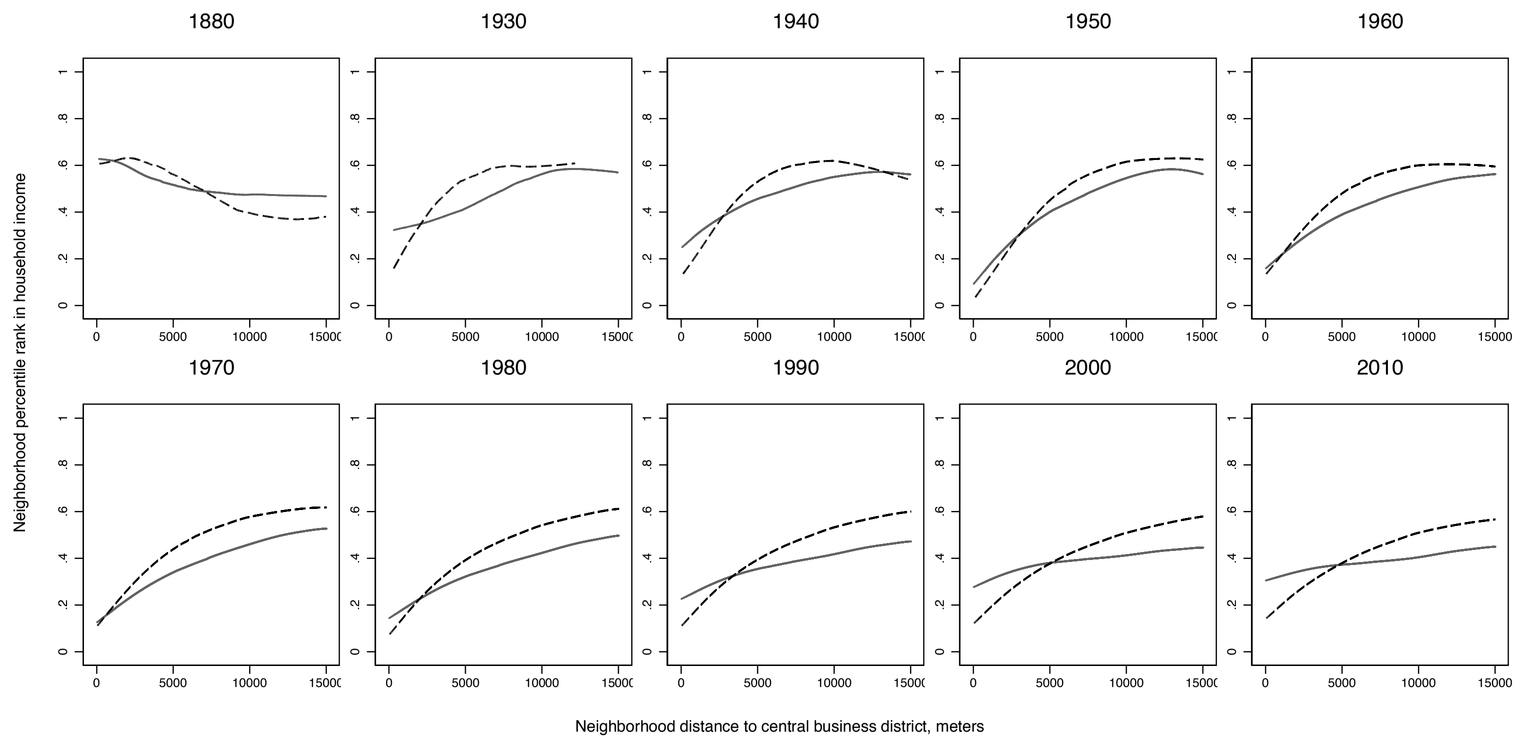


Figure 2: Income and residential location for coastal versus interior cities, 1880–2010

Each panel shows, for a different census year, the pattern of neighborhood average household income on the vertical axis versus neighborhood distance to the city center (up to 15km) on the horizontal axis. The smoothed lines are from lowess regressions. Two groups of cities are shown in each panel: Coastal cities are represented by a solid line and interior cities by a dashed line. The sample is a fixed group of 29 metropolitan areas with some missing city observations in 1930–1950. Coastal and interior cities are classified in Appendix Table A2.

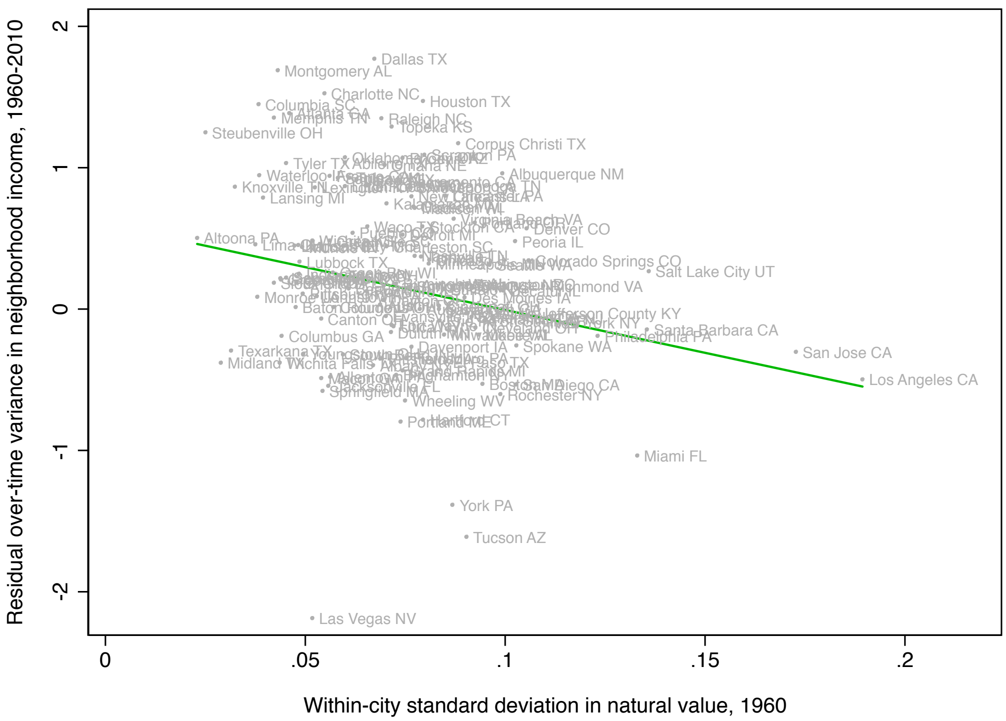


Figure 3: Neighborhoods in naturally heterogeneous cities experience smaller over-time fluctuations in income

The vertical axis measures the metropolitan-level residual from a regression of mean neighborhood 1960–2010 variance in percentile rank by income on within-metropolitan variance in neighborhood income and log changes in metropolitan population and land area over the same period (i.e., the same controls and weights as in Table 5, column 4). The horizontal axis measures the within-metropolitan variance (standard deviation) in natural value, using estimated hedonic weights as described in the text. The slope of the fitted line corresponds to the estimate in Table 6, column (5).

Table 1: Expected income change conditioned on initial income and natural feature

| Initial percentile rank | Beach | Desert |
|---|---|---|
| $r^H$ | $(r^L - r^H)Pr(S_2 \vert S_1)$ | $(r^L - r^H)Pr(S_1 \vert S_2)$ |
| $r^L$ | $(r^H - r^L)Pr(S_1 \vert S_2)$ | $(r^H - r^L)Pr(S_2 \vert S_1)$ |

Table 2: Number of neighborhood and metropolitan observations by census year

| **Year** | **Metropolitan areas** | **Neighborhoods** |
|---|---:|---:|
| 2010 | 308 | 60,757 |
| 2000 | 308 | 60,766 |
| 1990 | 308 | 60,299 |
| 1980 | 259 | 56,176 |
| 1970 | 229 | 49,888 |
| 1960 | 135 | 38,669 |
| 1950 | 51 | 17,681 |
| 1940 | 43 | 11,527 |
| 1930 | 10 | 1,962 |
| 1920 | 2 | 2,505 |
| 1910 | 1 | 1,748 |
| 1880 | 29 | 3,071 |

Table 3: Conditioned on current income, coastal neighborhoods increase in income versus other neighborhoods

| | (1) | (2) | (3) | (4)$^{\dagger}$ |
|---|---|---|---|---|
| 1(Ocean or Gr. L. within 1km) | -0.009$^{b}$ | 0.023$^{c}$ | 0.024$^{c}$ | 0.049$^{c}$ |
| | (0.003) | (0.003) | (0.004) | (0.004) |
| Current %ile rank by income ($r_{i,t}$) | -0.162$^{c}$ | -0.192$^{c}$ | -0.180$^{c}$ | -0.188$^{c}$ |
| | (0.005) | (0.004) | (0.008) | (0.007) |
| Neighborhood log population density | | -0.027$^{c}$ | -0.027$^{c}$ | -0.027$^{c}$ |
| | | (0.002) | (0.002) | (0.002) |
| Metro log change in population | | | -0.104$^{c}$ | -0.090$^{c}$ |
| | | | (0.015) | (0.015) |
| Metro log change in land area | | | 0.017$^{c}$ | 0.031$^{c}$ |
| | | | (0.006) | (0.007) |
| Extra controls$^{\ddagger}$ | | | Yes | Yes |
| *Neighborhoods* | 298,778 | 298,778 | 282,581 | 282,581 |
| $R^2$ | 0.081 | 0.175 | 0.172 | 0.174 |

Each column displays estimates from a separate regression. Dependent variable is change in percentile rank by income ($\Delta r_{i,t}$); mean 0, standard deviation 0.16. Standard errors, clustered on metropolitan area, in parentheses; $^{b}$—$p<0.05$, $^{c}$—$p<0.01$. Regressions use observations of 61,047 neighborhoods in 308 metropolitan areas, 1880–2010. †—Explanatory variable in column (4) is indicator for an Ocean or Great Lake within 1km *and* neighborhood is in top income decile. ‡—Unreported control variable in (3) and (4) is mean housing unit age in years, which is unavailable for some neighborhoods and census years.

Table 4: Conditional effect of natural amenity on neighborhood change

| | (1) Ocean or Gr. L. | (2) Lake | (3) River | (4) Hills | (5) Moderate and dry | (6) Low flood risk | (7) High hedonic natural value |
|---|---|---|---|---|---|---|---|
| Indicator for natural feature | $0.024^{c}$ | $0.017^{b}$ | $-0.007^{b}$ | 0.012 | $0.027^{c}$ | 0.006 | $0.029^{c}$ |
| | (0.004) | (0.007) | (0.003) | (0.008) | (0.007) | (0.004) | (0.002) |
| Indicator for natural feature *and* top decile by income† | $0.049^{c}$ | $0.041^{c}$ | $0.026^{c}$ | $0.043^{c}$ | $0.050^{c}$ | $0.036^{c}$ | $0.052^{c}$ |
| | (0.004) | (0.012) | (0.004) | (0.006) | (0.011) | (0.003) | (0.003) |

Each cell displays estimates from a separate regression. Dependent variable is change in percentile rank by income ($\Delta r_{i,t}$); mean 0, standard deviation 0.16. Standard errors, clustered on metropolitan area, in parentheses; $^{b}$—$p$<0.05, $^{c}$—$p$<0.01. Explanatory variable is indicator for proximity within 1km in columns (1)–(3), average slope greater than 15 degrees in column (4), mean January minimum temperature between 0 and 18 degrees Celsius and mean July maximum temperature between 10 and 30 degrees Celsius and mean annual precipitation less than 800mm in column (5), mean annual flood probability less than 1% in column (6), and above-average natural value estimated using hedonic weights as described in the text in column (7). Regressions use 282,581 observations of 61,047 neighborhoods in 308 metros, 1880–2010, except column (6), which uses 238,455 observations of 49,780 neighborhoods in 261 metros with valid floodplain data. †—Explanatory variable in second row is indicator for natural amenity *and* neighborhood is in top income decile.

Table 5: Effect of within-metropolitan heterogeneity in proximity to oceans and Great Lakes on over-time neighborhood fluctuations in income

| | $[\mu]$ $[(\sigma)]$ | (1) | (2) | (3) | (4) |
|---|---|---|---|---|---|
| Metro indicator for proximity to coast | [0.33] | -0.424$^{b}$ (0.191) | -0.430$^{b}$ (0.181) | | |
| Within-metro s.d. in log nbhd. distance to coast | [0.42] [(0.57)] | | | -0.327$^{b}$ (0.144) | -0.334$^{b}$ (0.139) |
| Within-metro s.d. in nbhd. income ($1,000s) | [1.91] [(0.43)] | | -0.157 (0.162) | | -0.159 (0.162) |
| Metro log change in population, 1960–2010 | [0.94] [(0.78)] | | 1.106$^{c}$ (0.221) | | 1.100$^{c}$ (0.220) |
| Metro log change in land area, 1960–2010 | [1.61] [(1.15)] | | -0.539$^{c}$ (0.113) | | -0.544$^{c}$ (0.115) |
| $R^2$ | | 0.070 | 0.481 | 0.071 | 0.483 |

Each column displays estimates from the second level of separate two-level regressions. First-level OLS regressions (unreported) use 38,293 neighborhood observations in census year 1960 to estimate 135 metropolitan area means and cluster-robust standard errors. Dependent variable is over-time variance in percentile rank $\times 100$, 1960–2010; mean 2.13, standard deviation 2.69 in balanced panel of 38,293 neighborhoods over six census years. Second-level WLS regressions use 135 metropolitan areas. Dependent variable is estimated metropolitan area means from first level and weights are inverse estimated variance from first level. Robust standard errors in parentheses; $^{b}$—$p<0.05$, $^{c}$—$p<0.01$.

Table 6: Effect of within-metropolitan natural heterogeneity on over-time neighborhood fluctuations in income

| | | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
|---|---|---|---|---|---|---|---|---|---|
| | Base year: | 1880 | 1930 | 1940 | 1950 | 1960 | 1970 | 1980 | 1990 |
| | [$\mu$]<br>[($\sigma$)] | [4.7]<br>[(0.94)] | [3.0]<br>[(0.82)] | [2.6]<br>[(0.64)] | [2.2]<br>[(0.79)] | [2.3]<br>[(0.89)] | [2.0]<br>[(0.82)] | [1.4]<br>[(0.65)] | [1.3]<br>[(0.61)] |
| Within-metro s.d. in log nbhd. distance to coast | [0.42]<br>[(0.62)] | -0.024<br>(0.334) | 0.475<br>(0.576) | -0.141<br>(0.203) | -0.079<br>(0.166) | $-0.334^{b}$<br>(0.139) | $-0.177^{c}$<br>(0.061) | 0.007<br>(0.041) | 0.062<br>(0.039) |
| Within-metro s.d. in nbhd. average slope | [0.40]<br>[(0.22)] | 1.08<br>(0.873) | 1.39<br>(2.40) | -0.600<br>(0.453) | $-1.19^{a}$<br>(0.649) | $-1.97^{c}$<br>(0.414) | $-1.14^{c}$<br>(0.267) | $-0.372^{c}$<br>(0.122) | -0.220<br>(0.137) |
| Within-metro s.d. in nbhd. natural value | [0.08]<br>[(0.03)] | 6.92<br>(5.77) | 1.21<br>(13.2) | -4.06<br>(2.83) | $-7.64^{b}$<br>(3.34) | $-9.03^{c}$<br>(1.80) | $-5.93^{c}$<br>(0.821) | $-1.73^{c}$<br>(0.565) | -0.424<br>(0.716) |
| *Metropolitan areas* | | 29 | 9 | 43 | 51 | 135 | 227 | 277 | 308 |
| *Neighborhoods* | | 3,002 | 1,935 | 11,352 | 17,420 | 38,293 | 49,660 | 55,911 | 60,063 |
| *Census years*† | | 7 | 9 | 7 | 7 | 6 | 5 | 4 | 3 |

Each cell displays estimates from the second level of separate two-level regressions. Regression specifications are analogous to Table 5, column (4). First-level OLS regressions (unreported) use neighborhood observations in base year to estimate metropolitan area means and robust standard errors. Dependent variable is over-time variance in percentile rank $\times 100$, between base year and 2010; metropolitan-level means and standard deviations in first row. Second-level WLS regressions use metropolitan areas. Dependent variable is estimated metropolitan area means from first level and weights are inverse estimated variance from first level. Robust standard errors in parentheses; $^{a}$—$p<0.10$, $^{b}$—$p<0.05$, $^{c}$—$p<0.01$. Metropolitan-level means and standard deviations (in census year 2010) of explanatory variables in first column. †—For each base year, we balance our neighborhood panel to calculate over-time variances. Thus, historical base years have fewer than expected time periods, since we drop years with missing data.

# A Appendix

## A.1 Full Model

This section presents the full model which allows a city to have more than two neighborhoods and the amenity shocks $\epsilon_{j,t}$ to be correlated over time. The full model differs from the simple model presented in section 2.1 in the following ways. First, the city has $J \in \mathbb{N}$ neighborhoods and $J$ unit measure of workers. Second, the aggregate amenity shock $\epsilon_{j,t}$ follows an AR(1) process: $\epsilon_{j,t+1} = \rho\epsilon_{j,t} + \nu_t$ where $\nu_t$ is independent and identically distributed. Third, we extend the equilibrium selection rule in Section 2.3 as follows. When there are multiple equilibria, we choose the one that is closest to the selected equilibrium in the previous period, in terms of the Euclidean distance in the vector of average incomes (i.e., endogenous amenities) across neighborhoods.

### A.1.1 Equilibrium Within a Period

Lemma 1, which states that higher income workers sort into superior aggregate amenity neighborhoods, holds with the full model, because it is driven solely by workers' preferences.

In order to precisely describe the sorting with $J$ neighborhoods, we introduce new notation. First, we partition the set of worker incomes $[\underline{\theta}, \bar{\theta}]$ into $J$ intervals $\{\Theta_1, \Theta_2, ..., \Theta_J\}$ so that each group has a unit measure of workers; $\Theta_1$ is the top income group, and $\Theta_J$ is the bottom group. $\hat{\theta}_{i,i+1}$ denotes workers who divide group $i$ from group $i+1$; i.e., $\Theta_J \equiv [\underline{\theta}, \hat{\theta}_{J-1,J}]$, $\Theta_{J-1} \equiv [\hat{\theta}_{J-1,J}, \hat{\theta}_{J-2,J-1}]$, ..., $\Theta_1 \equiv [\hat{\theta}_{1,2}, \bar{\theta}]$. Second, we define a neighborhood rank function $r_t : J \to J$ such that $r_t(j)$ is the rank of neighborhood $j$ in terms of aggregate amenities in period $t$. For example, suppose that neighborhood 1 is the third best neighborhood in terms of aggregate amenities in period 2. Then we have $r_2(1) = 3$. Note that its inverse function $r^{-1}$ maps back to neighborhood index numbers, given aggregate amenity rankings. For example, suppose that the second best neighborhood in period 3 is neighborhood 4. Then we have $r_3^{-1}(2) = 4$.

Since each $\Theta_j$ group of workers consumes one unit size of land and each neighborhood has one unit size land, each $\Theta_j$ group of workers occupies one and only one neighborhood. Further, since higher-skill workers select into better aggregate amenity neighborhoods, each group $\Theta_j$ occupies neighborhood $r_t^{-1}(j)$ in each period $t$.

We characterized workers' location choices as a function of the distribution of aggregate amenities across neighborhoods. In turn, workers' location choices have to generate a distribution of aggregate amenities across neighborhoods that is consistent with their location choices. In other words,

$$a_{r_t^{-1}(j)} + E(w|\theta \in \Theta_j) + \varepsilon_{r_t^{-1}(j)} \geq a_{r_t^{-1}(j+1)} + E(w|\theta \in \Theta_{j+1}) + \varepsilon_{r_t^{-1}(j+1)}. \quad (13)$$

Proposition 2, which states that an equilibrium exists in each period and there can be multiple equilibria, holds with the full model. First, an equilibrium always exists because condition (13) is satisfied if higher income workers choose to live in neighborhoods with greater *exogenous* amenities: $a_j + \varepsilon_{j,t}$. Second, there can be multiple equilibria because condition (13) is satisfied for any matching pattern between income groups and neighborhoods if *exogenous* amenities (i.e., $\alpha_j + \epsilon_{j,t}$) are identical across all neighborhoods.

Now we characterize how rents are determined in each period. We normalize rent for the least favored neighborhood to be 0, i.e.

$$R_{r_t^{-1}(J)} = 0.$$

For the other neighborhoods, equilibrium rent $R_{r_t^{-1}(j)}$ is recursively determined so that $\hat{\theta}_{j,j+1}$ workers (i.e., workers who divide $\Theta_j$ and $\Theta_{j+1}$) are indifferent between neighborhood $r_t^{-1}(j)$ and $r_t^{-1}(j+1)$.

$$A_{r_t^{-1}(j)} \cdot (\hat{\theta}_{j,j+1} - R_{r^{-1}(j),t}) = A_{r_t^{-1}(j+1)} \cdot (\hat{\theta}_{j,j+1} - R_{r_t^{-1}(j+1)})$$

This equation recursively pins down rent for each neighborhood. Note that neighborhood rents follow the same order as average incomes, as with the simple model.

### A.1.2 Equilibrium Selection and History Dependence

When there are multiple equilibria, we choose the one that is closest to the selected equilibrium in the previous period, in terms of the Euclidean distance in the vector of average incomes (i.e., endogenous amenities) across neighborhoods.

Partly because the number of possible location-choice patterns increases dramatically with that of neighborhoods (i.e., $J!$ with $J$ neighborhoods), we cannot analytically prove Propositions 4, 5, and 6 with the full model. Instead, we use numerical methods to demonstrate that the results are robust with more than two neighborhoods and serially correlated amenity shocks.

For various combinations of parameters, we calculate the equilibrium path for 100,000 periods and test whether the Propositions hold. The following list of parameters are used in the simulations. For the number of neighborhoods $J$, we use 3, 5, and 7 neighborhoods. For the natural amenity distribution across neighborhoods, we use $\xi \times (1, 2, ..., J)$ and vary $\xi$ to be 1, 3, 5, and 10. Note that the variance in natural amenity levels increases as $\xi$ increases. For the average income distributions across neighborhoods, we use $\psi \times (1, 2, ..., J)$ and vary $\psi$ to be 1, 3, 5, and 10. For amenity shocks, we assume that $\epsilon_{j,t}$ follows an AR(1) process $\epsilon_{j,t} = \rho\epsilon_{j,t-1} + \nu_{j,t}$, where $\nu_{j,t}$ follows a Normal distribution $(0, \sigma^2)$. We vary $\rho$ to be 0, 0.2, 0.6, 0.9, 0.95, 0.98, 0.99, and 1, and vary $\sigma$ to be 1,3,5, and 10. $\rho$ determines how much the amenity shocks are correlated over time. Note that the amenity shocks are stationary if $\rho$ is less than 1. $\sigma$ determines how volatile the shocks are. This grid of parameters generates 1,536 unique combinations of parameters.

We begin with Proposition 4. For each combination of all parameters with $\rho < 1$ (1,344 combinations in total), we obtain a number $J \times 100,000$ of simulated neighborhood level data. For each combination with $\rho < 1$, we regress change in percentile rank income of a neighborhood on its percentile rank natural amenity level and its current period percentile rank income. This is the base specification we use in our empirical analysis.

Proposition 4.(i) implies that the coefficient on natural amenity should be positive. Our simulation results confirm this prediction with nonstationary amenity shocks. With $\rho \leq 0.98$, the coefficients were weakly positive for all 1,152 combinations. With $\rho = 0.99$, only four parameter combinations out of 192 show small negative values. The small number of negative outcomes seem to be driven by numerical errors, as it becomes close to nonstationary unit-root process. With a unit-root process (i.e. $\rho = 1$), our predictions do not hold. 128 out of 192 cases show negative values.

Proposition 4.(ii) predicts that the coefficient on current percentile rank should be negative. Our simulation results confirm this prediction. The coefficient estimates are weakly negative for all parameter combinations.

Now we test Proposition 5. The effect of superior natural amenities is captured by the coefficient on percentile rank natural amenity level in the previous regressions used to test Proposition 4. We test if the coefficients tend to increase with $\xi$. We calculate the mean value of the coefficient estimates for each $\xi$=1, 3, 5, 10. Each $\xi$ group has 384 parameter sets. The results show that the mean coefficient increase monotonically with $\xi$.

Finally, we test Proposition 6. We calculate $E(Var(\bar{w}_{j,t}|j))$ for each parameter set. The Proposition implies that $E(Var(\bar{w}_{j,t}|j))$ decreases with $\xi$, and our simulation results show that $E(Var(\bar{w}_{j,t}|j))$ indeed decreases with a stationary amenity shock (i.e., $\rho < 1$).

## A.2 Data

### A.2.1 Census Data and Boundary Normalization

We use 2010 census tract data from the American Community Survey (ACS) 5-year summary file, via the National Historical Geographic Information System (NHGIS) (Minnesota Population Center, 2011). These data cover the entire geographic extent of the U.S., though we focus on metropolitan (core-based statistical) areas only. The ACS is the annual replacement for the decennial long-form data, and it includes much of the detailed information on population and housing (e.g., income) that is no longer reported in the decennial census. However, the ACS has one important limitation. Because of small annual sample sizes and privacy concerns, these data represent five-year averages of residents and houses located in each tract. Thus, though we refer to these data as coming from the year 2010 throughout the paper, they really represent an average over 2006–2010. Finally, since these data already follow 2010 census tract boundaries, no normalization is required.

Census data for 1970–2000 are from the Geolytics Neighborhood Change Database (NCDB) (Tatian, 2003). These data are already normalized to 2000 (n.b., not 2010) census tract boundaries. The NCDB methodology compares maps of 2000 census tract and block boundaries to earlier years. Then, 1990 census block information (each tract is composed of many blocks) is used to determine the proportion of people in each historic tract that should be assigned to each overlapping 2000 tract. These proportions are then used as weights to normalize the data to 2000 boundaries.[15]

To normalize the NCDB data to 2010 census tract boundaries, we use the Longitudinal Tract Database (LTDB) (Logan, Xu, and Stults, 2012). The LTDB uses the same block-weighting methodology as the NCDB. Thus, our analysis uses weights defined by 2000 census block populations to normalize all of the Geolytics data from 1970–2000 to 2010 census tract boundaries. It is important to note that in 1980 and earlier, the entire geographic extent of the U.S. was not completely organized into tracts, and missing data problems are more severe for earlier years. However, since we focus mostly on metropolitan areas, data quality is quite good as early as 1970. (We also drop tract observations in years when their respective metropolitan area is incompletely tracted. See more on sample selection below.)

For census years 1910–1960, we use decennial census information from the NHGIS. The 1940, 1950, and 1960 NHGIS extracts are collectively known as the Bogue files (2000a, 2000b, and 2000c), and they are also available from the Inter-University Consortium for Political and Social Research. These files contain tract information for selected cities and metropolitan areas. The 1910, 1920, and 1930 NHGIS extracts are known as the Beveridge files. Note that data availability is sparse, especially before 1950. Even for cities that are completely tracted, sometimes the data do not contain complete information on population, housing, or income. (For example, in 1910, tract information on household income is only available for New York City; in 1920, such information is only available for New York City and Chicago. Ten metropolitan areas have valid data in 1930, and 43 metropolitan areas have valid data in 1940.) We normalize these data to 2010 census tract boundaries ourselves using NHGIS map layers. For each decade, we compare historical tract boundaries to 2010 census tract boundaries. Since sub-tract or block information on population is unavailable for these historical years, we are unable to exactly follow the NCDB and LTDB methodologies of constructing weights using block populations. Instead, we normalize using a simple apportionment based on land area.

Finally, we draw 1880 census information from the Integrated Public Use Mi-

[15] We make a small adjustment to the 1980 Geolytics NCDB. The 1980 census prized identification of "places" (e.g., towns, villages, boroughs) over tracts when confidentiality restrictions were binding. The NCDB propagates this censoring in their normalization procedure, even if the proportion of households in the tract with suppressed income data is negligible. We restore this income information from the original 1980 census as long as the proportion of censored households in a census tract is less than 20 percent.

crodata Series (IPUMS) (Ruggles et al., 2010). We use both the 100 percent census and the 10 percent population sample; the 10 percent sample includes information on literacy while the 100 percent census does not. The IPUMS includes data on each person's place of residence, via the enumeration district variable. Enumeration districts were areas assigned to census enumerators to gather data, and they are comparable in population size to modern-day census tracts. (In fact, on average, they are slightly smaller than modern-day census tracts.) We use enumeration districts to normalize the historical 1880 data to 2010 census tract boundaries. First, we obtain maps on historical enumeration district boundaries from the Urban Transition Historical GIS Project (UTHGIS) (Logan et al., 2011). Maps are available for 32 present-day metropolitan areas (totaling 29 consolidated metropolitan areas). Second, using the same procedure as for 1910–1960, we compare historical enumeration district boundaries to 2010 census tract boundaries. We apportion to 2010 census tract boundaries using land area.

### A.2.2 Natural Amenity Data

We spatially match our consistent-boundary neighborhoods to a number of natural and persistent geographic features.

*Water features—coastlines, lakes, and rivers.* We use data on water features from the National Oceanic and Atmospheric Administration's Coastal Geospatial Data Project. These data consist of high-resolution maps covering (i) coastlines (including the Atlantic, Pacific, Gulf of Mexico, and Great Lakes), (ii) other lakes, and (iii) major rivers. For each 2010 census tract, we separately calculate the distance to each of the nearest water features (ocean, lake, river) from the centroid of the tract.

*Elevation and slope.* We use the elevation map included in the ESRI 8 package. These data have a 90-meter resolution. In ArcGIS, we use the slope geoprocessing tool to generate a slope map. Then, we use the zonal statistics tool to calculate average slope in each 2010 census tract.

*Floodplains.* The Federal Emergency Management Agency (FEMA) publishes National Flood Hazard Layer (NFHL) maps covering much of the U.S. The NFHL shows areas subject to FEMA's flood zone designations. We assign to tracts either a high-risk or low-risk indicator. High risk means that an area has at least a 1 percent annual chance of flooding (a 26 percent chance of flooding over a 30-year period), as determined by FEMA. Note that flood maps are unavailable for some metropolitan areas. In our data, 261 metropolitan areas have valid flood zone information.

*Temperature and precipitation.* We match tracts to temperature and rainfall data available from the PRISM Climate Group at Oregon State University. These data are 1971-2000 averages, collected at thousands of weather monitoring stations and processed at a spatial resolution of 30-arcseconds for the entire spatial extent of the U.S., of annual precipitation, July maximum temperature, and January minimum temperature.

### A.2.3 Sample Selection

We drop tracts in Alaska, Hawaii, and Puerto Rico. We exclude tracts with zero land area (these are typically "at sea" populations, i.e., personnel on ships) or zero population (e.g., airports or zones otherwise reserved for nonresidential uses).

We do not consider tracts outside of metropolitan areas defined in 2009. One problem with nonmetropolitan tracts it that many of them are not available before 1990, the first year that the U.S. was fully organized into census tracts. Another problem with rural tracts is the difficulty in grouping these tracts into units that share common labor, housing, product, and input markets. (One exception are the core-based statistical areas called micropolitan areas. However, many of these micropolitan areas feature a very small number of tracts, making them unsuitable for our analysis. The very small number of tracts means that the entry of even one new neighborhood can elicit a volatile response in within-micropolitan area rankings.)

We drop tracts in particular years that are clearly nonurban. This restriction is more salient in historical years, when tracts or enumeration districts on the urban fringe were not subject to urban land uses. We classify tracts as nonurban if (i) the entire tract population is classified by the census as "rural" or (ii) population density is less than 32 people per square mile, or 1 person per 20 acres. (Lowering this threshold to 1 person per 40–160 acres affects the number of excluded tracts minimally. Population densities of less than 32 people per square mile are already well short of standard definitions of urban population densities.) We reason that while these tracts are within counties that contain urban uses, at the time of observation they are likely to be outside of metropolitan areas and urban household location decisions. In this way, we also address concerns about changing metropolitan area boundaries over time.

We exclude tracts where our normalization procedure is likely to be poor. In some cases, especially for early census years and tracts on the urban fringe, historical tracts cover only a portion of 2010 census tract areas. This is more likely to be the case when historical city boundaries were much smaller than present-day extents. When historical tracts cover less than 50 percent of the land area of the present-day tract, we exclude these data from our analysis.

We also eliminate tract observations that disappear from one year to the next. This problem is partly mechanical; we cannot compute income changes for a tract that does not appear in the next period. It also is mostly limited to the transition between the 1880 UTHGIS data and the subsequent NHGIS data. The reason this problem arises is because the UTHGIS maps, which we use for our normalization procedure, typically cover entire counties, whereas the NHGIS data and maps used in the early 20th century are confined mostly to city boundaries. Thus, many of the UTHGIS tracts outside city boundaries are dropped anyway because they are nonurban (see above), but, to avoid the problem of contracting metropolitan boundaries, we exclude the remaining earlier tracts that do not appear in subsequent years.

A consequence of the unbalanced nature of the data is that forward lags vary by metropolitan area and year. For example, after 1880, it is only 30 years until our next observation of New York neighborhoods (in 1910), but it is 70 years until our next observation of Omaha neighborhoods (in 1950). Out of 1,684 metropolitan area-year groups in our data, 1,342 follow the standard 10-year gap between census year observations. As a result, the actual number of neighborhoods used in regressions varies according to whether the specification requires balancing across two subsequent census years or balancing over a large number of years. In addition, some variables, such as flood hazard or average housing unit age, are unavailable in some years, further affecting sample selection.

### A.2.4 Summary Statistics

Figure A2 shows the evolution of several New York neighborhoods over our sample period. Recall that each neighborhood corresponds to data normalized to 2010 census tract boundaries. The solid lines show the relative rankings of three neighborhoods—tracts corresponding to the Upper East Side, East Harlem, and Tribeca. (Levittown was unpopulated in 1880 and a corresponding solid line does not appear in the figure.) An interesting feature of this graph is variation in income dynamics across neighborhoods. For example, the Upper East Side has remained a high-income neighborhood throughout our sample period. East Harlem, which was a relatively high-income neighborhood in 1880, experienced decline and has been a low-income neighborhood since 1910. Tribeca saw a large increase in average household income in the 1980s.

The dotted lines show the relative rankings of these three neighborhoods after 1960 and also the relative ranking of a fourth neighborhood, Levittown, which first appeared in that census year. In comparing the solid to the dotted lines, note that we have changed the universe used to compute neighborhood ranks from 1880 to 1960 neighborhoods, but the dynamic patterns for the extant three neighborhoods remain qualitatively similar.

In our sample, most neighborhoods experience changes in percentile rank that are close to zero—that is, neighborhood income ranks are largely persistent over time, especially over the ten-year changes that are predominant in our sample. Few neighborhoods experience dramatic increases or declines in rank. The distribution of percentile rank changes has a mean zero and standard deviation of 0.164.

Finally, in Figure A3, we show a pattern of mean reversion in our sample data. The graph shows the results from a local polynomial regression of change in tract percentile rank on initial tract percentile rank. High-income neighborhoods—

neighborhoods with initial percentile ranks greater than 0.5—tend to decline in rank over the subsequent period. Low-income neighborhoods tend to improve in rank. One feature of interest is that despite using nonlinear techniques, the pattern of mean reversion is close to linear in initial rank. Later, we condition on initial rank linearly in our regressions.

### A.2.5 Robustness Checks

In Appendix Figure A4, we show that the effect of proximity to an ocean or Great Lake is consistent for definitions of proximity from 100 meters to 10 kilometers. The gray line connects regression estimates from 100 separate regressions of neighborhood change on proximity, varying the proximity indicator. (As we move to the right along the horizontal axis, each regression classifies more neighborhoods as being "close" to the natural amenity and fewer neighborhoods as being "far" from the natural amenity.) The black line displays regression estimates when we use only natural features near top-income neighborhoods, as in Table 3, column 4. (The intersection of these lines with the vertical dotted line are the estimates from our baseline estimates using a 1-kilometer definition, shown in Table 3.) Recall that we expect these features to be more likely to be positive, versus negative, amenities. As expected, the results using this variable are always stronger than the results using oceans and Great Lakes unconditioned on income.

Figure A4 also shows the same results for lakes, rivers, and hills. These results show consistent patterns. The important feature of this figure is that it shows that conditioning rivers on their proximity to high-income neighborhoods improves the estimated effect of (positive-amenity) rivers on neighborhood change. This is consistent with the view that, on average, rivers in our sample are a disamenity for households.

Our results are also robust to varying the time horizon over which we calculate rank changes (i.e., differences of 10, 20, 30 or more years) or estimating our regressions for each year separately, as shown in Appendix Figure A5. In this figure, each point shows the estimated conditional effect from a separate regression that varies the base year $t$ and the time horizon $\Delta t$. The vertical axis measures the estimated conditional effect of superior natural amenities conditioned on initial proximity to high-income neighborhoods (as in Table 3, column 4). The dependent variable is the change in percentile rank by income from $t$ to $t + \Delta t$. The initial year $t$ used in the regression can be read by tracing the line segment back to the beginning year, and the outcome year $t + \Delta t$ is measured on the horizontal axis. With few exceptions, the estimated conditional effects are positive and precisely estimated.

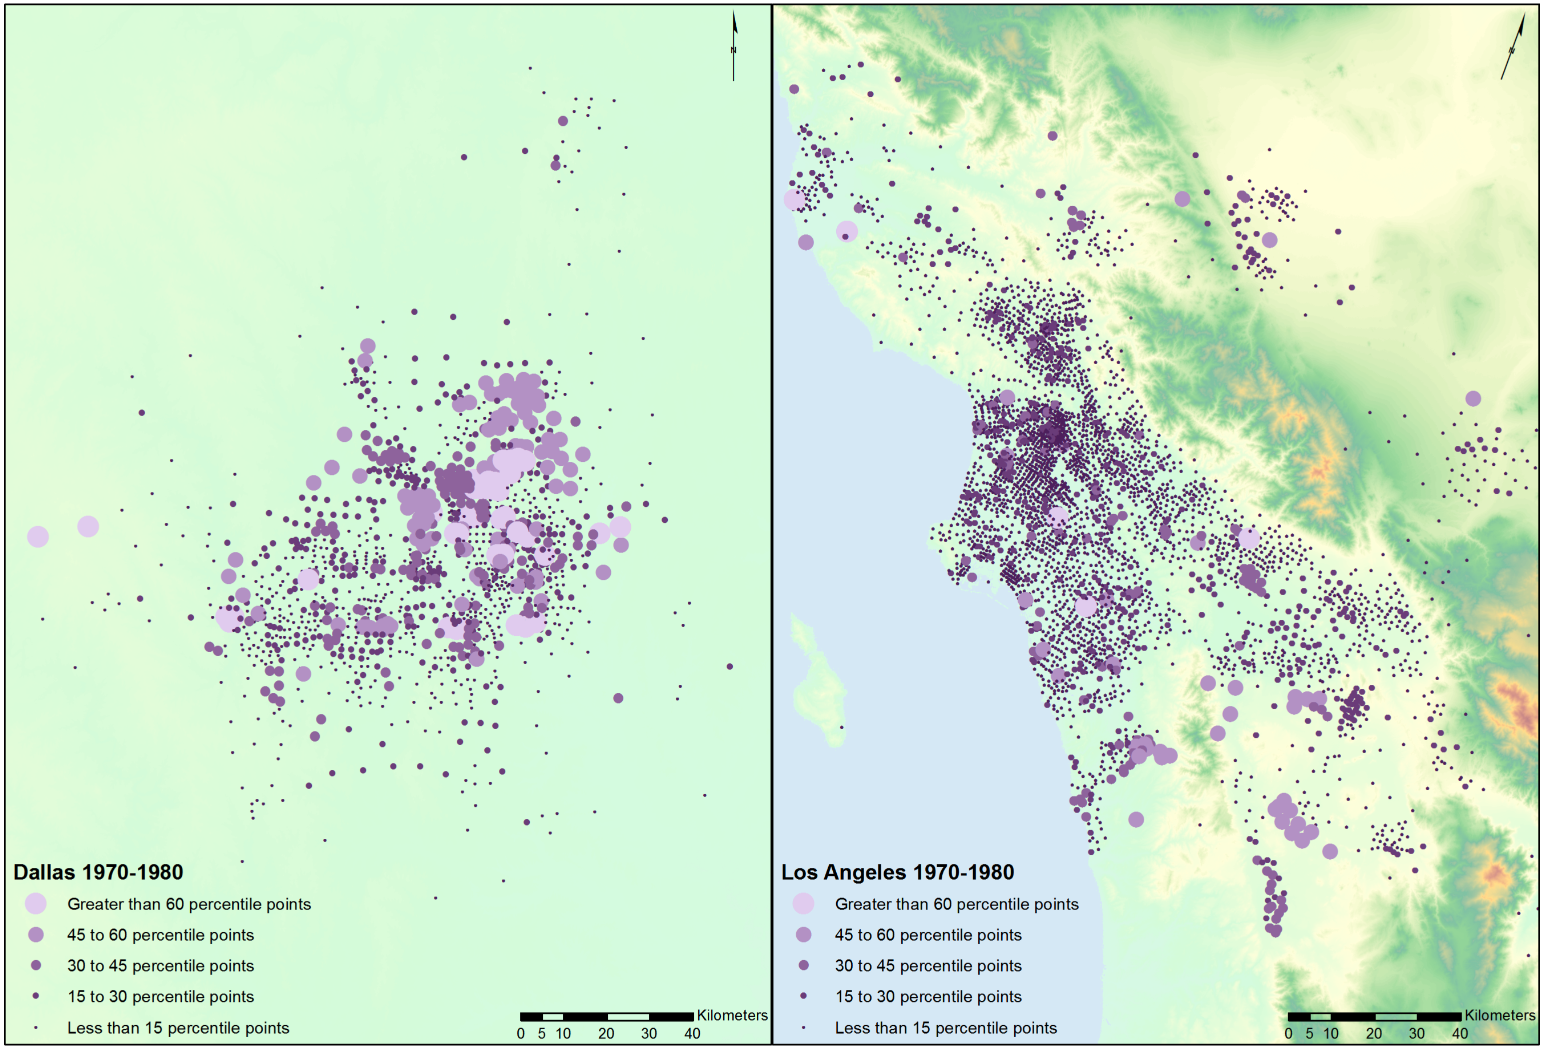


Figure A1: Churning in the spatial distribution of income

These maps show 1970 neighborhoods in the Dallas and Los Angeles metropolitan areas as dots. Dots are sized and colored according to the absolute value of change 1970–1980 in each neighborhood's percentile ranking within the city by average household income. The average absolute change for neighborhoods in Dallas was 21 percentile points; for neighborhoods in Los Angeles, it was 9 percentile points.

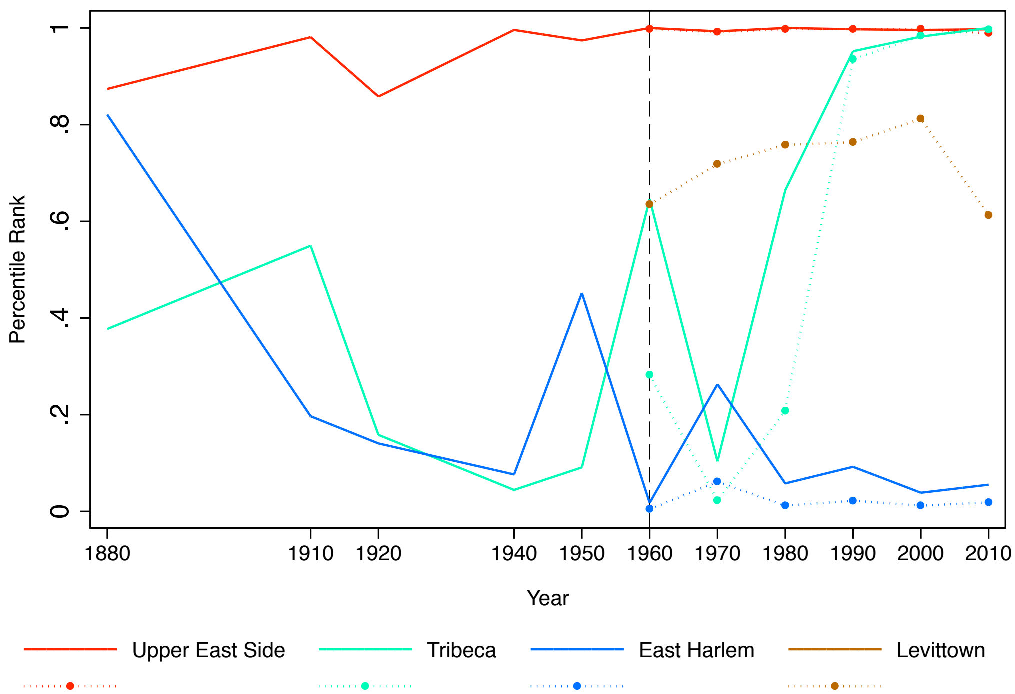


Figure A2: New York metropolitan area rankings over time, selected neighborhoods

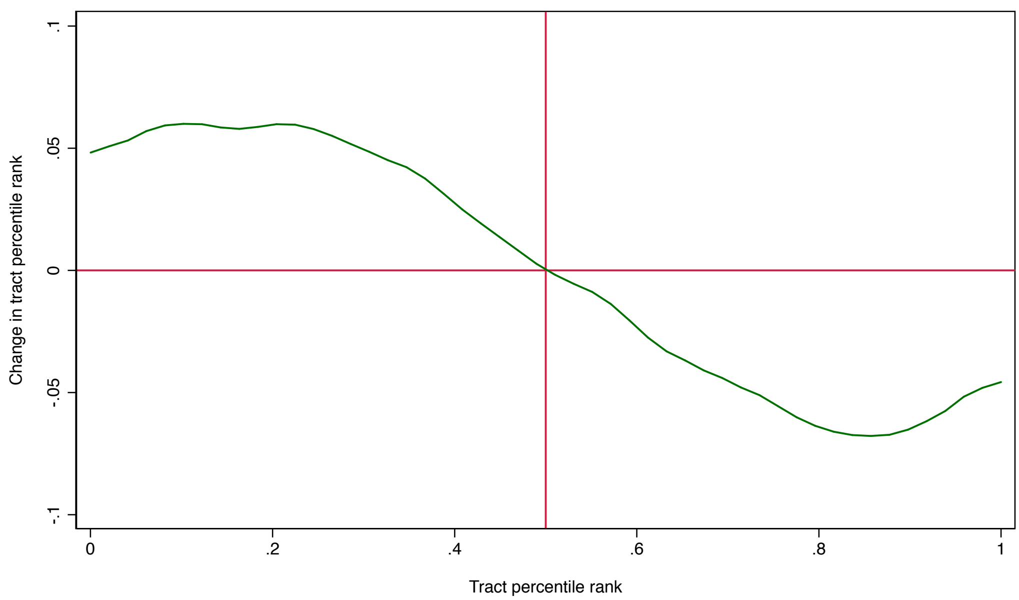


Figure A3: Mean reversion in neighborhood percentile rank

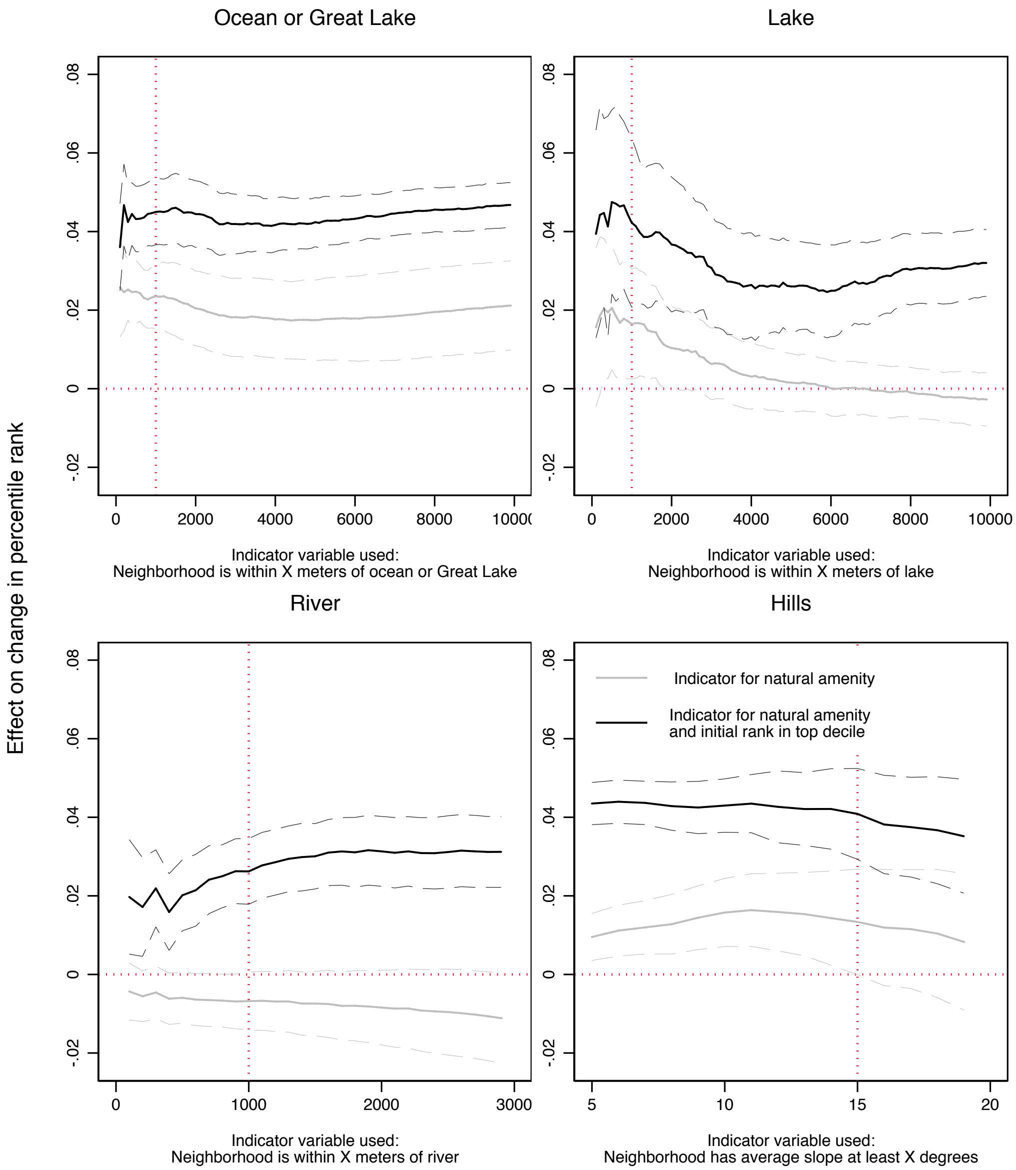


Figure A4: Robustness to indicator variable thresholds

These graphs show the conditional effect of natural features on neighborhood change for varying indicator definitions of proximity to natural features.

A. Ocean or Great Lake within 1km

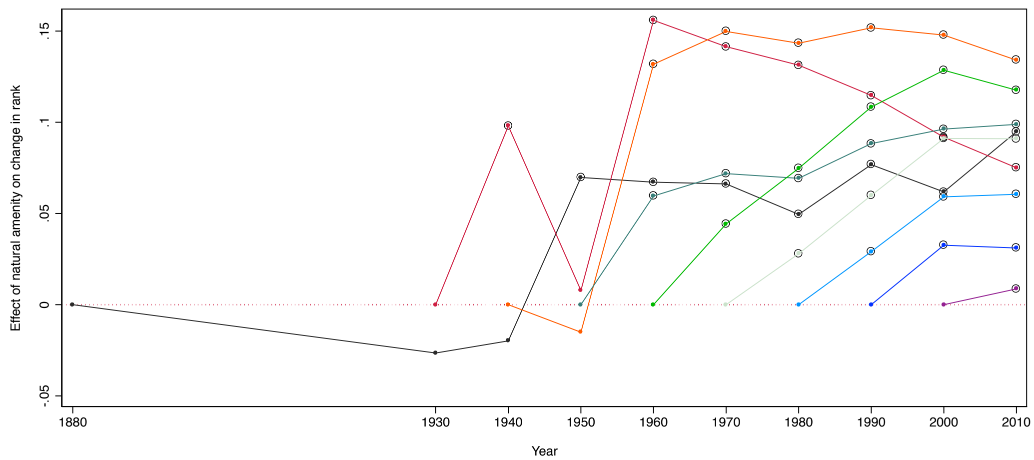


B. Average slope greater than 15 degrees

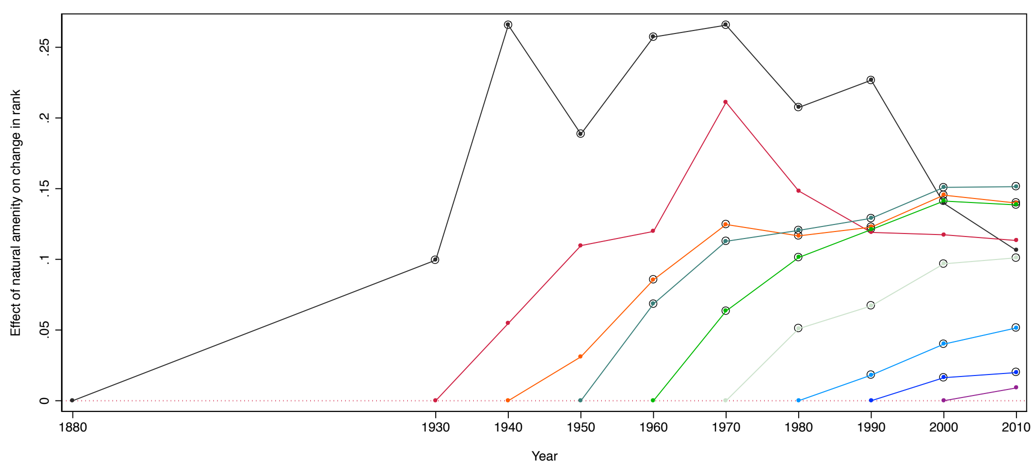


Figure A5: Robustness to year selection and time horizon

Each point shows results from a separate regression, that varies the base year $t$ and the time horizon $\Delta t$. The vertical axis measures the estimated conditional effect of indicator for natural amenity *and* neighborhood is in top income decile. The dependent variable is change in percentile rank by income from $t$ (the beginning year of its corresponding line segment) to $t + \Delta t$ (the year corresponding to the horizontal coordinate of the point). Circled points are significant at $p < 0.05$.

Table A1: Summary statistics

| | $\mu$ | $(\sigma)$ |
|---|---|---|
| *Neighborhoods* | | |
| $\Delta r$, Ten-year change in percentile rank | 0.00 | (0.16) |
| $\text{Var}(r_{i[m]}) \times 100$, Variance in 1960–2010 percentile rank | 2.13 | (2.69) |
| | | |
| Population, 2010 | 4,283 | (1,912) |
| Land area (km$^2$) | 27.5 | (73.5) |
| Persons per square km, 1880 | 5,940 | (12,406) |
| Persons per square km, 1960 | 2,901 | (6,159) |
| Persons per square km, 2010 | 2,335 | (4,807) |
| Housing units, 2010 | 1,796 | (786.5) |
| Mean age of housing units (years), 2010 | 37.3 | (14.1) |
| Distance from centroid to city center (km), 2010 | 29.9 | (27.2) |
| | | |
| *Share of 2010 neighborhoods* | | |
| ... with centroid within 1km of ocean or Great Lake | 0.092 | |
| ... with centroid within 1km of lake (ex. Great Lakes) | 0.012 | |
| ... with centroid within 1km of major river | 0.210 | |
| ... with average slope greater than 15 degrees | 0.069 | |
| ... with moderate temperatures* | 0.091 | |
| ... with less than 800mm average annual precipitation | 0.063 | |
| ... with less than 1% average annual flood risk† | 0.454 | |
| | | |
| *Metropolitan areas* | | |
| $\overline{\text{Var}(r_m)} \times 100$, Mean variance in 1960–2010 percentile rank | 2.3 | (0.89) |

*–Average January minimum temperatures between 0 and 18 degrees Celsius and average July maximum temperatures between 10 and 30 degrees Celsius. †–Flood information available for 49,517 neighborhoods.

Table A2: Coastal and interior cities

| Coastal Cities | Interior Cities |
|---|---|
| Boston, MA | Albany, NY |
| Buffalo, NY | Atlanta, GA |
| Charleston, SC | Cincinnati, OH |
| Chicago, IL | Columbus, OH |
| Cleveland, OH | Hartford, CT |
| Detroit, MI | Indianapolis, IN |
| Milwaukee, WI | Kansas City, MO |
| Mobile, AL | Louisville, KY |
| New Orleans, LA | Memphis, TN |
| New York, NY | Minneapolis, MN |
| Rochester, NY | Nashville, TN |
| San Francisco, CA | Omaha, NE |
| | Philadelphia, PA |
| | Richmond, VA |
| | Pittsburgh, PA |
| | St. Louis, MO |
| | Washington, DC |

These are the principal cities for consolidated metropolitan areas that have available data on neighborhood income in the census years shown in Figure 2.