Re-measuring gentrification

Urban Studies
XX(X):1–17

Re-Measuring Gentrification

devin michelle bunten[1], Benjamin Preis[1], and Shifrah Aron-Dine[2]

Abstract

We develop an expectations-based measure of gentrification. Property values today incorporate market participants' expectations of the neighborhood's future. We contrast this with present-oriented variables like demographics. To operationalize the signal implicit in property values, we contrast the percentile rank of a neighborhood's average house price to that of its average income, relative to its metropolitan area. When a neighborhood's house value percentile begins to rise above its income percentile, that is a signal of gentrification. We show that a gap between the house value and income percentiles predicts future income growth. We further validate our metric against existing approaches to identify gentrification, finding that it aligns meaningfully with qualitative analyses built on local insight. Compared to existing quantitative approaches, we obtain similar results but usually observe them in earlier years and with more parsimonious data. Our approach has several advantages: conceptual simplicity, communicative flexibility with graphical and map forms, and availability for small geographies on an annual basis with minimal lag.

Keywords

gentrification, neighborhood change, measurement, urban economic theory

Introduction

Gentrification scholarship is characterized by debate on its definition, causes, consequences—and measurement (Brown-Saracino 2013). Davidson and Lees (2005) argue that gentrification consists of capital reinvestment, "social upgrading" as high-earners arrive, landscape change, and displacement of low-income groups. Even with this conceptual clarity, Finio reports over 100 quantitative measures of gentrification in the literature, collectively utilizing over 3 dozen variables in combinations that are "often vague or arbitrary" (Finio 2021, 261). Finio follows a common dichotomy by classifying input variables as pertaining to either demand (e.g., income) or supply (e.g., tenure), and argues that measures should include both. Nevertheless, theoretically-grounded measures composed of similar variables have been shown empirically to produce very different classifications when applied to the same city in the same time period (Preis et al. 2020).

We intervene in these debates by proposing a classification of candidate measurement variables based on whether they reflect *expectations* of a neighborhood's future, or if they instead reflect its present conditions. Present-oriented variables are tethered to the current status of the neighborhood: e.g., incomes do not rise today because of expectations the rich will arrive tomorrow. Conversely, expectations-based variables respond to anticipated changes. For example, property values this year reflect anticipated changes next year: a future influx of the wealthy will raise resale values, and property purchasers who expect this will raise their willingness to pay now. Expectations-based variables include physical capital investment and city plans; along with property values, they are all generated through processes incorporating actors' assessments of the neighborhood's future.

We construct an expectations-based signal of gentrification by contrasting variables that *do* reflect expectations to variables that *don't*. Using insights from asset valuation theory (Fisher 1906), we show that property values are expectations-based: prices are generated by transactions involving market participants who make and apply assessments of the neighborhood's future when transacting. Accordingly, property values may rise in response to expectations of the four components of gentrification identified by Davidson and Lees (2005)—even *before* those components take hold. We operationalize property values using house prices, and we use income as our present-oriented variable. These choices are contextual and practical: in the US, house value and income data are annually available for small geographies. We convert each neighborhood's house value and income levels into percentile-ranks relative to its metropolitan statistical area (MSA). The most expensive and high-income neighborhoods of a given city will take values just under 1.0, while low-price and -income neighborhoods will take values close to 0. A sizeable gap between a neighborhood's house value and income percentiles is our empirical signal of gentrification.[1]

To test the strength of this signal, we study its relationship to income growth in gentrifiable US neighborhoods. The opening of a 25-percentile gap is associated with rising incomes within three years, and a 5% faster increase in neighborhood real income ten years later, after controlling for baseline socioeconomic and geographic characteristics.

[1]Department of Urban Studies and Planning, Massachusetts Institute of Technology, USA
[2]Department of Economics, Stanford University, USA

Corresponding author:
devin michelle bunten
Department of Urban Studies and Planning, MIT
77 Massachusetts Avenue, Cambridge, MA 02139, USA.
Email: bunten@mit.edu

The effect is larger in neighborhoods with more Black residents, those closer to downtown, and those that gained more housing units.

We validate the signal using qualitative and quantitative understandings of gentrification developed by researchers across four cities. We compare the house value and income percentile-ranks for Boston and Chicago neighborhoods to findings from qualitative studies. Next, we compare the percentile-ranks for Portland neighborhoods to a quantitative approach that uses a broader base of measurement inputs, and to a prospective approach implemented by Los Angeles planners to detect displacement threats. Our approach maps well onto the qualitative research while capturing many of the same patterns as existing quantitative approaches—in many cases, *before* alternative approaches, emphasizing the value of an expectations-based approach. Across these comparisons, we visualize our signal in three ways: charting percentile-ranks over time, mapping the gap across space, and mapping the year a gap first crossed a threshold, illustrating how policymakers and researchers can use the signal in their work.

Understanding Competing Conceptions of Gentrification

Our paper contributes to several conversations in the gentrification literature. In line with our empirical setting, we concentrate our discussion on US-based literature. First, we contribute to a long-running literature on quantitative measurement of gentrification spanning disciplines including geography (Hammel and Wyly 1996), planning (Freeman 2005), economics (Ellen and O'Regan 2011), and sociology (Rucks-Ahidiana 2021). Researchers in this tradition typically study gentrification by (1) identifying gentrifiable neighborhoods and (2) diagnosing a treated subset as gentrifying using changes in demographic and housing market characteristics, often using Census data. However, minor differences in variable selection can lead to substantial differences in the set of neighborhoods identified as gentrifying. A parallel literature—channeling Beauregard (1986) and Galster and Peacock (1986)—has troubled these approaches (Barton 2016; Finio 2021; Preis et al. 2020).[2] Academically, this diagnostic instability amounts to uncertainty in whether a neighborhood should be in the treatment or the control group. Practically, it limits planners' ability to tailor anti-gentrification policies, as well as their ability to learn *from* academic research.

We contribute to this literature by conceptualizing some variables as *expectations-based* and contrasting these with *present-oriented* variables to construct an expectations-based signal of gentrification. We use insights from asset valuation theory (Fisher 1906) to argue property values are expectations-based: they are generated by actors with knowledge of (or plans for) a neighborhood's future. Present-oriented variables reflect current conditions. We focus on house prices and incomes, which are widely available at high temporal frequency for small geographic areas in the US. Our signal is intuitive and enables easy communication between academics, planners, and other city residents. These features address several of Finio's (2021) criteria for better metrics. We term our measure a *signal* because we do not seek to overturn existing definitions of gentrification as such; instead, we offer an indicator the process is occurring.

Second, we address a literature investigating quantitative and qualitative assessments of gentrification (Brown-Saracino 2017) and connecting these insights (Easton et al. 2019; Goetz et al. 2019). Our expectations-based approach to measurement incorporates some of the insights from this strand of the literature by distinguishing variables grounded in the practices and beliefs of gentrifiers, sellers, and locally-informed market participants. By identifying how property values encode local knowledge into housing transactions, we are able to incorporate some local knowledge from essentially every neighborhood. In line with Brown-Saracino (2016) and Goetz et al. (2019), we validate our measure against the findings of qualitative research in Boston and Chicago.

More recent literature has developed novel approaches to gentrification identification. One thread uses data from technology platforms to identify ("nowcast") gentrification (Chapple et al. 2022; Glaeser et al. 2018; Jain et al. 2021). An overlapping thread applies machine learning and other statistical methods to both traditional and novel data, often training or baselining the models using traditional data (Jain et al. 2021; Liu et al. 2019; Reades et al. 2018). Our approach likewise uses data published frequently in near real time. Our signal doesn't depend on specific technology platforms and user behaviors, nor on machine learning that may not capture evolving (and out-of-sample) modes of gentrification. It has intuitive graphical representations, an aid for academic and practical communication.

An Expectations-Based Signal of Gentrification

Property values are expectations-based. Asset valuation theory, formalized by Fisher (1906), identifies the value of an asset with the appropriately-discounted stream of future income it produces.[3] Referring to the net income—or *returns*—earned in period $t+n$ as r_{t+n}, and the discount factor as $\delta < 1$, the *net present value* of an asset today is $NPV_t = r_t + \sum_{n=1}^{\infty} \delta^n r_{t+n}$, where the second term is the (potentially) infinite sum of future income generated by the asset. For predictable assets—like annuities, where the payouts are known in advance—the price of the asset, p_t, should approximate NPV_t. Because returns to property aren't known with certainty, we write $\mathbb{E}[r_{t+n}]$ to denote *expected* returns. Returns may take different forms. For landlords, returns are rent minus costs. For an owner-occupier, the returns may constitute the use-value of the property, including access to local services and amenities. Our key theoretical equation is

$$p_t = r_t + \sum_{n=1}^{\infty} \delta^n \mathbb{E}[r_{t+n}] \tag{1}$$

This equation captures the key insight we pull from asset valuation theory. If market participants begin expecting gentrification in several years, then landlords will anticipate being able to raise future asking rents by more than otherwise, while owner-occupiers anticipate an increase in their resale value.[4] They will incorporate this into their willingness to pay, causing expectations of gentrification to *directly increase* the value of property today, and the price at which informed market participants expect it to transact.

This insight rests on a few features of market participants. First, participants must have (some, imperfect) knowledge about the gentrification status and trajectory of the neighborhood. Qualitative studies identify the quite detailed insights residents have about change within their neighborhood. This knowledge may take a spatial form, with a tacit understanding of which neighborhood is *next* based on proximity to past gentrification and proximity to natural amenities or transportation infrastructure (e.g., Brown-Saracino 2009, 58). Second, participant knowledge, however imperfect, needs to inform their actions (again, Brown-Saracino 2009, 58). Third, it is not necessary that gentrification expectations be the key determinant of property values, only that they are sufficient to alter the value of property, all else equal.[5]

Given the US context, we use neighborhood-level house values as our expectations-based variable and income as our present-oriented variable; both are available annually for small geographies. The signal could be improved by incorporating additional expectations-based variables beyond house values, or it could be modified to adapt to data availability in other countries—such as multifamily property values, physical investment, property tax valuation, or comprehensive plans—and additional present-oriented measures like race or housing conditions.

To compare these variables, we convert neighborhood-level average house values and incomes into relative percentile-ranks within a metropolitan area: the neighborhood with the highest house prices will be around 1, and that with the lowest will be around 0. In general, the percentile-rank of house prices and income are highly correlated, with the most expensive places also among the richest. For gentrifying neighborhoods, we expect the house price percentile to be greater than the income percentile.

Figure 1 presents a stylized version of house value (solid) and income (dashed) percentile-ranks across three hypothetical neighborhoods: one rich (blue), one poor (green), and one that gentrifies during the period (orange). The blue lines cluster at the top of the city's distribution for both house value and income percentiles while the green lines cluster near the bottom. The orange lines begin near the bottom but, over time, market participants begin to expect gentrification. House values rise first, with the newly-opened gap between signaling expectations of future gentrification.[6] After several years, expectations become reality, and incomes rise too. Eventually, house values levels off at a high level, while income growth continues.

Measuring Gentrification: Relative Income and House Prices

In this section, we summarize construction of our signal and the dynamic difference-in-difference regression model.[7] For the signal, we use the smallest geographies—the census tract and ZIP code — with relatively frequent data releases. We use home price data from the Federal Housing Finance Agency (FHFA) House Price Index and income data Internal Revenue Service (IRS) Statistics of Income. Given limitations of FHFA, we also construct our measure using the Zillow Home Value Index (ZHVI). We rely on Census and American Community Survey (ACS) data, provided by the National Historical Geographic Information System (NHGIS) database (Manson, Steven et al. 2021). We examine historical gentrification using reweighted census data from Lee and Lin (2018) for census years from 1940–onward, harmonized to 2010 census tract boundaries.

FHFA provides an annual estimate of *changes* to single-family house values relative to the prior year—it does not provide absolute values. At the tract and ZIP level, we reconstruct values for each geography in each year by multiplying the relative changes from FHFA by the median house value from the 2000 Census. In some contexts, it is more appropriate to use Zillow's ZHVI at the ZIP level. The ZHVI provides black-box estimates of "typical" house prices at the ZIP level, inclusive of single-family, condo, and co-op typologies. We observe house prices from 1990 through 2020.

The IRS reports average income for each ZIP annually for 1998–2018. When doing analysis at the spatial geography of the ZIP, we use the average household income constructed directly from the IRS data. At the tract level, we take the rate of year-over-year income change from the IRS for the ZIP in which the tract is located and multiply it by the median household income from the 2000 Census to estimate tract-level income for 1998-2018.

Next, we calculate the percentile-rank of each neighborhood's income and house values vis-á-vis the distribution of income and house value within its MSA. For every tract and ZIP within an MSA, we calculate the percentile rank of the house value, weighted by housing unit counts. For income, we weight by population.

To conduct statistical analyses, we construct a binary measure of gentrification for 1998–2018. Gentrifiable neighborhoods are those in the inner third of an MSA based on distance to the central business district (CBD) with an initial income below the 25th percentile. These thresholds follow the gentrification literature in defining "gentrifiable" areas (Finio 2021). A neighborhood is classified as gentrifying when a 25-percentile gap opens between the house values and incomes. This threshold balances competing risks: a low threshold could be triggered by short-term fluctuations; a high threshold may never be triggered—especially in neighborhoods with little rental or social housing, where rising prices directly exclude new low-income buyers. In the qualitative comparison section, we highlight these tradeoffs by exploring alternative thresholds.

To test whether a gap predicts future income growth in a neighborhood, we use a dynamic difference-in-difference design (Sun and Abraham 2021).[8] We ask whether a house price/income gap in central low-income neighborhoods is associated with future income growth, and how income growth depends on neighborhood characteristics. Because we wish to include neighborhoods with few single-family homes, we use ZHVI data at the ZIP code level. Out of 3,329 centrally-located ZIPs, we identify 213 newly-gentrifying ZIPs between 1999 and 2018, and 97 already-gentrifying ZIPs. Our baseline estimating equation is

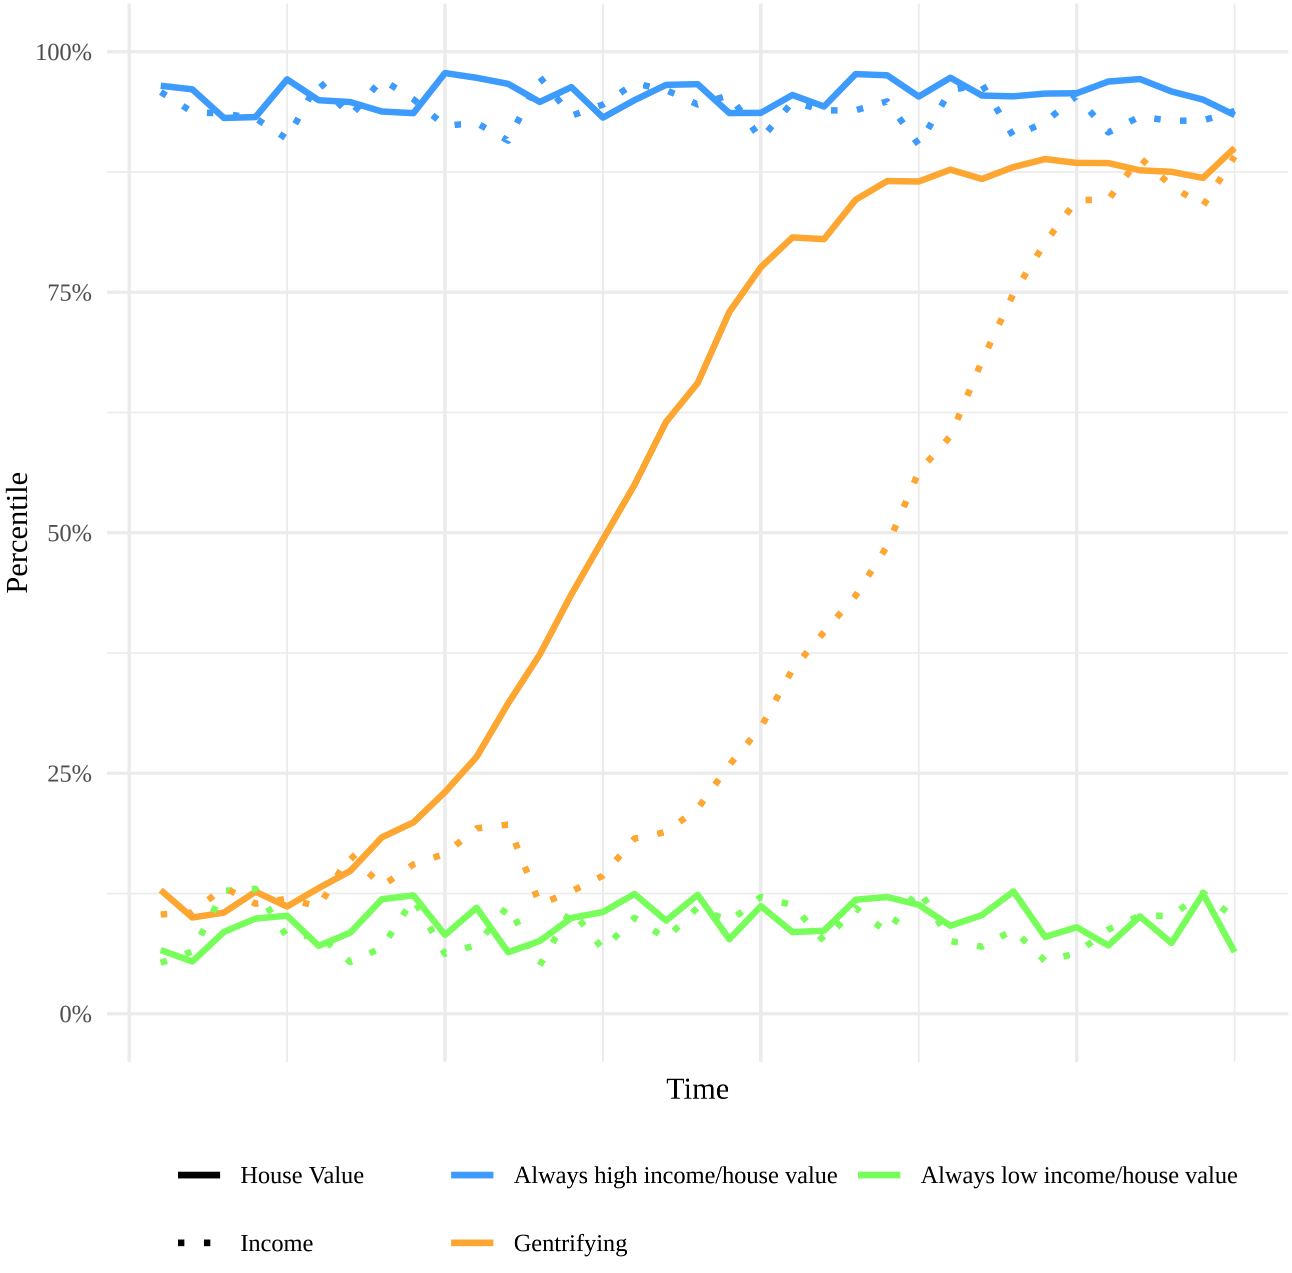

Figure 1. Conceptual example of house value and income percentiles in gentrifying and non-gentrifying neighborhoods.
Notes: This conceptual example does not use any real data, but illustrates how we might expect three different neighborhoods' house price and income percentiles to change over time.

$$y_{t,z,m} = \left[\sum_{j=-10}^{19} \mathbf{1}_{t-\tau_z=j} \times \alpha_j \right] + \mathbf{X}_{t,z,m}\beta + \gamma_m + \eta_t + \epsilon \quad (2)$$

Our dependent variable of interest, $y_{t,z,m}$, is the log of the average income of ZIP z in year t in MSA m. The variable τ_z is the year the ZIP first has a 25-percentile gap between house values and incomes. The variable $[\mathbf{1}_{t-\tau_z=j}]$ is an indicator variable that equals 1 if a neighborhood in year t is j years away from τ_z, and 0 otherwise. All non-gentrifying neighborhoods take a value of 0, as do already-gentrifying neighborhoods for which a gap opened before the start of our data. Our main coefficient of interest is α_j, which represents the relative income growth, in log-points, before or after gentrification onset in gentrifying neighborhoods. We include a vector of control variables, $X_{t,z,m}$: the log of neighborhood income in 1998, pre-1998 gentrification status, neighborhood, socio-economic variables (from NHGIS), natural amenities (from Lee and Lin 2018), and employment characteristics (from Manduca 2020). The variables γ_m and η_t are MSA and year fixed effects, respectively.

To test whether gentrifying neighborhoods experience income growth differently depending on neighborhood characteristics, we run additional regressions using equation 3. We interact gentrification status $[\mathbf{1}_{t-\tau_z=j}]$ with another binary variable, $\mathbf{1}_v$, where a neighborhood is either above

(1) or below (0) a threshold for the variable of interest. The variables of interest selected for comparison are: share of the neighborhood that is Black, the relative change in units in the neighborhood between 2017 and 2000, and the distance to the CBD. For example, $\mathbf{1}_v = 1$ if the neighborhood is in the inner sixth of the MSA. These variables were chosen because of their relevance to threads in the literature: how gentrification is shaped by racialization, new construction, and centrality (Davidson and Lees 2005; Rucks-Ahidiana 2021; Smith 1979).

$$y_{t,z,m} = \left[\sum_{v=0,1} \sum_{j=-10}^{19} \mathbf{1}_v \times \mathbf{1}_{t-\tau_z=j} \times \alpha_{j,v} \right] + \mathbf{X}_{t,z,m}\beta + \gamma_m + \eta_t + \epsilon \quad (3)$$

Gentrification, Neighborhood Context, and Income Growth

Figure 2 plots the estimates for α, obtained from our regression models, and re-expressed as percentage growth in average neighborhood income in the years before (to the left of 0) and after (to the right) a 25-percentile gap opens. These models control for the covariates described in the previous section. The shaded areas show the confidence intervals in the estimated income changes over time.

Panel A of Figure 2 shows the baseline results of Equation 2. Fifteen years after a gap opening, average income growth is 14% higher than would have been expected without a gap opening. This reflects rapid changes in neighborhood composition in the years following gap opening—for comparison, US median household income grew 5% from 1998–2018 (U.S. Census Bureau 2022). In contrast, Panel A shows relatively little change in income prior to gap opening: income growth is slightly negative, but the estimates are stable rather than trending upward—which would indicate our measure is "too late".[9]

The other panels present estimates of Equation 3, comparing income trajectories among gentrifying neighborhoods with different characteristics. Panel B shows that majority-Black neighborhoods see rapid and sustained income growth after gentrification onset, while others see slower income growth. Prior to gentrification onset, majority non-Black neighborhoods experience relatively low average income growth, matching the baseline figure but distinct from the experience of majority-Black neighborhoods. These dynamics contrast somewhat with the findings of Rucks-Ahidiana (2021), who finds increases in higher-educated and White residents—but not high earners—in majority-Black gentrifying neighborhoods. The different findings may be accounted for by differences in the time period under study, the gentrification measure, or the measure of income changes.

Panel C shows that neighborhoods with more housing growth experience greater income growth in the years following gentrification onset. This may reflect a few possible channels, which our approach cannot distinguish among: new construction may attract high earners, an influx of high earners may attract new construction, and the poor may be displaced through the construction process. This finding connects to Leguizamon and Christafore (2021), who show that neighborhoods in development-constrained cities are somewhat less likely to gentrify. Because Panel C shows income growth among neighborhoods that *do* gentrify, our finding is compatible with theirs.

Panel D reveals that gentrifying neighborhoods close to the CBD saw faster income growth, while neighborhoods further out saw no faster growth upon gentrification onset. Fifteen years after a gap opens, neighborhoods close to the CBD saw nearly 20% faster income growth, compared to essentially flat income growth in gentrifying neighborhoods further from downtown. Centrality helps shape gentrification (Smith 1979).

Collectively, these findings validate using house prices as an expectations-based signal for evaluating the onset of gentrification. The relationship between gap opening and income growth is mediated by other neighborhood attributes: income growth follows more quickly among neighborhoods that are closer to downtown, adding homes faster, and (initially) majority Black.

Validation: Qualitative and Quantitative Comparisons

In this section, we apply the signal to Boston and Chicago and compare our measure with extant qualitative studies in these cities using (variously) participatory, archival, ethnographic, and interview methods to establish gentrification status. We view systematic qualitative investigation as the most appropriate benchmark for validating a gentrification measure. While we don't conduct our own qualitative work, our investigation of "neighborhoods that qualitative researchers often highlight" responds to calls to "bridge methodological divides" (Brown-Saracino 2016) by benchmarking our quantitative signal against qualitative insights. We also compare our measure to two established quantitative measures, both of which were focused on planning applications: Bates (2013), whose work was used in Portland's comprehensive planning process (Bureau of Planning and Sustainability 2018), and Los Angeles's Index of Displacement Pressure, created by the Office of the Mayor's Innovation Team (Pudlin 2018).

Boston Region

Binet et al. (2021) uses survey and longitudinal interview methodologies within a participatory action research (PAR) process to study how gentrification affects caregiving relationships for residents in nine Boston-area neighborhoods. Binet collaborated with resident researchers from these neighborhoods to jointly develop hypotheses, research instruments, and analyses of the resulting data. The study selected sites based on four criteria: having a walkable urban center, a need for economic growth, early/mid-stage transformation, and significant population health challenges (Binet et al. 2021, 48). After identifying three such sites with major health equity-oriented development projects planned, each was paired with two comparable sites without such plans. These criteria ruled out the South End, a traditional site of gentrification research in Boston, instead targeting neighborhoods that began gentrifying more recently (e.g., Roxbury) as well as those that are experiencing other modalities of development (e.g., Brockton). We view the multi-site

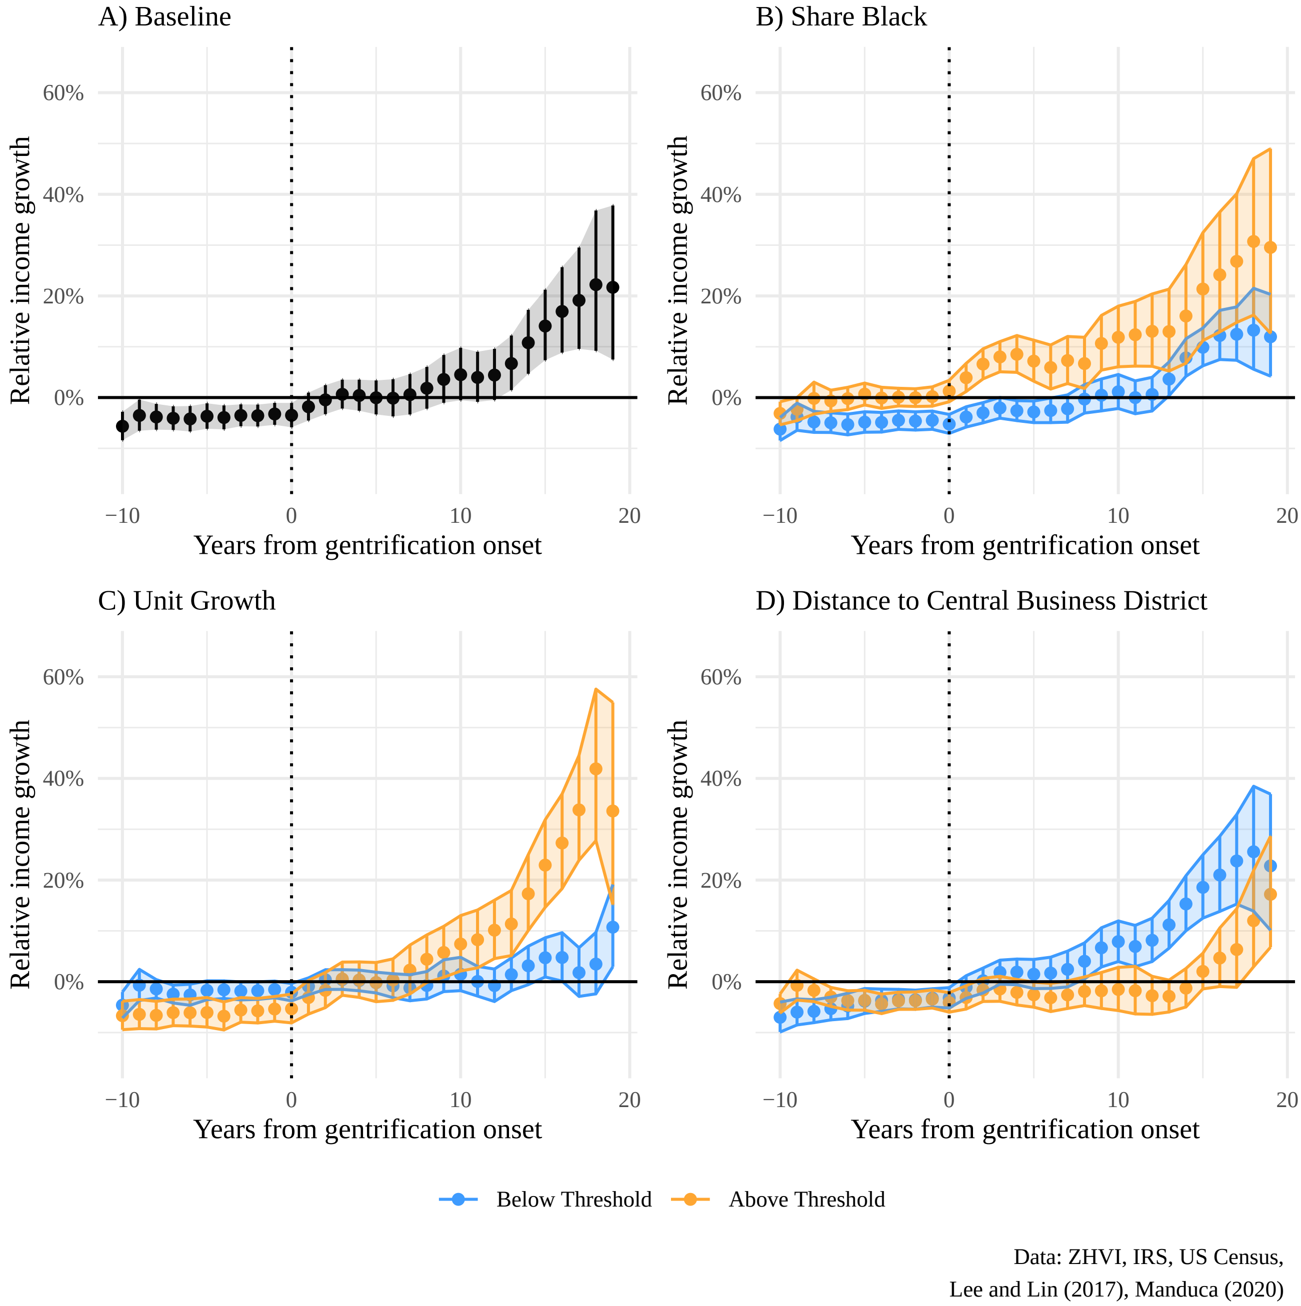

Figure 2. Effect of a 25-percentile gap between house values and incomes.
Notes: Regression results from equation 2 are shown in panel A. Regression results from equation 3 are shown in panels B-D. The thresholds for comparison groups are 50% Black for Panel B, a 10% increase in new units between 2000 and the 2015-2019 ACS for panel C, and within the inner sixth of the MSA for panel D. Dots are coefficients from the regression, lines and shaded areas represent the confidence intervals, either 95% for panel A or 83% for panels B-D.

comparative nature of the study—including neighborhoods in Boston proper, immediately adjacent communities, and more outlying places—as very useful for establishing a contemporaneous baseline of comparison to our quantitative signal.

In some neighborhoods—Roxbury, Dorchester, especially, as well as Mattapan and the nearby small cities of Chelsea and Everett—residents described strong community ties and social support coupled with threats to stability from new development priced beyond their reach and new businesses that didn't serve their needs. In contrast, residents of Brockton were as likely to describe the lack of investment, services, and social connections as major challenges—features common in more outlying places in the study. Based on their analyses, we expect to see strong signals of gentrification in the core neighborhoods of Roxbury and Dorchester as well as Mattapan, Chelsea and Everett, but not in Brockton.

Figure 3 applies our signal to these neighborhoods using our ZIP dataset. Brockton is the clear outlier: house prices and incomes remain among the lowest in the MSA. By contrast, gaps have opened in every neighborhood in which residents describe development pressures as a threat to caregiving responsibilities, with larger (and earlier) gaps in Roxbury and Dorchester. Our method provides a quantitative

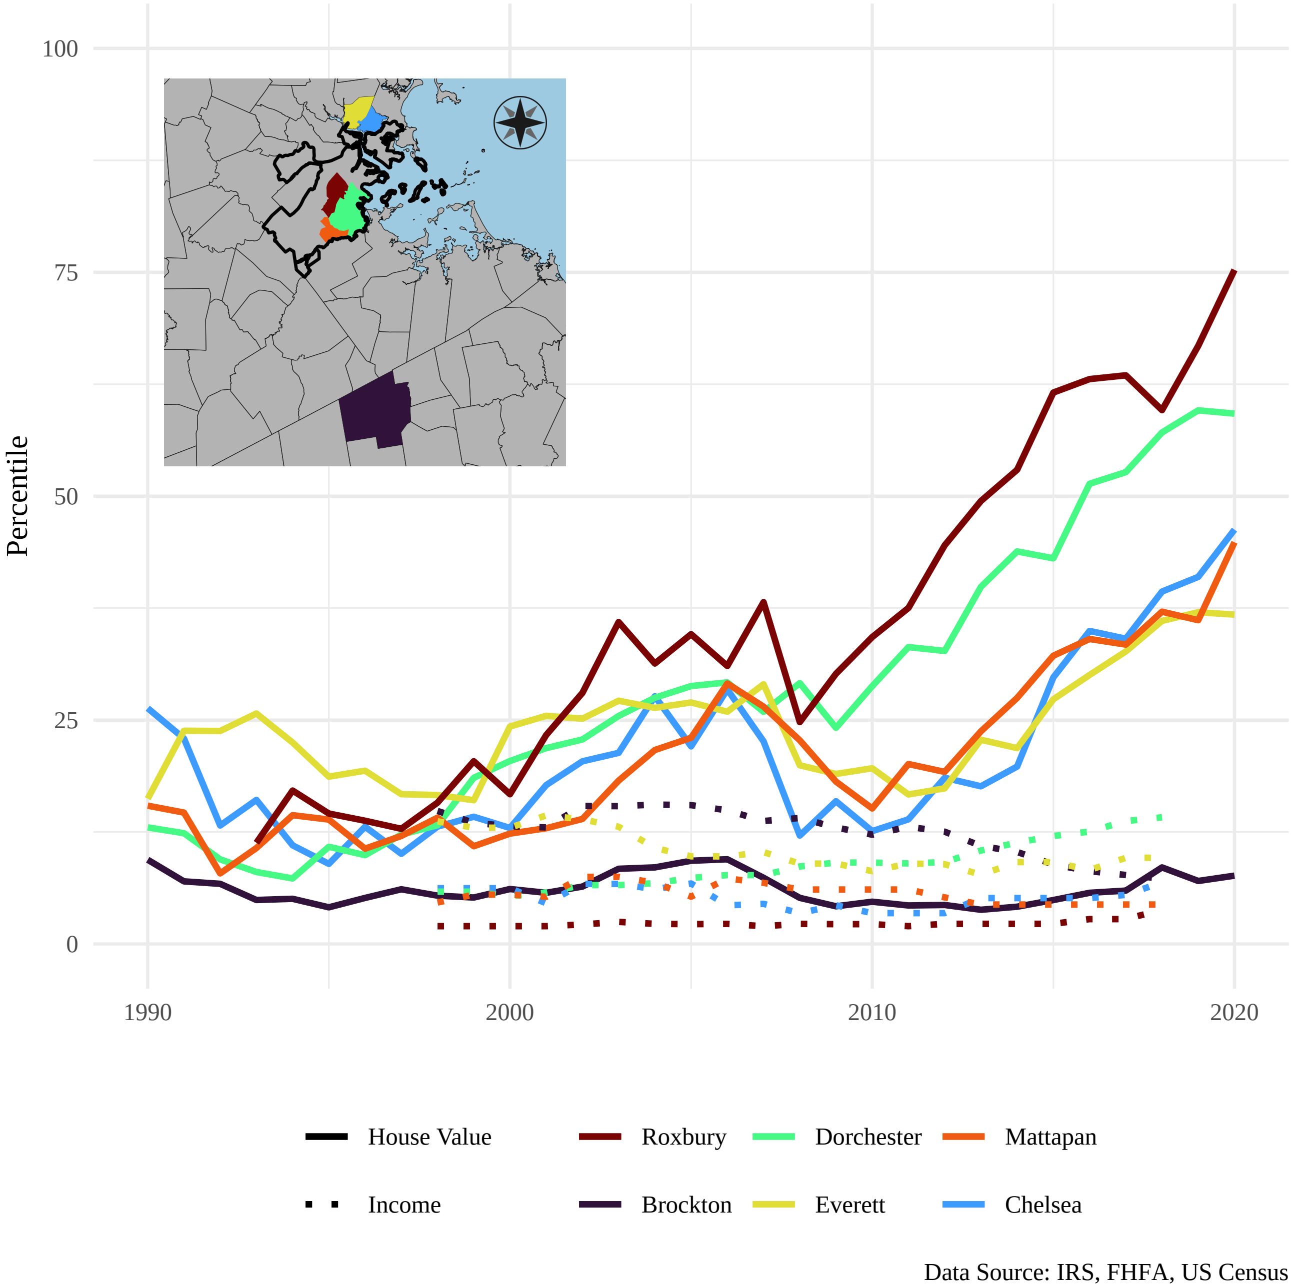

Figure 3. Greater Boston Gentrification.
Notes: House price and income percentiles in greater Boston. Colors in map match colors in line chart. Boston is highlighted in a bold, black outline in the map.

signal of the local knowledge Binet captures through PAR-based surveys and interviews. Using contemporaneous data, we see what's happening on the ground shortly after it takes place.[10] However, our MSA-based operationalization misses two places in Binet's study that lie in southern Massachusetts, beyond the borders of the Boston MSA. Those places could be included by recalculating the percentiles inclusive of this area, reflecting the necessity of accounting for boundary effects.

Chicago

In this subsection, we apply a threshold-based signal to study the history of gentrification in Chicago and compare our findings to the qualitative work of Perez (2004), Pattillo (2008), Hyra (2008), Brown-Saracino (2009), and Hertz (2018). Figure 4 maps the first year a gap opens between house value and income percentiles, using data from 1940 to 2019. To show the flexibility of our measure, we present two binary thresholds: a 20-percentile gap on the left panel and a 30-percentile gap on the right. We use two alternative neighborhood definitions: the top panels use Census tracts while the bottom panels aggregate tracts into city-defined community areas. Some tracts are missing price/income data in some years. For community areas, we only show results in years where data is reported for over three-fourths of the population.

Figure 4 enables a cartographic reading of Chicago's history of gentrification. Old Town was an early exemplar

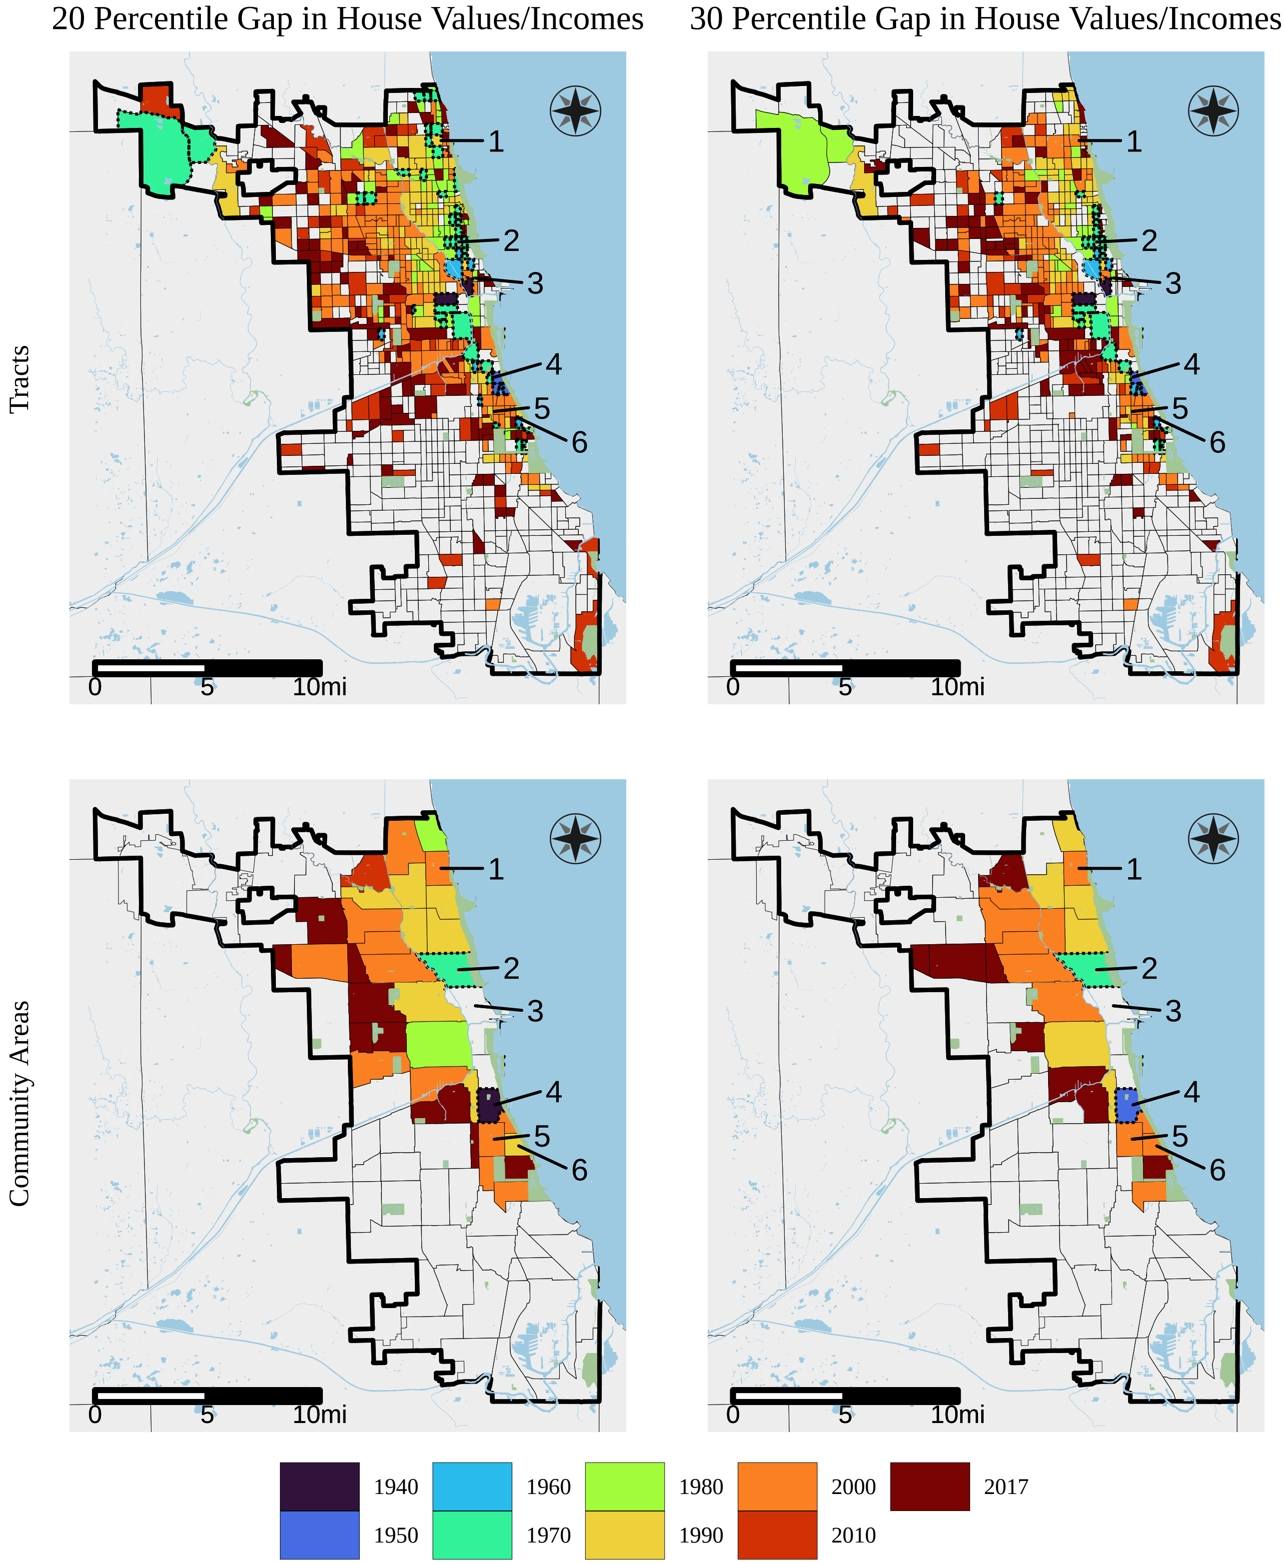

Data: Census via Lee and Lin (2017), ACS

Figure 4. Gentrification in Chicago since 1940.
Notes: Neighborhoods that are discussed in the text are labeled: 1, Edgewater; 2, Lincoln Park; 3, Old Town; 4, Douglas; 5, Bronzeville, 6, Kenwood. A dotted line indicates an area where a gap opened in 1970 or earlier.

of gentrification (Hertz 2018). The neighborhood at its commercial heart saw gentrification as early as 1960. By the early 1970s, rising rents had pushed the bohemians north towards Lincoln Park where extensive gentrification throughout the community area is recorded as of 1970. Parts of Lincoln Park still had gaps open in recent decades despite having incomes well above the median, suggesting advanced gentrification.

Farther north, Edgewater is shown as gentrifying by 2000. Looking to its constituent census tracts, we can see substantial heterogeneity. Gaps opened in the sub-neighborhoods of Andersonville during the 1980s and 1990s and Argyle by 1990 or 2000, in line with (Brown-Saracino 2009). West and southwest of Lincoln Park, Puerto Rican and Ukrainian neighborhoods show as gentrifying by 1990 or 2000, consistent with Perez (2004).

Pattillo (2008) and Hyra (2008) document gentrification in 1990s Kenwood/Oakland and Bronzeville, respectively. Unlike the north and northwest-side neighborhoods discussed above, these Southside neighborhoods were home to mostly Black residents at the onset of the processes, and Pattillo's book documents a process of Black gentrification. In the context of racialized housing markets, gentrification may not generate expectations of rapid house price appreciation in Black neighborhoods. In our maps, a single tract of Kenwood is gentrifying by 1990, and the community areas cross the 20-percentile threshold by 2000. Two Bronzeville tracts are shown as gentrifying by 1990, and more cross the threshold by 2000. Despite the different nature of gentrification in Black neighborhoods, the signal works: the house price/income gap is significant in several tracts and opens in line with the processes described in their work.

However, the Douglas community area—overlapping Bronzeville—registers as gentrifying by 1940 or 1950. Douglas was not gentrifying in the 1940s; it was the core of the intensively segregated Black South Side. Why was there a gap? Intense segregation may have been directly responsible: the limited supply of housing available to Black families pushed prices up while labor markets segregation held down Black workers' earnings (Boustan 2016). The Douglas example emphasizes the importance of combining any metric with local knowledge, and the simplicity of doing so with our metric.

Comparing across panels reveals tradeoffs of using different neighborhood boundaries and gap thresholds. Community areas are larger than the neighborhoods qualitative researchers generally study, and mask substantial variation across tracts. At the same time, some spatial variation is statistical noise, which aggregating smooths. The 20-percentile threshold results in a very advanced gentrification frontier in recent years. By contrast, the larger threshold misses some places with rising incomes and house prices that never see a 30-percentile gap—including many surrounded by gentrifying places. These tensions are inherent to quantitative measurement, and our signal cannot avoid them. Our use of a 25-percentile threshold elsewhere in the paper aims to balance these competing risks.

Portland

Bates (2013) studies gentrification in Portland, Oregon, between 1990 and 2010. She draws definitional characteristics from Freeman (2005), and her approach has been taken up since, e.g. by Chapple et al. (2022), thus offering a practice-engaged and academically-representative example of quantitative gentrification measurement. Bates classifies tracts based on the presence of a "vulnerable" population, housing market factors, and demographic change. (Most tracts lack these features and were coded *NA*.)

Figure 5 presents our measure for Portland tracts during the period 1998–2018, with separate panels for each of Bates's tract types. For clarity, we bold low-income tracts after a 25-percentile gap has opened. Our measures largely agree. Many tracts undergoing "early" gentrification see sizeable house value/income gaps open, and tracts classified as "Dynamic," "Late," or "Continued Loss," have rapidly rising house prices with trailing, but increasing, incomes. However, our measure picks up likely gentrification Bates's approach misses. The bold, purple, tract in the "NA Tract, High Vulnerability" panel appears to be experiencing post-industrial gentrification: house values increased from near the median to the top quartile by 2003, while incomes increased from the 3rd to the 15th percentile by 2018.

There is overall alignment between neighborhoods Bates classifies as undergoing gentrification, and those tracts where we see rising house values and lagging (but rising) incomes. Against a popular quantitative measure of gentrification, our measure performs similarly.

Los Angeles

The Los Angeles Innovation Team developed the Los Angeles Index of Displacement Pressure (LAIDP) to map gentrification and influence planning (Pudlin 2018). This prospective measure identifies neighborhoods with *future* displacement risks by integrating the Los Angeles Index of Neighborhood Change (Pudlin 2016)—a retrospective index akin to Bates (2013)—with displacement risk factors: expiring affordable housing units, transit facilities, rental market factors, and a proprietary forecast of house price growth from Esri.

Figure 6 maps the LAIDP (top left) and the house value/income gap in five-year increments for 1998–2018 for LA ZIPs (other panels). The city borders are outlined in black. There is substantial concordance between the approaches. Very high- and high-risk areas in Central LA have sizable gaps, while smaller gaps are visible in South and Northeast LA that are medium risk in the LAIDP. Areas in the distant edges of the San Fernando Valley have modest *negative* gaps and are largely classed as low risk by the LAIDP. Differences arise in the large gaps of late-stage gentrifying neighborhoods like Venice Beach and in college-adjacent areas like University Park, labeled low risk in the LAIDP. Further, the LAIDP shows higher risks in other parts of downtown than our measure. This reflects extremely rapid income growth downtown, surpassing the median and closing the gap by 2010. The displacement risk warned of by the LAIDP was already visible in the rearview mirror.

Mapping the gap over time offers some unique insights. Changes across panels are subtle—and for much of the city, the panels are nearly identical (and the gaps are near zero). These subtleties reveal variation in how far in advance house values anticipate future projected displacement risk: Central LA has large gaps open by 1998, while Northeast LA only sees a gap open more recently. Relative to the maps of Chicago, this approach reveals gaps closing, as in Venice Beach, a (now) wealthy coastal enclave. For gentrifying places, a constant gap does not imply stasis in the neighborhood measured; instead it could reflect rising house values *and* incomes.

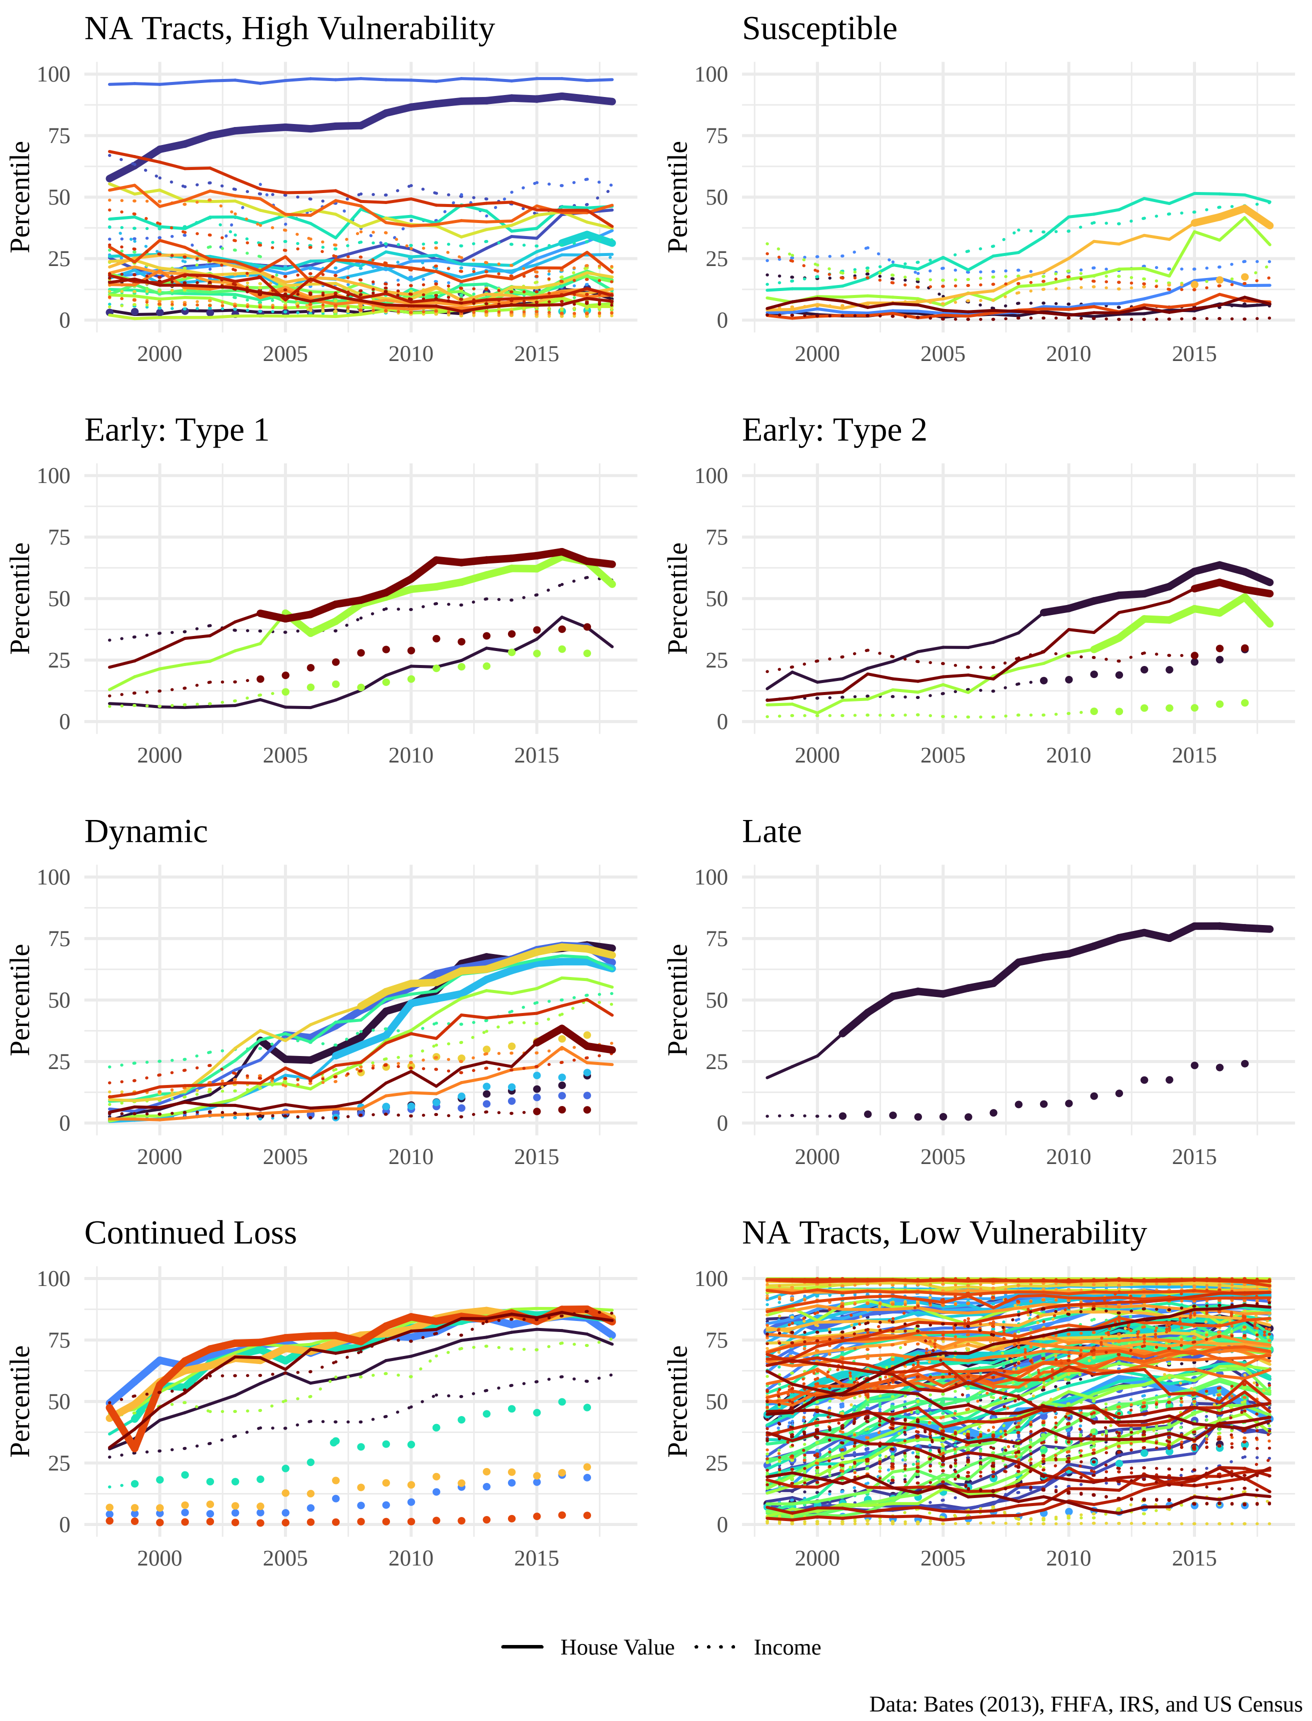

Figure 5. Gentrification in Portland, comparing the Bates (2013) findings to our measure.
Notes: Bold lines indicate a gentrifiable tract has had a gap larger than 25 percentiles open between house prices and income. Here as in our regression analysis, we define “gentrifable” as a tract in the inner third of the MSA, with an income percentile-rank in the bottom quartile.

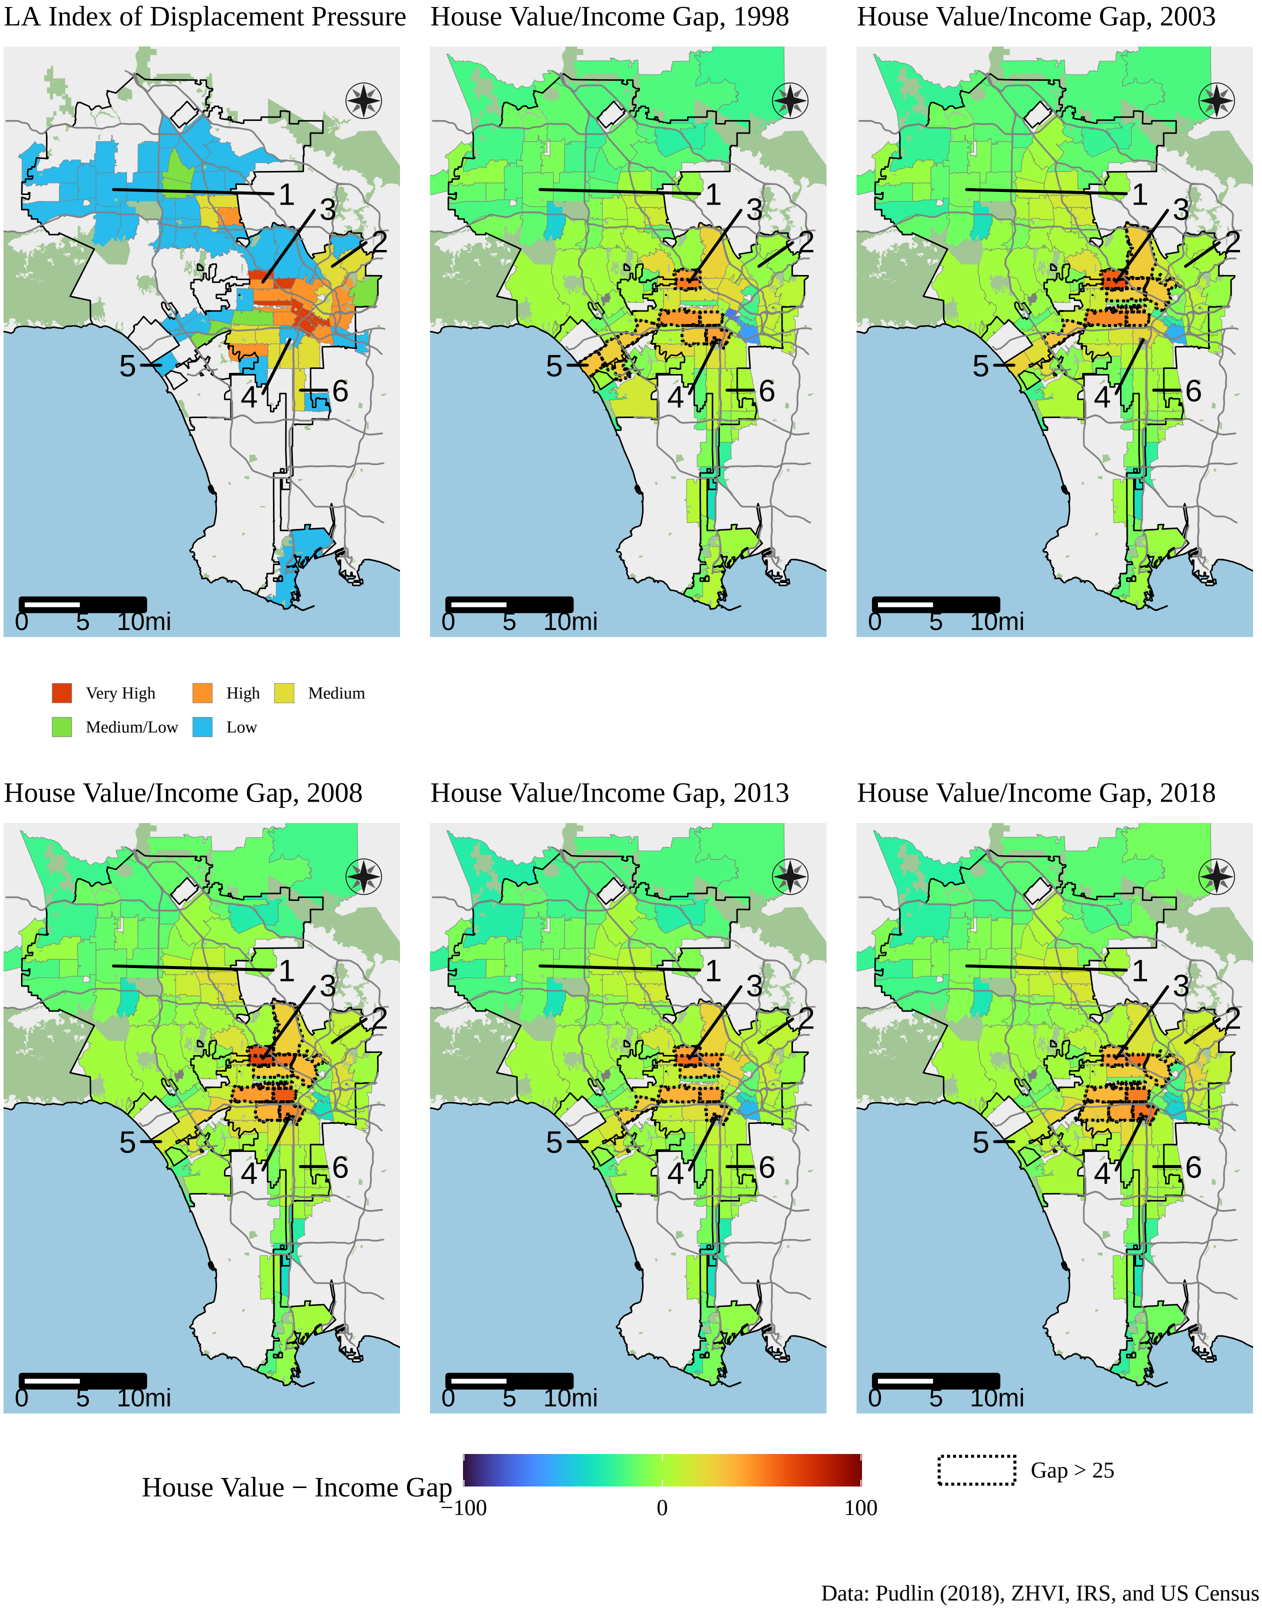

Figure 6. Los Angeles Index of Displacement Pressure and House Value/Income Gap Over Time.
Notes: Approximate location of areas that are discussed in the text are labeled: 1, San Fernando Valley; 2, Northeast LA; 3, Central LA; 4, University Park; 5, Venice Beach, 6, South LA.

An Expectations-Based Signal Improves Understanding

In this paper, we developed an expectations-based measure of gentrification. Asset valuation theory shows that property values incorporate the expectations of market participants. We use this theory to interpret property values as incorporating local market participant knowledge about a neighborhood's future. If their expectations are correct, the future holds rising incomes, capital investment, landscape change, displacement, and other changes characterizing gentrification.

We operationalize this insight by comparing the percentile-rank of a neighborhood's house prices and incomes. In the US, these components are released on at least an annual basis, enabling rapid identification of expected gentrification. We interpret a sizeable gap between the two as a signal of gentrification. Using annual data and a dynamic difference-in-difference framework, we demonstrate that incomes rise rapidly following the opening of a substantial gap. Our signal overlaps empirically with existing measures of gentrification and improves upon them by offering easy application to time-series, cross-sectional, and panel contexts. The signal can be plotted over time (as we demonstrate for Boston and Portland) and mapped cross-sectionally (as for Los Angeles) or by mapping gentrification's path through a city over time (as for Chicago).

We note several limitations. Our emphasis on the convenience of house prices and incomes costs us nuance. House prices may proxy poorly for property values in areas with mostly rental or social housing. We may miss marginal gentrification that doesn't translate immediately into house prices, as well as interventions like state-led gentrification. In these cases, house prices may be a lagging indicator. Other variables incorporating expectations of the future include multifamily property values, investment decisions, and city plans. Income doesn't fully characterize vulnerability to gentrification, and without including (e.g.) a direct racial component, we may misstate risks. Our percentile-based measure may flatten meaningful differences. Brooklyn Heights, in the period Lees (2003) studies, has a small house value/income gap, but the fractal nature of top income inequality means "super-gentrification" may nevertheless push house prices beyond the reach of the merely rich. In Section 5, we only test income growth, not other relevant outcomes. Empirically, we identify some neighborhoods as gentrifying that do not have established records of research, raising the possibility of false positives.

Set against these limitations are the significant benefits of a timely, well-understood, and readily-available measure of gentrification. Our approach can be used by practitioners and researchers alike to track gentrification at the local level. Practitioners implementing policies to mitigate negative effects of gentrification can only do so if they have accurate, timely measures of on-the-ground changes. Our signal meets those needs, while providing interpretability and flexibility allowing for its deployment in planning contexts. For researchers, the annual signal and difference-in-difference implementation offer a new way of studying diverse outcomes in gentrifying places (e.g., Kavanagh-Smith 2021). Beyond the gap, plotting house price and income percentiles over time offers insight into gentrification by revealing how a gentrifying neighborhood has moved through its city's socioeconomic hierarchies—even in cases where a gap doesn't open. Our conceptual distinction of expectations-based variables offers a new approach to identifying gentrification, and we hope further variables can be brought into this framework.

Acknowledgements

We are deeply indebted to the research assistance of Madeleine Daepp (for her exceedingly thorough and diligent data analysis) and Emily Moss (for her support in tracking the literature and copy editing). We are grateful to insightful comments from an array of readers, including Vicki Been, Japonica Brown-Saracino, Mike Lens, Paavo Monkkonen, and Justin Steil.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

This paper draws in part on work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1745302. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Notes

1. Our gap extends Smith's (1979) directly: a neighborhood's house price/income gap emerges when market participants begin valuing real estate according to its *potential* ground rent—i.e., when the rent gap begins closing.
2. Some scholars study definitional variation itself (Hwang and Shrimali 2021; Rucks-Ahidiana 2021), although quantitative data has limited ability to validate how these distinctions correspond to experiences on the ground (Goetz et al. 2019).
3. See Glickman (2014) and Kaplan et al. (2020) for the ongoing relevance of asset valuation theory to practical and academic work, respectively.
4. Landlords may also take actions to increase the likelihood of gentrification; gentrification enables rent increases even absent such investment.
5. Contexts where expectations are reflected in property values include changes to flood risk (Fonner et al. 2022) or transit investment (Golub et al. 2012).
6. Zapatka and Beck (2021) argue that gentrifiers lead one year-ahead house price growth, although this doesn't necessarily conflict with our multi-year window.
7. Detailed methods are available in the supplementary online materials.
8. Sun and Abraham (2021) caution against causal interpretation when the "treatment" effect—here, gentrification—can be anticipated; we interpret our results as correlations.
9. A downward trend would have been worrisome too, as it could imply the gap opened because of declining incomes rather than growing house values.
10. Incomes do not rise much over this period. In Roxbury, the first to see a gap, just over half of the housing stock is income-restricted, slowing the manifestations of gentrification. Boston has the highest rate of income-restricted housing among major US cities (of Neighborhood Development 2021).

References

Barton M (2016) An exploration of the importance of the strategy used to identify gentrification. *Urban Studies* 53(1): 92–111. DOI:10.1177/0042098014561723. URL https://doi.org/10.1177/0042098014561723.

Bates L (2013) Gentrification and Displacement Study: Implementing an Equitable Inclusive Development Strategy in the Context of Gentrification. *Urban Studies and Planning Faculty Publications and Presentations* DOI:10.15760/report-01. URL https://pdxscholar.library.pdx.edu/usp_fac/83.

Beauregard RA (1986) The chaos and complexity of gentrification. In: Smith N and Williams P (eds.) *Gentrification of the City*. Routledge. ISBN 978-1-315-88909-2. Num Pages: 21.

Binet A, Rio GZd, Arcaya M, Roderigues G and Gavin V (2021) 'It feels like money's just flying out the window': financial security, stress and health in gentrifying neighborhoods. *Cities & Health* 0(0): 1–16. DOI:10.1080/23748834.2021.1885250. URL https://doi.org/10.1080/23748834.2021.1885250. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/23748834.2021.1885250.

Boustan LP (2016) *Competition in the Promised Land: Black Migrants in Northern Cities and Labor Markets*. Princeton University Press. ISBN 978-0-691-15087-1. URL https://www.nber.org/books-and-chapters/competition-promised-land-black-migrants-northern-cities-and-labor-markets. Backup Publisher: National Bureau of Economic Research Type: Book.

Brown-Saracino J (2009) *A Neighborhood That Never Changes*. Fieldwork Encounters and Discoveries. The University of Chicago Press. URL https://doi.org/10.7208/9780226076645.

Brown-Saracino J (2013) *The Gentrification Debates*. 0 edition. Routledge. ISBN 978-1-134-72564-9. DOI:10.4324/9781315881096. URL https://www.taylorfrancis.com/books/9781134725649.

Brown-Saracino J (2016) An Agenda for the Next Decade of Gentrification Scholarship. *City & Community* 15(3): 220–225. DOI:10.1111/cico.12187. URL http://journals.sagepub.com/doi/10.1111/cico.12187.

Brown-Saracino J (2017) Explicating Divided Approaches to Gentrification and Growing Income Inequality. *Annual Review of Sociology* 43(1): 515–539. DOI:10.1146/annurev-soc-060116-053427. URL http://www.annualreviews.org/doi/10.1146/annurev-soc-060116-053427.

Bureau of Planning and Sustainability (2018) 2035 Comprehensive Plan. Technical report, City of Portland. URL https://www.portlandoregon.gov/bps/2035-comp-plan.pdf.

Chapple K, Poorthuis A, Zook M and Phillips E (2022) Monitoring streets through tweets: Using user-generated geographic information to predict gentrification and displacement. *Environment and Planning B: Urban Analytics and City Science* 49(2): 704–721. DOI:10.1177/23998083211025309. URL https://doi.org/10.1177/23998083211025309. Publisher: SAGE Publications Ltd STM.

Davidson M and Lees L (2005) New-Build 'Gentrification' and London's Riverside Renaissance. *Environment and Planning A: Economy and Space* 37(7): 1165–1190. DOI:10.1068/a3739. URL https://doi.org/10.1068/a3739.

Easton S, Lees L, Hubbard P and Tate N (2019) Measuring and mapping displacement: The problem of quantification in the battle against gentrification. *Urban Studies* : 0042098019851953DOI:10.1177/0042098019851953. URL https://doi.org/10.1177/0042098019851953.

Ellen IG and O'Regan KM (2011) How low income neighborhoods change: Entry, exit, and enhancement. *Regional Science and Urban Economics* 41(2): 89–97. DOI:10.1016/j.regsciurbeco.2010.12.005. URL http://linkinghub.elsevier.com/retrieve/pii/S0166046211000044.

Finio N (2021) Measurement and Definition of Gentrification in Urban Studies and Planning:. *Journal of Planning Literature* DOI:10.1177/08854122211051603. Publisher: SAGE PublicationsSage CA: Los Angeles, CA.

Fisher I (1906) *The Nature of Capital and Income*. New York: The Macmillan Company.

Fonner R, Izón G, Feist BE and Barnas K (2022) Capitalization of reduced flood risk into housing values following a floodplain restoration investment. *Journal of Environmental Economics and Policy* URL https://www.tandfonline.com/doi/abs/10.1080/21606544.2022.2136765. Publisher: Routledge.

Freeman L (2005) Displacement or Succession?: Residential Mobility in Gentrifying Neighborhoods. *Urban Affairs Review* 40(4): 463–491. DOI:10.1177/1078087404273341. URL https://doi.org/10.1177/1078087404273341. Publisher: SAGE Publications Inc.

Galster G and Peacock S (1986) Urban gentrification: Evaluating alternative indicators. *Social Indicators Research* 18(3): 321–337. DOI:10.1007/BF00286623. URL https://doi.org/10.1007/BF00286623.

Glaeser EL, Kim H and Luca M (2018) Nowcasting Gentrification: Using Yelp Data to Quantify Neighborhood Change. *AEA Papers and Proceedings* 108: 77–82. DOI:10.1257/pandp.20181034. URL https://www.aeaweb.org/articles?id=10.1257/pandp.20181034.

Glickman EA (2014) Property Valuation. In: *An Introduction to Real Estate Finance*. Elsevier. ISBN 978-0-12-378626-5, pp. 129–151. DOI:10.1016/B978-0-12-378626-5.00005-X. URL https://linkinghub.elsevier.com/retrieve/pii/B978012378626500005X.

Goetz EG, Lewis B, Damiano A and Calhoun M (2019) The Diversity of Gentrification: Multiple Forms of Gentrification in Minneapolis and St. Paul. Technical report, Center for Urban And Regional Affairs, University of Minnesota. URL http://gentrification.umn.edu/sites/gentrification.dl.umn.edu/files/media/diversity-of-gentrification-012519.pdf.

Golub A, Guhathakurta S and Sollapuram B (2012) Spatial and Temporal Capitalization Effects of Light Rail in Phoenix: From Conception, Planning, and Construction to Operation. *Journal of Planning Education and Research* 32(4): 415–429. DOI:10.1177/0739456X12455523. URL http://journals.sagepub.com/doi/10.1177/0739456X12455523.

Hammel DJ and Wyly EK (1996) A Model for Identifying Gentrified Areas with Census Data. *Urban Geography* 17(3): 248–268. DOI:10.2747/0272-3638.17.3.248. URL https://doi.org/10.2747/0272-3638.17.3.248.

Hertz DK (2018) *The Battle of Lincoln Park: Urban Renewal and Gentrification in Chicago*. Belt Publishing.

Hwang J and Shrimali BP (2021) Constrained Choices: Gentrification, Housing Affordability, and Residential Instability in the San Francisco Bay Area. Technical Report 2021-2, Federal Reserve Bank of San Francisco. URL https://www.doi.org/10.24148/cdrb2021-02.

Hyra DS (2008) *The New Urban Renewal: The Economic Transformation of Harlem and Bronzeville*. Chicago, IL: University of Chicago Press. ISBN 978-0-226-36604-3. URL https://press.uchicago.edu/ucp/books/book/chicago/N/bo5764141.html.

Jain S, Proserpio D, Quattrone G and Quercia D (2021) Nowcasting Gentrification Using Airbnb Data. *Proceedings of the ACM on Human-Computer Interaction* 5(CSCW1): 38:1–38:21. DOI: 10.1145/3449112. URL https://doi.org/10.1145/3449112.

Kaplan G, Mitman K and Violante GL (2020) The Housing Boom and Bust: Model Meets Evidence. *Journal of Political Economy* 128(9): 3285–3345. DOI:10.1086/708816. URL https://www.journals.uchicago.edu/doi/10.1086/708816.

Kavanagh-Smith D (2021) *Does Gentrification Contribute to Homelessness? A Study of Family Shelter Entry across New York City, 2005–2015*. Dissertation, The New School, New York. URL https://www.proquest.com/pagepdf/2557509612?accountid=12492.

Lee S and Lin J (2018) Natural Amenities, Neighbourhood Dynamics, and Persistence in the Spatial Distribution of Income. *The Review of Economic Studies* 85(1): 663–694. DOI: 10.1093/restud/rdx018. URL https://academic.oup.com/restud/article/85/1/663/3091867. Publisher: Oxford Academic.

Lees L (2003) Super-gentrification: The Case of Brooklyn Heights, New York City. *Urban Studies* 40(12): 2487–2509. DOI:10.1080/0042098032000136174. URL https://doi.org/10.1080/0042098032000136174.

Leguizamon S and Christafore D (2021) The influence of land use regulation on the probability that low-income neighbourhoods will gentrify. *Urban Studies* 58(5): 993–1013. DOI: 10.1177/0042098020940163. URL http://journals.sagepub.com/doi/10.1177/0042098020940163.

Liu C, Deng Y, Song W, Wu Q and Gong J (2019) A comparison of the approaches for gentrification identification. *Cities* 95: 102482. DOI:10.1016/j.cities.2019.102482. URL https://linkinghub.elsevier.com/retrieve/pii/S0264275118315440.

Manduca R (2020) The spatial structure of US metropolitan employment: New insights from administrative data. *Environment and Planning B: Urban Analytics and City Science* DOI:10.1177/2399808320934821. URL https://journals.sagepub.com/doi/10.1177/2399808320934821. Publisher: SAGE PublicationsSage UK: London, England.

Manson, Steven, Schroeder, Jonathan, Van Riper, David, Kugler, Tracy and Ruggles, Steven (2021) National Historical Geographic Information System: Version 16.0. DOI:10.18128/D050.V16.0. URL https://www.nhgis.org/. Version Number: 16.0 Type: dataset.

of Neighborhood Development D (2021) Income Restricted Housing in Boston 2020. Technical report, City of Boston. URL https://www.boston.gov/sites/default/files/file/2021/03/Income%20Restricted%20Housing%202020_0.pdf.

Pattillo M (2008) *Black on the Block: The Politics of Race and Class in the City*. Chicago, IL: University of Chicago Press. ISBN 978-0-226-64932-0. URL https://press.uchicago.edu/ucp/books/book/chicago/B/bo4149945.html.

Perez G (2004) *The Near Northwest Side Story: Migration, Displacement, and Puerto Rican Families*. University of California Press. ISBN 978-0-520-23368-3.

Preis B, Janakiraman A, Bob A and Steil J (2020) Mapping gentrification and displacement pressure: An exploration of four distinct methodologies:. *Urban Studies* 58(2): 405–424. DOI:10.1177/0042098020903011. URL https://journals.sagepub.com/doi/full/10.1177/0042098020903011?casa_token=8gA9rAdLbDgAAAAA%3ALzpVApZyGzxOaD-OII_VNsmM4uD_9R_L5AbALOqUuvPpnCuCsnPJScVCpHNgqdlSJ8IaOXLS-S1c. Publisher: SAGE PublicationsSage UK: London, England.

Pudlin A (2016) Los Angeles Index of Neighborhood Change. URL http://geohub.lacity.org/datasets/57e9231c3bd34d44ae49b309b0cb440e_1.

Pudlin A (2018) Los Angeles Index of Displacement Pressure. URL http://geohub.lacity.org/datasets/70ed646893f642ddbca858c381471fa2_0.

Reades J, De Souza J and Hubbard P (2018) Understanding urban gentrification through machine learning. *Urban Studies* : 0042098018789054DOI:10.1177/0042098018789054. URL https://doi.org/10.1177/0042098018789054.

Rucks-Ahidiana Z (2021) Racial composition and trajectories of gentrification in the United States. *Urban Studies* 58(13): 2721–2741. DOI:10.1177/0042098020963853. URL https://doi.org/10.1177/0042098020963853. Publisher: SAGE Publications Ltd.

Smith N (1979) Toward a Theory of Gentrification A Back to the City Movement by Capital, not People. *Journal of the American Planning Association* 45(4): 538–548. DOI:10.1080/01944367908977002. URL https://doi.org/10.1080/01944367908977002.

Sun L and Abraham S (2021) Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225(2): 175–199. DOI:10.1016/j.jeconom.2020.09.006. URL https://www.sciencedirect.com/science/article/pii/S030440762030378X.

US Census Bureau (2022) Real Median Household Income in the United States. URL https://fred.stlouisfed.org/series/MEHOINUSA672N. Publisher: FRED, Federal Reserve Bank of St. Louis.

Zapatka K and Beck B (2021) Does demand lead supply? Gentrifiers and developers in the sequence of gentrification, New York City 2009–2016. *Urban Studies* 58(11): 2348–2368. DOI: 10.1177/0042098020940596. URL http://journals.sagepub.com/doi/10.1177/0042098020940596.

Supplemental material

Construction of House Value and Income Quantiles

To construct a spatially compact, temporally useful signal of gentrification, the two main data sources for this paper are (1) home price data from the Federal Housing Finance Agency (FHFA) and (2) Internal Revenue Service (IRS) Statistics of Income. We use 2010-vintage census tracts as our primary unit of analysis. The census tract is the smallest geographic unit for which house price data is available. At times, we instead rely on zip codes, alternatively known as ZIP Code Tabulation Areas (ZCTAs). Throughout the paper, we identify when we are using zip codes and when we are using tracts. This choice is largely dependent on data quality and availability, and we tend towards the smallest geography possible. The other relevant geographic unit is the Metropolitan Statistical Area (MSA), which we use as the overall distribution from which we construct percentile rankings for each neighborhood's house price and income. We use the 2018 vintage MSA definitions provided by the Office of Management and Budget. MSAs are agglomerations of counties, largely defined on bilateral commuting flows. They are quite large, relative to core cities. For instance, New York City has a land area of 300 square miles (777 sq km), while the MSA encompassing New York City has a land area of 6,685 square miles (17,314 sq km). Because census tracts are nested within counties, we are able to assign all census tracts to a specific MSA. We assign ZCTAs to MSAs based on a crosswalk between counties an ZCTAs provided by the Department of Housing and Urban Development.[1]

FHFA provides an annual estimate of *changes* to the housing values in a geographic area relative to the prior year — it does not provide absolute values. To construct values for each year, we use 2000 Census median house value to calculate the house value quantile for each tract within its city. For earlier and later years, we grow out these base-year prices using the FHFA tract-level indices. There are some tracts that have FHFA in some years, but are missing in others. In those years, we use impute the FHFA values using Kalman Smoothing as implemented in the "imputeTS" package in R (Moritz et al. 2022). We use those imputed values for the purpose of creating accurate percentiles among all geographies in an MSA, but we drop the imputed values from regressions and plots to ensure that we do not draw inferences from imputed results. We do not impute values beyond the last year of the provided FHFA data, however, to avoid extrapolation (e.g., if a tract has data 1998-2005, and 2007-2013, we'll impute the value for 2006 but leave the values beyond 2013 as NA.) When tract-level indices are wholly unavailable, we compute changes using the index from the tract's zip code.[2]

From this constructed dataset of house price levels, we calculate the percentile rank of every tract within an MSA for its house price in every year for which data is available, 1998 through 2020. In constructing the percentiles, we weight each tract by 2000 counts of the the number of housing units provided by the US Census and the National Historical Geographic Information System (NHGIS).

There are some zip codes and tracts that are generally absent from the FHFA indices. Because the FHFA relies on repeat-sales data of single-family homes, downtowns in central cities with many condominiums and cooperatives, but few single-family homes, are generally absent. As such, we also construct our same measure using the Zillow Home Value Index (ZHVI). The ZHVI includes all housing types and is available at the zip-code level. We find broad consistency across FHFA and ZHVI data: a correlation coefficient of approximately 0.88 for zip code quantiles. ZHVI is provided at monthly intervals from 1996 through 2019; we average across all months within each calendar year.

For annual income data, the IRS reports annual gross income at the zip code level. When doing analysis at the spatial geography of the zip code, we take the IRS data as-in. At the tract level, we calculate the rate of year-over-year income change from the IRS at the zip code level and thus calculate the changes from the 2000 income based on the US Census median income estimates, to calculate expected income at the tract level for 1998-2018. This allows us to compute a percentile ranking and gap for each year. This process at the tract level is akin to how we "grow out" house value data from the FHFA HPI.

Our main two variables of interest are the percentile rank of neighborhoods' income and house values vis-à-vis the distribution of income and house value across an entire MSA. For every tract and zip code within an MSA, we calculate the percentile rank of the median house value, weighted by housing unit counts. For income, we similarly calculate the percentile rank, weighted by population. We use the "wtd.rank" function from the "Hmisc" package in R (Harrell Jr and Dupont 2022).

When we map historical gentrification in Chicago, we use a data set from Lee and Lin (2018). Lee and Lin (2018) reweighted census data dating back to 1880, providing consistent measures of population, house value, and income for most censuses over the last 150 years, harmonized to 2010 census tract data. Lee and Lin (2018) report *average* house price and household income, rather than median, differentiating the estimates from our more contemporary data. We calculate house value and income quantiles the same way as in the annual, contemporary case.

Dynamic Difference-in-Difference Event Study

Variable Selection and Creation Other than House Value and Income Quantiles Census and American Community Survey (ACS) data was downloaded from the National Historical Geographic Information System (NHGIS) database. The NHGIS has constructed crosswalks to facilitate re-weighting of prior-year census data onto 2010 tract geographies. We follow their weights. We use 1990 and 2000 census data in a few capacities. First, the files contain latitude/longitude coordinates for each tract, which we use to construct various spatial indicators. Second, we use census data to calculate initial neighborhood percentile rankings—as described above. Third, we use 1990 income and house value percentiles to determine if a tract was gentrifiable in 1990 for the purposes of our regressions, which determines gentrification status from 1998 onwards. Finally, we use various demographic and housing unit information as control variables in some regression models, described in more detail in the next section.

NHGIS also provides time-series consistent data for zip codes for 1990, 2000, and 2010 for some demographic and housing variables. We use these to the greatest extent possible. However, in some of our regression analysis, we wish to use some demographic indicators, such as poverty status or college attainment, that are not in the time series data. For these, we only have cross-sectional (rather than panel) data.

To calculate the distance to downtown, operationalized as distance to the central business district (CBD) we use a dataset for 383 MSAs compiled by Manduca (2020). Manduca (2020) uses estimates of job densities, housing densities, and gross number of jobs to estimate the main business district across MSAs. His estimate of the central business district almost always matches with a historical business district, though there are a few instances where the main job center today is a hospital or university, not in downtown.

Explanation of DiD Approach Difference-in-differences is a technique to estimate the effect of a treatment (in this case, gentrification onset) by comparing outcome variables of a control group and a treatment group before and after the treatment. In the current case, neighborhoods (operationalized as ZCTAs) are considered treated if they began to gentrify, and they are in the control group if they do not gentrify, are not gentrifiable. (The latter group includes neighborhoods that were *already* gentrifying as of 1998, as well as other high-income neighborhoods. Such neighborhoods are susceptible to advanced gentrification, but this regression is studying income growth at gentrification onset, not advanced stages). Gentrifiable neighborhoods are those neighborhoods in the bottom quartile of incomes in an MSA, while gentrifying neighborhoods have 25-percentile gap between house values and incomes. We only consider neighborhoods within the inner third of the MSA, given the traditional focus of gentrification scholarship.

A standard difference-in-difference approach has one "pre" period and one "post" period, resulting in an estimating equation:

$$y_{t,z,m} = \alpha tg + \delta g + \lambda t + \mathbf{X_{t,z,m}}\beta + \gamma_m + \epsilon \tag{1}$$

Where $y_{t,z,m}$ is the outcome variable, log of the average income of ZIP *z* in time period *t* in MSA *m*. A dummy variable, *t*, represents whether observations are is post-treatment, while a dummy variable *g* indicates whether a neighborhood gentrified (both variables take the value of 1 if the statement is true and 0 otherwise). In this basic setup, one interacts the time dummy and the treatment dummy to see the impact of being treated on the housing price, making α the main variable of interest. X is a matrix of independent covariates, β is the vector of coefficients to be estimated for the covariates. The variable γ_m is for MSA fixed effects, and ε is the unobserved error term. The above equation, thus, would estimate the impact of a neighborhood being gentrified, comparing it both to its pre-gentrification status, as well as to other neighborhoods, given the independent variables.

Difference-in-differences can be used when treatment takes place at different times, or to extend beyond a binary "pre" and "post" treatment and instead to focus on multiple periods before and after treatment (Callaway and Sant'Anna 2021; Sun and Abraham 2021). This describes our case: gentrification is not simultaneous across neighborhoods, and the immediate vs. later income effects may be quite different. In this case, our estimating equation is as follows:

$$y_{t,z,m} = \left[\sum_{j=-10}^{19} \mathbf{1}_{t-\tau_z=j} \times \alpha_j \right] + \mathbf{X}_{t,z,m}\beta + \gamma_m + \eta_t + \epsilon \tag{2}$$

Thus, rather than estimating the singular variable α, we estimate 30 different values for α_j, from j = -10 (10 years before gentrification onset) to j = 19 (19 years after gentrification onset). In this second equation, we have a binary indicator $\mathbf{1}_{t-\tau_z=j}$ to account for whether a neighborhood *z* is *j* years away from gentrification; it is in lieu of the two dummy variables *t* and *g* in the simple pre/post case.

The following variables are included in $\mathbf{X}_{t,z,m}$:

- Log of income in 1998, from the IRS
- Gentrification status in 1998, from our measure, based on gentrifiability in 1990
- Fraction of the ZCTA that is White, fraction Black, fraction of population older than 25 with a college degree, fraction of the population with income under the poverty line, fraction of owner-occupied housing units, fraction of units that are vacant, log of the number of units. All of these variables are calculated at the ZCTA level and come from the 1990 Census as provided by the NHGIS
- Distance quantile from the CBD, number of jobs within two miles, taken to the fourth root, both of which come from Manduca (2020)[3]
- A dummy variable as to whether the ZCTA is within 1 mile of a shoreline, from Lee and Lin (2018)

Lastly, we include fixed effects for MSAs (γ_m) and years (η_t). The former controls for the average income level in each metropolitan area while the latter controls for the average income by year for the whole sample. Standard errors are clustered at the MSA level. For our three regressions where we interact gentrification status with neighborhood characteristics, we create a dummy variable $\mathbf{1}_v$, to indicate whether the neighborhood is above the threshold ($\mathbf{1}_v = 1$) or below the threshold ($\mathbf{1}_v = 0$), resulting in the following estimating equation:

$$y_{t,z,m} = \left[\sum_{v=0,1} \sum_{j=-10}^{19} \mathbf{1}_v \times \mathbf{1}_{t-\tau_z=j} \times \alpha_{j,v} \right] + \mathbf{X}_{t,z,m}\beta + \gamma_m + \eta_t + \epsilon \tag{3}$$

The thresholds for $\mathbf{1}_v$ are as follows:

- $\mathbf{1}_v = 1$ if the share Black in a ZCTA is greater than 50%
- $\mathbf{1}_v = 1$ if the unit growth is greater than the median unit growth among gentrifying ZCTAs, translating to a cut point of 10%
- $\mathbf{1}_v = 1$ if the ZCTA has a distance quantile greater than 1/6th of the CBD.

Notes

1. https://www.huduser.gov/portal/datasets/usps_crosswalk.html
2. Tracts do not map perfectly to zip codes; we use the dominant zip code by land area.
3. We take the fourth root to transform the number of jobs to a smaller scale, in lieu of taking a log. This is because many ZCTAs have no or few jobs within two miles. See Gelman (2007)

References

Callaway B and Sant'Anna PH (2021) Difference-in-Differences with multiple time periods. *Journal of Econometrics* 225(2): 200–230. DOI:10.1016/j.jeconom.2020.12.001. URL https://linkinghub.elsevier.com/retrieve/pii/S0304407620303948.

Gelman A (2007) The 1/4-power transformation. URL https://statmodeling.stat.columbia.edu/2007/01/02/the_14power_tra/.

Harrell Jr FE and Dupont C (2022) Hmisc: Harrell Miscellaneous. URL https://CRAN.R-project.org/package=Hmisc.

Lee S and Lin J (2018) Natural Amenities, Neighbourhood Dynamics, and Persistence in the Spatial Distribution of Income. *The Review of Economic Studies* 85(1): 663–694. DOI: 10.1093/restud/rdx018. URL https://academic.oup.com/restud/article/85/1/663/3091867. Publisher: Oxford Academic.

Manduca R (2020) The spatial structure of US metropolitan employment: New insights from administrative data. *Environment and Planning B: Urban Analytics and City Science* DOI:10.1177/2399808320934821. URL https://journals.sagepub.com/doi/10.1177/2399808320934821. Publisher: SAGE PublicationsSage UK: London, England.

Moritz S, Gatscha S, Wang E and Hause R (2022) imputets: Time series missing value imputation. URL https://CRAN.R-project.org/package=imputeTS.

Sun L and Abraham S (2021) Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225(2): 175–199. DOI:10.1016/j.jeconom.2020.09.006. URL https://www.sciencedirect.com/science/article/pii/S030440762030378X.