The Price of Housing in the United States, 1890–2006*

Ronan C. Lyons† Allison Shertzer‡ Rowena Gray§ David Agorastos¶

May 2025

Abstract

We construct the first annual market rent and home sale price series for American cities over the twentieth century using over 2.7 million newspaper real estate listings. Our findings revise several stylized facts about U.S. housing markets. Real market rents did not fall over the postwar period in most cities and rose nationally by 60% from 1890 to 2006. We also document higher sale price growth between 1953 and 1987 relative to previous series, but price trajectories varied considerably across cities with demand and new housing supply. We find that rent-to-price ratios fell from about 8% in the early twentieth century to 3% by 2006, consistent with declines in the cost of owning housing relative to renting. We close by discussing the return to owning housing over the 20th century United States.

Keywords: Housing prices; rental indices; hedonic analysis; housing markets.
JEL codes: E3, N1, O18, R3

*The authors gratefully acknowledge the financial support of the National Science Foundation (SES-1918554), the Lincoln Institute of Land Policy, and Trinity College Dublin. We also acknowledge support from The Center for the Restoration of Economic Data (CREED) at the Philadelphia Federal Reserve and thank Grace Kneidel for her support. Many excellent postdoctoral, graduate, undergraduate and professional research assistants contributed to this project; see Appendix A for complete acknowledgments. We thank seminar audiences at Trinity College Dublin, Bonn, UIUC, UC Davis, UC Irvine, Caltech, NBER (DAE), NYU, Paris School of Economics, Philadelphia Fed, Stanford (Cities, Housing, and Society Workshop), Wharton, Wisconsin, Yale, the FHFA, and the Board of Governors. We benefited from conversations with Justin Contat, Morris Davis, Rebecca Diamond, William Doerner, Gilles Duranton, Barry Eichengreen, Fernando Ferreira, Daniel Fetter, Price Fishback, Ed Glaeser, Adam Guren, Jessie Handbury, Walker Hanlon, Philip Lane, Jeffrey Lin, Bob Margo, Raven Molloy, Tim McQuade, Thomas Piketty, Sarah Quincy, Jonathan Rose, Moritz Schularick, Todd Sinai, Bryan Stuart, Alan Taylor, James Vickery, and Maisy Wong. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

†Department of Economics and Centre for Economics, Policy & History (CEPH), Trinity College Dublin; email: ronan.lyons@tcd.ie

‡Research Department, Federal Reserve Bank of Philadelphia; email: allison.shertzer@phil.frb.org

§Department of Economics, University of California, Merced; email: rgray6@ucmerced.edu

¶email: dna17@pitt.edu

1 Introduction

Housing is both a key service and an important asset in the United States, with housing the largest single component of consumer expenditure and household balance sheets. Despite the centrality of housing to the American economy, existing long-run housing price series are very limited, particularly prior to the 1970s for home sales and prior to the 1990s for current market rents.[1] Scholars interested in rental prices have used the Bureau of Labor Statistics (BLS) Rent of Primary Residence (RoPR) series, which is based on contract rents and used to construct consumer price indices. Scholars interested in sale prices have primarily relied on the pioneering work of Robert Shiller, who spliced together several data sources to obtain a national housing sale price series beginning in 1890 (Shiller, 2015). Together, the RoPR and Shiller series underpin important recent work studying the U.S. housing market (Jordà et al., 2019; Knoll et al., 2017).

On the other hand, consistently constructed annual housing price series that span the 20th century do not exist for American cities, for either market rents or sales. Annual, quality-adjusted price series for rents in the open market for U.S. cities do not begin until the end of the 20th century or beginning of the 21st.[2] Annual, quality-adjusted sale price series for cities are not available until 1975, when the Federal Housing Finance Agency (FHFA) indices begin. Economists thus know relatively little about how the market prices of rented and owned housing have evolved within and across U.S. cities prior to the last few decades. As a consequence, we cannot measure the long-run return to owning housing — either the rental return or capital gain — across American cities over the 20th century. It has also been impossible to ask how market rent-to-price ratios have evolved with changes in credit conditions over the long run, and scholars have been limited in their ability to systematically investigate how or why housing prices have increased in some cities more than others. These questions are fundamental to our understanding of housing markets.

[1] We use the term "market rents" to mean the prices facing households signing a new lease. "Contract" or "average" rents from the BLS instead measure rental prices paid by all current renters. Because of stickiness in rental terms, contract rents typically lag market rents.

[2] For instance, the "Repeat Rent Index" of Ambrose et al. (2015) begins in the early 2000s.

This paper introduces new price series for both rents and sales constructed from archival newspaper real estate sections for 30 cities over the period from 1890 to 2006. These "Historical Housing Prices" (HHP) series draw on over 2.7 million listings that include a price as well as measures of size, type, and location within the city. Using hedonic methods, we construct new housing price indices for both the rented and owned segments of the housing market. We aggregate using population weights to construct national indices. These series are unique because they are constructed consistently across the period and reflect market conditions in each city at an annual frequency. In addition, and in contrast to many other datasets on housing sales or rental prices, the HHP series are based on publicly available sources and can be freely distributed.[3]

The first part of this paper compares the HHP series to existing sources of information on U.S. housing prices. Beginning with market rents, we find substantially more cyclical behavior in market rental price levels compared with the RoPR, which smooths market fluctuations by design and historically incorporated new construction with substantial lags. In addition, and in contrast to the RoPR series, we find no evidence that real rents fell during the 20th century in most cities. Rather, national rental prices in 2006 were 60% higher relative to 1890 and 36% higher relative to 1914. Furthermore, rents in the U.S. were cyclical within a relatively narrow band compared to sale prices but exhibited an upward trend over time. Our national series substantially agrees with adjustments to the RoPR series proposed by Gordon & van Goethem (2007) and Crone et al. (2010) that have been adopted by economists in recent work on national U.S. housing price levels (Knoll et al., 2017).

For sale prices, the HHP series aligns well with the Case-Shiller-Weiss and FHFA repeat-sales series at both the national and city level after 1987, despite the differences in methodology. However, the series diverge substantially in earlier decades precisely where historians of housing markets have noted limitations of the data sources underlying the Shiller index, particularly the reliance on a single retrospective survey for the 1890–1934

[3]The HHP indices are available at https://www.philadelphiafed.org/surveys-and-data/regional-economic-analysis/historical-housing-prices.

period (Fishback & Kollmann, 2014), the use of truncated loan data for the 1953–1975 period (Greenlees, 1982), and the inclusion of appraisals in the 1975–1986 period. Importantly, we find evidence of a sizable interwar housing cycle that is not present in the Shiller series. In addition, we find that U.S. housing prices began increasing in real terms earlier than in the Shiller index, which reports relatively little housing price growth from 1950-1995. Our series show inflation-adjusted sales price increases of more than 50% between the late 1960s and the late 1970s, in line with the work of Davis & Heathcote (2007) and Davis et al. (2007), who use census housing values.

Having established the usefulness of the HHP series, we next turn to trends in the price of housing for rent and sale over the 20th century. We categorize the housing market of 1890–2006 into three phases. The first phase from 1890–1940 saw a largely flat long-term trend in housing prices with significant cycles in both sale and rental prices around the Panics of 1893 and 1896. We also provide new evidence on the housing market dynamics around the Great Depression. Sales prices surged 70% over the 1920s and then fell over the 1930s, in line with the work of Fishback & Kollmann (2014) and Rose (2022). On the other hand, rents fell from 1925–1934 and began to rise thereafter. The second phase of the housing market from 1940–1970 was characterized by greater stability of prices and the beginning of real price growth in sales while rental price levels remained mostly flat. The final phase of the housing market from 1970–2006 was marked by faster and more volatile price growth across both segments. Much of the growth in sale prices occurred in the late 1970s, the 1980s, and the decade from 1996–2006. Rents increased much more slowly and at different periods, with growth in the late 1960s and the decade from 1975–1985 followed by periods of decline. A key aspect of this final phase of the 20th century housing market is the permanent divergence of the sale and rental prices of housing.

The second part of the paper uses the HHP series to document new facts about housing markets across the 20th century United States. We begin by observing that the aggregate growth in the sale price of housing did not occur evenly across cities, and in fact some cities saw little price growth between World War II and 1995. We investigate this pattern

from the perspective of housing supply and construction costs using long-run cost indices from RSMeans and newly digitized data on historical housing permits. We find that nearly every city in the U.S. responded to the 1920s housing boom by constructing new housing. However, the trajectory of cities diverged after World War II. Some cities, like Philadelphia and St. Louis, entered the postwar period with sufficient housing to meet much of the demand over the rest of the century. Other cities, such as Atlanta, continued to build to meet demand through the 1990s, consistent with work on the post-1985 period by Glaeser & Gyourko (2018). These two groups of cities exhibit relatively flat housing price series until the end of the period. The rest of the cities in the sample ceased building housing at some point after 1970, with San Francisco one of the earliest examples. We discuss the cities in this group and link the slowdown in housing construction relative to demand to zoning regulations (Saks, 2008).

We next investigate the divergence of market rents from sale prices in a user cost framework (Himmelberg et al., 2005; Poterba, 1984) and compute, for the first time, a rent-to-price ratio across the 20th century United States. The ratio was largely flat from 1890 to 1920 at around 8% before following a classic bubble-crash cycle during the interwar period. The rent-to-price ratio then fell unevenly from around 10% at the end of World War II to 3% in 2006. We discuss how this trend matched developments in housing finance, from the short-term home loans of the pre-1920 period to the emergence of the modern mortgage during the Great Depression. Years to maturity and loan-to-value ratios increased dramatically between 1935 and 1960, reducing the user cost of owned housing. The period after 1970 was characterized by fluctuations in real interest rates and falling down payment requirements within the framework of a 30-year, fully amortized mortgage that further reduced the costs of owning housing relative to renting.

Finally, we produce for the first time annual series reporting the rental return, capital gain, and total return to owning housing at both the national and city level for the full 1890–2006 period. In aggregate, the average annual real housing return across the 1890–2006 period was 8.6%, with gross rental returns accounting for 7.3% of this and capital

gains just 1.3%. Capital gains were largely unimportant until the 1940s, after which they averaged 1.7% in real terms for the rest of the 20th century. Capital gains were far more volatile and cyclical relative to rental returns. We discuss the implications of these series for our understanding of the rate of return associated with owning housing versus equities (Jordà et al., 2019).

The rest of the paper is structured as follows. Section 2 describes the HHP dataset. Section 3 outlines the construction of our new housing price indices. Section 4 compares the HHP indices to existing series and documents trends in the price of housing over the 20th century. Sections 5-7 present new facts on housing markets before the final section concludes.

2 HHP Newspaper Data

The price series for both rents and sales studied in this paper were constructed from archival newspaper real estate sections. The "Historical Housing Prices" (HHP) series draw on over 2.7 million listings from the 1890–2006 period. Before 1890, the secondary housing market was too small to yield a sizable set of listings in all but the largest U.S. cities and, after 2006, newspapers were increasingly eclipsed by the internet as the primary medium for advertising housing. However, during the twelve decades covered by the dataset, newspapers contain the most consistently collected information on the price of owned and rented housing that can be feasibly obtained for a large set of cities in the United States.[4] We chose 30 cities that represented a diverse sample in terms of region and economic trajectory and that had a complete and accessible newspaper repository. Appendix Table B1 reports our sample cities, newspapers, and access sites.

[4]Scholars focused on individual cities have been able to collect series of sales transactions, for instance Nicholas & Scherbina (2013) for Manhattan in the 1920s and 1930s. Collecting historical transactions requires scanning surviving records from local archives, a process that is infeasible for many cities. We are not the first scholars to instead rely on newspapers, as Rees & Jacobs (1961), Shiller (2015), and Fetter (2016) did so in their work covering multiple cities at various points in the 20th century. However, the HHP dataset is much broader in scope and, unlike Rees-Jacob and Shiller, adjusts for the mix of properties by within-city location, size, and type.

Our sampling procedure targeted 150 valid rental and 250 valid sale listings from each sampled newspaper, typically the last Sunday of the month of interest. To ensure that every area covered by the newspaper would be included in the sample, research assistants sampled across all columns in the real estate section. We collected listings only if they contained (1) a price, (2) a measure of size, (3) a property type (house or apartment), and (4) an indication of location. The location could be an address, an intersection, or an area within a city. Further details on the sampling strategy can be found in Appendix B.

Most major cities enter the sample in 1890, but some enter only when the secondary housing market became sufficiently dense. For example, due to its small size, Las Vegas did not have a substantial volume of listings until after World War II. Major differences in observation counts across cities typically reflect sampling effort rather than systematic features of the data, in particular whether we sampled one newspaper per year or four.[5]

Our dwelling size measure is generally total rooms before World War II and number of bedrooms thereafter, while early New York listings sometimes mention number of stories. The summary statistics for the sample can be found in Appendix Tables B3 and B4 for the rent and sale segments, respectively. The modal rental dwelling had two bedrooms and the modal home for sale had three bedrooms in most cities and periods. For location, we used a simple machine learning classification algorithm to standardize the location from the geocoded location information. We use a set of twenty standardized areas for each city to control for location; however, our indices are robust to alternative area controls. Intuitively, we allow the newspaper to define the housing market boundaries in each year, and so the geographic area covered by the housing market for each city grows over time as the associated metro area expands. Details on geocoding can be found in Appendix C.

Although transacted prices are the ultimate object of interest, we rely on listed prices for both rented and owned housing as a proxy. A limitation of our approach is that listing

[5]Half of the sample was collected prior to the Covid-19 outbreak, and we collected data from four newspapers per year for these cities. The other half of the sample was collected in 2020 and 2021 while navigating campus closures. For this half of the sample, we were able to collect data from one newspaper per year only, typically from May. New York has the largest dataset, reflecting other research being undertaken by two of the authors of this paper on sale and rental housing prices in that city.

prices may diverge from transacted rents and sale prices, particularly across business cycles (Han & Strange, 2014). Few empirical studies explore the relationship between these prices over the long run, but we expect that the same strategic considerations would have applied in the past, particularly that sellers cared about both the transacted price and time on the market and set listing prices with the goal of balancing these objectives (Haurin et al., 2010; Yavas & Yang, 1995). In more recent settings, a substantial share of houses transact at their exact listing price (Han & Strange, 2016) or close to it, even during periods of volatility (Lyons, 2019), which is consistent with real estate as a setting where learning valuations is costly (Lester et al., 2017). We leave it to future work to investigate further the relationship between listed and transacted prices over the long run and across business cycles.

We argue that our listing-based indices are nonetheless valuable for at least three reasons. Firstly, to date, there existed no annual city-level housing price series for owned housing in U.S. cities before 1975 or market rents for any city before the 1990s, so even a dataset based on listing prices is a significant resource for researchers. Notably, long-run housing price series have been assembled for many major European cities, beginning with the work of Eichholtz (1997) on Amsterdam, and the HHP series provide a useful counterpart for scholars interested in U.S. cities. Secondly, our sale price and rental series align well with city-level series from the FHFA after 1975 and Case-Shiller after 1987, lending credence to the earlier decades of the HHP dataset and to our overall hedonic methodology. We discuss these comparisons in the next section. Finally, cross-sections of our housing price data match up well with the distribution of housing prices from the census in many cities and years irrespective of the business cycle, particularly once we control for number of rooms. While the census data do not contain transaction prices either, the self-reported values have been widely used. We discuss these comparisons in detail in Appendix D.1.

3 Price Index Construction

Measuring the historical performance of housing markets is challenging because of the difficulty of observing property characteristics back in time. In this section, we discuss how we construct price indices from the newspaper data and compare them to existing sources of information on both national and city-level housing markets for the United States. In more recent settings, scholars are often able to observe the same property more than once, allowing the construction of "repeat sales" indices that minimize bias associated with unobserved quality changes (Bogin et al., 2019; Contat & Larson, 2022). However, constant-quality assumptions can bias these indices (Nowak & Smith, 2020) and the requirement that a housing unit be observed more than once means that indices are often based on a very small portion of the overall market (Nagaraja et al., 2014). In any event, a repeat sale index is not possible using newspaper listings, as the same property is unlikely to appear more than once across sample years.

We instead use a hedonic model with controls for both observed and unobserved housing unit attributes. Shen & Ross (2021) caution that hedonic methods may be biased upwards due to unobserved quality changes. To address unobserved quality changes over the long run, we use a "rolling windows" methodology (Silver, 2016). Intuitively, this approach avoids a regression where all years of the dataset are combined, which imposes a fixed vector of coefficients on size, type and location variables. This is undesirable because the relative price of a housing unit's size or location is likely to change substantially over time. We instead employ a rolling windows (RW) approach using the listings from two years in sequential regressions to allow for changes in coefficients for size, type and location over longer periods. This approach better measures the true like-for-like change in prices. We use two-year rolling windows as our baseline since it is the most conservative specification our annual data allows, with three- and five-year rolling windows as robustness checks.

In both sale and rental specifications, the regressors include measures of size, hous-

ing type, and standardized location within the metro area as well as the year of listing, which is our principal regressor used to construct the price indices. For rents, we standardize rental prices so that they are expressed in monthly terms, where necessary, and also include the rental frequency as an additional regressor.[6] We control flexibly for each size measure using dummies, allowing for any individual measure to be missing, with bathrooms rounded to the nearest half. Type is standardized to house or apartment.

The hedonic pricing model is thus the following regression equation, where we create an index for one city at a time. For each city c and for a particular base year BY we run:

$$\ln(\text{Price})_{ict} = \alpha + \underbrace{\sum_{\min(Y),y\neq BY}^{\max(Y)} \beta_{cy}\ 1_{(y=t)}}_{\text{Coefficients of Interest}} + \underbrace{\mathbf{X\Gamma}}_{\text{Controls}} + \varepsilon_{ict} \tag{1}$$

The rolling windows approach means that we run up to 116 separate regressions for each city, with regressions each covering the base year plus the rolling window length. For example, to create an index with rolling windows of size two years, we would run the above regression for 1890–1891 to obtain the coefficient for 1891, 1891–1892 to obtain the coefficient for 1892, and so on.[7] Controls include location within the city, size (rooms, bedrooms, bathrooms, stories), dwelling type (house or apartment), and rental payment frequency.

Across the almost 7,000 city-segment-year combinations, there were inevitably some cases of observation counts too low for the two-year RW to generate reliable results. To

[6] Rental frequencies are stated often but not always: approximately 0.77m of our 1.23m rental listings do not have a stated rental frequency. Of approximately 0.46m stated frequencies, the vast majority (0.44m) are monthly with most of the remainder weekly (over 22,000, while the remaining 6,000 observations are annual). We use city-year thresholds to identify frequencies where not stated. In most cases, this is straightforward as monthly rentals dominate, particularly after WWII. Weekly rentals were more likely to occur in the 1930s than in other decades while annual rentals were most common in the 1890s. To impute rental frequency when it was missing, we compared the three-year rolling average of the 90th percentile for weekly rents with the 5th percentile for monthly rents, and similarly the 95th percentile for monthly rents with the 10th percentile for annual rents. Values less than 60% of the 5th percentile of monthly rents were classed as weekly, values more than 5 times the 95th percentile were classed as annual, and values between the 5th and 95th percentile were classed as monthly. This rule reduced the number of rental listings with unknown frequency to just 33,000.

[7] If the missing city-year segment was at the start or end of the series, we assigned a weight of zero so that the city contributes to the national index for a particular segment only in years where the city index is defined.

address cases of insufficient listings in some cities, particularly rental listings during World War II and sale listings during the Great Depression, we use longer window lengths to obtain more stable year coefficients where necessary.[8] All cases of missing observations and exceptions to the two-year rolling windows specification are listed in Appendix Table B2. Starting in each city's earliest year, we take the coefficients for subsequent years from the relevant regression to build a city-specific index over time.

$$\iota_t = \iota_{t-1} \exp(\beta_t) : t \in \{1891, 2006\} \tag{2}$$

In our baseline, we aggregate city-level percentage changes each year by the city's population share that year to create the corresponding national index, effectively a Paasche Price Index. Population weights are constructed using the metropolitan area population from the Census Bureau and interpolated between census years.[9] It would be preferable to weight by households, rather than individuals. However, as households in 1890 cannot easily be measured, we use persons in the baseline. Appendix Figure E1 shows, for the sample from 1900 on, that using population rather than households as the source of weights does not significantly affect our overall results.

We subject our indices to a number of other robustness controls, including for rolling window length and area definitions. Figure E2 shows the series using alternative rolling window lengths. In general, these display very similar trends, with the main difference being between the two- and three-year rolling windows for the rental sector. This difference is driven overwhelmingly by the period 1944-1947. Given the allowances made for smaller samples due to WWII rent controls in this period in our baseline specification, we prefer the adjusted two-year RW over an unadjusted three-year rolling window. The hous-

[8] In a few cases and despite our best archival efforts, we were unable to locate a newspaper for a given year and city. In these cases we interpolated across the missing year to obtain index values.

[9] We match our city to the corresponding metro area as defined by the census and allow MSAs to expand over time with population growth. The components of census metro areas we used to construct population counts for the sample period can be found at https://usa.ipums.org/usa/volii/county_comp1a.shtml.

ing price indices with and without adjustments for area are shown in Appendix Figure E3. The inclusion of controls for area modestly reduces price inflation, especially for rents after WWII.

Indices of real sale and rental prices for all 30 cities are presented in Figure 1.[10] The city-level series show local booms and busts, including in the earlier decades when measures of local economic activity are scarce. For example, the 1893 Panic and the resulting economic downturn that affected in particular cotton- and grain-exporting states are visible in the data. Sale prices fell by at least a third in a number of Southern cities (and by more than half in Nashville and Dallas) over the course of the 1890s. Secondly, the impact of the 1906 San Francisco earthquake is also obvious. Rents in that city were over 50% higher in 1907 than in 1905 while, in contrast, rents in Los Angeles actually fell slightly in the same two-year period. And, perhaps most prominently, the 1925 Florida land boom is immediately evident in the Tampa and in particular Miami series (Calomiris & Jaremski, 2023; Eichengreen, 2015). Rents in Miami in 1925 were over three times their 1922 level, while sale prices almost doubled between 1924 and 1925 alone, before falling back down.

The city-level figures also show that cities had different long-run trends over the course of the long 20th century. Some cities saw rents rise since World War I (San Francisco), others had largely flat rents over the whole period (Atlanta), and still others saw rents decline since World War II (Detroit). The extent of sale price growth in the postwar era is also very different across local housing markets. For instance, the sale price index of Louisville is largely flat from 1948 to 2000 while sale prices have increased consistently in San Diego since 1970. While the extent of housing price growth varies across cities, the divergence of sale prices and rents in the second half of the period is common across most cities in the sample. Section 4 goes into more detail on broad patterns in the data, while Sections 5 and 6 discuss the changes in housing supply and credit conditions, respectively, that have accompanied these long-run trends. Section 7 investigates how these trends have shaped

[10]Throughout the paper we put our series in real terms to facilitate meaningful comparisons over time. We use an extended CPI-U to deflate rather than alternatives such as the PCE deflator, as it is available throughout our period on an annual basis. Specifically, we use Officer-Williamson (https://measuringworth.com/).

Figure 1: HHP Real HPI and RI by City, 1890–2006

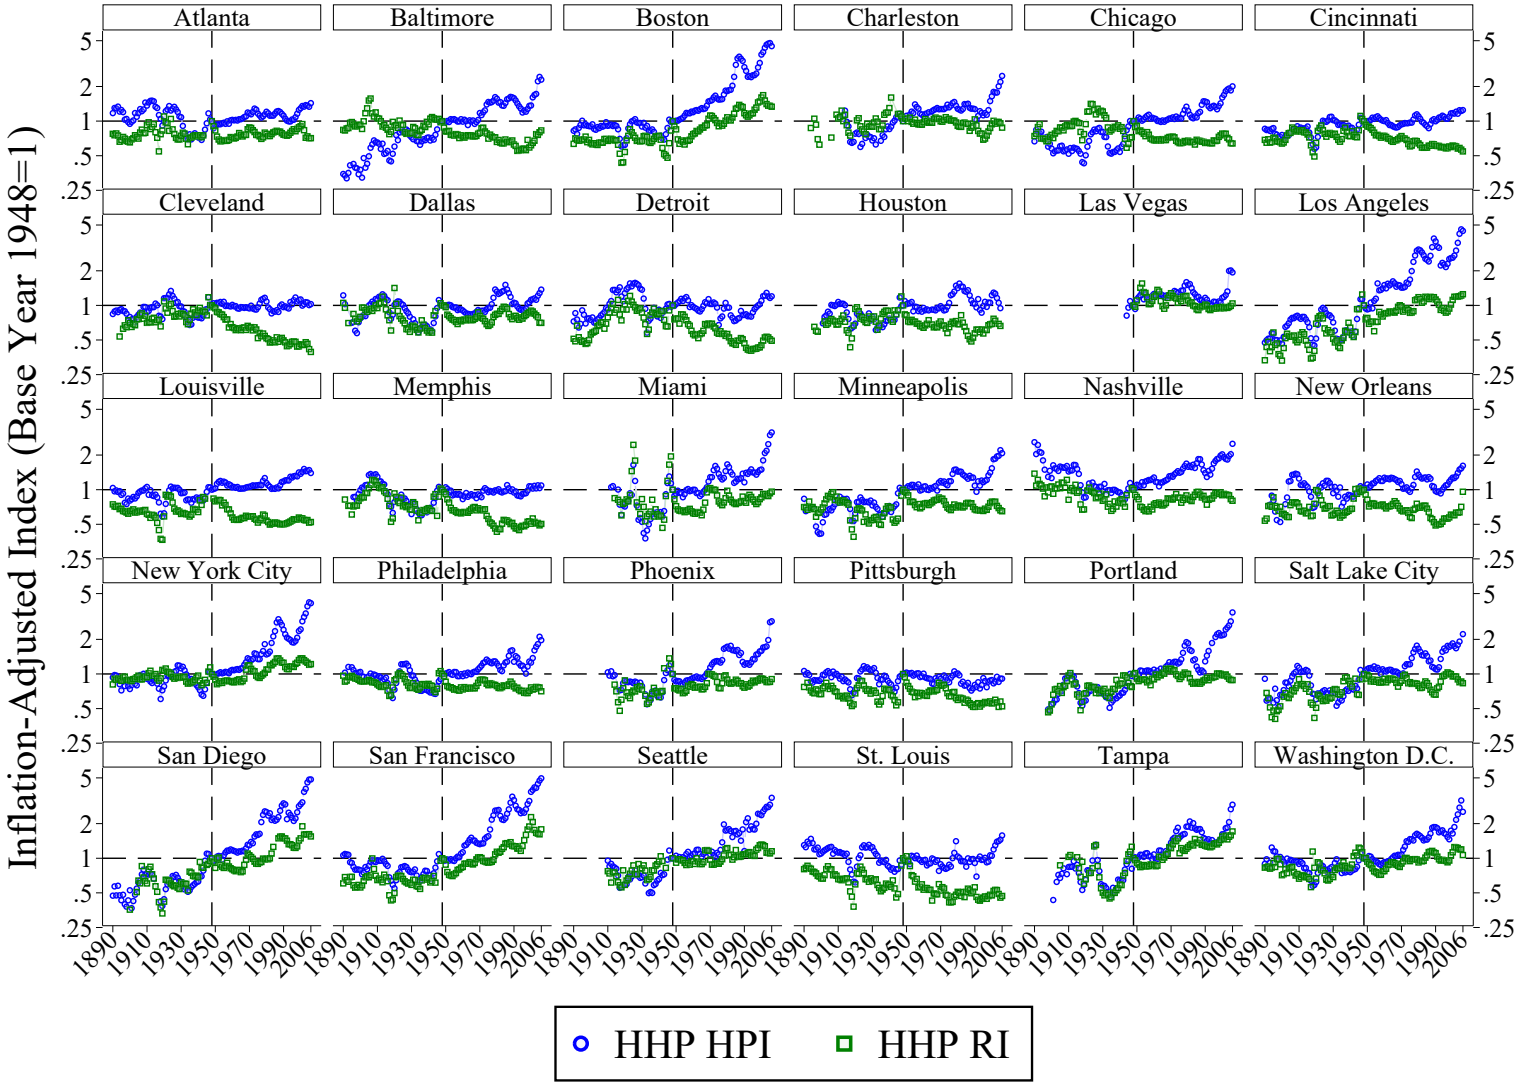

Note: This figure shows the baseline rental and sale price indices by city in real terms and with a base year of 1948. All series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

the returns associated with owning housing over time.

4 The HHP Rent and Price Indices

In this section, we first compare the HHP indices to the widely-used indices that underlie much of the previous work on U.S. housing markets and explain how the consistent newspaper hedonic methodology underlying the HHP addresses many of the limitations of existing series. We then identify and discuss three phases in the evolution of sale and rental prices between 1890 and 2006.

4.1 Comparisons to Existing Series

Rental Series In Figure 2 we compare the HHP national rental price series to existing national rental price series that have been widely used in economics, with 1915 set as the base year to match the RoPR series. Throughout this section, we include recession shading on each figure corresponding to the NBER recession indicators.[11] For the period prior to 1914, we compare our series to the only source for rental prices, Rees and Jacobs (1961), a series based on the unweighted, unadjusted average rental price from six city newspapers covering the 1890–1914 period.[12] Both series show similar cycles over the 25-year period, with peaks in the early to mid-1890s and again around 1906. Nonetheless, there are important differences: while the Rees-Jacobs (RJ) index implies real rents were roughly 5% lower in 1914 than in 1890, the HHP series implies real rents were about 17% higher and largely stable in the decade leading up to World War I.

We next compare our market rental index with the BLS RoPR series, which has been collected since 1914.[13] The comparison is not strictly like-for-like: the RoPR series is a

[11] Federal Reserve Bank of St. Louis, NBER based Recession Indicators for the United States from the Period following the Peak through the Trough [USREC], FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/USREC.

[12] Tables 22 and 32 in Chapter 4 of Rees & Jacobs (1961) provide estimates of market rent indices for six cities and a national index based on the unweighted average of their levels.

[13] The BLS RoPR series can be accessed from the St. Louis Federal Reserve Bank FRED website: U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: Rent of Primary Residence in U.S. City Average [CUUR0000SEHA].

Figure 2: HHP Rental Series Compared with BLS and Rees-Jacobs

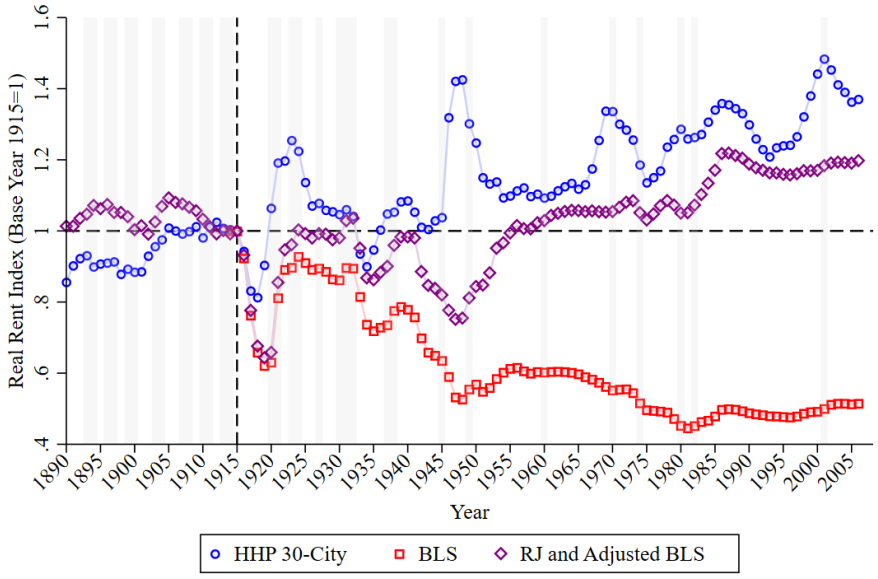

Note: The figure shows the HHP 30-city rental price series, the BLS Rent of Primary Residence series, and a rental series spliced together from Rees & Jacobs (1961) and the Rent of Primary Residence series with adjustments similar to those proposed by Crone et al. (2010) and Gordon & van Goethem (2007). All series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

measure of average rents paid by renters used to compute the Consumer Price Index. The RoPR series is not intended to capture market conditions, meaning the prices facing renters signing new leases. We are nonetheless interested in this comparison since average rents should converge to the market rate over the long run, and the overall trajectory of rental prices is of interest. The HHP series points to an upward trend in real rents, with 2006 real rents 60% higher than in 1890 and 36% higher than in 1914. The upward trend of 0.5% per year is in sharp contrast to the RoPR series, which reports that real rents fell by nearly half between 1914 and 2006 (or -0.6% per year, on average).

Interestingly, some cities *do* exhibit falling real rental price levels in the HHP data, including Detroit and Cleveland (we report MSA-level comparisons to the RoPR, where available, in Appendix Figure D4). However, for most cities, the fall in rents in the BLS series is not evident in the HHP data. The difference likely relates to well-known limitations of the methods for collecting data on rents over time, as described in detail by Crone et al. (2010), in particular non-response bias caused by vacant units. This was not an issue prior

to 1942, when the BLS obtained its information on rents through what was effectively a dwelling-level survey (Humes & Schiro, 1948, 1949). The switch to a household-level survey meant that, from 1942 onward, the Bureau of Labor Statistics (BLS) largely omitted rent increases when tenants changed or when units were vacant, leading to a downward bias in measured inflation. Index compilers were aware of this vacancy bias and attempted to address it in numerous ways, making several substantive changes between 1953 and 1994.[14]

The HHP series, through use of listings, does not suffer from non-response bias. The risk from listings is instead a "non-listing" bias, which may occur during periods of market scarcity (particularly wartime rent controls), when properties listed may be systematically different from both the market as a whole and listings in other periods. To investigate this possibility, we compare our series to adjustments to the RoPR proposed by Gordon & van Goethem (2007) for 1915-1939 and Crone et al. (2010) for 1940–1995. As shown in Figure 2, we find reasonably close agreement with the Gordon & van Goethem (2007) series, with both series seeing similar real rental price levels in 1914 and 1939. For the period since 1940, by construction, the Crone et al. (2010) series matches the cycles in the original BLS series, resulting in substantial differences in the cycles. But, it is the overall level of growth that is more of interest and on this point there is agreement between the two series with real rents 15%–25% higher in 2000 than in 1940. Appendix Figure D5 benchmarks the HHP series to each adjustment discussed here.

Our finding of rising real rental prices over the 20th century thus accords with the most influential proposed revisions to the RoPR series. In aggregate, the radically different assessment of the path of rents over the long run in HHP compared to BLS – with rents 60% above, rather than 50% below, their 1890 level in 2006 – has implications not just for the measurement of housing prices but also for the wider cost of living. Given the weight

[14]Earlier changes include a shift from quarterly to semiannual rent collection in 1953, which reduced the number of instances of no change in rents, and an increased reliance on personal visits and telephone surveys (1964). Later changes include one in 1978 aimed at reducing non-response rates by new tenants that effectively allowed responses from the landlord or building manager. Further adjustments were made to correct for vacancy bias with final changes made in 1988 (correcting for aging bias) and in 1994 (when one-month changes were dropped, eliminating the risk of recall bias).

of shelter in household expenditure throughout the period, and the role of market rents in calculating housing costs for owner-occupiers (Ozimek, 2013), substantially greater price growth in rents would mean faster overall inflation and thus slower growth in real living standards over the 20th century.

Figure 3: HHP Sale Series compared with the Case-Shiller Index

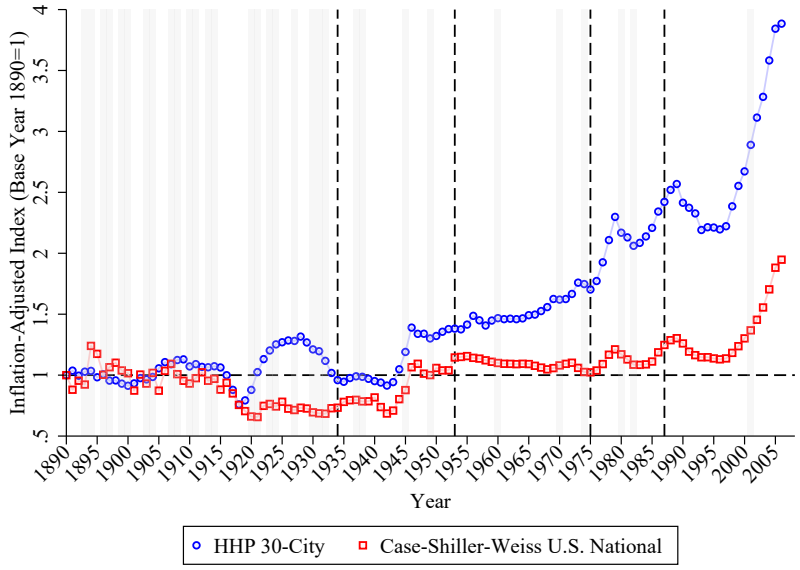

Note: This figure shows the HHP 30-city sales price series against the Case-Shiller-Weiss U.S. National sales price index obtained from Robert Shiller's website, http://www.econ.yale.edu/~shiller/data.htm. The dashed lines denote different data sources used by Shiller, as described fully in Appendix D. Both series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

Sale Series Next, in Figure 3, we compare the HHP sale price series to the index proposed by Shiller (2015), with vertical lines indicating when the underlying source data for the Case-Shiller index changes.[15] The period 1890–1934 relies on a single retrospective survey, the period 1935-1953 on a five-city newspaper sample, and the period 1953–1975 on the home purchase component of the CPI. The period from 1975 onward is based on two repeat sale indices: the OFHEO (now FHFA) index was used from 1975 to 1986 and the Case-Shiller-Weiss (CSW) repeat sales index from 1987 onward.

[15] We obtained these data from Robert Shiller's website, http://www.econ.yale.edu/~shiller/data.htm.

We discuss each of these components in turn in Appendix D and benchmark the HHP series by component in Appendix Figure D6. To summarize this exercise, we find strong evidence that the Grebler et al. (1956) survey that comprises the 1890–1934 period misses much of the interwar housing cycle, consistent with the work of Fishback & Kollmann (2014) using the census and New Deal data sources. We also find that the data Shiller relied on for the 1953–1987 period understates housing price growth, mainly due to the exclusion of many sales that did not meet the FHA conforming loan limit from 1953-1975 and the use of appraisals in the repeat sales index for 1975–1986.[16] This missed housing price growth in the Case-Shiller index is consistent with limitations of the underlying data noted by Greenlees (1982) and Davis & Heathcote (2007). Overall, we find real sale prices grew 142% between 1890 and 1987 compared with just 37% in the Case-Shiller index. On the other hand, the HHP series aligns very closely with the CSW portion of the Shiller index for 1987–2006, with increases of 60% and 65%, respectively, in real housing prices.

Appendix D contains a wealth of other benchmarking figures for the interested reader, including city-level comparisons to the FHFA indices from 1975–2006 in Appendix Figure D7, the S&P Case-Shiller indices from 1987–2006 in Appendix Figure D8, and the early Case-Shiller indices from 1970–1986 in Appendix Figure D9. These figures generally show substantial agreement in spite of the different methodologies.[17] The high degree of agreement in later years provides evidence of the robustness of our method earlier in the 20th century.

[16]The impact of including appraisals in repeat-sales indices is ambiguous. Appraisals may be positively biased (Leventis, 2006) or may lag market conditions as appraisers seek to ground their valuation using past transactions. Consistent with Shiller's own assessment of this data (Shiller, 2015, p. 235), we find that the use of appraisals appears to have lowered measured price growth between 1975 and 1986. For more information on the OFHEO data used in the Case-Shiller index during this period, see Leventis (2008).

[17]For instance, the Case-Shiller indices are arithmetic indices while the FHFA indices are geometric with a correction to make them more comparable to arithmetic indices (see Calhoun (1996)). Our log-linear hedonic specification means that our index is closer in spirit to a geometric mean, as it works in log points (similar to percentage changes) rather than dollar changes in prices.

Figure 4: HHP Rent and Sales Series

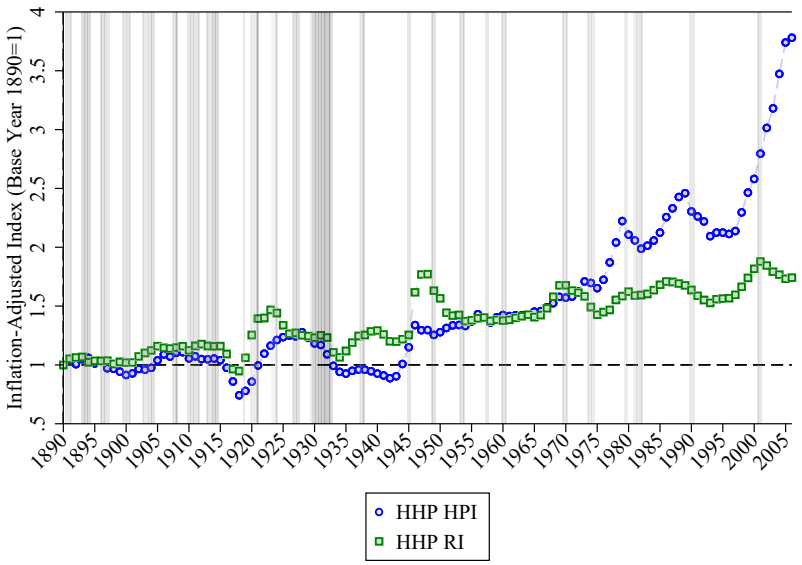

Note: This figure plots the baseline real HPI and RI aggregated by population. Both series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

4.2 Housing Prices, 1890–2006: Three Phases

Having established the usefulness of the HHP series, we now turn to what we can learn about the evolution of housing prices over the 20th century. We proceed by categorizing the housing market of the 1890–2006 period into three phases: (1) the era of early booms and busts around a largely flat trend before WWII, (2) the era of stability and slowly increasing prices after WWII and before 1970, and (3) the era of sharply increasing prices and volatility after 1970.

Phase 1: Early Cycles During the first phase from 1890–1940, the long-term trend rate for real housing prices in the U.S. was stable or even slightly negative, with significant cycles in both sale and rental prices. As shown in Figure 4, sale prices peaked before both the 1893 and 1906 panics, with a trough around 1900. Largely due to wider inflation, real prices fell sharply in the late 1910s. The picture for rents before the 1940s is similar. The panics of 1893 and 1906 are also evident in this series, with peaks in the early 1890s and mid-1900s.

While real rents were largely stable for the next decade, they fell sharply between 1916 and 1918 before surging in the following five years, peaking in 1923 one quarter above their 1914 level and almost 50% above their 1890 level. Overall, rents in the late 1930s were very similar to rents in the early 1890s.

One of the most important historical contributions of the HHP series is providing new evidence on the trajectory of housing prices around the Great Depression. Real sale prices rose by almost three quarters over the Roaring Twenties, peaking in 1928 at 23% above 1890 levels. Sale prices fell consistently over the Great Depression, and in 1940 prices were almost one-fifth below their 1890 level. Sale prices would not return to their 1928 peak until 1946. Importantly, we find that sale prices were 20% lower in real terms in 1940 relative to 1930 (rather than 10% higher as suggested by the Shiller index), consistent with evidence from the census (Fishback & Kollmann, 2014). Meanwhile, real rents peaked five years earlier than sale prices and then fell by almost one third over the early Depression, bottoming out in 1934. The final rental market cycle of this era occurred in the later Depression years, with rents rising from 1934–1940 and falling in real terms from 1940–43.

Phase 2: Steadier Growth and Stability The second phase of the U.S. housing market lasted from 1940 through 1970. This phase was marked by greater stability of prices and a positive trend rate of growth. In the case of sale prices, after increasing by over half in real terms between 1942 and 1946, prices rose by just 17% over the following 25 years, implying an average growth rate (AGR) of just 0.6%. Further, while prices fell slightly in real terms in the late 1940s and again in the late 1950s, in general growth was steady during this period. In only three of these 25 years was the change in real prices in either direction greater than three percent.

For rents, the initial impact of World War II was similar, with real rents rising 40 percent 1943–1947, a period largely under federal rent control. As noted above, our indices rely on dwellings listed for rent during these years. Given the prevalence of rent controls, this is unlikely to be a representative sample of all rental dwellings. Thus, rent controls

themselves may have both moderated rent increases for sitting tenants and increased market rents by reducing mobility. Nonetheless, in most city/year-pair combinations there are sufficient listings to reliably estimate like-for-like changes in market rents. Unlike with sales, which were not subject to government controls, market rents immediately fell (by 20%) once restrictions were removed. Thereafter, rents were largely stable from the mid-1950s to the mid-1960s before a substantial increase in the late 1960s. Overall, real rents in the early 1970s were only about 15% above their 1914 level. Rents experienced their flattest and least volatile period of the 20th century from 1950 to 1970.

Phase 3: Growth & Volatility The final phase of the U.S. housing market began around 1970. This period was marked by faster but more volatile growth in housing prices, especially sale prices. While the average growth rate of sale prices was -0.1% during 1890–1939 and 0.6% during 1946–70, it was 2.5% for the period 1971–2006. The equivalent AGRs for market rents were 0.5%, -0.1%, and 0.1%, respectively. The very strong overall growth in sale prices after 1970 occurred in three cycles, the first in the late 1970s, with real sale prices rising 35% between 1975 and 1979. A modest downturn over the following three years gave way to another run-up in prices: the real price in 1989 was 25% higher than in 1982. A more protracted downturn then followed, with prices bottoming out in 1996. The final decade (1996–2006) saw prices rise by 77%. Between 1945 and 1979 prices almost doubled and increased a further 70% by 2006.[18] This was in sharp contrast to the first half of the 20th century, where real sale prices were similar in 1945 to levels in 1932, 1914 and the mid-1890s.

The growth in housing prices from 1970 and 2006 was unprecedented in U.S. history. The other remarkable aspect of housing prices during this period was the decoupling of sale and rental prices. While sale prices were 140% higher in 2006 than 35 years earlier, rents were just 5% higher. We discuss this divergence further in Section 6. Alongside the growth in sale prices was a notable increase in volatility. While the change in sale prices

[18]It is worth noting the scale of the downturns in the early 1980s and early 1990s, however. Adjusting for inflation, prices in 1998 were only slightly above the 1979 peak (and still below the 1988–89 peak).

was greater than three percent (in either direction) in just three of the 25 years leading up to 1971, a change of at least this magnitude happened in 18 of the 35 years between 1971 and 2006. This was unique to the sale segment, though. In the rental segment, there were more substantial year-on-year changes of at least three percent in either direction in the second phase (eight out of 25 years) relative to the third (six out of 35 years).

Nonetheless, there were substantial market cycles in this third phase in market rents, which enjoyed strong growth 1966–69 and then a fall in the early 1970s. Between 1975 and 1987, market rents rose by almost one fifth, before falling back 11% from 1987–1993. Then, in the run-up to the Great Recession, and similarly to the Great Depression, rents peaked five years earlier than sales. While sale prices peaked in 2006 at the end of our sample, rents peaked in 2001, having risen 23% in the previous eight years. Overall, real market rents rose by 31% between the bottom of the cycle in 1975 and the top of the cycle in 2001. Sale prices, despite the different timing of their cycles, rose by twice as much as rents (67%) 1975–2001. To summarize this phase, the rental and sale segments of the housing market showed both different trend rates and greater cyclicality in this third and final phase of the U.S. housing market.

We have thus far demonstrated the usefulness of the HHP series for rents and sales and used them to document three eras of housing price movements over the 20th century. In the second part of the paper, we use the HHP data to revisit several significant questions about the long-term evolution of housing markets in American history.

5 Housing Prices and New Housing Supply

In this section we explore one of the starkest facts evident in Figures 1 and 3: while national sale price levels began to climb in earnest by the late 1960s, this growth varied substantially across cities. Our goal is to compare the price of housing for sale relative to the costs of supplying new housing over the long run. Following Glaeser & Gyourko (2018), we consider a housing production function with three components: land (L), construction

Figure 5: Housing Prices Relative to Costs and New Housing Permits

Note: The city-level BLS building permit reports for the 1920-1959 period were accessed from the FRASER website. The Building Permit Surveys were accessed from HathiTrust for 1959-1979, the Census Bureau's website for 1980-1987, and the FRED database at the St. Louis Fed for each city from 1988-2006. See Appendix B for details. The HPI relative to construction cost is the HHP baseline HPI scaled by the city-level costs indices from the RSMeans "Square Foot Cost" volumes (1950=1).

costs (CC), and a measure of entrepreneurial profit (EP) for the builder. The production cost (PC) of housing in period t is thus:

$$PC_t = (L_t + CC_t) \times EP_t \tag{3}$$

Developing consistent measures of housing production costs over such a long time period is complicated. The price of vacant land in the U.S. is particularly difficult to measure. Economists who have developed methods to infer land values using construction costs for recent decades (Davis & Heathcote, 2007; Davis & Palumbo, 2008) argue that land prices have increased in most cities, with substantial volatility, and at a faster rate than structure prices. Descriptive work from earlier in American history suggests that land markets were also highly volatile before World War II, particularly in the late 19th century (Hoyt, 2000). We leave the estimation of land values from the HHP data to future work and focus directly on construction costs.

Some of the earliest, consistently measured information on city-level construction costs comes from the home building industry, specifically the cost index from RSMeans that begins in 1940.[19] This index is available every five years from 1940 to 1980 and annually thereafter. Appendix Figure D10 shows the construction cost series relative to real housing price levels for each city from 1940 onward. During the immediate postwar period (1945–1970), housing prices grew more slowly than construction costs in every city except Boston and Los Angeles. However, in the last few decades of the 20th century (1970–2003), housing prices grew faster than construction costs in every city in the sample, suggesting a rising role for land scarcity and regulatory constraints. We return to this idea below.

To investigate how housing supply responded to changes in the price of new housing relative to costs, we also require information on changes in the housing stock each year. Starting in 1959, annual construction permits at the level of the metropolitan statistical

[19]These historical costs indices can be found in the back of the RSMeans "Square Foot Cost" volumes. We thank Raven Molloy for providing us with a transcription of this index.

area (MSA) are available from the Census Bureau's Building Permit Survey. To estimate new housing construction on an annual basis earlier, we use surveys of housing permits conducted by the Bureau of Labor Statistics (BLS) for cities from 1920 to 1950.[20] We scale the BLS housing permit data by the corresponding city population and the Building Permit Survey by the MSA population, interpolating between each census year.

We construct a housing price index relative to construction costs by scaling our baseline sales price index by the construction cost series obtained from RSMeans. We plot the adjusted index for each city in our sample (with base year 1950=1) in Figure 5. We also plot permits by population measures (for city or MSA) as our measure of housing construction intensity. The cost and construction series in this figure thus show the long run trends in housing prices relative to costs and new housing construction relative to population, complementing research focused on the period after 1980 (Glaeser et al., 2005).

Although our cost-adjusted housing price series does not go back to the 1920s, cities in the United States experienced a surge in demand and price growth during this decade of industrialization and urbanization, with national housing price levels increasing by 70%. We see that virtually every major city in the United States undertook new housing construction in response to the booming demand of the 1920s, which was perhaps the only decade of the last century in which every major city could be classified as a growing, elastically supplied market (in the parlance of Glaeser & Gyourko, 2018). New York City built more housing in 1927 alone than in the entire period from 1965–1987. The Florida cities of Miami and Tampa saw particularly dramatic housing construction booms in the mid-1920s.[21] Housing construction fell steeply everywhere during the Great Depression, however. Most cities saw relatively flat construction levels during World War II, although some cities with substantial wartime economic activity saw growth, including San Diego and Washington D.C.

[20] The city-level BLS building permit reports for the 1920-1959 period were accessed from the FRASER website. The Building Permit Surveys were accessed from HathiTrust for 1959-1979, the Census Bureau's website for 1980-1987, and the FRED database at the St. Louis Fed for each city from 1988-2006. See Appendix B for details on the sources for the building permit series.

[21] The only sample city for which we do not have housing market information is Las Vegas, which was a hamlet of just over 2,000 residents in 1920.

It is in the postwar era that the housing supply trajectories of cities in the United States began to diverge, principally between those that entered the period with a sizable housing stock relative to population growth and those that still needed to build. The "Low Demand" metropolitan areas in our sample include the northeastern and midwestern industrial centers of Baltimore, Chicago, Cincinnati, Cleveland, Detroit, Louisville, Minneapolis, Philadelphia, Pittsburgh, and St. Louis. All of these cities saw very low levels of per capita housing permitting activity alongside flat or modestly increasing housing prices relative to costs. The data do not reveal a uniform experience across this group. For instance, Chicago saw more price growth relative to costs after 1970 while both Baltimore and Louisville saw some housing construction after World War II. Detroit is remarkable in the extent to which housing prices relative to construction costs declined in the postwar period, reflecting the durability of the prewar housing stock in the presence of industrial decline (Glaeser & Gyourko, 2005).

The cities in the "High Demand" group that still needed to build substantial housing to accommodate their postwar populations looked reasonably similar in the 1950s and 1960s. Housing demand was met with new housing construction, leading to a volatile permitting series alongside relatively flat sale prices relative to construction costs. New York City is an exception here, as the country's largest city saw both limited construction and little price growth in the immediate postwar decades. However, New York City became a High Demand metropolitan area by 1980. What differentiates the cities in this High Demand category is whether they continued building housing to meet demand. Some of the Sun Belt cities continued building housing throughout the sample period, including Atlanta, Charleston, Dallas, Houston, Las Vegas, Memphis, New Orleans, Nashville, Phoenix, and Tampa. These cities saw relatively stable prices relative to construction costs but substantial construction activity. Again, there is some heterogeneity here, as Phoenix saw more cost-adjusted price growth relative to the other cities. However, it is clear that there are two different types of cities with stable housing prices over the postwar period: those that experienced low demand for housing and those that built housing to meet higher demand.

The other cities in the High Demand category vary by the year in which housing construction in response to demand came to a halt. Los Angeles was one of first cities to largely stop building, with booming housing prices and relatively little construction as early as 1970. San Francisco and Miami also built minimal new housing in the 1970s, although demand pressures were substantially larger along the California coast relative to Florida. San Diego still built substantial housing in response to increasing demand in the 1970s, but construction slowed relative to price growth by the 1980. The 1980s also saw the end of new housing construction in response to demand in northeastern metropolitan areas including Boston, New York City, and Washington D.C. Mountain and northwestern cities including Portland, Salt Lake City, and Seattle met demand with new housing construction until closer to 1990, when sale prices relative to costs began to climb while permitting activity slowed.

Why did some cities in the United States stop meeting demand for housing with new construction? One possible explanation is land use regulation. To explore this relationship, we follow a literature relating capital gains to land use regulation (Saks, 2008) over a longer time horizon. We take our measures from the seminal work of Saiz (2010), which reports regulatory constraints on housing construction based on the Wharton Residential Urban Land Regulation Index (Gyourko et al., 2008) for all of the cities in the HHP dataset. This regulatory index measures constraints on development from a survey done near the very end of our sample period. However, social scientists have pointed out that high and growing home prices incentivize the adoption of restrictive zoning (Baum-Snow & Han, 2024; Fischel, 2001; Trounstine, 2018). It is thus informative to see how earlier home price appreciation varied with the future regulatory environment.

In Figure 6 we present simple scatterplots of the average annual real capital gain at the city level plotted against the Saiz measure of zoning stringency for the 1940–1969 and 1970–2006 periods. There is already a positive relationship between real average annual capital gains and zoning restrictiveness measured in 2008 from 1940 to 1969, when there was still significant construction in cities across the United States. It is possible both that

Figure 6: Capital Gains and Zoning Stringency

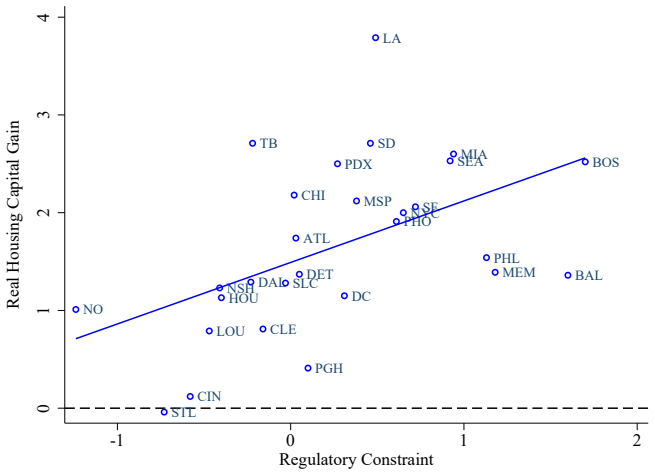

(a) 1940-1969

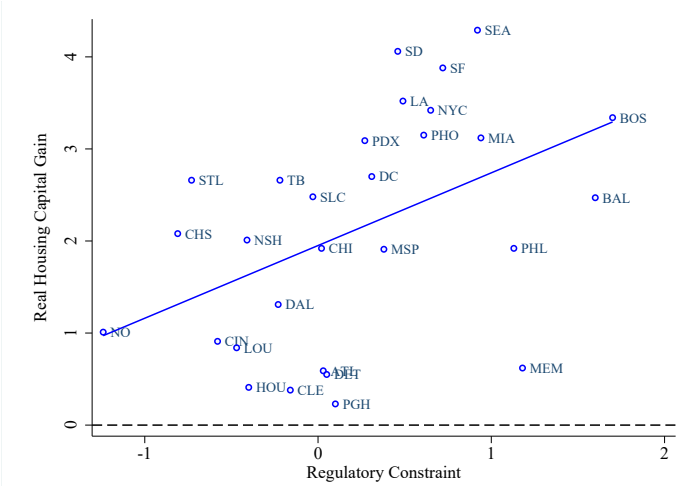

(b) 1970-2006

Note: The figures plot the average annual real capital gain from the HHP baseline HPI against the Saiz (2010) measure of regulations on housing construction based on the Wharton Residential Urban Land Regulation Index (Gyourko et al., 2008).

the zoning rules already in place were causing higher price growth and that home price appreciation was incentivizing cities to adopt even more restrictive measures, particularly by the 1970s (Fischel, 2015; Molloy et al., 2020). Indeed, we see an even more positive correlation[22] between capital gains and the same 2008 zoning measure across the last phase

[22]The regression coefficient increases from .63 from 1940–1960 to .79 from 1970–2006.

covered in the paper from 1970 to 2006, suggesting a closer relationship between zoning and home price appreciation towards the end of the 20th century. The fundamental endogeneity here is difficult to unpack without better information on historical zoning regimes. However, that land use regulation has become a more important correlate of across-city differences in housing price growth over time is consistent with recent scholarship (Gyourko & Krimmel, 2021).

6 Changes in Mortgage Conditions over Time

The second major fact about U.S. housing prices evident in Figure 4 is the evolving relationship between rental and sale prices. During the cycles of the prewar era and the period of stability thereafter, the two series largely moved together, at least within a few years. After 1970, the sale price and rental price series permanently diverged, with sale prices increasing by far more than rents. We use the HHP data to shed new light on the relationship between sale and rental prices over the long run. In equilibrium the annual cost of owning a home should be equivalent to renting it. This "no arbitrage" condition (Poterba, 1984) can be expressed in terms of the user cost of owning housing:

$$R_t = P_t u_t \tag{4}$$

where

$$u_t = i_t + \tau_t + \delta_t - g_{t+1} + \gamma_t \tag{5}$$

and i_t is the risk-free interest rate, τ_t is the tax benefit or cost associated with homeownership, δ_t is maintenance costs, g_{t+1} is expected appreciation (or depreciation) over the year, and γ_t the risk associated with homeownership relative to renting. That is, the user cost of housing is equal to the ratio of rents to prices in equilibrium. A challenge with computing that ratio is that we do not observe rental prices for owned housing, so rents are usually

imputed from sale prices using a measure of user costs (Himmelberg et al., 2005). We take a different approach here and compute the rent-to-price ratio using our national HPI and RI series. We base our rent-to-price ratio index to a 2006 benchmark value (3.16%) based on the Case-Shiller-Weiss index taken from Davis et al. (2008).[23]

We visualize the rent-to-price ratio for our baseline specification HPI and RI for the full sample beginning in 1890 in Figure 7a. This is the longest annual rent-to-price ratio series that has been assembled for the United States to date.[24] Two striking aspects of this figure are the substantial cycle in the ratio between 1920 and 1940 and the long, uneven decline in the ratio during the postwar period. Below we discuss this series and the associated changes in the user cost of housing in three phases, echoing our approach in Section 4.

Before examining the ratio over time, we pause to consider the extent to which the rents and sales indices reflect the "same" property, a theoretical assumption underlying the user cost equation 5. By basing our ratio to the contemporary benchmark value in our terminal year of 2006, we can remove the concern that systematic differences in sale and rental datasets (for example in size or location of units) affects the level of rental returns over time. One could also be concerned that there are differential trends across different segments of the housing markets, for instance because rental homes have concentrated over time into smaller and more centrally located neighborhoods. However, our hedonic approach controls for such differential trends by holding the mix of observed housing attributes constant to estimate year-to-year changes in price levels. Furthermore, the use of short rolling windows limits the scope for unobserved housing attributes that could change over time.

To provide evidence that our hedonic approach is working well, in Appendix Figure F1 we compute rent-to-price ratios for the full sample and for the top five areas for each city by rental share with a relaxed specification.[25] Intuitively this sample restricts attention

[23] Note that the ratio here, of contemporaneous sale and rental prices, is subtly different from the concept of the rental return to housing, used in the next section, where rental prices are expressed as a fraction of sale prices in the previous period.

[24] We compare our series to the widely-used rent-to-price ratio of Davis et al. (2008) for the period from 1960 onward in Appendix D.

[25] In order to estimate rental and sales housing price indices for this relatively small subset of our data, we

Figure 7: Rent-to-Price Ratio, Homeownership Rate, and Mortgage Terms

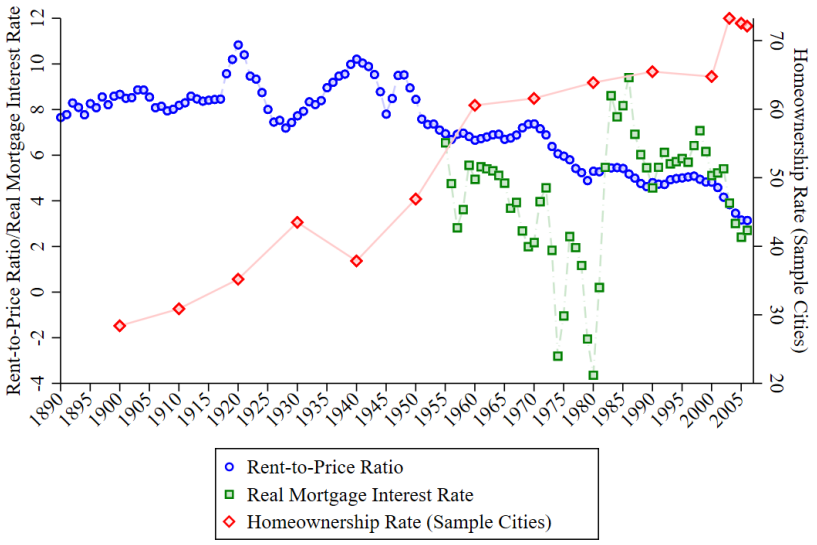

(a) Changes in Rent-to-Price Ratio, Homeownership, and Interest Rates

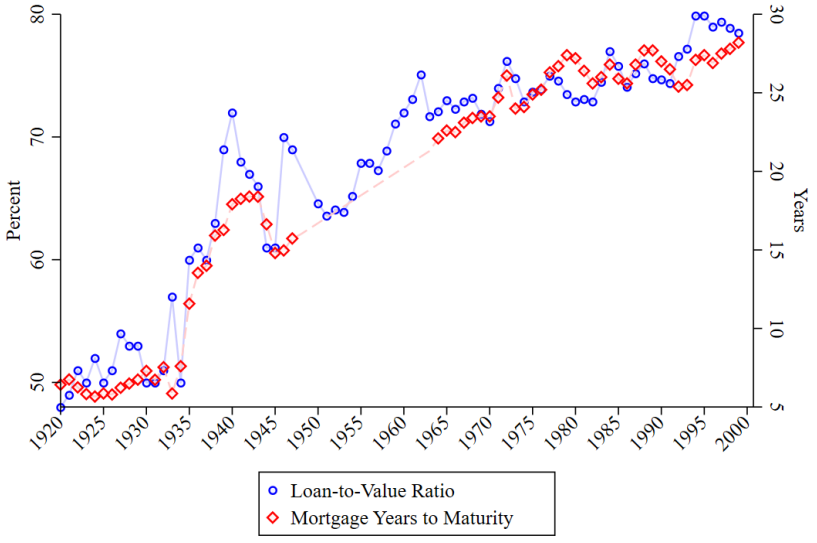

(b) Mortgage Terms

Note: Panel (a) shows (1) the rent-to-price ratio from the HHP dataset benchmarked to the Davis et al. (2008) value in 2006, (2) the population-weighted homeownership share from the census for the 30 cities in the HHP dataset, and (3) the real mortgage interest rate on a conventional first mortgage taken from Drehmann et al. (2024). Panel (b) shows (1) the loan-to-value (LTV) on conventional mortgage loans from volumes of the Historical Statistics of the United States used in Fetter (2013) and (2) average years to maturity series across all loan types (VA, FHA, and conventional) from Drehmann et al. (2024).

to areas where similar homes for sale and rent are available and the marginal consumer is more likely to be indifferent between the two segments. The high-rental-stock rent-to-price ratio is slightly higher relative to the full sample but follows a very similar trajectory over time, lending credence to our approach.

Phase 1: Abundant Construction and Limited Financing (1890–1935) The decades from 1890 to 1920 saw a stable rent-to-price ratio of 8–9%, as substantial new housing was built in U.S. urban areas to accommodate both immigrants from Europe and internal migrants from rural areas. During this period of minimal land use regulation and abundant land, sale and rental prices were closely tied to construction costs and followed a similar path. We show the population-weighted homeownership share from the census for our sample cities in Figure 7a as well. Homeownership rates increased moderately throughout this period from just under 30% in 1900 to over 40% in 1930.

Before World War I, families financed home purchases using building and loan associations, life insurance companies, commercial banks, or directly from the seller of the house (Snowden, 2010). Loan terms were generally short and loans were much smaller than the property value. In Figure 7b we show the conventional mortgage loan-to-value (LTV) ratio and years to maturity (YTM) for the typical home loan in the United States.[26] LTV ratios hovered around 50% and loan duration was about five years for the entire 1920–1935 period. Loan terms were even shorter before 1920, often just two or three years (Fetter et al., 2018). This first phase closes with the dramatic rise and fall of the rent-to-price ratio associated with changes in sales prices and rents during the Roaring Twenties and Great Depression. The turmoil in the housing market in the early 1930s prompted the federal government to intervene, permanently changing housing finance in the United States.

relax the specification by removing rolling windows and controls for area and set the start year of 1920, by which point we have enough observations for most cities to compute indices for this subsample.

[26]The years to maturity series was taken from Drehmann et al. (2024) and reflects the average across all loan types (VA, FHA, and conventional). The LTV series is from volumes of the Historical Statistics of the United States used in Fetter (2013) and reflects the LTV on a conventional mortgage. LTV ratios were typically even higher for VA and FHA mortgages. We thank Sarah Quincy and Daniel Fetter for their assistance with the data.

Phase 2: The Emergence of the Modern Mortgage (1935-1970) We demarcate the first and second phases in this section at 1935 to reflect the establishment of the Homeowners' Loan Corporation (HOLC) in 1933 and the Federal Housing Administration (FHA) in 1934. The creation of these agencies was a watershed moment for mortgage structure. The HOLC bought distressed loans and refinanced them into 15-year, fully-amortized mortgages with a subsidized interest rate of 5%. Between 1933 and 1936, the HOLC refinanced loans for about 10% of the nonfarm housing stock.[27]

The FHA provided mortgage insurance backed by the federal government, which reshaped mortgage terms for borrowers who did not already have HOLC loans. The establishment of the Federal National Mortgage Association (Fannie Mae) in 1938 further increased the supply of longer-term, higher-LTV mortgages. Fully amortized mortgages with a duration of 25 years (30 by the mid-1950s) for 80% or even 90% of the purchase price became widely available. The GI Bill in 1944 provided Veterans Administration (VA) loans at generous terms to returning veterans. As a result of these programs, the typical YTM of a mortgage began to climb by 1934, increasing from 5 to 18 years by 1940.[28]

The widespread availability of fully-amortized, long-duration mortgages greatly reduced the risk to homeowners associated with rolling over their home loans. This reduction in γ_t should have reduced the user cost of housing and the rent-to-price ratio, all else equal. This is precisely what we see over the 1935–1960 period, with the ratio falling from 9.4% in 1935 to 7.4% in 1970. Meanwhile, the supply of housing in most of these metro areas expanded dramatically after World War II (see Figure 5). The rapidly expanding housing stock coupled with favorable changes in mortgage structure led to a sharp increase in homeownership. The rate in our sample cities increased from 40% to over 60% by 1960. Some of this increase was due to rent control during World War II, which shifted some of the housing stock from renter to owner occupied (Fetter, 2016). The transforma-

[27] The book *Well Worth Saving* provides an insightful history of New Deal housing programs (Fishback et al., 2019).

[28] Some of these programs were generally available to both black and white borrowers (HOLC loans) while others largely excluded black borrowers and neighborhoods, particularly before the Fair Housing Act was passed in 1968 (FHA-insured loans). For a discussion of racial discrimination by New Deal housing programs, see Fishback et al. (2023) and Fishback et al. (2024).

tion of mortgage structure had largely concluded by 1970, with only modest increases in the LTV ratio and YTM thereafter.

Although we do not attempt to directly measure differences in the tax benefit of housing ownership across space here, we note an important change in τ_t in this period. The expansion of the income tax base during World War II greatly increased the share of the population filing taxes with itemized deductions including home mortgage interest. Loan interest had been deductible since the inception of the federal income tax in 1913, but it was not until the postwar period that this provision of the tax code affected the average American's housing costs. Tax benefits were thus another factor that pushed user costs of owning down relative to renting and contributed to the postwar decline in the rent-to-price ratio.

Phase 3: Changing Credit Conditions (1970-2006) The final phase of 20th century housing finance was characterized by changing credit conditions for borrowers within a relatively stable framework of a 30-year, fully-amortized mortgage. Before 1950, nominal mortgage interest rates fluctuated little with economic conditions, with most borrowers holding loans with rates between 7% in the 1920s and 5% by the 1940s (Drehmann et al., 2024). However, mortgage interest rates became far more volatile in the latter part of this phase, particularly after 1970. We plot the real mortgage interest rate for the typical conventional loan in the United States in Figure 7a.

Real interest rates plummeted over the 1970s, at times dropping below zero, as inflation soared. Long mortgage terms meant that housing equity could enjoy the uplift, even as housing debt was eroded by inflation. This fall in the user cost component i_t was reflected in another drop in the rent-to-price ratio over the same decade. At the same time, we see a far more muted shift into homeownership, perhaps reflecting the slowdown in construction across some U.S. cities. Interest rates and the rent-to-price ratio continued to move together, increasing as real mortgage interest rates spiked in the 1980s and then fell dramatically starting in the late 1990s.

Our sample ends in the mid-2000s, at the peak of a substantial boom in sale prices. Duca et al. (2011) document important changes in mortgage market conditions between the 1970s and the 2000s. In particular, they show that sale prices (adjusted for inflation or relative to rental prices) responded to the steady relaxation of conditions in the mortgage market between the late 1980s and the mid-2000s, with the typical down-payment by first-time-buyers falling from about 17% in 1988 to about 5% in 2006. Related work by Greenwald & Guren (2021) estimates that changes in credit standards explain between one third and one half of the fall in the rent-to-price ratio during the pre-Great Financial Crisis boom.

7 The Return to Owning Housing

In this section, we revisit the debate on the long-run return to housing, especially when compared to other assets. A growing body of research has examined the long-run returns to residential housing, both in the United States and internationally. Jordà et al. (2019) find that, over the past century, real annual returns on housing have averaged around 7%, comparable to equities, though driven more by rental income than capital appreciation. Importantly, they find that across countries housing returns have exhibited lower volatility than stocks, leading to superior risk-adjusted performance in many contexts.[29] Economists have particularly debated the average annual capital gain associated with homeownership in the United States. One prominent reference point is the real annual return implied by the national Shiller index of 0.7% AGR from 1890–2006.

We contribute to this debate by using our series on sale and rental prices to calculate housing returns. For each city c, we compute the total return to housing, R, as the sum of

[29] Making these comparisons requires many methodological choices and assumptions, including whether and how to account for maintenance costs, taxes, leverage, and the treatment of owner-occupied versus rental properties. Jordà et al. (2019) do not attempt to account for taxation or leverage but do estimate costs of maintenance, insurance and management. For this reason, they estimate net rental returns and compare to dividends (rather than earnings) of stocks.

the capital gains, H, and the rental return (or income), Y:

$$\underbrace{R_{c,t}}_{\text{Total Return}} = \underbrace{H_{c,t}}_{\text{Capital Gain}} + \underbrace{Y_{c,t}}_{\text{Rental Return}} \tag{6}$$

$$= \frac{\text{HPI}_{c,t} - \text{HPI}_{c,t-1}}{\text{HPI}_{c,t-1}} + \frac{\text{RI}_{c,t}}{\text{HPI}_{c,t-1}} \tag{7}$$

To study long-run asset returns, we align with Shiller (2015) in two ways. First, to measure annual income, we focus on asset earnings rather than net income. In practice this means using earnings rather than dividends for equities and the gross return for housing. Second, we deflate capital gains by the CPI to better understand true asset price growth. We also align with Jordà et al. (2019) in our specific measure of returns: rather than calculating the annual rent as a fraction of contemporaneous capital value (in time t, the "rent-to-price ratio" from Section 6), we calculate rent in t as a fraction of the price in $t-1$ (the "rental return"), the last period when an agent could take ownership of the asset in time to earn the income in period t. Thus, where $\pi_t = \frac{\text{CPI}_t - \text{CPI}_{t-1}}{\text{CPI}_{t-1}}$, and lower-case refers to inflation-adjusted series, we have the following formula for real capital gains:

$$h_{c,t} = \frac{1 + H_{c,t}}{1 + \pi_t} - 1$$

The total return to housing r for city c and year t can therefore be calculated straightforwardly from our baseline indices of sale prices $\text{HPI}_{c,t}$ and rental prices $\text{RI}_{c,t}$, and the consumer price index.

As before, we can look at city-specific returns or, using population weights to aggregate, returns at the national level. We define the average annual return for each component of total returns as the arithmetic mean for the purpose of considering the average year in any given period. For capital gains, where $t \in T = \{1891, 2006\}$, we have:

$$\bar{h}_c = \frac{1}{|T|} \sum_{t=1891}^{2006} h_{c,t} \tag{8}$$

Figure 8: Real Gross Return to Housing

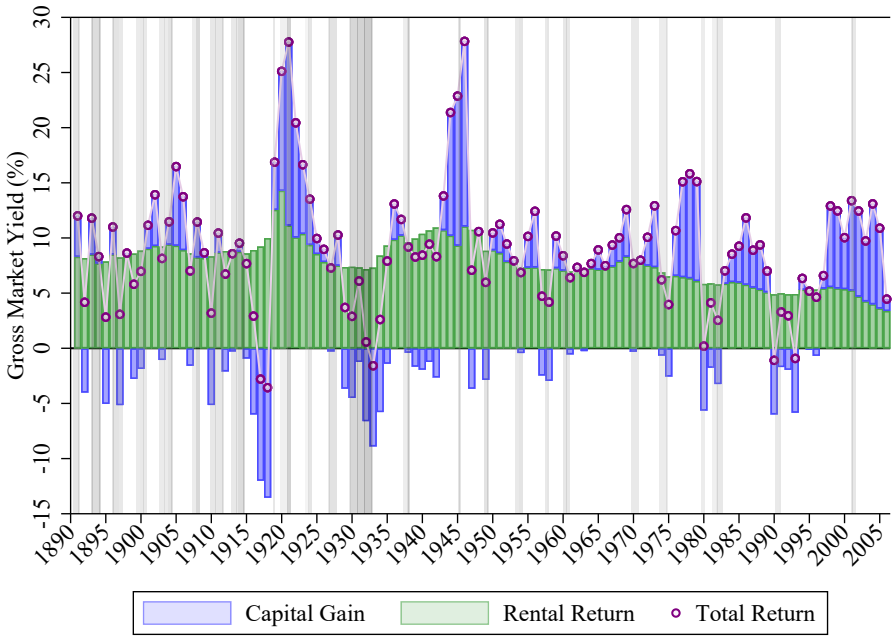

Note: The figure plots the total gross return to housing in real terms for the full HHP sample. The total return is broken down into the inflation-adjusted capital gain and rental return components from Equation 6.

Analogs for the arithmetic mean of the real rental return are $\bar{y}_c$ and for the real overall housing return $\bar{r}_c$.

Stylized Facts of Housing Returns Figure 8 presents our estimate of real housing returns, by year and component, for the U.S. from 1890 to 2006. We also provide city-level returns series in Appendix F. Below we describe five principal stylized facts about total housing returns before comparing them with long-run equity returns.

Total real returns to owning housing were rarely negative: Firstly, we find that total real returns to U.S. housing were rarely negative during our sample period. In fact returns were negative in only five out of 116 years (1917, 1918, 1933, 1990, and 1993). This stability is largely due to the persistence of rental income, mirroring a finding of Jordà et al. (2019) that with steady rental income and less volatile capital gains, overall housing returns are more stable than equity returns.

The average annual return to housing was 9%: Second, the average annual total return to housing over the full period was 9.0%, with gross rental return accounting for the bulk at 7.7% and real capital gains contributing 1.3%. This composition is broadly consistent with Jordà et al. (2019), who report a similar structure of housing returns as dominated by income rather than capital gains. Our estimate of capital gains, however, does represent a significant upward revision of about one third to their long-run figure (1.0%). Much of this revision is driven by far greater capital gains in the post-WWII decades, consistent with Davis et al. (2008), with prices growing by a full percentage point more than Shiller's index from 1940 to 2000 (1.7% vs. 0.7% AGR).

Capital gains were more volatile than rental returns: Third, capital gains were significantly more volatile and cyclical than rental returns. With an average real capital gain of 1.3% and a standard deviation of 5.0%, capital gains had a coefficient of variation (CV) of 3.8. For rental earnings, however, the mean was 7.7% and the standard deviation just 1.9%, giving a CV of 0.2. Such consistent earnings are central to the "puzzle" of housing having returns that were, relative to other assets, both stable and high. However, there is a marked secular decline in gross rental returns over time, from an average of over 8% in the early decades to a low of just 3.4% in 2006. In other words, the more stable portion of housing returns has declined over time.

Prewar capital gains were close to zero: Fourth, with the notable exception of the immediate aftermath of World War I, capital gains were largely unimportant in overall housing returns until the 1940s. While there was an average capital gain of 1.7% from 1940 to 2006, the real capital gain for the average year during 1890–1940 was effectively zero (an arithmetic mean of +0.04%) and alternated between periods of decline (such as the 1890s, 1910s and 1930s) and periods of growth (the 1900s and 1920s). There were differences in postwar capital gains by decade, with strongest growth in the 1940s and 1970s as well as in the period 2000-2006. Nonetheless, for each decade from the 1940s, the average year saw capital gains, not losses.

Rental returns declined and capital gains rose: Lastly, with rental returns declining over

time and capital gains rising, capital gains became more important as a source of returns on housing as an asset, particularly after 1970. Capital gains did not contribute at all to housing returns for the typical year in the 1890–1935 period (rental returns of 8.8%; total returns of 8.9%). During the second phase (1935–1970), a capital gain of 1.6% in the average year made up approximately one-sixth of total returns (compared to an 8.6% rental return). In the final period, 1970–2006, capital gains made up almost one third of housing returns in the average year (2.6% of 8.1%).

Housing and Equities Compared We now compare our new estimates of long-run returns to housing with those of equities, relying in particular on the work of Shiller (1992) in assembling series on prices and earnings for the S&P 500, but also factoring in the stylized facts of Jordà et al. (2019) (JKKST) on relative returns across countries. As noted above, we use earnings (rather than dividends) for equities and gross (rather than net) rental returns for housing to provide a consistent comparison of total returns across assets. Figure 9 compares housing and equities in terms of total real returns, earnings, and capital gains.

We confirm the finding in Jordà et al. (2019) that equities offer greater total returns than housing, but at the cost of greater volatility. Adjusting for inflation, stock prices in 2006 were over ten times their 1890 level, while housing prices had increased by a factor of just under four. Using the HHP series on housing returns for 1890–2006, we find that the average annual total real return on housing was 9% compared to 11.5% for equities. With the standard deviation of equity returns over three times that for housing (17.6% compared to 5.4%), the coefficient of variation for returns to equity was 2.5 times that of housing returns (1.5 vs. 0.6).

Our new figures do not alter the ranking across asset classes seen in JKKST. While we revise upwards the real capital gain associated with housing for the U.S. over the long run, from 0.7% in Shiller and 1.0% in JKKST to 1.3%, this is not enough to offset the differential in capital gains between asset classes. For equities, the real capital gain in the U.S. in the average year was 3.6% annually for our period. Figure 9a suggests a change in the corre-

lation across housing and equities, consistent with Jordà et al. (2019). We present ten-year annual average growth rates in this panel to move beyond the volatility in the annual series and examine more systematic patterns over time. In the pre-World War II period, housing and equity capital gains moved in tandem while afterwards, the correlation weakened. However, our findings suggest that the U.S. experienced unusual relative performance of housing and equities before World War II. Unlike the typical country in the JKKST sample, we do not see housing outperforming equities before the 1940s. Indeed, there are only three brief periods across our period — the early 1920s, the early 1940s, and the late 1970s to early 1980s — when housing returns exceeded those of equities.

Figure 9b shows the pattern of annual earnings for equities and housing. There are clear similarities, including the overall level (7.9% for equities and 7.7% for housing). However, while rental returns fell across the three major periods (as noted above), average stock earnings was largely similar over time: 7.8% before 1935, 8.3% between 1935 and 1970, and 7.5% after 1970. Spikes in rental income were less extreme than those seen in earnings on equities. Altogether, with earnings for stocks stable and those for housing falling, housing lost its initial earnings advantage over equities. Rental returns for housing were higher than earnings for stocks for 22 of the first 24 years (1890–1914), almost as many as in the final six decades of the sample (25 years out of 60, 1947–2006).

The greater volatility is driven principally by capital gains: real capital gains associated with equities were almost three times as large as those on housing in the average year (3.6% compared to 1.3%). But, as Figure 9a shows, those gains were volatile and significantly more so for housing. This finding is not wholly a product of aggregation: Figure 9c gives total returns for three sample cities and, while the returns are somewhat more volatile than for the national average, they remain in general significantly less noisy than those for equities (Figure 9d).

Figure 9: Earnings, Prices, and Returns by Asset Class

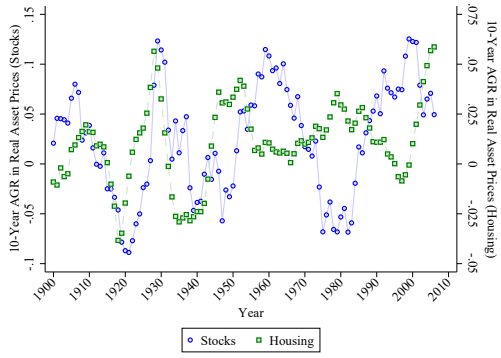

(a) 10-Year AGR in Real Asset Prices

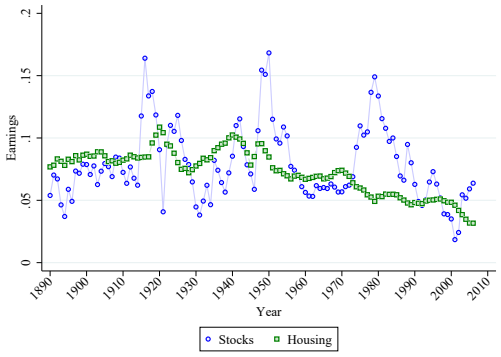

(b) Earnings by Asset Class

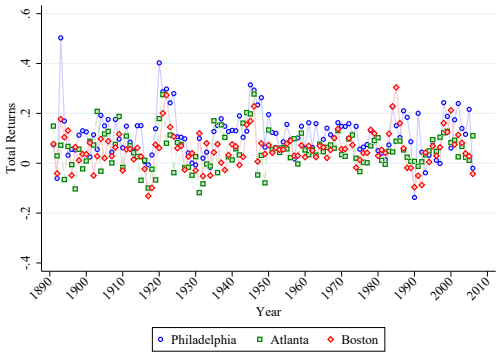

(c) Total Returns by City

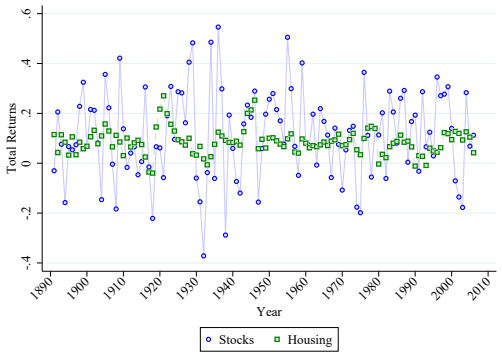

(d) Total Returns by Asset Class

Note: Panel (a) shows the real 10-year average growth rate for the prices of stock vs. housing, panel (b) the annual earnings on stocks vs. rental return on housing, panel (c) the total real return on housing in three sample cities, and panel (d) the average annual total return to owning stocks vs. housing in real terms. The source of the housing data is the HHP series for rent and sale prices. The stock price and earnings data series were used in Shiller (2015) and downloaded from Robert Shiller's website https://shillerdata.com/. All series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

8 Conclusion

In this paper we examined housing price dynamics in American cities over the long run. We constructed, for the first time, annual market price indices for both rented and for-sale housing for 30 cities and the nation covering the period 1890–2006. Our work extends our understanding of housing market conditions within and across cities and over the long run and documents several new stylized facts about U.S. housing markets. In doing so, we updated what is known about key debates in the housing literature, including how housing performs relative to other assets over the long run, why returns differ across space and time, and how changes in housing construction and finance have influenced trends in prices and returns. The data series produced by this paper include rental prices, sale prices, capital gains, rental returns, total housing returns, and rent-to-price ratios for cities and the U.S. as a whole. Our hope is that these new series will be of great use to future researchers and spur new inquiry into the evolution of housing markets in the United States.

References

Ambrose, B. W., Coulson, N. E., & Yoshida, J. (2015). The repeat rent index. *Review of Economics and Statistics*, *97*(5), 939–950.

Baum-Snow, N., & Han, L. (2024). The microgeography of housing supply. *Journal of Political Economy*, *132*(6), 1897–1946.

Bogin, A., Doerner, W., & Larson, W. (2019). Local house price dynamics: New indices and stylized facts. *Real Estate Economics*, *47*(2), 365–398.

Calhoun, C. A. (1996). Ofheo house price indexes: Hpi technical description. *Office of Federal Housing Enterprise Oversight*, *20552*, 1–15.

Calomiris, C. W., & Jaremski, M. (2023). Florida (un) chained. *Journal of Financial Intermediation*, *55*, 101043.

Case, K. E., & Shiller, R. J. (1987). Prices of single family homes since 1970: New indexes for four cities.

Contat, J., & Larson, W. D. (2022). *A flexible method of house price index construction using repeat-sales aggregates* (tech. rep.). Working Paper.

Crone, T. M., Nakamura, L. I., & Voith, R. (2010). Rents have been rising, not falling, in the postwar period. *Review of Economics and Statistics*, *92*(3), 628–642.

Davis, M. A., & Heathcote, J. (2007). The price and quantity of residential land in the United States. *Journal of Monetary Economics*, *54*(8), 2595–2620.

Davis, M. A., Lehnert, A., & Martin, R. F. (2008). The rent-price ratio for the aggregate stock of owner-occupied housing. *Review of Income and Wealth*, *54*(2), 279–284.

Davis, M. A., Ortalo-Magné, F., & Rupert, P. (2007). What's really happening in housing markets? *Economic Commentary*, (7/1/2007).

Davis, M. A., & Palumbo, M. G. (2008). The price of residential land in large us cities. *Journal of Urban Economics*, *63*(1), 352–384.

Drehmann, M., Juselius, M., & Quincy, S. (2024). Aggregate debt servicing and the limit on private credit.

Duca, J. V., Muellbauer, J., & Murphy, A. (2011). House prices and credit constraints: Making sense of the us experience. *The economic journal*, *121*(552), 533–551.

Eichengreen, B. J. (2015). *Hall of mirrors: The great depression, the great recession, and the uses-and misuses-of history*. Oxford University Press.

Eichholtz, P., Korevaar, M., & Lindenthal, T. (2019). 500 years of housing rents, quality and affordability. *Quality and Affordability (July 11, 2019)*.

Eichholtz, P., Korevaar, M., Lindenthal, T., & Tallec, R. (2021). The total return and risk to residential real estate. *Review of Financial Studies*, *34*(8), 3608–3646.

Eichholtz, P., Straetmans, S., & Theebe, M. (2012). The Amsterdam rent index: The housing market and the economy, 1550–1850. *Journal of Housing Economics*, *21*(4), 269–282.

Eichholtz, P. M. (1997). A long run house price index: The herengracht index, 1628–1973. *Real estate economics*, *25*(2), 175–192.

Fetter, D., Rose, J., & Snowden, K. (2018). Housing in american economic history. *The Oxford Handbook of American Economic History Volume 2*.

Fetter, D. K. (2013). How do mortgage subsidies affect home ownership? Evidence from the mid-century GI bills. *American Economic Journal: Economic Policy*, *5*(2), 111–147.

Fetter, D. K. (2016). The Home Front: Rent control and the rapid wartime increase in home ownership. *Journal of Economic History*, *76*(4), 1001–1043.
Fischel, W. A. (2001). Homevoters, municipal corporate governance, and the benefit view of the property tax. *National Tax Journal*, *54*(1), 157–173.
Fischel, W. A. (2015). *Zoning rules! The economics of land use regulation*. Lincoln Institute of Land Policy.
Fishback, P., & Kollmann, T. (2014). New multicity estimates of the changes in home values. *Housing and Mortgage Markets in Historical Perspective*, 203.
Fishback, P., Rose, J., Snowden, K. A., & Storrs, T. (2024). New evidence on redlining by federal housing programs in the 1930s. *Journal of Urban Economics*, *141*, 103462.
Fishback, P. V., LaVoice, J., Shertzer, A., & Walsh, R. P. (2023). The holc maps: How race and poverty influenced real estate professionals' evaluation of lending risk in the 1930s. *The Journal of Economic History*, *83*(4), 1019–1056.
Fishback, P. V., Rose, J., & Snowden, K. (2019). *Well worth saving: How the new deal safeguarded home ownership*. University of Chicago Press.
Glaeser, E., & Gyourko, J. (2005). Urban decline and durable housing. *Journal of political economy*, *113*(2), 345–375.
Glaeser, E., & Gyourko, J. (2018). The economic implications of housing supply. *Journal of economic perspectives*, *32*(1), 3–30.
Glaeser, E., Gyourko, J., & Saks, R. E. (2005). Why have housing prices gone up? *American Economic Review*, *95*(2), 329–333.
Gordon, R., & van Goethem, T. (2007). Downward bias in the most important component of the CPI: The case of rental shelter, 1914–2003. *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches, edited by Ernst R. Berndt and Charles R. Hulten, University of Chicago Press: Chicago*.
Grebler, L., Blank, D. M., Winnick, L., et al. (1956). Capital formation in residential real estate: Trends and prospects. *NBER Books*.
Greenlees, J. S. (1982). An empirical evaluation of the CPI home purchase index, 1973–1978. *Real Estate Economics*, *10*(1), 1–24.
Greenwald, D. L., & Guren, A. (2021). *Do credit conditions move house prices?* (Working Paper No. 29391). National Bureau of Economic Research.
Gyourko, J., & Krimmel, J. (2021). The impact of local residential land use restrictions on land values across and within single family housing markets. *Journal of Urban Economics*, *126*, 103374.
Gyourko, J., Mayer, C., & Sinai, T. (2013). Superstar cities. *American Economic Journal: Economic Policy*, *5*(4), 167–199.
Gyourko, J., Saiz, A., & Summers, A. (2008). A new measure of the local regulatory environment for housing markets: The Wharton residential land use regulatory index. *Urban Studies*, *45*(3), 693–729.
Han, L., & Strange, W. C. (2014). Bidding wars for houses. *Real Estate Economics*, *42*(1), 1–32.
Han, L., & Strange, W. C. (2016). What is the role of the asking price for a house? *Journal of Urban Economics*, *93*, 115–130.

Haurin, D. R., Haurin, J. L., Nadauld, T., & Sanders, A. (2010). List prices, sale prices and marketing time: An application to US housing markets. *Real Estate Economics*, *38*(4), 659–685.

Himmelberg, C., Mayer, C., & Sinai, T. (2005). Assessing high house prices: Bubbles, fundamentals and misperceptions. *Journal of Economic Perspectives*, *19*(4), 67–92.

Hoyt, H. (2000). *One hundred years of land values in chicago: The relationship of the growth of chicago to the rise of its land values, 1830-1933*. Beard Books.

Humes, H., & Schiro, B. (1948). The rent index: Part i, concept and measurement. *Monthly Labor Review*, *67*(6), 631–637.

Humes, H., & Schiro, B. (1949). The rent index: Part ii, methodology of measurement. *Monthly Labor Review*, *68*(1), 60–68.

Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., & Taylor, A. M. (2019). The rate of return on everything, 1870–2015. *Quarterly Journal of Economics*, *134*(3), 1225–1298.

Jordà, Ò., Schularick, M., & Taylor, A. M. (2017). Macrofinancial history and the new business cycle facts. *NBER Macroeconomics Annual*, *31*(1), 213–263.

Knoll, K., Schularick, M., & Steger, T. (2017). No price like home: Global house prices, 1870–2012. *American Economic Review*, *107*(2), 331–353.

Lester, B., Visschers, L., & Wolthoff, R. (2017). Competing with asking prices. *Theoretical Economics*, *12*(2), 731–770.

Leventis, A. (2006). Removing appraisal bias from a repeat-transactions house price index: A basic approach.

Leventis, A. (2008). Revisiting the differences between the ofheo and s&p/case-shiller house price indexes: New explanations. *Research Paper. Office of Federal Housing Enterprise Oversight*.

Lyons, R. C. (2019). Can list prices accurately capture housing price trends? Insights from extreme markets conditions. *Finance Research Letters*, *30*, 228–232.

Molloy, R., et al. (2020). The effect of housing supply regulation on housing affordability: A review. *Regional Science and Urban Economics*, *80*(100), 1–5.

Nagaraja, C., Brown, L., & Wachter, S. (2014). Repeat sales house price index methodology. *Journal of Real Estate Literature*, *22*(1), 23–46.

Nicholas, T., & Scherbina, A. (2013). Real estate prices during the Roaring Twenties and the Great Depression. *Real Estate Economics*, *41*(2), 278–309.

Nowak, A. D., & Smith, P. S. (2020). Quality-adjusted house price indexes. *American Economic Review: Insights*, *2*(3), 339–356.

Ozimek, A. (2013). *Sticky rents and the CPI for owner-occupied housing*. Temple University.

Poterba, J. M. (1984). Tax subsidies to owner-occupied housing: An asset-market approach. *Quarterly Journal of Economics*, *99*(4), 729–752.

Quigley, J. M., & Raphael, S. (2004). Is housing unaffordable? Why isn't it more affordable? *Journal of Economic Perspectives*, *18*(1), 191–214.

Rees, A., & Jacobs, D. P. (1961). The cost-of-living index. In *Real wages in manufacturing, 1890-1914* (pp. 74–119). Princeton University Press.

Rose, J. D. (2022). Reassessing the magnitude of housing price declines and the use of leverage in the depressions of the 1890s and 1930s. *Real Estate Economics*, *50*(4), 907–930.

Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics*, *125*(3), 1253–1296.
Saks, R. E. (2008). Job creation and housing construction: Constraints on metropolitan area employment growth. *Journal of Urban Economics*, *64*(1), 178–195.
Shen, L., & Ross, S. (2021). Information value of property description: A machine learning approach. *Journal of Urban Economics*, *121*, 103299.
Shiller, R. J. (1992). *Market volatility*. MIT Press.
Shiller, R. J. (2015). Irrational exuberance. In *Irrational exuberance*. Princeton University Press.
Silver, M. (2016). *How to better measure hedonic residential property price indexes*. International Monetary Fund.
Snowden, K. A. (2010). The anatomy of a residential mortgage crisis: A look back to the 1930s. In *The panic of 2008*. Edward Elgar Publishing.
Trounstine, J. (2018). *Segregation by design: Local politics and inequality in American cities*. Cambridge University Press.
Vandell, K. D. (1995). FHA restructuring proposals: Alternatives and implications.
Yavas, A., & Yang, S. (1995). The strategic role of listing price in marketing real estate: Theory and evidence. *Real Estate Economics*, *23*(3), 347–368.

The Price of Housing in the United States, 1890–2006
Online Appendices

A Acknowledgments

The project was made possible by the combined efforts of many incredible people whose contributions we wish to acknowledge here.

We are deeply grateful for the work of our pre-doctoral, doctoral, and post-doctoral scholars: Michael Coury (University of Pittsburgh, now SUNY Buffalo); Grace Kneidel (Philadelphia Federal Reserve); Maximilian Guenniwig-Moenert, Lea Hauser, Sean Moran, Johann Ohler, Conor O'Malley, Carla Revilla Diez, Sameer Shaikh and Sam Watling (all Trinity College Dublin).

We thank the following undergraduate research assistants at the University of Pittsburgh: William Lang, Ella Donley, Derek Glausser, Peter Martin, Sammi Zheng, Emily Xu, Victoria Beck, Brady Groves, Caroline Layding, Rebecca Markowitz, and Pawarisa Bunsophak.

We thank the following undergraduate research assistants at the University of California Merced: Tracy Nguyen, Clare Recinos, Souen Reyes-Sor, Magali Torres, and Lipi Vallabhaneni.

We thank Jason Barr (Rutgers) for his assistance on data relating to New York City.

We thank the following research assistants, mostly undergraduates, who worked on the project at Trinity College Dublin (TCD): Noah Andriatch, Melissa Barrett, Frida Bowe, Liam Brady, Mikaela Blount, Gavin Cahill, Aislinn Carty, Alex Crean, Mateo Crowley, Sam Daly, Ruairi Devereux, William Dickey, Chris Doughty, Christian DiAntonio, Eoin Dignam, Sophie Donnelly, Katie Duffy, Stephen Garavan, Danny Gray, Luca Graziadei, Ryan Grunwell, Luke Harjo, Ryan Hayes, Josh Healy, Mark Hennessy, Luke Hosford, David Kavanagh, Thomas Khadoo, Brittney Kidwell, Philip King, Matt Kolovsky, Anselm Kuesters, Priyanka Lodhia, Jeffrey Lui, Michael Mahony, Lucille McKnight, Finn McLaughlin, Amelia Melanson, Jackson Miller, Patrick McDonagh, Christian McKeever, Conor Murray, Sean Murray, Henry O'Brien, Gillian O'Connell, Eoin O'Donnell, Seamus O'Donovan, Nicole O'Sullivan, Gian Palleschi, Stefano Paparoni, Yera Patel, Sarah Pubins, Peter Quigley, Ciaran Ramdeo, Giorgio Ravalli, Ben Riseberg, Fionn Roche, Bryan Rohan, Corey Russell, Michael Scheriger, Emma Schloegel, Sabrina Schoenfeld, Iseult Sheehy, Octavia Smirren, Robert Stern, Chris Swords, Madina Tembulatova, Melanie Wallskog, Nina Zhao, and class members of the following modules at TCD: EC4150 (2016/17, 2017/18 and 2018/19) and ECU44081/3 (2019/20 and 2020/21).

We thank our dedicated data entry partners: Elmer Devilleres, Hazel Lyne Fianza,

Matthew Kimuyu, Geniva Pongasi Amion, Lavenia Tampus, Jeriford Devilleres, Dinabeth Maranga, Kate Giljang, Rica Mae Sombilon, Marie Mar Maranga, Ritse Gimeno, Jo-an Gador, Patricia Maranga, Angelyn Maranga, Sheryll Jandayan, Naomi Kathure, Perpetual Kariuki, Hellen Nyambura, and Nicole Ochilo.

We also thank Elisa Pepe (NBER), Semida Munteanu and Joan Youngman (Lincoln Institute of Land Policy), and Naoise McNally (now University College Dublin), for their tireless support with the grant administration of this project and for all their assistance on matters financial and logistical.

B Data Construction

Newspaper Data

In this appendix we provide further detail on the dataset covering our 30 cities. Listings were obtained by sampling the residential real estate sections of newspapers published in each city. Our sampling procedure aimed to identify 150 rental and 250 sale listings from each sampled newspaper edition, typically the last Sunday of the month of interest. To ensure that all areas covered by the newspaper would be proportionally included in the sample, research assistants sampled across columns. This sampling approach was particularly important after 1950, when the real estate section frequently spanned many pages and contained hundreds of listings. Listings were then chosen at random to be digitized as long as they contained (1) a price, (2) size as measured by number of either rooms or bedrooms, and (3) an indication of location within the city. The location could be either an address, an intersection, or an area. If the research assistants could not identify at least 150 rentals or 250 sales that met these criteria, they would consult the newspaper from the last Sunday of adjacent months. This is particularly relevant earlier in the sample, particularly in the 1890s, when valid listings were more scarce.

Figure B1: Example of a Valid Listing

Note: A sample real estate listing from the LA Times in 1980.

Figure B1 is an illustration of acceptable listings from the *LA Times*. The neighborhood assigned to these listings would be "Huntington Harbour," based on the column heading. The first and third listings (in yellow blocks) are valid because they contain a price, a measure of size (bedrooms), and are clearly houses for sale. They also have an address in addition to the neighborhood. These are the types of listings we use to train our city area algorithm (see Appendix C for details). The second listing is not valid because it lacks a measure of size.

The newspapers, access websites, and years for each newspaper are listed in Appendix Table B1. We also list the start date for each city segment in the final two columns. We started the series for a particular city and segment if we could consistently obtain about 30 valid listings per year. As a result, while most major cities start in 1890 for both sales and rentals, some cities enter the dataset later: Miami and Phoenix sale listings start in 1910 while Las Vegas sales start in 1943.

Table B1: Newspapers and Access Sites

City	Newspapers Sampled	Access Site	Start Year (City)	
			Rent	Sales
Atlanta	*Atlanta Constitution* (1890–2006)	Newspapers.com; NYPL	1890	1890
Baltimore	*Sun* (1890–2006)	Newspapers.com; NYPL	1890	1890
Boston	*Boston Globe* (1890–2006)	ProQuest	1890	1890
Charleston	*Post-Courier* (1893-2006)	Genealogybank.com; NYPL	1893	1911
Chicago	*Chicago Tribune* (1890–2006)	ProQuest	1890	1890
Cincinnati	*Cincinnati Enquirer* (1890–2006)	Newspapers.com	1890	1890
Cleveland	*Plain Dealer* (1890–2006)	Genealogybank.com	1894	1890
Dallas	*Dallas Morning News* (1890–2006)	dallasnews.com; Genealogybank.com; NYPL	1890	1890
Detroit	*Detroit Free Press* (1890–2006)	Newspapers.com	1890	1890
Houston	*Post* (1896-1924)	Newspapers.com	1896	1900
	Chronicle (1924-2005)	Genealogybank.com; HPL		
Las Vegas	*Review Journal* (1943-2006)	Genealogybank.com	1948	1943
Los Angeles	*LA Times* (1890–2006)	Newspapers.com; NYPL	1890	1890
Louisville	*Courier-Journal* (1890–2006)	Newspapers.com	1890	1890
Memphis	*Commercial Appeal* (1890–2006)	Genealogybank.com	1890	1895
Miami	*Miami Metropolis/Daily News* (1910-1940)	Newspapers.com	1914	1910
	Miami Herald (1920-2006)	Newspapers.com		
Minneapolis	*Star Tribune* (1890–2006)	Newspapers.com	1890	1890
Nashville	*Tennessean* (1890–2006)	Newspapers.com	1890	1890
New Orleans	*Times-Picayune* (1890–2006)	Newspapers.com; NYPL	1890	1893
New York City	*New York Times* (1890–2006)	ProQuest	1890	1890
	New York Daily News (1980-2006)	Newspapers.com		
	Brooklyn Daily Eagle (1890–1940)	Newspapers.com		
	Bronx Homes News (1907-1940)	BCA		
Philadelphia	*Philadelphia Inquirer* (1890–2006)	Newspapers.com	1890	1890
Phoenix	*Arizona Republic* (1910-2006)	Newspapers.com	1910	1910
Pittsburgh	*Post Gazette* (1890–2006)	Newspapers.com	1890	1890
Portland	*Oregonian* (1898-2006)	Genealogybank.com	1898	1898
Salt Lake City	*Tribune* (1890–2006)	Genealogybank.com; NYPL	1891	1890
San Diego	*Union* (1890–2006)	Genealogybank.com	1900	1890
San Francisco	*San Francisco Chronicle* (1890–2006)	Genealogybank.com	1890	1890
Seattle	*(Daily) Times* (1910-2006)	Genealogybank.com; NYPL	1910	1910
St. Louis	*Post Dispatch* (1890–2006)	Newspapers.com	1890	1890
Tampa	*Tribune* (1900-2006)	Newspapers.com	1905	1900
	Bay Times (1900-1950)	Newspapers.com		
Washington D.C.	*Washington Post* (1890–2006)	ProQuest; NYPL	1890	1890

Note: The table lists the newspapers we consulted for each city and years in cases where multiple newspapers were consulted for a single city. NYPL refers to New York Public Library; HPL refers to Houston Public Library; BCA refers to Bronx County Archives. We were unable for find a newspaper for Houston for 2006.

We digitized only purely residential listings and disregarded any listings that made mention of commercial or industrial uses. We also disregarded listings that appeared to mix owned and rented housing as well as listings that described an entire building or a single room for rent, including any sublet units. We report the observation counts for each city in Table B2. The difference in observation counts across cities reflects sampling effort rather than systematic differences across newspapers. Specifically, half of the sample was collected prior to the Covid-19 outbreak, and we collected data from four newspaper editions per year for these cities. The other half of the sample was collected in 2020 and 2021 while navigating campus closures. For this half of the sample, we were able to collect data from one newspaper per year only, typically from May, before running into resource constraints. We also collected extra data for New York City for a related book project by two authors of this paper.

In Section 3 we discuss our estimation strategy, which relies primarily on two-year "rolling windows" to estimate yearly changes in price levels. However, there are several periods during the sample where we had difficulty obtaining our targeted number of listings. This was most likely to happen in the 1930s for sales (due to the Great Depression), the 1940s for rents (due to rent control), and the 1890s for both segments (emerging secondary housing markets). In the cases of low observation counts, we relaxed our specification and used longer windows of time to estimate yearly changes in price levels. Specifically, we ran a single pooled regression with year fixed effects across each set of years with insufficient listings to obtain estimates of the year-over-year change in price levels. We then spliced these "static window" estimated changes into our rolling window price index. Table B2 also lists the very rare city-year pairs for which we interpolated (year in bold, usually due to missing data) and each set of years over which we used a static window estimation (year ranges in italics).

Table B2: Observation Counts and Relaxed Specification Years

City	Total Number of Observations	Relaxed Specification Years
Atlanta	119,728	*1942-1952* (rent)
Baltimore	40,065	*1891-1905* (sale)
Boston	153,223	–
Charleston	31,319	*1917-1922*, *1940-1946*, *1962-1964* (rent)
Chicago	168,153	–
Cincinnati	162,406	–
Cleveland	46,422	–
Dallas	123,885	*1890–1900*, **1894**, *1942-1948* (rent) ; *1890–1900*, **1891**, **1894** (sale)
Detroit	103,678	*1942-1949* (rent)
Houston	38,379	*1896-1904*, **1925** (rent) ; *1900-1905*, *1924-1930*, **2003** (sale)
Las Vegas	22,531	*1943-1948* (rent)
Los Angeles	156,136	*1944-1947* (rent)
Louisville	125,290	*1918-1922*, *1943-1948* (rent)
Memphis	40,916	*1891-1897*, *1945-1948* (rent)
Miami	83,068	*1914-1924* (rent)
Minneapolis	44,805	*1944-1948* (rent)
Nashville	39,737	*1890–1900*, *1909-1914*, *1943-1949* (rent)
New Orleans	129,719	*1914-1922* (rent)
New York	279,066	–
Philadelphia	145,913	*1890–1900* (sale)
Phoenix	102,375	*1910-1920*, *1942-1947* (rent) ; *1910-1920* (sale)
Pittsburgh	49,186	*1945-1948* (rent) ; *1890–1900* (sale)
Portland	42,042	*1898-1903*, *1943-1948* (rent)
Salt Lake City	38,015	*1941-1944*, *1945-1948*, *1948-1951* (rent)
San Diego	37,401	*1900-1910*, *1941-1948* (rent) ; *1890–1905* (sale)
San Francisco	142,874	*1943-1948* (rent)
Seattle	31,486	*1940-1944*, *1944-1948*, *1948-1951* (rent) ; *1910-1920* (sale)
St. Louis	45,922	*1944-1947* (rent)
Tampa	34,791	*1905-1915* (rent) ; *1905-1920* (sale)
Washington D.C.	123,307	–

Note: The table lists the sample city, the total number of observations for the city, and the interpolated (bold) and static window (italicized) years for the city.

Table B3: Rent: Real Median Price and Modal Size

	Rent					
	Median Price (2006 dollars)			Modal Size		
	1890-1945	1946-1970	1971-2006	1890-1945	1946-1970	1971-2006
City						
Atlanta	582	552	843	6	2	2
Baltimore	489	544	609	3	2	3
Boston	456	609	1,281	5	2	2
Charleston	399	418	655	6	2	2
Chicago	766	812	963	6	2	2
Cincinnati	527	593	575	4	2	1
Cleveland	564	695	590	5	2	2
Dallas	469	609	938	5	2	3
Detroit	588	492	556	5	2	2
Houston	480	555	814	5	2	2
Las Vegas	.	754	795	.	2	2
Los Angeles	530	770	1,239	4	1	2
Louisville	440	439	527	4	2	2
Memphis	526	451	439	4	2	2
Miami	464	573	1,036	4	1	2
Minneapolis	485	707	749	5	1	1
Nashville	491	483	700	5	2	2
New Orleans	476	520	635	5	1	2
New York City	1,000	1,032	1,436	4	2	1
Philadelphia	476	556	764	6	2	1
Phoenix	354	544	795	3	2	2
Pittsburgh	556	605	631	5	2	1
Portland	431	523	572	5	2	2
Salt Lake City	397	466	650	3	2	2
San Diego	386	564	740	5	1	1
San Francisco	499	812	1,330	4	2	1
Seattle	438	607	750	3	1	1
St. Louis	509	456	506	3	2	1
Tampa	358	419	650	5	2	2
Washington D.C.	696	749	1,058	6	1	1

Note: The table shows the median unit rental prices in 2006 dollars and the modal number of rooms (fourth column) or bedrooms (fifth and sixth column) in each of three periods.

Table B4: Sale: Real Median Price and Modal Size

	Sale					
	Median Price (2006 dollars)			Modal Size		
	1890-1945	1946-1970	1971-2006	1890-1945	1946-1970	1971-2006
City						
Atlanta	74,300	107,199	170,079	6	3	3
Baltimore	46,781	80,106	163,728	6	3	3
Boston	78,140	109,279	273,253	6	3	3
Charleston	51,745	84,225	143,570	6	3	3
Chicago	113,223	146,181	232,885	6	3	3
Cincinnati	83,901	111,750	156,559	6	3	3
Cleveland	92,606	134,004	170,005	6	3	3
Dallas	67,269	135,784	220,852	5	3	3
Detroit	84,134	102,381	145,551	8	3	3
Houston	64,340	107,425	183,686	5	3	3
Las Vegas	72,164	116,946	174,614	4	3	3
Los Angeles	84,234	205,643	430,839	5	3	3
Louisville	58,304	82,538	129,102	5	3	3
Memphis	66,224	87,979	155,187	5	3	3
Miami	86,041	112,702	197,217	5	3	3
Minneapolis	55,902	106,506	181,182	5	3	3
Nashville	56,842	90,359	167,981	5	3	3
New Orleans	69,261	111,005	162,650	5	3	3
New York City	147,368	188,372	325,161	6	3	3
Philadelphia	58,007	103,091	191,669	6	3	3
Phoenix	53,617	94,915	177,366	5	3	3
Pittsburgh	85,909	111,176	147,634	6	3	3
Portland	54,723	92,715	153,334	5	3	3
Salt Lake City	53,476	104,758	167,130	5	3	3
San Diego	56,798	112,799	309,145	5	3	3
San Francisco	84,604	164,930	456,947	5	3	3
Seattle	53,146	111,176	234,173	5	3	3
St. Louis	70,774	89,238	140,216	6	3	3
Tampa	44,800	75,349	153,008	5	2	3
Washington D.C.	108,792	144,268	269,176	6	3	3

Note: The table shows the median housing prices in 2006 dollars and the modal number of rooms (fourth column) or bedrooms (fifth and sixth column) in each of three periods.

Housing Permits

Our measure of housing permits relates to dwellings for which permission to build was granted during the year. For the period from 1921 to 1959, we use tabulations for the city only. From 1960 onward, we use data for the full metropolitan statistical area (MSAs). Prior to 1980, all series were taken from either the U.S. Census Bureau sources made available on the HathiTrust website and were processed either manually or via OCR. From 1980 onward, all series were taken from digital datasets from various U.S. government sources:

- **Central Cities, 1921-40 and 1946-1948**: United States, Congress. *Handbook of Regional Statistics*. US Government Printing Office, 1950.
- **Central Cities, 1949-60**: United States, Census Bureau. *Statistical Abstract of the United States*. US Government Printing Office, 1953-1961.
- **MSAs and Central Cities, 1961-1980**: United States, Census Bureau. *Annual Construction Reports*. US Government Printing Office, 1967-1980. Annual summaries (13th editions) were consulted for yearly building permit data.
- **MSAs, 1980-1987**: United States, Census Bureau. *Census Text Files*. US Census Bureau, 1980-1987.
- **MSAs, 1988-2020**: United States, Census Bureau. *Electronic Permit Data and Census Text Files*. US Census Bureau, 1988-2024.

C Geocoding

The geocoding process uses a random forest algorithm to create a consistent set of areas for each city using the disparate geographic information we have from newspaper listings. A valid listing requires an indication of location within the city. However, this information could be either an address, an intersection, or a neighborhood from either the listing or a column heading such as "West Philadelphia." For some southern cities such as New Orleans (and prior to the Fair Housing Act), a common column heading was "Colored" or similar to indicate housing open to African Americans. We took whatever information was listed in newspapers, so our "area" definitions are not necessarily anchored to mutually exclusive geographies. "Downtown" or cardinal directions are also common.

Our process is as follows: upon receiving the digitized listings, we first formed an address or general location using all available information. For instance, we would combine a street address from the listing with a neighborhood from the newspaper heading. In some early cases, the most detailed geographic information we could get was a quadrant, for example "Southwest Chicago." We next ran the address or general location through Google Maps to get a latitude and longitude. There are two types of inputs. In the first case, we had an address or intersection along with an area from the newspaper. In the second case, we had only an area, and the geocoder returned the centroid of some area or neighborhood from Google Maps.

We then took observations from the case where we had both a neighborhood and a latitude and longitude from an address or intersection to train our random forest algorithm to generate a set of 20 consistent areas. The areas were chosen to be the top twenty most frequently occurring in the data for each city, plus a missing category. We input the coordinates, the price (listing price or capitalized rent), the total rooms, and the segment (rent or sales) to predict neighborhood. We used this model to add the neighborhood to observations without one – for instance, in the case of listings that had only an address. The algorithm places listings with only an area in the listed area if it is in the top twenty and in the most similar neighborhood according to the algorithm if not.

We chose twenty areas so that there was a sufficient set of listings to obtain credible identification of the area fixed effects given our relatively small samples in each year. The top twenty areas are based on the entire 116 year period for each city, so in early years most cities do not yet span all twenty neighborhoods. For instance, New York City has twenty areas over the whole sample while Salt Lake City has ten in 1890; many listings were assigned to the missing category in early years. We also repeated the algorithm using ten areas to address the concern that our sample size was insufficient for twenty neighborhoods. The indices are similar whether we use ten or twenty areas.

Figure C1: Example of Geocoding and Area Classification, Las Vegas

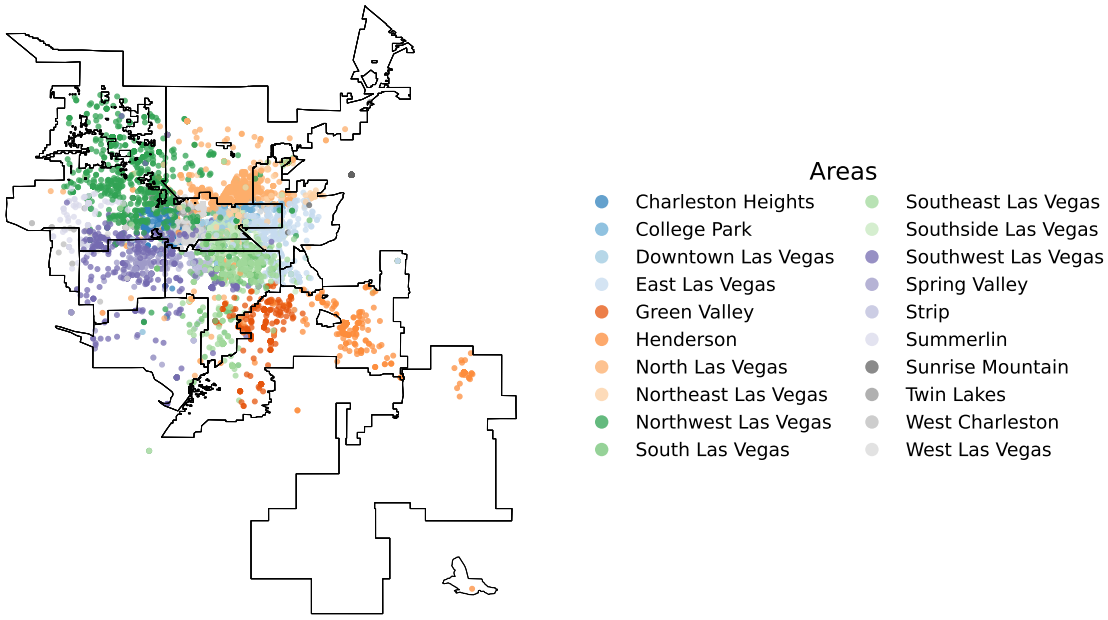

Note: The figure shows the geographic distribution of HHP listings for both sale and rental segments over the 1948-2006 period along with the corresponding area assigned by the machine learning approach.

Figure C1 shows a representative example of the geocoding and classification of listings into twenty areas. These twenty areas include neighborhoods, cardinal directions within Las Vegas, and surrounding cities within the Clark County/Las Vegas-Henderson Paradise MSA. This figure makes clear that the areas are not perfectly delineated by geographic area because we also use price to predict area assignment.

D Comparisons to Existing Housing Price Data and Series

In this appendix, we review the existing sources of information on the price of housing in U.S. cities and compare them to the HHP data. We begin by benchmarking the HHP to the census to compare the respective distributions of prices and number of rooms, which were recorded in both datasets. We then compare the HHP rental index to both the BLS Rent of Primary Residence (RoPR) series as well as commonly used revisions to this series at both the national and city level. Finally, we compare the HHP sale price index to indices produced by Case-Shiller-Weiss and the FHFA at both the national and city level.

D.1 Benchmarking HHP to the Census

One first source for information on home values and rents paid by American households that we consider is the census, which first asked these questions in 1930. The census is the most commonly used source for scholars needing information on the distribution of rents for a specific city or metro area (for instance, see Quigley & Raphael, 2004). We first compare the cross-sectional distribution of the HHP data to the census at the city level for both prices (first asked by the census in 1930) and the number of bedrooms (first asked by the census in 1960). Figure D1 shows the distribution of sale prices or rents from HHP compared with owner valuations or contract rents from the census using IPUMS microdata samples for each year.[30] The city-year pairs were chosen to illustrate the typical cases of distributional overlap. The top four panels show sales versus owner valuations and the bottom four panels show market rents versus census contract rents.

We begin with sales. A common case is that the distribution of HHP sales aligns well with the census owner valuations, which we see in many city-year pairs. For instance panel *a* shows very close alignment for Philadelphia in 1940, a year in which the market was close to the 1938 trough. Another common case is shown for Philadelphia in 1970 in panel *b*, when our HHP sales data are shifted to the right relative to the censored census distribution (at $50,000 in 1970). Topcoding is even more of an issue in expensive coastal cities, for instance see panel *c* with San Francisco in 1970. Our data thus contain more information on the upper end of the housing market relative to the census, particularly in years with binding topcodes like 1970.

What is driving the rightward shift of our data relative to the census in cases such as Philadelphia in 1970? One explanation is limited homeowner awareness of nominal price inflation, which was greater in the 1960s relative to earlier decades. However, owner-

[30] Steven Ruggles, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Renae Rodgers, and Megan Schouweiler. IPUMS USA: Version 15.0 [dataset]. Minneapolis, MN: IPUMS, 2024. https://doi.org/10.18128/D010.V15.0

Figure D1: Census Benchmarking: Prices

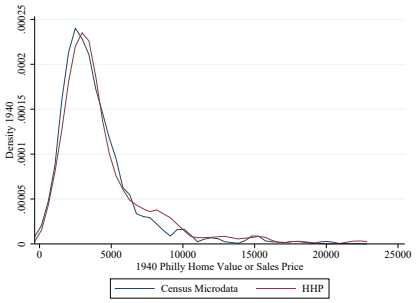

(a) Philadelphia Sales in 1940

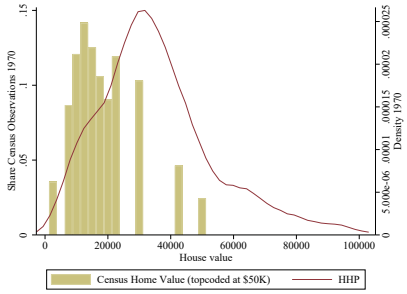

(b) Philadelphia Sales in 1970

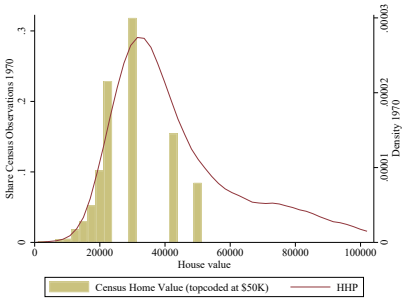

(c) San Francisco Sales in 1970

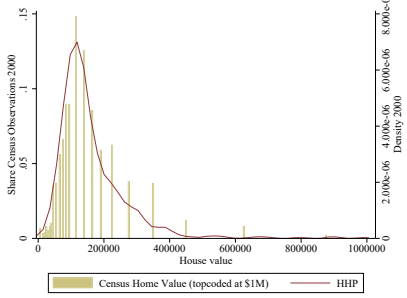

(d) Louisville Sales in 2000

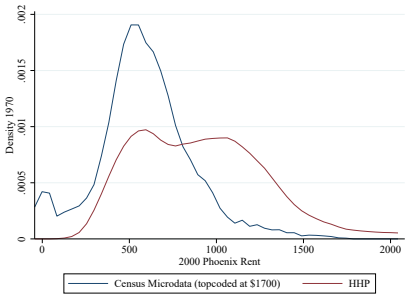

(e) Phoenix Rents in 1970

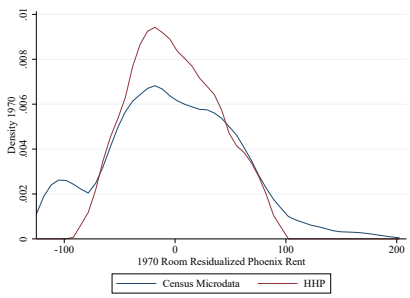

(f) Phoenix Residual Rents in 1970

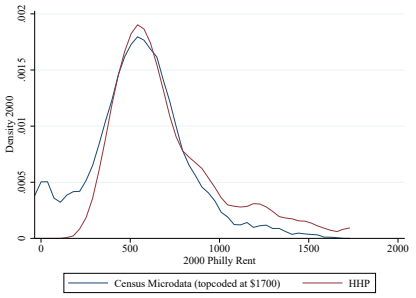

(g) Philadelphia Rents in 2000

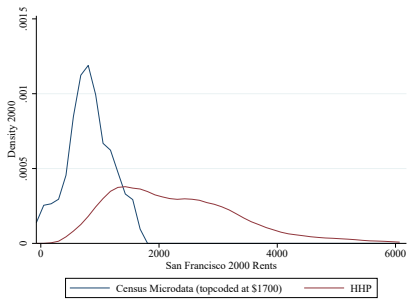

(h) San Francisco Rents in 2000

Note: This figure shows the distribution of census owner valuations against the corresponding year's HHP sale prices (panels a-d) and the distribution of census contract rents against the corresponding year's HHP rental prices (panels e-h). Each panel shows nominal prices in that year's dollars.

reported values and HHP sales data are back in alignment by 2000 in many cities, as seen in panel *d* for St. Louis, suggesting that homeowners were more aware of housing price inflation during the 1990s. Another explanation is positive selection into newspaper listing. The binned format of home value data in the census makes it hard to adjust for housing attributes; however, below we residualize the continuous measure of rents in the census by number of rooms (which is available starting in 1960) to compare with our data, also residualized by number of rooms, and find a much closer alignment in the distributions. Thus, to the extent that our sample of homes for sale is likely to be positively selected in some cities or years, controlling for size and area of the city should address a substantial portion of bias.

We similarly explore the relationship between listed rents and rents reported by households in the census in the final four panels of Figure D1. Although renters know their contract rent with more accuracy than homeowners know their home's market value, these distributions could diverge because census rents lag market conditions, particularly during periods of high inflation. There could also again be positive selection into newspaper listings. We generally found a close alignment or a rightward shift relative to the census in most city-year pairs. Panel *e* shows the relationship between market rents in Phoenix in 1970 from the HHP data and reported values from the census. It is clear that the HHP distribution is shifted to the right relative to the census. To understand this shift, we residualize the rents from both sources by number of rooms and repeat the density plots for Phoenix in 1970 in panel *f*. The residualized rents show a much closer alignment, suggesting that our controls in the hedonic specification should be effective. We can also close the gap by using previous years of HHP rents, such as comparing 1968 HHP market rents with 1970 contract rents (not shown). It is not generally possible to know how much of the gap is driven by selection versus sticky rents, as we don't know when leases were signed in the census. However, in many city-year pairs we find a close alignment between census and HHP rents even without adjusting for housing unit size, for instance see panel *g* for Philadelphia in 2000.

When we look at a city with a longstanding rent control policy, such as San Francisco in 2000, we see a major divergence in market rents from HHP and reported rents from the census. We note that the census topcoded rents at $1,700 in 2000 while a large share of market rents from our data are above $2,000. Our data in this case cover a very selected portion of the rental housing stock in the city of San Francisco, namely units that were available to rent at market prices, plus listings from expensive Bay Area counties like Marin where no rent control policy was in place. Rent control introduced during World War II resulted in fewer market-rate units being advertised in the 1940s. We collected listings using the same process in these years but typically ended up with fewer observations. We relaxed our

rolling window specification for cities with relatively few observations during the 1940s to obtain more robust year coefficients, as discussed in Appendix B.

To summarize the benchmarking of HHP rents to the census, rent control policies create the greatest wedge between market conditions and reported rents in the census. In normal years, we found either a close correspondence or a rightward shift in the HHP data relative to the census. To the extent that these shifts are driven by sticky rents or inaccurate homeowner valuations, our data are likely more indicative of market conditions. To the extent that these shifts are driven by positive selection into newspaper listing, our controls for housing unit size and area should address the most important sources of bias.

Figure D2: Census Benchmarking: Room Size of Sale Observations

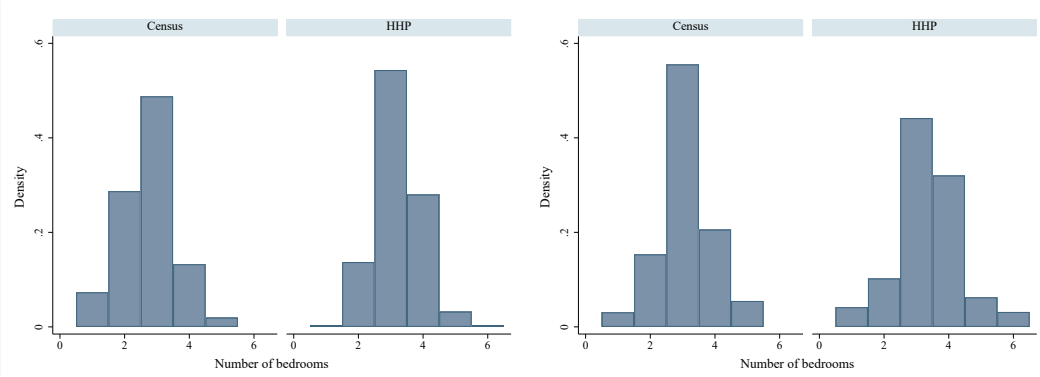

(a) Phoenix Sales in 1970

(b) Philadelphia Sales in 1980

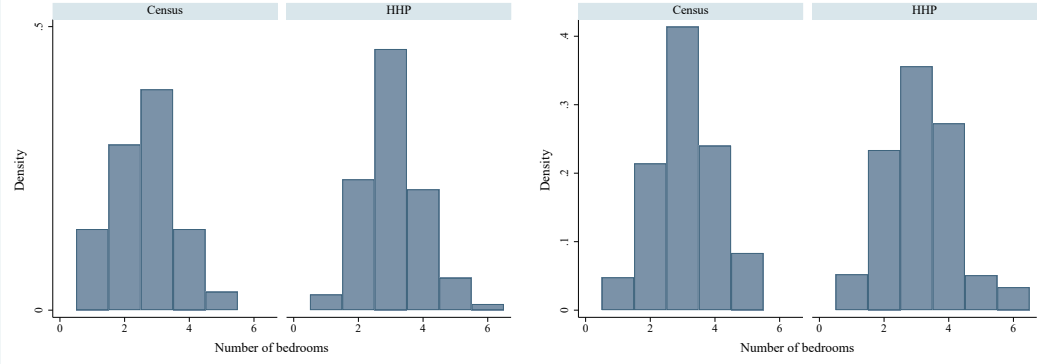

(c) Miami Sales in 1990

(d) Boston Sales in 2000

Note: This figure shows the distribution of the bedroom distribution of owner-occupied housing from the census versus the corresponding year's HHP bedroom distribution from the same metro area.

As noted above, the rightward shift of the HHP price distribution relative to the census suggests positive selection into newspaper listing. To investigate this notion further, we compare the distribution of number of bedrooms across the two datasets. This question

Figure D3: Census Benchmarking: Room Size of Rental Observations

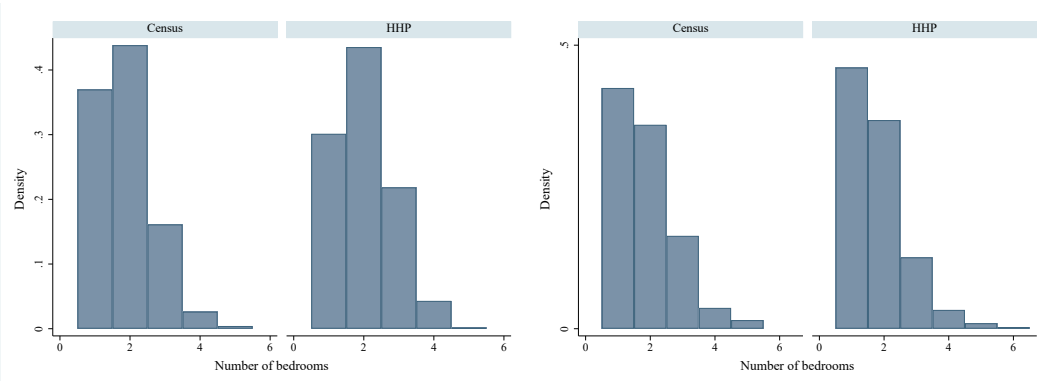

(a) Phoenix Rents in 1970

(b) Philadelphia Rents in 1980

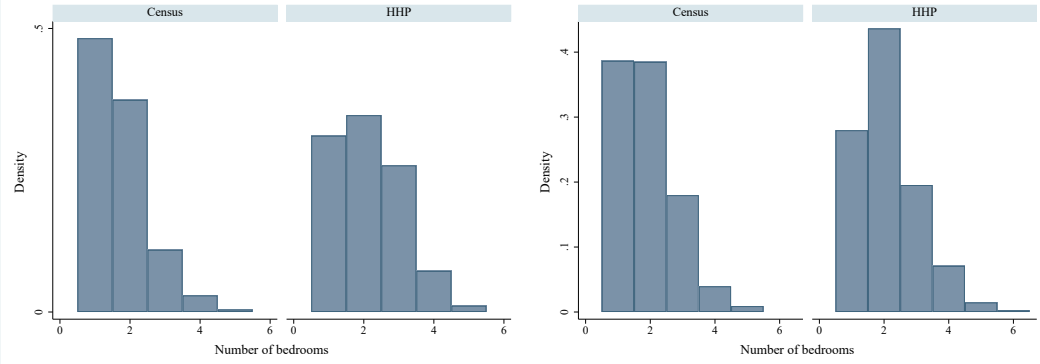

(c) Miami Rents in 1990

(d) Boston Rents in 2000

Note: This figure shows the distribution of the bedroom distribution of owner-occupied housing from the census versus the corresponding year's HHP bedroom distribution from the same metro area.

was first asked in 1960, so we can only do this exercise for the second half of the sample. Figure D2 shows the relationship between the number of bedrooms in owned housing units for four cities in four years, again with cases chosen to reflect typical cases of overlap. For every city shown (Phoenix in 1970, Philadelphia in 1980, Miami in 1990, and Boston in 2000), we see that the modal home for sale (HHP) or occupied by owner (census) has three bedrooms. However, the HHP dataset has relatively more large housing units (with four bedrooms) while the census reports more housing units with two bedrooms. This positive selection of housing into newspaper listing is consistent with the rightward shift of the price distribution that can be corrected by residualizing by number of bedrooms as discussed above.

Figure D3 shows a similar correspondence for rental units. In both the census and HHP, the modal rental unit typically has two bedrooms, with one-unit bedrooms the second-most common. The census typically reports more one-bedroom apartments rel-

ative to the HHP dataset, however. The cases of Miami in 1990 and Boston in 2000 illustrate this difference. It seems that larger rental units were more likely to be advertised in the newspaper relative to one-bedroom units. We are able to adjust for room size in our hedonic specification to address this selection. The fact that the indices computed from the HHP data align very well with the FHFA and Case-Shiller indices from the 1975-2006 and 1987-2006 periods, respectively, lend credibility to our indices covering the full 20th century.

D.2 Benchmarking HHP to Existing Rental Series by Data Source

Figure D4: Comparison to BLS Real Rent Indices, 1914-2006

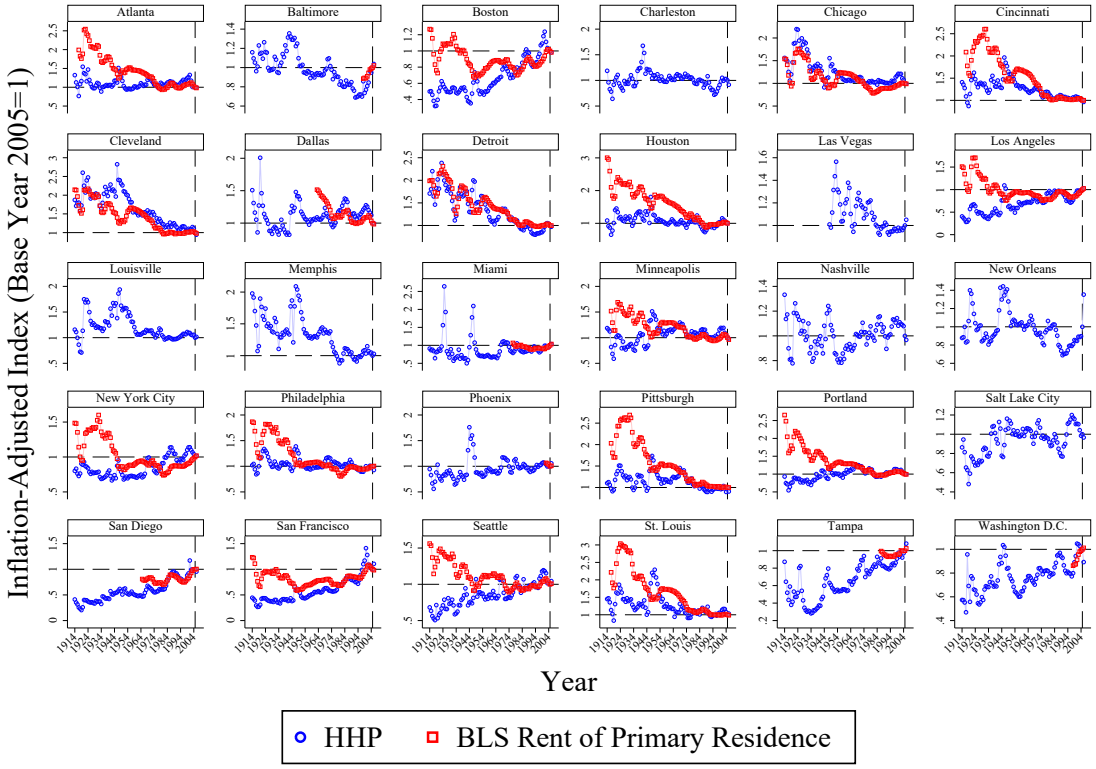

Note: This figure compares the baseline HHP rental series for each sample city to the BLS Rent of Primary Residence series for the corresponding MSA (when available) with a base year of 2005. The MSA-level series were accessed from FRED at the Federal Reserve Bank of St. Louis, for instance the New York series was obtained from U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: Rent of Primary Residence in New York-Newark-Jersey City, NY-NJ-PA (CBSA) [CUURA101SEHA], retrieved from FRED, Federal Reserve Bank of St. Louis: https://fred.stlouisfed.org/series/CUURA101SEHA. All series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

In this subsection we compare the HHP rental index to existing rental series at the na-

tional and city level. Since the seminal paper by Eichholtz (1997) on Amsterdam housing prices since the 1620s, a growing literature has assembled very long-run rental series for major European cities (Eichholtz et al., 2019; Eichholtz et al., 2021; Eichholtz et al., 2012). However, very little information on market rents is available for U.S. cities until the 1990s or 2000s (Ambrose et al., 2015). Scholars interested in the history of American rental housing markets have instead relied on Bureau of Labor Statistics (BLS) surveys of current rents undertaken to compute the CPI. These surveys, which were first done in 1914, form the basis of the rental series underlying the Jordà-Schularick-Taylor (JST) Macrohistory Database for the U.S. (Jordà et al., 2017).[31]

In the paper, we discuss the main findings of this comparison, but in this appendix we review the benchmarking exercise in more detail. In Figure D5a we focus on the 1890-1914 period and compare HHP with the Rees and Jacobs (1961) series, which is based on the unweighted, unadjusted average rental price from six city newspapers.[32] While the cyclical behavior across the two series align well, considering the methodological and sample differences, the overall change in rents 1890–1914 (adjusting for inflation) is very different: while in the Rees-Jacobs series, real rents were 4% lower in 1914 than in 1890, in the HHP series, they are 17% higher.

We next compare our market rental index with the RoPR series, which has been collected since 1914.[33] This is not a straightforward exercise, as the RoPR series is not intended to capture market conditions. We are nonetheless interested in this comparison since market conditions should converge on average rents over the long run, and the overall trajectory of rental prices is both in dispute and of interest. Figure D5b plots the RoPR against the HHP market rents series. The paths of rents across the two series are strikingly different. The HHP series has a broadly upward trend, albeit with six substantial market cycles, with real rents peaking in 1923, 1948, 1969, 1986 and 2001 and bottoming out in 1918, 1934, 1960, 1975 and 1993. As discussed in the main text, market rents in 2006 were 36% higher than in 1914.

In contrast, the RoPR series indicates steadily falling real rents, in particular after World War II, with real rents in 2006 effectively at effectively half their 1914 level. As scholars have previously expressed skepticism at this finding, proposing corrections to the RoPR series, we also compare our rental series to adjustments proposed by Gordon & van

[31] See https://www.macrohistory.net/database/.

[32] Tables 22 and 32 in Chapter 4 of Rees & Jacobs (1961) provide estimates of market rent indices for six cities and a national index based on the unweighted average of their levels.

[33] The BLS RoPR series can be accessed from the St. Louis Federal Reserve Bank FRED website: U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: Rent of Primary Residence in U.S. City Average [CUUR0000SEHA]. The CPI we use in this paper is the BLS Consumer Price Index for All Urban Consumers (CPI-U) Database, series [CUUR0000SA0], with the ultimate source being quarterly reports on the Consumer Price Index for the urban U.S.

Figure D5: Benchmarking Against Rees-Jacobs and the BLS, 1890–2006

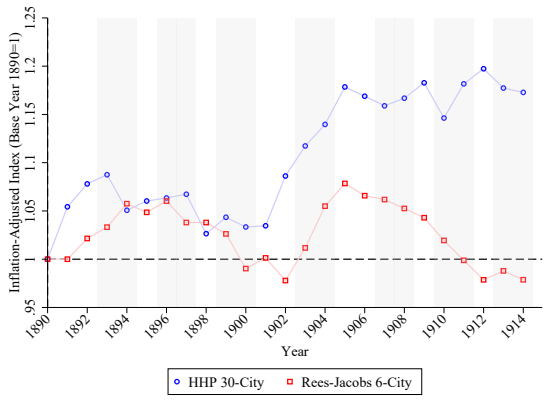

(a) Rees-Jacobs, 1890–1914

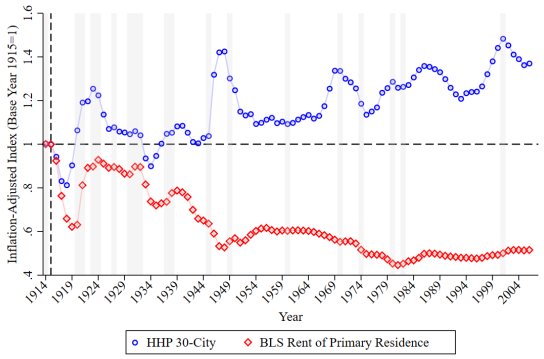

(b) BLS RoPR, 1914-2006

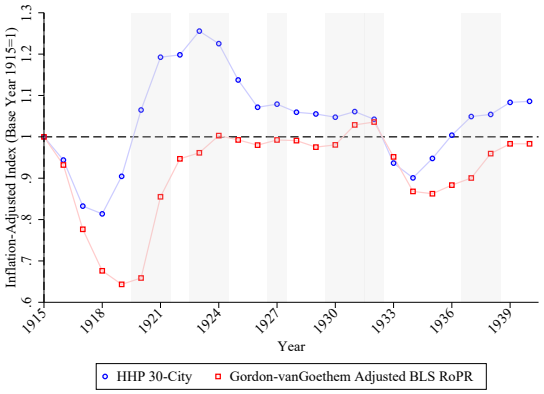

(c) Gordon-van Goethem, 1915-1939

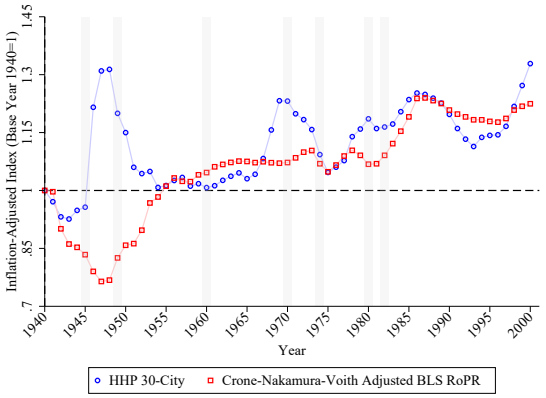

(d) Crone-Nakamura-Voith, 1940-1995

Note: This figure shows the baseline HHP national rental price series against the Rees-Jacobs series in panel (a), the BLS Rent of Primary Residence in panel (b), the adjusted BLS Rent of Primary Residence series proposed by Gordon & van Goethem (2007) in panel (c), and the adjusted BLS Rent of Primary Residence series proposed by Crone et al. (2010) in panel (d). All series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

Goethem (2007) for 1915-1939 and Crone et al. (2010) for 1940-1995 in Figures D5c and D5d, respectively. We find reasonably close agreement with the Gordon & van Goethem (2007) series, with both series seeing similar real rental price levels in 1914 and 1939. For the 1940-2000 period covered by Crone et al. (2010), there is similarity in the overall change in rents between HHP and CNV, although – by construction – the CNV series cannot exhibit any cyclical behaviour not in the original RoPR series.

D.3 Benchmarking HHP to Existing Sale Price Series

For the sale price of housing, city-level information prior to 1975 is primarily limited to the owner valuations reported in the census (Davis & Heathcote, 2007; Gyourko et al., 2013). Beginning in 1975, the FHFA provides MSA-level sale price indices, and the S&P/Case-Shiller MSA-specific indices begin in 1970 at the earliest, with other markets added in 1987. Scholars needing information on the price of housing in the U.S. at an annual frequency before 1975 have typically relied on the pioneering work of Shiller (2015), who spliced together several sources of data to create a national housing price index from 1890 to the present:

- For the period 1890–1934, the index is based on a survey of owner valuations taken in 1934 using the owner's recollection of the transaction price and their assessment of its value in that year (Grebler et al., 1956).
- For the period 1934-1953, the index is based on median listing prices from newspaper advertisements from five major cities. An average of 30 newspaper listings was used for each city-year.
- For the period 1953-1975, the index uses data from government-backed mortgage programs (the Home Purchase Component of the CPI). The underlying data are a truncated sample of housing market transactions with the price ceiling for mortgages covered by the programs changing over time.
- For the period 1975-1987, the index uses the home price index created by the Office of Federal Housing Enterprise Oversight (OFHEO, now the FHFA). This is a repeat-sales index but includes valuations based on appraisals, as well as open-market transactions, with the relative share of appraisals unclear.
- From 1987 to the present, the Shiller index is the same as the Case-Shiller index, also based on a repeat sales method.

The data sources underlying the Shiller index at different points in time are thus very different from each other. Economists have argued that the Shiller index is particularly difficult to reconcile with other sources around the Great Depression (Fishback & Kollmann, 2014) and that the index appears to understate inflation-adjusted increases in the value of homes from the census, particularly in the 1970s (Davis & Heathcote, 2007).

We begin with the 1890–1934 portion of the index in Figure D6a. While our series aligns to a large degree before World War I, the HHP series reports a strikingly different interwar housing price cycle relative to the 1934 survey used by Grebler et al. (1956). This

Figure D6: Benchmarking Against Components of the Shiller Index, 1890–2006

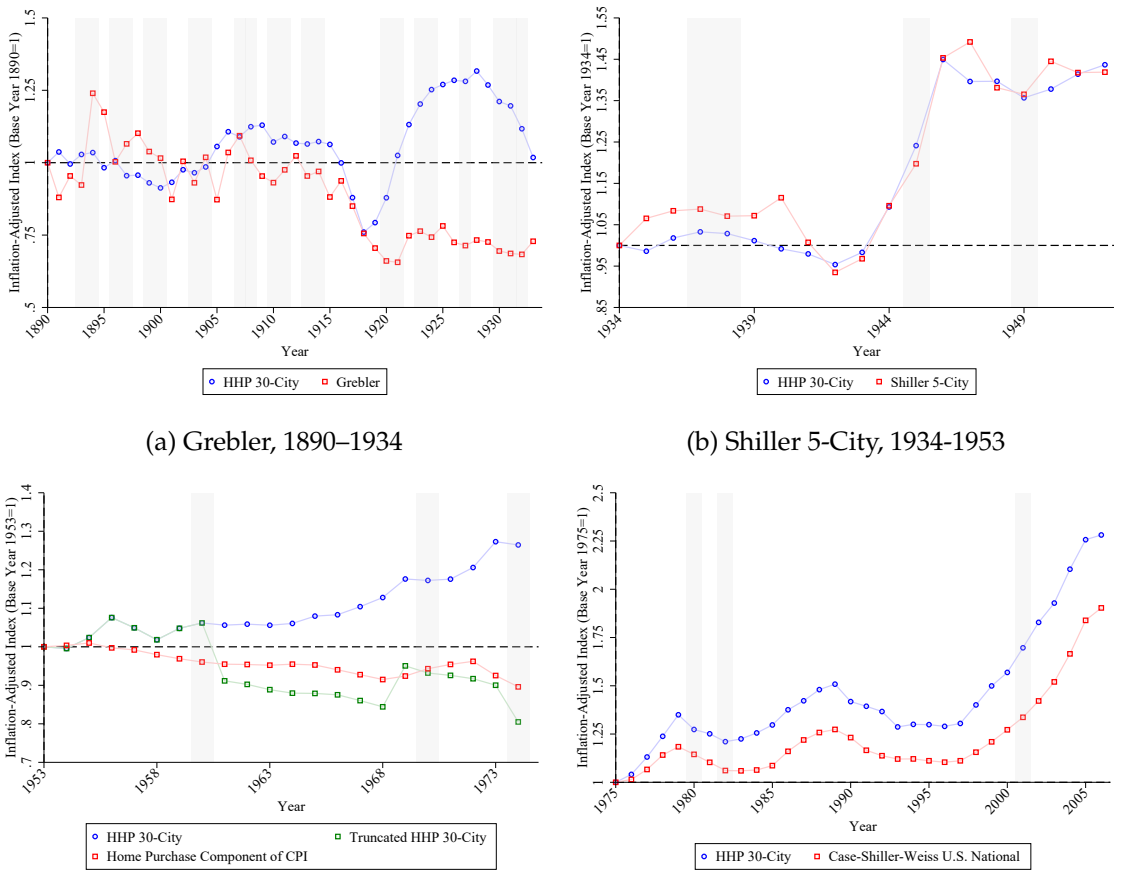

(a) Grebler, 1890–1934

(b) Shiller 5-City, 1934-1953

(c) CPI Home Purchase, 1953-1975

(d) Case-Shiller-Weiss, 1975-2006

Note: This figure shows the baseline HHP national sale price series against the following components of the Shiller index: the Grebler et al. (1956) survey in panel (a), newspaper medians for five cities in panel (b), the Home Purchase component of the CPI in panel (c), and the Case-Shiller-Weiss repeat sales index in panel (d). All series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

divergence is likely due to homeowners' lack of awareness around the changing value of their homes over the previous decades. Importantly, we find that real housing prices were lower in 1940 than in 1930, consistent with what is reported in the census, Nicholas & Scherbina (2013)'s study of New York City, and Fishback & Kollmann (2014)'s study of New Deal reports. Overall, we find that real prices rose by 73% 1918–1928 (rather than 5%, in Shiller) and then fell by 28% 1928–1935 (rather than rising a further 4.5%, per Shiller). While in the Shiller index, prices in 1935 were already above their 1928 level, in HHP, prices did not pass their 1928 peak until 1946.

Figure D6b shows the relationship between the HHP and Shiller series for the next portion of the latter index, which is based on a simple average of 30 newspaper listings

for each of five cities. Overall, these two sources give very similar trajectories for these two decades, with prices roughly 40% higher in 1953 than two decades before. This is notwithstanding the differences in sample size, city selection and methodology.

Third, we benchmark to the 1953-1974 portion of the Shiller index, which is the Home Purchase Component of the CPI constructed from a truncated sample of government-backed mortgages. The two series, reported in Figure D6c, suggest very different trajectories for housing prices during this period. In particular, the Shiller index suggests a decline in housing prices of almost 10% while the HHP series shows a 23% increase. Scholars such as Greenlees (1982) have criticized the downward bias of the Home Purchase Component of the CPI due to the exclusion of higher-valued houses as a result of loan limits imposed by the FHA.[34] To explore the impact of FHA requirements, we truncate our data using the same statutory limits imposed by the FHA: $25,000 from 1961 to 1968 and $33,000 from 1969 and 1975 (as reported in Vandell (1995)). Roughly 40% of our sale listings across 30 cities (105,130 of 246,154 listings) from 1961-1975 are excluded if the FHA limits are used. Our truncated data match the Shiller index more closely. Our data thus support the Greenlees critique and suggest real housing prices began rising earlier than 1997, as also argued by Davis & Heathcote (2007).

Finally, we benchmark to the last portion of the Shiller index (1975–2006) in Figure D6d. This period in the Shiller index uses two main sources, both of which are repeat-sales indices (Shiller, 2015, p.235): the 1975–1986 portion is the OFHEO repeat sales index, while the 1987–2006 portion is the Case-Shiller-Weiss (CSW) index. For the entire period, the HHP series shows more real price growth than CSW (2.7% per year rather than 2%), with most of the difference stemming from the 1970s. Between 1981 and 2006, the Shiller series shows 80% growth in real prices, while the HHP series shows 82%. However, for 1975–1981, the HHP series indicates growth in prices of one quarter, compared to just 4% in Shiller. Why does the HHP series indicate more growth in the 1970s? As Shiller (2015, p.235) notes, the OFHEO series uses both sales and appraised values and appraised values may only reflect changes in market conditions with a lag.

To further explore the performance of the most recent decades of the HHP dataset, we benchmark our sample at the MSA level to the FHFA All-Transaction home price indices (starting in 1975), the S&P/Case-Shiller home price index (starting in 1987 when available for an MSA), and the early Case-Shiller indices (1970-1986) in Figures D7, D8, and D9, respectively.[35] The MSA-level indices reveal a close agreement when both series are avail-

[34]These criticisms contributed to the abandonment of the home purchase approach and the adoption of the rental equivalence method in 1983.

[35]The source for the FHFA series is the FHFA All-Transactions Index for MSAs (base year of 1975). These data were accessed from https://www.fhfa.gov/data/hpi/datasets. The S&P Case-Shiller series (when available for the MSA) were accessed from FRED at the Federal Reserve Bank of St. Louis: for instance the

Figure D7: Benchmarking Against FHFA Real House Price Indices, 1975-2006

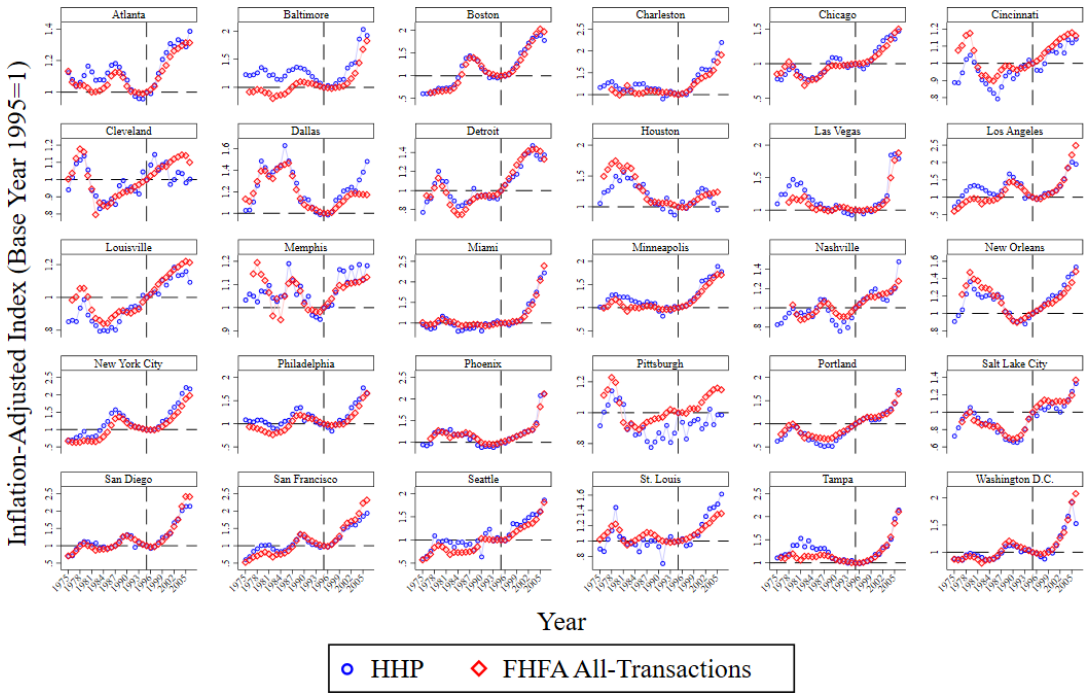

Note: This figure shows the comparison between the baseline HHP city sales series and the FHFA All-Transactions Index for MSAs (base year of 1995). These data were accessed from https://www.fhfa.gov/data/hpi/datasets. All series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

able for most cities; however, for some cities, such as Cleveland and Pittsburgh, there is greater volatility in the HHP series, likely due to smaller sample sizes.

In sum, the HHP series aligns well with the Case-Shiller-Weiss and FHFA series at both the national and city level. Divergences for earlier periods correspond closely to known limitations of the data underlying those series, as highlighted by previous scholars.

S&P Dow Jones Indices LLC, S&P CoreLogic Case-Shiller GA-Atlanta Home Price Index was accessed from https://fred.stlouisfed.org/series/ATXRNSA. The source for the 1970-1986 series is Case & Shiller (1987).

Figure D8: Benchmarking Against S&P/Case-Shiller Real Home Price Indices, 1987-2006

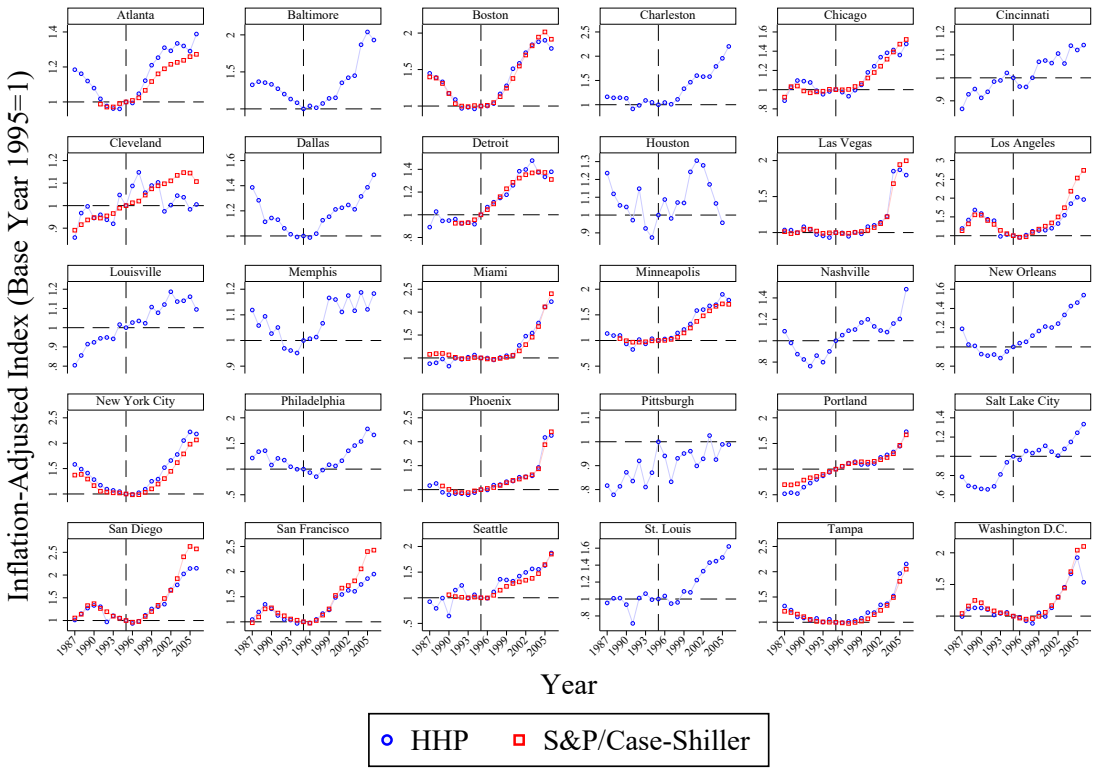

Note: This figure shows the comparison between the baseline HHP sales series for each sample city against the S&P Case-Shiller series for the corresponding MSA (when available) with a base year of 1995. The S&P Case-Shiller series were accessed from FRED at the Federal Reserve Bank of St. Louis; for instance the S&P Dow Jones Indices LLC, S&P CoreLogic Case-Shiller GA-Atlanta Home Price Index was accessed from https://fred.stlouisfed.org/series/ATXRNSA. All series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

Figure D9: Case-Shiller 1970-1986

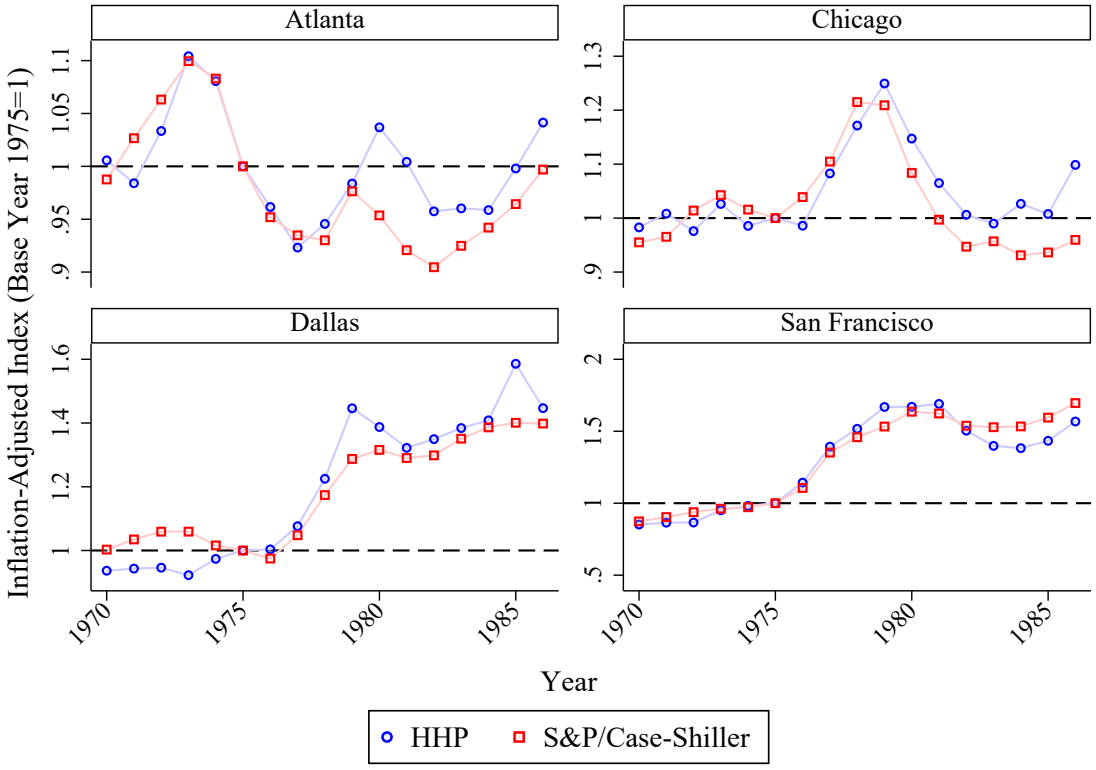

Note: This figure shows the comparison between the baseline HHP sales series for each city against the series for Atlanta, Chicago, Dallas, and San Francisco from Case & Shiller, 1987.

D.4 Construction Costs

In this appendix we visualize the nominal construction cost indices from RSMeans alongside the nominal HHP index.

Figure D10: Real Housing Prices and Construction Costs

Note: The baseline HHP sales price index is based to year 1950 and in real terms. The city-level cost indices are from the RSMeans "Square Foot Cost" volumes and we thank Raven Molloy for her transcription of this data.

D.5 Benchmarking Rental Returns

In this section we examine our estimates of rental returns. Figure D11 presents the estimated gross market rental yield, back to 1890, benchmarked to the Davis Lehnert Martin (DLM) series, which runs quarterly from 1960 onward (Davis et al., 2008). By construction, the national DLM and HHP values in 2006 are the same; unlike in the main paper, we use here for consistency with DLM their benchmark value of the rental yield in 2006 (with contemporaneous sales prices) rather than the rental return (with sales prices lagged one period). The series show broadly similar cycles and trends between 1960 and 2006, but the HHP yield in the 1960s is higher. This reflects rents having risen more slowly than sale prices to a greater extent in HHP than in DLM. In part, this reflects the constraints of the DLM method, which connects between rents in census years using other economic information. As a result, the late 1960s surge in rents is not reflected in DLM, whose yield is lower in 1970 than in 1960. Similarly, the greater sale price growth in the late 1970s/early 1980s in HHP compared to OFHEO, as described above, means that the HHP yield rises by one tenth 1980-1985 while the DLM yield falls by one tenth. The overall pattern of yields across time in the HHP data, including in the pre-1960 period, is described more fully in the main text.

Figure D11: National Rental Yields in HHP vs. Davis Lehnert Martin

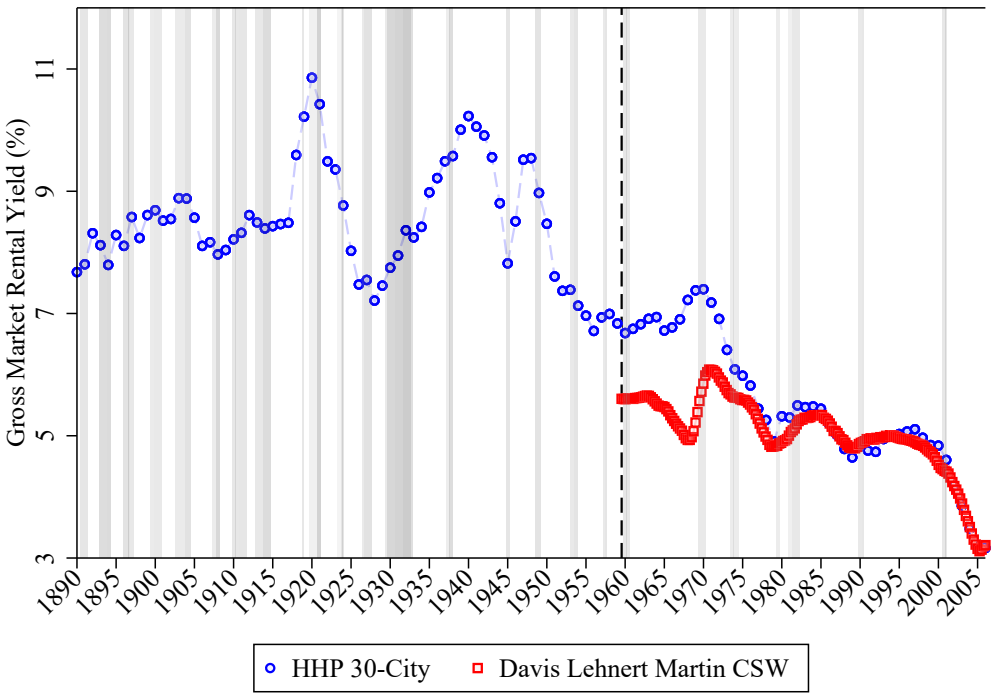

Note: This figure plots the HHP gross rental yield against the Davis Lehnert Martin series (Davis et al., 2008) with base year 2006 = 3.2%.

E Robustness Checks

Here we present additional robustness checks. In the two panels of Figure E1, we show that our results are robust to the use of household shares, rather than population shares, to aggregate from city- to national-level indices. We can only compute the household-weight index from 1900 onward due to difficulties with the 1890 census.

In Figure E2, we provide a comparison of our rental and housing price indices with and without rolling windows. The "No Rolling Window" specification is an all-in-one regression, pooling data from all years. This means we impose fixed price differentials – across areas and property sizes, for example – over time. For rental prices, there is a clear difference between the index from an all-in-one regression and those from either two- or three-year rolling windows. In particular, our preferred specification, two-year rolling windows, shows the smallest increase of the three specifications, highlighting that we have chosen the most conservative approach to generate our baseline results. The lower panel shows the equivalent three indices for sale prices although, in this segment, the differences between the three are much smaller.

Figure E3 shows the national housing price indices, both sale and rental, with and without adjustments for area. While there are start-to-end differences in both rental and sale indices, with greater estimated inflation when there are no geographic controls included, overall the pattern in both segments is strikingly similar whether these location controls are included or not. This is not surprising, given that our baseline specification relies on changes across two years, a span of time short enough that it is unlikely that there was enough differential appreciation across neighborhoods to affect the aggregate indices.

Figure E1: Household Weights versus Individual Weights

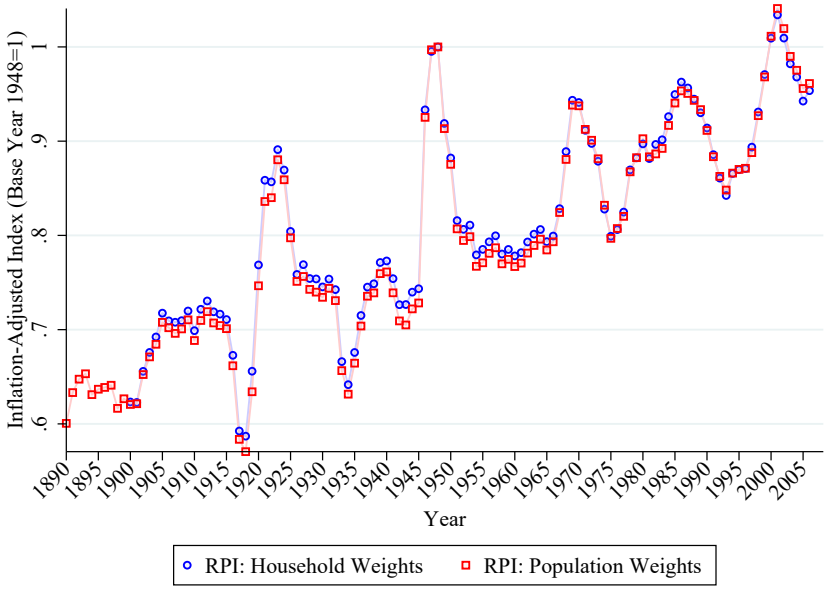

(a) Rents

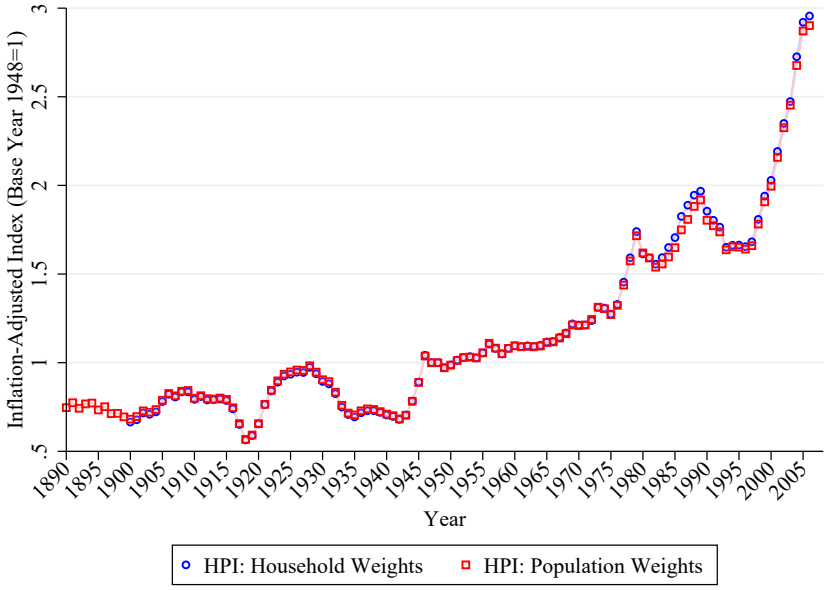

(b) Sales

Note: This figure shows the baseline real HHP rental series in panel (a) and sales series in panel (b) with and without the control with individual vs. household weights. These weights are constructed from the census and interpolated between census years. These indices are based to 1948=1 because household weights cannot be readily constructed using surviving records from the 1890 census. All series were put in real terms using the Officer-Williamson (https://measuringworth.com/) CPI.

Figure E2: HHP 30-City Real Indices by Rolling Window Size, 1890–2006

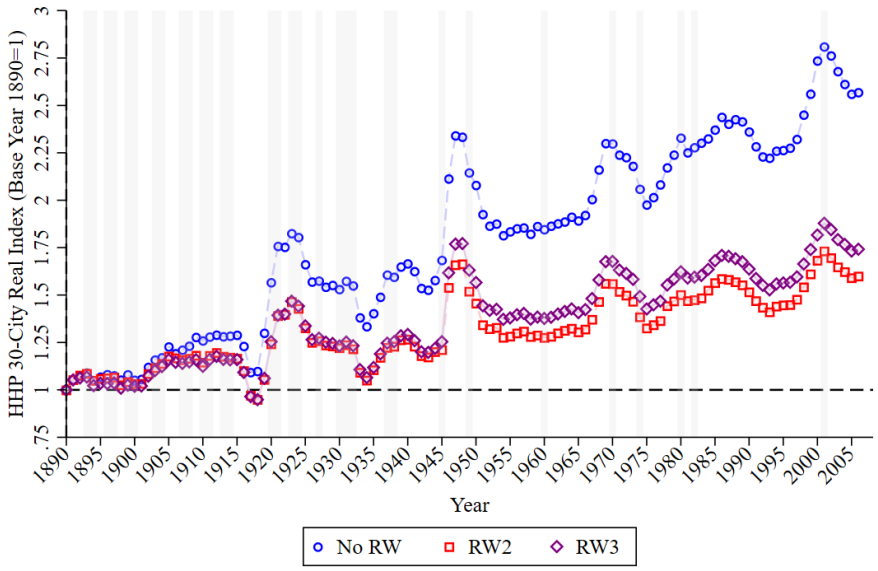

(a) Rents

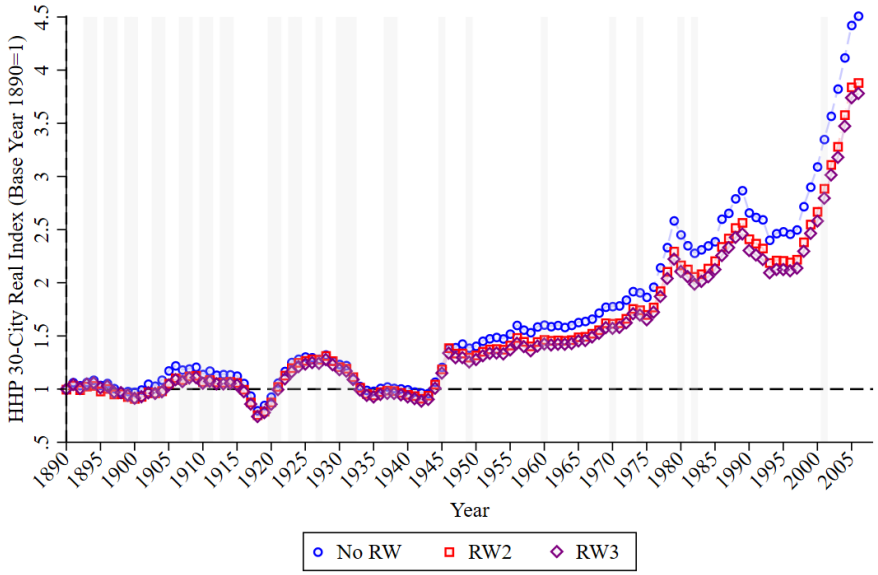

(b) Sales

Note: This figure shows the baseline real HHP rental series in panel (a) and sales series in panel (b) with various rolling window specifications. The "No RW" specification is an all-in-one regression with data pooled across all years in the sample.

Figure E3: HHP 30-City Real Indices by Geography Control, 1890–2006

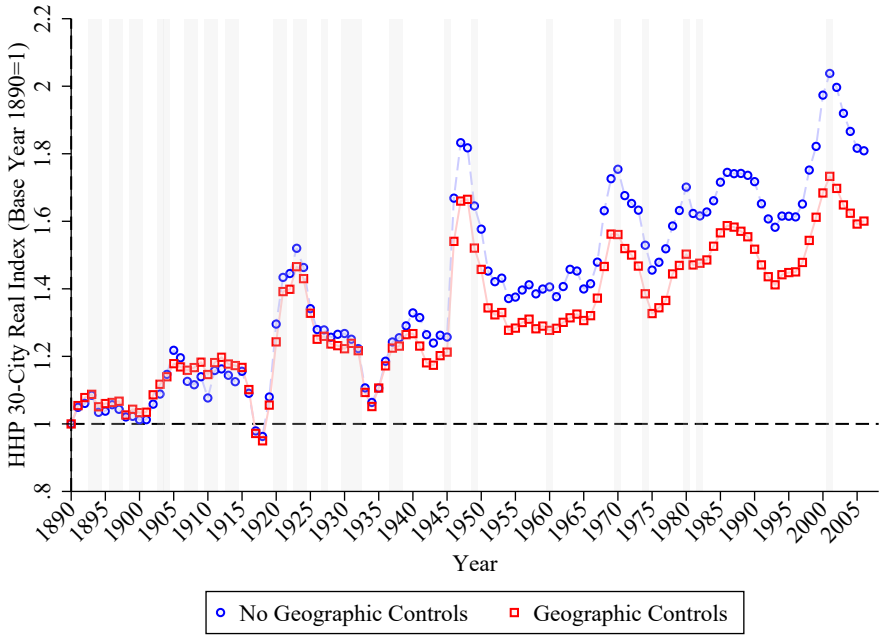

(a) Rents

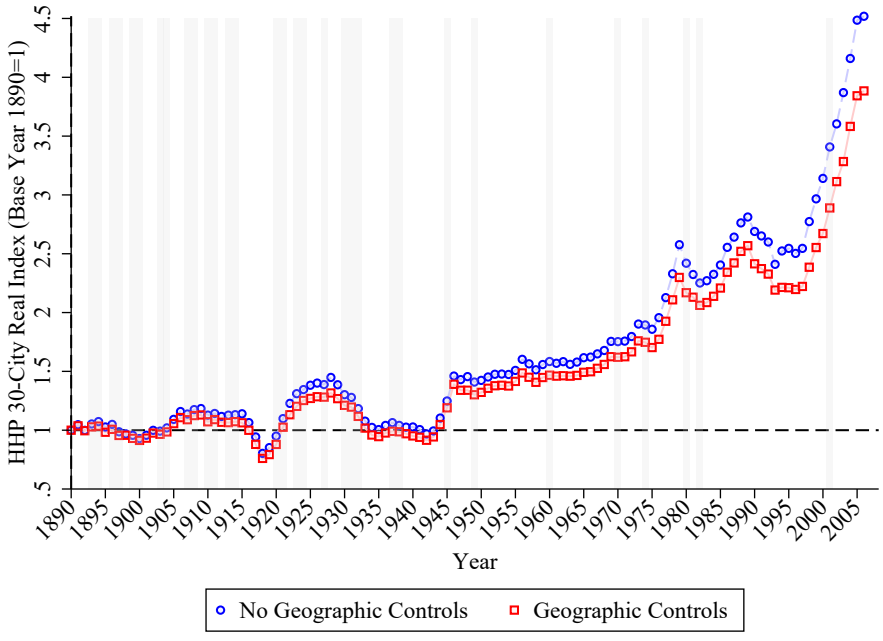

(b) Sales

Note: This figure shows the baseline real HHP rental series in panel (a) and sales series in panel (b) with and without controls for area.

F Housing Returns by City

Our final robustness check examines whether the path of returns over time is affected in any way by a combination of differential growth rates in housing prices across areas and a shifting composition of rental homes to sales homes; for example, rental homes concentrating over time, in more central locations with lower housing price growth rates. Provided that there is some coverage over time across segments with differential growth rates, the two-year rolling hedonic methodology employed will not be susceptible to such a bias. This is confirmed by calculating the rent-to-price ratio from 1920 solely using listings from neighborhoods that have a high share of rental homes throughout our period. This is shown in Figure F1: the path of returns 1920–2006 is effectively identical to a full sample (using here, due to data constraints, an all-in-one specification), with returns falling during the 1920s, rising in the 1930s and then falling again from the 1950s through the 2000s.

Figure F1: Rental Returns in High Rental Stock Neighborhoods

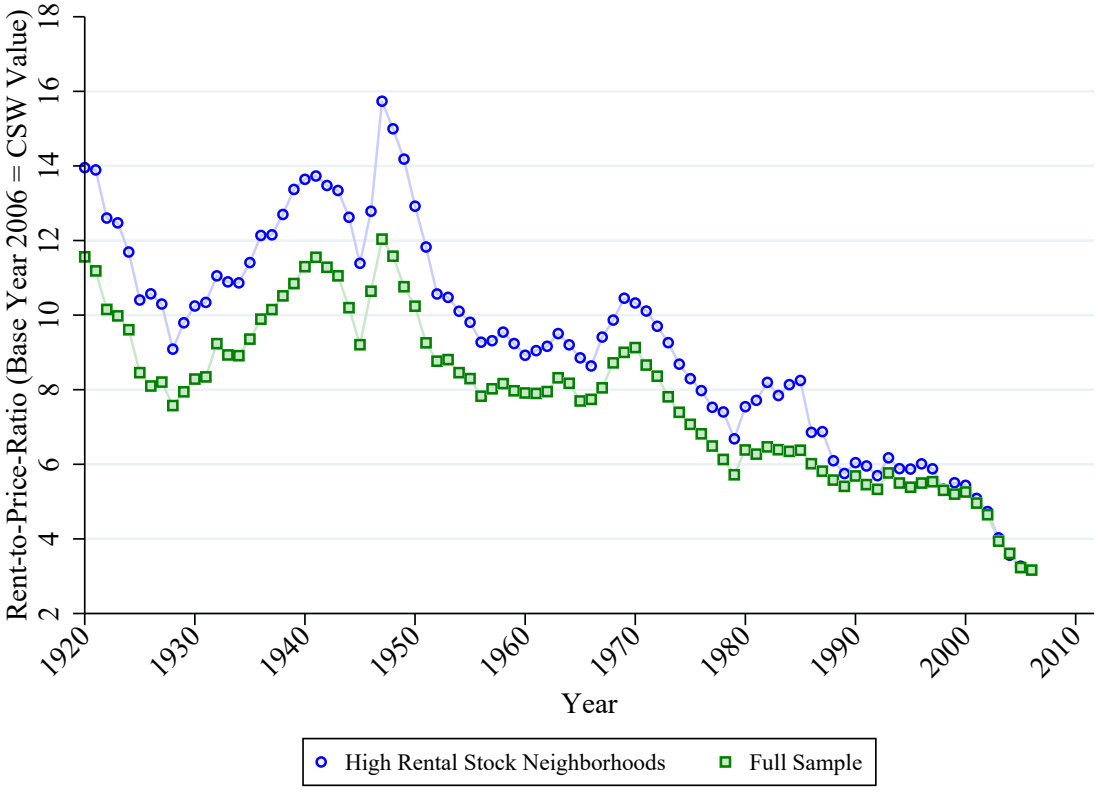

Note: The rental return shown here is computed using the top five areas for each city by rental share with a relaxed specification that removes the rolling windows and controls for area and sets the start year of 1920.

Table F1: Average Return Associated with Housing Ownership

	1890-1945			1946-1970		
	Cap. Gain	Rental Ret.	Total Ret.	Cap. Gain	Rental Ret.	Total Ret.
City						
Atlanta	-0.02	4.75	4.73	1.01	4.98	5.99
Baltimore	2.29	13.11	15.40	1.49	6.03	7.52
Boston	0.36	4.51	4.86	2.28	4.16	6.44
Charleston	0.61	13.48	14.09*	0.87	9.33	10.20
Chicago	0.95	18.19	19.14	0.71	9.34	10.05
Cincinnati	0.35	11.05	11.40	0.13	10.24	10.36
Cleveland	0.59	11.48	12.02*	0.30	9.96	10.26
Dallas	0.04	15.26	15.30	-0.15	13.91	13.76
Detroit	1.38	8.12	9.50	0.18	7.27	7.45
Houston	0.86	10.55	11.41*	0.63	8.49	9.12
Las Vegas	.	.	.	1.54	10.64	12.18
Los Angeles	1.78	8.82	10.60	2.18	7.28	9.46
Louisville	0.16	12.92	13.08	0.47	11.21	11.67
Memphis	0.64	13.83	14.47*	-0.29	11.73	11.43
Miami	2.93	16.06	18.49*	0.74	11.38	12.12
Minneapolis	0.74	8.08	8.82	1.17	6.47	7.64
Nashville	-1.00	10.93	9.93	1.14	10.36	11.49
New Orleans	1.27	8.40	9.67	0.70	7.22	7.92
New York City	0.15	14.37	14.52	2.01	11.91	13.92
Philadelphia	-0.18	12.50	12.32	1.63	11.31	12.94
Phoenix	0.65	14.89	15.54*	0.36	13.99	14.35
Pittsburgh	-0.09	6.88	6.79	0.31	7.33	7.64
Portland	1.45	16.40	17.85*	1.57	13.68	15.25
Salt Lake City	0.65	11.76	12.41*	0.57	10.98	11.55
San Diego	2.27	10.20	12.88*	0.76	7.93	8.69
San Francisco	0.05	5.89	5.95	1.53	5.90	7.42
Seattle	0.75	9.90	10.64*	0.66	7.94	8.60
St. Louis	-0.50	11.25	10.75	0.78	12.88	13.66
Tampa	3.69	11.64	13.24*	1.95	9.86	11.82
Washington D.C.	0.52	7.56	8.08	0.56	7.08	7.63
National	0.47	9.08	9.55	1.31	7.97	9.28

Table F1, Continued: Average Return Associated with Housing Ownership

	1971-2006			1890-2006		
	Cap. Gain	Rental Ret.	Total Ret.	Cap. Gain	Rental Ret.	Total Ret.
City						
Atlanta	0.65	4.49	5.15	0.41	4.72	5.13
Baltimore	2.12	3.61	5.73	2.07	8.64	10.70
Boston	3.40	2.96	6.37	1.72	3.95	5.67
Charleston	2.13	7.42	9.55	1.25	10.09	11.35
Chicago	1.98	6.83	8.81	1.22	12.76	13.98
Cincinnati	0.87	7.63	8.50	0.46	9.81	10.28
Cleveland	0.42	6.66	7.08	0.47	9.61	10.06
Dallas	1.41	12.15	13.56	0.43	14.00	14.43
Detroit	0.82	5.87	6.69	0.95	7.24	8.19
Houston	0.29	6.48	6.77	0.61	8.70	9.32
Las Vegas	1.94	8.54	10.48	1.78	9.36	11.14
Los Angeles	3.75	4.46	8.21	2.48	7.13	9.61
Louisville	0.85	8.01	8.86	0.44	11.03	11.47
Memphis	0.71	8.10	8.81	0.45	11.50	11.95
Miami	2.96	8.36	11.31	2.37	11.82	14.00
Minneapolis	1.87	4.56	6.43	1.18	6.64	7.82
Nashville	2.05	8.24	10.29	0.41	9.97	10.38
New Orleans	1.21	5.85	7.07	1.13	7.33	8.45
New York City	3.56	8.14	11.71	1.61	11.91	13.52
Philadelphia	1.88	8.35	10.23	0.85	10.95	11.80
Phoenix	2.99	9.47	12.46	1.45	12.62	14.08
Pittsburgh	0.33	6.52	6.85	0.13	6.87	6.99
Portland	3.38	9.24	12.62	2.12	13.38	15.50
Salt Lake City	2.53	8.08	10.61	1.22	10.45	11.67
San Diego	4.20	5.40	9.60	2.54	8.05	10.80
San Francisco	4.03	4.02	8.04	1.60	5.31	6.92
Seattle	4.60	5.69	10.29	2.17	7.81	9.98
St. Louis	2.33	9.11	11.44	0.66	10.93	11.59
Tampa	2.51	8.93	11.44	2.88	10.25	12.25
Washington D.C.	2.73	5.12	7.86	1.22	6.70	7.91
National	2.56	5.59	8.14	1.30	7.76	9.06

Note: This table reports the capital gain, rental return, and total return to housing at the level of the city and nationally, with the city-level series aggregated by population weight to construct the national series. All values are arithmetic means of gross returns in real terms. Cities that start after 1890 are indicated with a star on the total return over the 1890–1929 period in the third column.