

Rental Prices and the Cost of Living in the United States, 1914–2006*

Ronan C. Lyons[†]

Allison Shertzer[‡]

Rowena Gray[§]

March 2026

Abstract

The Rent of Primary Residence (RoPR) series constructed by the Bureau of Labor Statistics (BLS) implies that nominal rental prices increased by just 2.6% per year from 1914 to 2006 while overall prices grew by 3.3%. We show that this “falling real rents” puzzle can be explained by the evolving treatment of shelter in the Consumer Price Index (CPI). In this paper we construct a new, methodologically consistent shelter price series using the Historical Housing Prices (HHP) Project rental index. We also construct a revised set of shelter weights going back to 1914 and combine them with the price series to create an alternate CPI that applies the owners’ equivalent rent (OER) concept of shelter consistently across time. The HHP shelter price series increases by a factor of 28.4 (compared with the 10.7 increase in RoPR) and lifts average CPI growth from 3.3% to 3.6% per year. The revised series eliminates the long-run decline in real rents in the CPI and provides a new benchmark for assessing trends in the cost of living and real income in the U.S. over the twentieth century.

Keywords: Housing prices; rental indices; CPI; housing markets; cost of living.

*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. We appreciate the insightful comments from Robert Martin and thank David Agorastos, Grace Kneidel, Johann Ohler, and Dahyun Park for their assistance with this project. We also appreciate comments from Thesia Garner, Jessie Hanbury, Miriam Steurer, Randal Verbrugge, and participants at CRIW (July 2025). 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

1 Introduction

Housing is one of the most important elements of consumer expenditure for American households, comprising more than 40% of the core Consumer Price Index (CPI) produced by the Bureau of Labor Statistics (BLS).¹ Understanding trends in the cost of housing is thus essential for accurately measuring changes in the overall cost of living over time. However, measuring changes in the price of shelter and its share within the overall consumption basket over the long run is challenging. Existing series often rely on inconsistent methodologies and sometimes omit parts of the housing market, making their trends difficult to interpret. While the share of spending allocated towards housing has generally risen over the 20th century (Albouy et al., 2016), inflation-adjusted housing costs as measured by the BLS have fallen substantially. In nominal terms, the BLS' Rent of Primary Residence (RoPR) component of the CPI grew by just 2.6% per year from 1914 to 2006. Meanwhile, the CPI excluding shelter increased by an average of 3.6% per year over the same period.² The differential paths of the shelter and non-shelter components of the CPI present a puzzle: why did the cost of housing rise by substantially less than other prices over the 20th century?

In this paper, we construct a revised CPI for 1914–2006 that makes progress on two of the challenges in measuring shelter over the long run in the United States. The first challenge is addressing the well-documented bias of the RoPR series related to vacancy and survey nonresponse. Specifically, BLS's household-based rent collection methodology under-captured rent resets between tenants in favor of within-tenancy observations, particularly prior to the mid-1980s (Ambrose et al., 2015, 2023; Ball & Koh, 2025). For continuing tenants, nominal rents are sticky.³ When units turn over, landlords often reset

¹Core CPI excludes more volatile food and energy costs, which are included in the headline CPI measure.

²For both the overall index and for the Rent of Primary Residence, the CPI we refer to in this paper is the official BLS Consumer Price Index, covering all urban consumers (CPI-U) from 1978 (FRED series CUUR0000SA0) and all wage-earners before this. CPI ex-shelter is given from 1935 onwards by FRED Table CUUR0000SA0L2 and extended back using the estimated expenditure shares of shelter and non-shelter and trends in overall CPI and rents.

³Genesove (2003) showed that 29% of units in the late 1970s did not change rent from year to year, and that fraction was higher among continuing tenancies.

to market conditions. However, such changes were observed only if the vacant or newly let unit entered the BLS sample; nonresponse on vacant units compounded the problem (Rivers & Sommers, 1983). Vacancy and nonresponse bias worsened in 1942 with the shift to tenant mail surveys instead of the direction collection of prices from the files of agents and owners. This bias was comprehensively addressed in survey changes implemented from 1978–1985. Existing scholarship argues that the RoPR series likely understates rent inflation before these changes were made (Crone et al., 2010; Gordon & van Goethem, 2007). For legal and contractual reasons, subsequent methodological improvements were not applied retrospectively, so the RoPR has never been comprehensively revised.⁴

The second challenge is how to consistently measure housing services for homeowners, whose share in the population has grown from about 48% in 1930 to over 66% in 2000.⁵ Policymakers have struggled to disentangle asset considerations from the services element of owner-occupied housing (Eiglsperger et al., 2022). The original CPI (1914–1952) measured shelter using tenant rents only. In 1953 the BLS added a home purchase/financing (HP&FTI) component, but concerns about conceptual consistency led to its 1983 replacement with owners' equivalent rent (OER), which imputes a rent from comparable dwellings to approximate the cost of the flow of housing services consumed by homeowners.⁶ The treatment of shelter in the CPI thus varied considerably from its inception until the 1980s.

We bring new data and a novel approach to address these challenges. First, we use new estimates of rental prices from the Historical Housing Prices (HHP) market rental index.⁷ The HHP rental price series is constructed consistently from historical newspaper listings using a flexible hedonic approach that adjusts for observable housing characteristics (size, type, and location) and allows the implicit prices of these characteristics

⁴See BLS CPI Methods Overview: <https://www.bls.gov/cpi/methods-overview.htm>.

⁵U.S. Census Bureau <https://www.census.gov/data/tables/time-series/dec/coh-owner.html>.

⁶The Stigler Commission first recommended using an imputation method in 1961 to capture owner-occupied housing costs, because the relevant article is the flow of services from durable goods, rather than the cost of acquiring those goods. This is especially true for a slow-depreciating durable good like housing.

⁷The HHP indices are available at <https://www.philadelphiafed.org/surveys-and-data/regional-economic-analysis/historical-housing-prices>.

to vary over time through a rolling-window specification. By estimating the model over short two-year windows, we minimize the extent to which slowly evolving unobserved attributes contaminate measured price changes. While this approach does not eliminate all bias from unobserved factors, the assumption that such attributes are approximately stable within a two-year window is substantially more plausible than in a full-sample regression spanning many decades (Randolph, [1988](#)). Second, we standardize the treatment of shelter and use the notion of owners' equivalent rent (OER) for the entire 1914–2006 period. To impute rents for owner-occupied housing before 1983, we use estimates of the ratio of tenant rents to imputed rents for homeowners from the Bureau of Economic Analysis (BEA). We combine these new data to construct an alternative shelter index and CPI for comparison with existing BLS measures.

Using the HHP rental series, we find that nominal rents grew by about 3.7% per year rather than 2.6% per year as the RoPR suggests, with the higher growth driven by the years after both World Wars as well as the 1965–1980 period. We also find greater inflation in market rents relative to either the RoPR measure or homeowners' costs for the 1953–1983 period, when the CPI included a home purchase component. The home purchase component of the CPI was also biased downward due to the exclusion of higher-valued houses as a result of loan limits imposed by the FHA (Greenlees, [1982](#); Lyons et al., [2026](#); Vandell, [1995](#)). Overall shelter prices rose slightly more than broader price levels, contradicting the downward real trends found in the RoPR series. Combining the HHP shelter price and and consistent share series, we construct an alternate CPI for 1914–2006 that suggests significantly greater inflation in the U.S. over the long run relative to the official estimate from the BLS. Specifically, our alternate CPI shows overall prices growing from 1914 to 2006 by 3.6% per year rather than 3.3% per year, with that difference driven by the 1914–1987 period (3.7% vs. 3.4%). This difference driven by the revised price of shelter rather than the revised share. If the growth in rental prices, especially to 1987, is understated in official data, then there has been less of an increase in the standard of living in the U.S. since World War I than was previously understood.

We contribute to the growing literature identifying and addressing issues with CPI measurement and its housing components. We are not the first to suggest alternative measures, but our effort combines consistently-collected and analyzed market rent data with a more systematic approach to measuring shelter share and owner-equivalent rent to produce a usable alternative CPI going back to the start of the BLS series. Contemporary work that closely relates to ours is Ambrose et al. (2015), which uses leases for new tenants from large multi-family buildings to construct a repeat rent index (RRI) for the US from 1998 to 2010. Ambrose et al. (2023) use this RRI to construct an alternative to the official CPI, the Penn State Alternative Inflation Rate, similar to our approach here. Their alternate measure is more volatile compared to the CPI, with larger upward swings in the years up to the Great Recession.

An important conceptual consideration in our approach is that we use market rents – that is, the prices of rental units coming onto the market and advertised in newspapers – as the measure of rents for both tenant housing and the imputation of owners’ equivalent rent (OER). The literature distinguishes between market rents and average contract rents, and the appropriate choice depends on the object being measured. For tenants, the theoretically correct price concept in a CPI is the rent actually paid by the average renter – that is, contract rents. Our concern is not with that conceptual target, but with the historical implementation of contract-rent measurement in the CPI. In particular, prior to methodological changes introduced between the late 1970s and mid-1980s, the BLS survey design tended to under-capture vacancy resets and new tenancies, and nonresponse among vacant units further artificially limited the pass-through of market conditions into the measured index (Crone et al., 2010; Genesove, 2003). As a result, measured rental inflation in certain periods likely understated true changes in contract rents, especially during episodes of rapid adjustment.

The conceptual case for using market rents is stronger for owner-occupied housing. The rental equivalence approach aims to capture the opportunity cost homeowners incur by occupying their homes rather than renting them out. As the BLS noted when intro-

ducing OER in 1983, “although homeowners’ rent is implicit, it is not hypothetical; it is a cost all homeowners actually incur” (U.S. Bureau of Labor Statistics, 1983, p.8). Because this opportunity cost is defined at current market prices, market rents provide a natural benchmark for OER.⁸ Indeed, from a cost-of-living perspective, the relevant question is the cost, *at prevailing market prices*, of maintaining a given standard of living (Poole et al., 2005). However, a CPI should reflect the purchasing power of a typical household, which will not transact at market rents every period (Adams et al., 2024). Nevertheless, the user-cost and opportunity-cost considerations embedded in owner-occupied housing are more closely aligned with market rents than with average contract rents (Ozimek, 2013).

In this paper, we use market rents because there are no consistent measures of average tenant rents extending back to 1914, other than the BLS RoPR series itself. While it would be ideal to observe both the rents faced by new tenants and those paid by the average tenant, over long horizons the two series should move together, differing primarily in timing rather than trend. Our objective in this paper is not to measure the inflation experienced by a representative sitting tenant in each period, but rather to recover the long-run trajectory of housing service costs relevant for intertemporal cost-of-living comparisons. From this perspective, market rents provide a coherent and empirically tractable basis for measuring both tenant rents and the imputed rental value of owner-occupied housing over the twentieth century.

The rest of the paper is structured as follows. Section 2 describes the construction of our HHP index of market rents and presents its main trends for the period 1914–2006, while section 3 outlines our new estimates of the shelter share for the same period. In Section 4, we combine these new series together with the existing ex-shelter CPI to generate an alternate measure of CPI extending back to 1914. The final section concludes with suggestions for avenues for future research.

⁸See Triplett (2001) for an excellent summary of the advantages of the OER methodology and discussion of the treatment of housing in general in the CPI. More recently, there has been discussion on the limits of OER methods, including that they do not allow the CPI to reflect the changing credit conditions under which housing is acquired (Chodorow-Reich et al., 2025).

2 New Estimates of Rental Prices

In this section we present new annual estimates of the nominal price of market rents in each of 30 U.S. cities over the period 1914–2006. We constructed rental price indices from historical newspaper real estate listings as part of the Historical Housing Prices (HHP) Project.⁹ Here, we provide background on the data used to construct the nominal rental price indices used in this paper; for additional information on the historical newspaper data used in the HHP dataset, interested readers should consult Lyons et al. (2026). For this paper, we draw on just over one million rental listings from the 1914–2006 period, covering 30 large and mid-size cities that represent a range of regions and economic trajectories and that also had a complete and accessible newspaper repository. Our sample represents 36% of the U.S. urban population in 1920. This compares to the original 39 cities in the BLS cost of living sample, which represented about 42% of the U.S. urban population, and to the 51-city expanded BLS sample which was in use from 1920, representing about 45% of the urban population.

The BLS sample contains more cities that remained small throughout the twentieth century (such as Little Rock) or that shrank as older New England manufacturing towns were replaced in the economic order by larger cities across the country (Fall River). Our sample provides greater coverage outside of New England, with some cities in the Southwest absent from the BLS sample (Las Vegas, Phoenix) and more coverage in the New South (Tampa) compared to the BLS use of Birmingham and Richmond. But, 24 cities are common to the two indexes and in both samples approximately two thirds of urban residents are renters (66% in the expanded BLS, 68% in our sample), compared to 61% in the urban U.S. overall. We thus conclude that our sample offers comparable coverage of the U.S. population over the long run.

In terms of our data work, we follow the pioneering efforts of Rees & Jacobs (1961), who used newspapers to construct rental price indices for several U.S. cities for the pre-

⁹The HHP indices are available at <https://www.philadelphiafed.org/surveys-and-data/regional-economic-analysis/historical-housing-prices>.

1914 period. The Rees-Jacobs index presented unadjusted averages. We improve upon their approach by controlling for the housing unit mix in our rental indices so that we are measuring like-for-like changes as closely as possible. The intellectual groundwork for quality adjustment in price indices was laid by Zvi Griliches and his coauthors in the 1960s (for example, see Adelman & Griliches, 1961). The application of hedonic methods they recommended has proven challenging for housing because property characteristics must be observable throughout the period. We use a flexible hedonic approach to control for changing housing characteristics over time.

2.1 HHP Rental Dataset

We drew a large sample of newspaper rental listings to construct our rental index. The procedure targeted 150 valid rental listings from each sampled newspaper, typically the last Sunday of the month of interest. Valid listings contain (1) a price, (2) a measure of size, (3) a property type (house or apartment), and (4) an indication of location. 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. Most major cities had sufficiently dense secondary housing markets by 1914 to provide adequate listings for constructing the index. However, some smaller cities did not have a substantial volume of listings until later; for instance our Las Vegas rental series begins in 1948.

A natural concern with newspaper listings is that they may represent a positively selected subset of the housing market, potentially skewed toward larger or higher-quality units. In earlier work using the same dataset, we benchmarked the distribution of observed housing characteristics against census data across cities and decades (Lyons et al., 2026, Appendices B–D). While advertised units are modestly larger on average than the full housing stock, the distribution of bedrooms and dwelling types tracks census patterns closely, and we find no evidence of systematic drift toward higher-quality segments over time or that long-run price or yield dynamics are driven by shifts in the composition of listed units. Importantly, any level differences arising from positive selection do not me-

chanically imply bias in estimated inflation rates, provided that the composition of listed units does not trend systematically toward higher-quality segments over time. We do not find support for this hypothesis in the benchmarking exercises. The summary statistics for the sample can be found in Appendix Table A.

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 storeys. To develop controls for neighborhood, we used a simple machine learning classification algorithm to geolocate each observation in one of 20 standardized neighborhoods for each city.¹⁰ 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. We standardize rental prices so that they are expressed in monthly terms, where necessary, and include the rental frequency as an additional regressor.¹¹ 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.

2.2 Rental Price Index Construction

Newspaper real estate sections do not allow us to use “repeat sales” approaches that have become the standard for rental indices in more recent years (Ambrose et al., 2015; Greenwald & Guren, 2021). Given the structure of our data, we use a hedonic model that controls for observed housing attributes recorded in listings. To allow the valuation of these char-

¹⁰Full details can be found in Appendix C of Lyons et al. (2026).

¹¹Rental 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 95th 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.

acteristics to evolve over time, we estimate the model using short rolling windows, which has become the standard outside the U.S. recently (Hill et al., 2022; Silver, 2016). In our baseline specification, each regression spans two adjacent years, so that year-to-year price changes are identified within a narrow time frame. This design reduces the risk that gradual changes in unobserved housing attributes – such as modernization, aging, or improvements in amenities – are mechanically attributed to price growth. Although unobserved characteristics may still vary within a given window, the assumption that such attributes are approximately stable over two years is considerably more plausible than assuming stability over the full 1914–2006 sample. As a result, the rolling-window approach mitigates bias arising from unobserved quality drift.

The hedonic pricing model is given by the following regression equation. For rolling window of size s , with base year b , we have:

$$\ln(\text{price})_{ict} = \alpha_{bc} + \sum_{y=b+1}^{b+s-1} \beta_{cy} \cdot 1_{(y=t)} + \mathbf{X}_{ict} \mathbf{\Gamma}_{bc} + \varepsilon_{ict} \quad (1)$$

where the regression is estimated separately for each city c using only observations i from years $t \in \{b, b + s - 1\}$, β_{cy} is the coefficient capturing price change from year y to $y + 1$ for city c , $1_{(y=t)}$ is an indicator function equal to 1 when observation is from year t , and $\mathbf{X}_{ict}$ is a vector of property characteristics. In the case of a rolling window length of two years, there is only one β term in the above expression, which captures the change in prices between the base year b and the next year. As an example, for our baseline rolling window of two years, we run the above regression on the years 1914–1915 to obtain the β coefficient for 1915, 1915–1916 to obtain the coefficient for 1916, and so on. Controls include location within the city, size (rooms, bedrooms, bathrooms, stories), dwelling type (house or apartment), and rental payment frequency. Our identification relies on the assumption that the observed characteristics and neighborhood controls proxy sufficiently for quality differences within each window.

Despite our efforts to meet our listing sampling target, there were inevitably some

cases of observation counts too low for the two-year rolling window to generate reliable results. To address cases of insufficient listings in some cities, particularly during the rent control policy around World War II, we use longer window lengths to obtain more stable year coefficients where necessary.¹² Starting in each city’s earliest year, we take the beta coefficient capturing the increase in the price level from the base year to the next year from the relevant regression to build a city-specific price index over time. For a city c , the value of the index in year t is defined as:

$$\iota_{ct} = \prod_{y=1915}^t \exp(\beta_{cy}). \quad (2)$$

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, using contemporaneous metropolitan population shares as aggregation weights. Population weights are constructed using the metropolitan area population from the U.S. Census Bureau and interpolated between census years.¹³ Lyons et al. (2026) contains a wealth of additional details on this specification, including robustness checks and city-level benchmarking exercises. We use the baseline specification from this paper with two-year rolling windows and area controls to proceed with an analysis of the nominal rental index in the remainder of this paper.

2.3 The HHP Rental Index versus the Rent of Primary Residence Index

The BLS rental sample and survey procedures have evolved substantially over time, including shifts in geographic coverage, sampling frames, and vacancy handling.¹⁴ Of particular relevance to our paper are the transition from agent-based reports to tenant mail

¹²In 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. 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.

¹³We match our city to the corresponding metro area as defined by the census and allow MSAs to expand over time with population growth.

¹⁴For more, see the BLS CPI Methods Overview: <https://www.bls.gov/cpi/methods-overview.htm>.


surveys in 1942 and methodological changes between 1978 and 1985 aimed at reducing vacancy nonresponse. Our goal is to provide an independent market-based measure of rental inflation against which the published RoPR series can be evaluated.

Previous work has argued that there is substantial understatement of rental inflation in the RoPR arising from vacancy nonresponse and within-tenancy rent rigidity (e.g. Crone et al., 2010; Gordon & van Goethem, 2007). To explore this issue further, we compare the BLS RoPR series and the HHP series. A graphical comparison, using a log-scale, is given in Figure 1a, where both series are indexed to 1914=1. Cumulatively over the period 1914–2006, the HHP index of market rents rose by a factor of 28.4 compared to just 10.7 for the RoPR series. All of this difference dates from before 1987: the HHP index rose by a factor of 15.8 from 1914 to 1987 compared to 5.9 in RoPR; both indices rose by a factor of 1.8 from 1987 to 2006. For context, the CPI excluding shelter index rose by a factor of 24.7 over the full period. In annual growth terms, our rent index is over one percentage point greater than the existing RoPR index: 3.7% per year versus 2.6%. The year-over-year change version of this figure can be found in Appendix Figure A1.


Using the HHP measure implies that rental inflation was slightly above general inflation, rather than well below it, as the annual average increase in CPI excluding-shelter over the same period was 3.6% per year. Reassuringly, the overall change in nominal rents seen in this series for 1940–2000 aligns with the estimate given by Crone et al. (2010). Further, as shown in Figure 1b, over the period 1985–2006 when the BLS implemented methodological changes, cumulative growth in the two series is very similar, with only modest differences in levels by the end of the sample. The principal distinction lies in short-run dynamics: the HHP index exhibits more pronounced cyclical movements whereas the RoPR series evolves more smoothly. This pattern is consistent with the CPI’s sampling design, which incorporates new tenancies gradually and thereby dampens short-run volatility.

With these validation exercises in place, the annual nature of both HHP and RoPR series allows investigation of the timing of potential mismeasurement of rental inflation. Three periods contribute most to the difference between the two series: 1914–1920, the

Figure 1: HHP Rental Index vs. BLS RoPR Series



(a) 1914-2006



(b) 1985-2006

Note: This figure plots the baseline HHP rental price series against the BLS RoPR series with base year of 1914 (panel a) and 1985 (panel b).

1940s, and 1965–1980; these differences are also summarized in Table 1 below. The first two periods are associated with both World Wars and significant general inflation at a time of expanded rent controls. One would expect an *a priori* difference between a series capturing open-market rents (as per HHP) and one capturing average contract rents (as per RoPR), especially during a period that combines rent controls and high inflation. However, these differences should disappear as controls are lifted and as tenants move to new market-rate units. Generally, the trajectory of market rents paid by new tenants and the average rents paid by all tenants should converge over time.

Evidence from the BLS series on new-tenant and all-tenant rents supports this notion in more recent years, namely that the difference between market and average contract rents is largely one of timing rather than long-run trend. Since 2005, the research New Tenant Rent (R-CPI-NTR) index has tended to lead the All Tenant Rent (R-CPI-ATR) index by approximately two to four quarters at turning points, most notably during the Great Recession and the post-pandemic rent surge. We plot the R-CPI-NTR and R-CPI-ATR with a three-quarter lead in Figure A2 to visualize how closely these series move together. From 2005 to 2025, cumulative growth differs minimally, with an average annual growth rate (AGR) of 3.0% for NTR and 2.9% for ATR. This pattern is consistent with contract rents incorporating market conditions with a lag rather than diverging permanently from them.

For the historical period we study in this paper, the divergence between the HHP and RoPR indices in the 1914–1920 period saw an increase in the RoPR series (35%) that was less than one third of the increase seen in the HHP series (119%). This is indeed partly offset by the RoPR increasing over 1920-1925 while market rents in HHP fell slightly. Given the nature of BLS measurement at that time, where the underlying data were collected from surveys of agents, this finding is indeed consistent with tenant rents reflecting market rents with a lag.

A similar pattern occurs after World War II, although substantial differences remain between the series. According to our HHP measure, market rents rose by more than one third in 1946 alone (38%) and this increase was followed by a further 23% increase in 1947.

Table 1: 5-Year Differences in the CPI and RPI by Shelter Series

	CPI (RoPR)	CPI (HHP)	RPI (RoPR)	RPI (HHP)
1915 - 1920	14.67	16.82	5.36	16.99
1920 - 1925	-2.63	-3.58	4.78	-1.35
1925 - 1930	-0.98	-1.16	-2.05	-2.59
1930 - 1935	-3.85	-3.72	-7.26	-5.74
1935 - 1940	0.41	1.17	2.06	3.20
1940 - 1945	5.13	5.35	0.91	4.18
1945 - 1950	5.99	7.36	3.67	9.97
1950 - 1955	2.16	1.21	3.69	-0.42
1955 - 1960	2.01	2.04	1.68	1.93
1960 - 1965	1.27	1.56	1.11	1.74
1965 - 1970	4.26	5.61	2.60	8.02
1970 - 1975	6.75	6.21	4.52	3.34
1975 - 1980	8.90	10.03	6.88	11.62
1980 - 1985	5.47	5.56	6.68	6.36
1985 - 1990	3.97	3.82	4.36	3.31
1990 - 1995	3.12	3.03	2.66	2.16
1995 - 2000	2.48	3.56	3.11	5.61
2000 - 2005	2.55	2.02	3.39	1.40

Note: This table summarizes the five-year average differences between the rental price and CPI series from the BLS and HHP.

By 1948, markets rents were 84% higher relative to 1945 while the RoPR increased by just 11% over these three years. The RoPR measure of nominal rents continued to rise in subsequent years, while the HHP measure stagnated as the postwar spike, driven by rent controls and limited construction, eased. Nonetheless, a decade after the introduction of rent controls in 1942, the HHP measure of rents had risen by 82%, compared to an increase of 30% in the RoPR.

The final period during which the two series differ, 1965–1980, is also one of high inflation in the wider economy. For the fifteen year period in total, our estimate is that market rents tripled while the RoPR doubled. While measured inflation across the two series was similar for 1969–1977 at roughly 45%, there are substantial differences in estimated rental inflation in the mid 1960s and the late 1970s. Between 1965 and 1969, the HHP measure increased by almost 40%, compared to less than 10% in the BLS measure. Similarly, between 1977 and 1980, the HHP measure saw twice the increase (50% vs 25%) of the BLS measure.

The timing of the divergence is consistent with the existing research that studies nominal rigidities in rents within tenancies and documents that the path of vacant and occupied rental dwelling prices is different. Such “vacancy effects” resulted in a major change in CPI methodology between 1978 and 1985. Specifically, the number of units surveyed was reduced, with effort redistributed to reducing the non-responses, including surveying the landlord or owner. An analysis by Rivers & Sommers (1983) of 18 months of responses from both existing and new tenants from October 1979 to March 1981 found that new tenants faced significantly higher inflation (18.6% annualized compared to 8.1%). It is most certainly the case that a share of true inflation was unmeasured when vacant units were omitted prior to these methodological improvements in the late 1970s and early 1980s.

One concern with this approach is that newspaper listings may reflect improvements in the quality of the housing stock over time, potentially inducing an upward bias in estimated rental inflation if higher-quality units become increasingly represented in advertised listings. While our rolling-window hedonic methodology controls for observable characteristics and allows their implicit prices to evolve, unobserved quality change re-

mains a general challenge for long-run price measurement. To explore how relevant this concern for our approach, we benchmark the HHP rental indices against the CBRE/Tortowheaton “same-store” rental indices, which are available starting in 1994 and is constructed from new and existing leases in large, professionally managed multifamily buildings and tracks the same properties over time. Because the same-store methodology holds property identity fixed, it provides a useful check against composition-driven bias for at least the last decade of our sample. Appendix Figure A3 compares the HHP and CBRE same-store indices for overlapping cities and years. The two series exhibit similar long-run trends for most cities, although the HHP index displays somewhat greater short-run volatility, consistent with its broader market coverage and reliance on new leases. The absence of systematic divergence in cumulative growth provides reassurance that the HHP series is not driven by progressive quality upgrading in newspaper listings.

The post-1985 convergence also provides indirect validation of the HHP methodology: when the CPI’s rent measurement improved, the two independently constructed indices evolve similarly in the long run (see Figure 1b). Together with the CBRE same-store benchmark, this evidence suggests that the long-run growth in the HHP index reflects underlying market rent dynamics rather than progressive changes in the composition of advertised units. Having provided support to the notion that our series reflects the true long run path of rents, and thus that rents increased in line with prices in the wider economy, we now turn to the measurement of housing weights in the CPI.

3 Constructing the Shelter Share Series

Since its inception the BLS has conducted surveys to estimate the relative importance of each consumption category in the the CPI.¹⁵ In this section, we briefly retrace the history of

¹⁵Some of the earliest surveys of household consumption were conducted in 1888 by the Bureau of Labor, and again in 1901, when the Bureau of Labor was temporarily an independent agency. In 1913, the renamed Bureau of Labor Statistics was transferred to the newly created Department of Labor, where it remains today. The BLS surveyed the income and consumption of urban households in the benchmark years of 1919, 1936, and 1941. From 1952 onward, consumption surveys were conducted at a nearly annual frequency, with spending shares taken from the preceding December. Exceptions include 1962 and 1963. For 1962, spending shares are

measuring owner-occupied housing in the CPI, noting the relevant shifts in the treatment of shelter, and outline our approach to standardizing the treatment of this category going back in time.

3.1 The Treatment of Owner-Occupied Housing in the CPI

The earliest household expenditure surveys conducted by the BLS included only urban renters. Spending relating to owner-occupied housing was excluded from the CPI until 1953, when a home purchase component was added. The weight for Rent of Primary Residence (RoPR) fell from 11.1% in 1952 to 5.5% the following year, while the new category of “Home Purchase and Upkeep” (of which the bulk comprised “Home Purchase & Financing, Taxes, and Insurance,” HP&FTI) accounted for 11.6%. The net effect of the methodological change was a jump in the shelter share, from 11.1% in 1952 to 17.5% a year later.¹⁶ This sizable jump in shelter share was thus driven by methodological changes rather than underlying shifts in household spending. Between 1953 and 1983, owner-occupied housing was included in the CPI through expenditures relating to home purchases, financing, taxation and insurance, and maintenance and repairs. During this period, one other methodological change significantly shifted the shelter share. The 1978 transition in the default CPI from CPI-W (wage earners) to CPI-U (all urban households) was accompanied by another sharp increase in shelter’s share. HP&FTI went from a share of 12.6% in 1977 to 19.2% a year later, with the overall shelter share rising from 21.3% to 29.2%.

The acquisitions approach was ultimately abandoned in favor of a “flow-of-services” approach because many of the elements of home purchase are not expenditures related to consumption (Rippy, 2014). The mortgage principal, for example, is strictly speaking investment, not consumption. Mortgage interest is considered a financing cost, and property taxes are a government transfer. The intellectual foundation of the shift away from

only available for top headings, while no information exists for December 1962 with which to construct 1963 spending shares.


¹⁶The full shelter category in 1953 included: 5.5% for tenant rents, 11.6% for home purchase/upkeep, and 0.4% for “Other Shelter.”

an acquisitions approach go back at least to the Stigler Commission, which had argued in 1961 that the ideal object to capture homeowners' housing consumption was the imputed or owners' equivalent rent (OER). The OER approximates the cost of the flow of housing services consumed by homeowners by estimating a hypothetical rent for owner-occupied homes.

OER was formally introduced into the CPI in January 1983, replacing the HP&FTI treatment of owner-occupied housing. As mortgage principal and financing components were removed and replaced with an imputed rent concept, the measured shelter share declined sharply between 1982 and 1983. The 1987 comprehensive CPI revision, based on updated Consumer Expenditure Survey (CE) weights and a survey redesign, introduced further changes. As in all periods, OER weights were derived from CE data; the increase in OER's relative importance in 1987 reflected updated expenditure weights, rising homeownership, improvements in the housing stock captured in the CE sample (Mason & Butler, 1987), and reclassification of certain components, including vacation homes (Marcoot & Bahr, 1986). The 1987 revision also altered the method used to measure OER price changes, namely using a sample of owner units linked to rents of comparable rental units rather than relying solely on reweighted renter samples (Marcoot, 1985). This change affected measured OER inflation but not its expenditure weight. The BLS returned to a renter-based framework for OER price estimation in 1998. Between 1986 and 1987, OER's relative importance rose from 14.4% to 19.7%, and total shelter from 22.3% to 27.8%. These shifts reflect CE-based reweighting and classification revisions rather than sharp changes in household behavior.

We plot the shelter share in the official CPI in Figure 2. Notwithstanding the shifts in shelter's share due to methodological changes, the overall trend for the period 1953–2006 is clear. Between 1953 and 1977 (a period of methodological consistency), shelter's share rose from 17.5% to 21.3%, while between 1987 and 2006 (under a different consistent methodology), shelter's share rose from 27.8% to 32.8%. Over the earlier of these two periods then, shelter's share rose by 0.16 percentage points per year on average, while in

Figure 2: Official and HHP Shelter Share



Note: This figure plots the revised shelter share constructed as described in Section 3 against the official shelter share used by the BLS to construct the CPI.

the latter period it rose by 0.26 percentage points per year. We note that the BLS “Shelter” category also includes lodging away from home and household insurance. However, these are, quantitatively, significantly less important than tenant rent or owner-equivalent rent: their combined weight in the 2018 Handbook was 1.2%, while rent (both tenant and owner) was roughly 26 times the share (31.5%). Thus, we assume they track rent.

3.2 Creating a Consistent Shelter Share

We follow the CPI’s lagged-weighting convention: for year t , expenditure shares are those from December $t - 1$, consistent with BLS practice. As in the official CPI, annual weights between major reference revisions reflect underlying reference-period expenditure shares updated using observed price changes. When official annual weights are available, we

adopt them directly rather than reconstructing intermediate updating paths using HHP price movements. Differences between the official and revised CPI therefore arise from differences in shelter prices and the construction of shelter shares, not from changes in the aggregation formula itself. Because our results are driven primarily by differences in shelter price growth rather than relative-weight dynamics, adopting a fully replicated CPI reference-period updating procedure would have only minor quantitative effects on the long-run inflation differential.

We take the following approach for constructing a consistent shelter share, starting from the most recent period and working backwards:

1987–2006 For 1987 onward, we take the shelter share directly from the BLS. This period reflects the modern OER methodology and consistent CPI treatment of owner-occupied housing.¹⁷ The year 1987 therefore serves as a level anchor for the extrapolation to earlier years we discuss below.

1929–1986 For 1929–1986, owner-occupied housing was either excluded from the CPI (pre-1953) or included using the acquisitions approach (1953–1983). To construct a consistent shelter share for this period, we turn to the Bureau of Economic Analysis (BEA) national accounts, which provide annual series beginning in 1929 for “Housing: Rental of tenant-occupied nonfarm housing” and “Housing: Imputed rental of owner-occupied nonfarm housing.” These components allow us to calculate the shelter share within total personal consumption expenditures (PCE) annually from 1929 onward. Our identifying assumption is that year-to-year movements in shelter’s share of total PCE provide a reasonable proxy for movements in shelter’s share within the CPI basket. The levels of shelter shares differ between PCE and CPI due to conceptual differences in coverage (e.g. third-party healthcare payments included in PCE but not CPI), but proportional changes in shelter’s importance over time are informative.

¹⁷Relative importance (expenditure weight) data were obtained from the Bureau of Labor Statistics, “Relative Importance and Weight Information for the Consumer Price Indexes,” available at <https://www.bls.gov/cpi/tables/relative-importance/home.htm>.

To anchor levels, we exploit the fact that 1987 is the first year in which both BEA and CPI reflect conceptually comparable OER-based shelter measures. In 1987, the ratio of OER to tenant rent in the CPI (2.5) is very similar to the BEA ratio of imputed rent to tenant rent (2.6), despite differences in overall scale. We therefore use 1987 as the benchmark level and allow earlier shelter shares to vary proportionally with the BEA shelter share. Formally, for each year $t \in [1929, 1986]$:

1. We take BEA shelter share in year t ;
2. We compute ratio of this share to the BEA shelter share in 1987;
3. We then multiply the 1987 CPI shelter share (27.8%) by this ratio.

For example, in 1950 the BEA shelter share was 10.5% compared to 15.4% in 1987 (BEA basis), implying a ratio of 0.683. Applying this ratio to the 1987 CPI shelter share yields a revised CPI-consistent shelter share of approximately 19% in 1950.

1914–1928 For 1914–1928, neither BEA nor CPI provides annual information on homeowners’ imputed rents. We therefore use information from the BLS household expenditure surveys for 1914–16 and 1934–36, which report rent-only shares (13.6% and 17.5%, respectively), specifically, in Table VIII of the BLS Report 0699, “Changes in the Cost of Living in Large Cities in the United States, 1913–41.” We then backcast in three steps:

1. Using the movement in total shelter share (rent + OER) between 1929 and 1934–36 in the BEA-based series, we infer a rent-only share in 1929 consistent with the BLS survey level in 1934–36.
2. We then compute the ratio of rent-only shares in 1914–16 relative to 1929 and apply this ratio to the 1929 rent+OER share (25.2%) to infer a total shelter share for 1914–16 (17.8%).
3. Finally, we linearly interpolate between 1916 and 1929 to obtain annual shelter shares for 1917–1928.

The total shelter share for 1914-1916 is estimated to be 17.8%, while the share for 1929 is estimated at 25.2%. This 7.4 percentage point change over 13 years translates into a linear increase of 0.57 points per year. Thus to compute the total shelter share for 1920, we interpolate between 1916 and 1929 to obtain a value of 20.1%. Aside from this linear interpolation prior to 1929, all variation in the revised shelter share from 1929 onward is driven by observed movements in the BEA national accounts.

Figure 2 plots the resulting HHP shelter share. By construction, the revised and official series coincide from 1987 onward. Between 1929 and 1986, the revised series follows movements in the BEA national accounts, anchored to the 1987 CPI level. Prior to 1929, the series is inferred from expenditure surveys and interpolated smoothly to connect with the BEA-based estimates.

4 Nominal Rental Prices and the CPI

In this section, we explore the implications of our revised rental price and shelter share series on the measurement of changes in consumer prices over the 20th century.

For our point of comparison, we continue to use the official CPI (Wage-earners until 1978 and All Urban Consumers thereafter), both headline and, where needed, CPI excluding (ex-) shelter.¹⁸ We construct the index using the same lagged annual expenditure shares as the official CPI. Specifically, for year t , we use the December $t - 1$ Consumer Expenditure Survey (CES) weights for shelter and non-shelter components, consistent with BLS practice. Thus, the aggregation mirrors the CPI's Laspeyres-type structure with lagged expenditure weights, allowing us to isolate the impact of alternative shelter prices and revised shelter shares.

We apply our market-rent measure to both principal components of shelter – tenant


¹⁸These series can be accessed through the Federal Reserve Bank of St. Louis FRED website, for instance series CUUR0000SEHA (Consumer Price Index for All Urban Consumers: Rent of Primary Residence in U.S. City Average) for the RoPR and the BLS Consumer Price Index for All Urban Consumers (CPI-U), series [CUUR0000SA0] for the CPI. Further, where needed as described below, we use the BLS CPI excluding shelter, obtained from the Federal Reserve Bank of St. Louis FRED website; series CUUR0000SA0L2 (Consumer Price Index for All Urban Consumers: All Items Less Shelter in U.S. City Average).

rents and owners' equivalent rent (OER) – as a common measure of the market price of housing services. We believe that this approach can address two well-documented limitations of the historical CPI: the understatement of rental inflation arising from vacancy nonresponse and within-tenancy rent rigidity, and the absence of an appropriate price basis for imputing housing services to owner-occupiers prior to the modern OER methodology. While market rents may at times move more quickly than the rents paid by sitting tenants, this potential bias is small relative to the gains from capturing vacancy resets and from using a conceptually appropriate opportunity cost measure for homeowners. Moreover, especially in later decades, OER accounts for the majority of shelter expenditure (2.7 times the tenant component for 1953–2006), making a market-based rent series useful for measuring aggregate shelter costs over the long run. We use HHP rental data through 2006 for consistency but note that the results do not differ substantially if we used BLS shelter prices from 1987 onward given the similarity of the trends across sources for this later period.¹⁹

There are limitations associated with using tenant rents for an OER as we do here, in part because renters and homeowners occupy different segments of the housing stock. In particular, owner-occupied units are more likely to be single-family structures and located in different neighborhoods than rental units, so that shifts in relative prices across structure types or locations can generate differences between rent and OER indices even absent measurement error. We believe our approach is nonetheless valid for two reasons. First, the HHP rental index is constructed from a broad set of market listings that include both houses and apartments, and our rolling-window hedonic specification controls flexibly for observable characteristics including dwelling type, size, and location. By allowing the implicit prices of these characteristics to evolve over time, the index measures like-for-like changes in market rents rather than shifts in the composition of advertised units. Second, in robustness exercises reported in earlier work using the same dataset, we found no evi-

¹⁹The BLS continued to make modifications to its rent survey methodology, including in 1988 and 1994 (see Crone et al. (2010)), so we used the HHP rental series throughout because it is consistently constructed over time. None of the results we report in this paper change significantly if we use the BLS OER and RoPR series for the period from 1987 onward.

Figure 3: Official CPI vs. HHP CPI




Note: This figure shows the official Consumer Price Index and two alternate CPIs: one that replaces only the shelter price component in each year with rental series from Section 2; and the second where both shelter prices and the revised shelter share (as described in Section 3) are replaced. For the official CPI, we use Wage-earners until 1978 and All Urban Consumers thereafter, both headline.

dence that the long-run dynamics of rental yields or price growth were driven by changes in the relative representation of smaller apartments versus larger houses. While we do not construct separate rent indices by tenure or structure type due to sample limitations, the design of the index mitigates concerns that differences between tenant rent and OER inflation in the official CPI reflect composition effects that would mechanically carry over to our series.

The resulting CPIs are plotted in Figure 3, where 1914 takes a value of 1. The differences between the HHP and BLS CPIs are concentrated over the period from 1914 to 1987 and so we separately report these magnitudes below. The official CPI increased by a factor of 20.1 between 1914 and 2006 (11.3 for 1914–1987), which is a weighted average of the 24.7-fold increase in ex-shelter prices and the smaller 10.7-fold increase in RoPR. To create an alternate CPI using the HHP series we have developed in this paper, we keep

the ex-shelter CPI the same throughout this exercise. To understand the respective roles of the HHP rental prices and shelter series, we first create a CPI using the HHP rental prices 1914–2006 and the official CPI shelter weights. The HHP rental price series increased by a factor of 28.4 over the 1914–2006 period (15.8 over 1914–1987), and updating the prices series alone gives a revised CPI value for 2006 of 26.1 (14.3 in 1987). We then also apply our HHP shelter weight series for 1914–2006, along with the HHP rental price series, and obtain a CPI with a value for 2006 of 25.5 (13.9 in 1987). We use the series with both revised prices and shares as the baseline for the remainder of the paper.

Figure 4: Five-Year Inflation Rate in Official and HHP CPIs




Note: This figure plots the five year average rate of inflation in the official CPI and in the HHP CPI.

Figure 3 implies that a richer and more consistent treatment of shelter in the consumer price index over time would substantially increase estimated inflation over time, with the change driven mostly by the alternate measures of prices rather than changes in shares. Overall, the estimated average annual growth rate (AGR) of consumer prices in the pub-

lished CPI was 3.3% for 1914–2006. Our new estimates suggest that the true change in consumer prices over this period was closer to 3.6% per year, on average, an increase in inflation of almost one tenth.

Figure 5: Ratio of the HHP CPI to Official CPI



Note: This figure shows the ratio of the HHP CPI to the official CPI.

To better understand the timing of the divergence between the official CPI and our revised CPI, we plot the five-year average rate of inflation in Figure 4. This figure highlights that the second half of the 1940s and the late 1970s to early 1980s were the principal periods during which the two series diverge. By contrast, there is little difference overall in inflation rates just after World War I. The two episodes of divergence coincide with periods in which known limitations of the official treatment of shelter were most acute: in the late 1940s, rapid increases in market rents following rent controls were only partially captured in tenant-based measures. Meanwhile, in the late 1970s and early 1980s, the divergence reflects both the delayed pass-through of market rents to sitting tenants and contempora-


neous changes in the measurement and weighting of shelter associated with the transition to OER.

Another way of visualizing the difference between the official and HHP CPI series is the ratio of the indices. This is shown in Figure 5, with 1914=1 by construction. By the 1980s, the HHP CPI is between one fifth and one quarter higher than the official CPI. While the same periods of greater difference between the series are evident, this figure shows that the overall divergence results from a steady accretion of slightly larger increases in shelter costs, in addition to those larger divergences in the 1940s and in the late 1970s to early 1980s. This chart shows, in other words, that the understatement of overall inflation due to the limitations of the measurement of shelter costs was not driven solely by two periods of high inflation.


For ease of comparison, we also summarize the decadal differences between the two CPIs in Table 1. The first column shows the average growth rates of the official CPI measure based on the RoPR, while the second gives the changes from an HHP CPI (with both prices and shelter shares revised). Consistent with the above discussion of the differences in the two rental series, we see the greatest divergence in the decades of 1935–1940, 1945–1950, 1965–1970, and 1975–1980. It is worth noting, also, that inflation in the wider economy was rising during the late 1960s and 1970s, peaking in 1980 at over 11%. The calculations presented here suggest that the long-run increase in living costs has been understated in standard CPI measures, and thus the long-run increase in the standard of living has been overstated for the urban United States.

Two final caveats to our approach are worth noting. First, our national series is constructed from 30 cities and thus may differ from a sample that covered a broader set of cities, particularly if the cities in our sample experienced systematically different housing market dynamics or regulatory environments than non-sampled areas. Second, the shelter-share series before the availability of consistent modern CPI weights inevitably relies on splicing and, in the earliest years, interpolation across sparse historical benchmarks and national accounts movements.

Figure 6: City-Level Differences in Rental Mismeasurement



(a) Houston



(b) Detroit

Note: The figure shows the RoPR series compared with the HHP rental price series for (a) Houston and (b) Detroit.

We close our analysis by noting that the divergence between the RoPR and our HHP rental price series varies by city, as does the mismeasurement of the CPI. Here, we restrict our attention to the rental series from the BLS and the HHP Project and ignore other city-level differences in the CPI. Figure 6 shows the RoPR series and HHP rental price series for Houston (Panel a) and Detroit (Panel b). There is a striking difference across the two cities. In the HHP data, nominal Houston rents increased by a factor of 19 during our period but by only 7 in the city-level RoPR series. On the other hand, Detroit rents increased by a factor of close to 11 in both the RoPR and HHP series.

This comparison suggests the difference may stem from city-level housing market dynamics interacting with the limitations of the BLS methodology in history. Detroit saw one of the smallest increases in rents across the 30 cities during this period, while Houston was close to the national average. In markets with flat or falling rents, issues around vacancy are likely to be less prominent as tenancy length is likely to be shorter, given the benefits of staying in the same tenancy are less pronounced. Ambrose et al. (2015) also found that there were large differences in the trends in rents across cities over time using their newly developed RRI index for the years from 2000 to 2010, as cities experienced the early century boom and subsequent Great Recession differently. They attribute the greater volatility in their repeat rent index compared to the BLS series for those decades as a consequence of the smoothing nature of the BLS sampling procedures.

5 Conclusion

This paper revisits a longstanding puzzle in U.S. price measurement: why the official Consumer Price Index suggests that rents rose far more slowly than other prices over the twentieth century. We address this question by constructing new and conceptually consistent series for shelter prices and the shelter share for the 1914 to 2006 period. Using the Historical Housing Prices rental price index derived from over one million newspaper rental listings in 30 cities and a rolling-window hedonic approach, we find that nominal market

rents grew by about 3.7 percent per year 1914–2006, substantially faster than the 2.6 percent implied by the BLS Rent of Primary Residence series, with that difference driven by 1914–1987 (3.9% vs. 2.5%). Combining this with a shelter-share series that standardizes the treatment of owner-occupied housing over time, we obtain an alternative CPI that implies average inflation of roughly 3.6 percent per year from 1914 to 2006, compared with 3.3 percent in the official CPI. The gap is concentrated in specific historical episodes – especially the postwar years and the late 1970s and early 1980s – consistent with the mechanisms emphasized in the existing literature. However, our findings also reflect a steady accumulation of slightly higher shelter inflation over the long run.

The findings have direct implications for policymakers and institutions that rely on inflation measurement. Because shelter is the largest component of the CPI and because rents inform costs for both tenants and owners, persistent understatement of shelter inflation can materially affect measured real wage growth, real interest rates, and long-run assessments of living standards. If shelter costs rose by more than previously thought, enough to meaningfully affect overall inflation across nine decades, then in real terms living standards also rose by less. These issues are not merely historical: debates about rent stickiness, vacancy resets, and the conceptual treatment of owner-occupied housing remain central to how inflation is perceived and how policy is communicated. Our results suggest that periods of rapid adjustment in housing markets are precisely when traditional survey-based rent measures are most vulnerable to missing marginal price movements. A revised CPI that applies a consistent shelter concept over time and that uses market-based rent information when appropriate can provide a more reliable benchmark for historical comparisons, distributional analyses, and the calibration of macroeconomic narratives about the evolution of U.S. living costs.

The paper also opens several avenues for future research. One path would be to integrate richer microdata on lease structures and tenant turnover, where available, to quantify more precisely the wedge between market and average contract rents. This would enable alternative CPI constructs targeting different welfare concepts. A second direction would

be to study systematically why mismeasurement varies across cities, linking divergence to housing supply elasticity, local regulation, and changes in survey design. This could be in the existing 30-city sample or indeed across a broader sample of cities. Third, the methods used here could be applied to other hard-to-measure components of historical CPIs, particularly where quality change and shifting market composition are central concerns such as other durable goods.

More broadly, our results reinforce the value of combining modern price-index methods with newly digitized historical sources to revisit foundational measures of inflation and living standards. We view the revised shelter series and CPI presented here as a step toward a more coherent long-run price benchmark and as an invitation for further work that improves historical measurement and deepens our understanding of the economic past.

References

- Adams, B., Loewenstein, L., Montag, H., & Verbrugge, R. (2024). Disentangling rent index differences: Data, methods, and scope. *American Economic Review: Insights*, 6(2), 230–245.
- Adelman, I., & Griliches, Z. (1961). On an index of quality change. *Journal of the American statistical Association*, 56(295), 535–548.
- Albouy, D., Ehrlich, G., & Liu, Y. (2016). *Housing demand, cost-of-living inequality, and the affordability crisis* (tech. rep.). National Bureau of Economic Research.
- Ambrose, B. W., Coulson, N. E., & Yoshida, J. (2015). The repeat rent index. *Review of Economics and Statistics*, 97(5), 939–950.
- Ambrose, B. W., Coulson, N. E., & Yoshida, J. (2023). Housing rents and inflation rates. *Journal of Money, Credit and Banking*, 55(4), 975–992.
- Ball, L. M., & Koh, K. W. (2025). *Market rents and cpi shelter inflation* (tech. rep.). National Bureau of Economic Research.
- Chodorow-Reich, G., Glaeser, E., Jaravel, X., & Lipton, A. (2025). Homeownership and the cost of living. *Manuscript*.
- 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.
- Eiglsperger, M., Arioli, R., Goldhammer, B., Gonçalves, E., & Kouvavas, O. (2022). Owner-occupied housing and inflation measurement. *Economic Bulletin Articles*, 1.
- Genesove, D. (2003). The nominal rigidity of apartment rents. *Review of Economics and Statistics*, 85(4), 844–853.
- 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.
- 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.
- Hill, R. J., Scholz, M., Shimizu, C., & Steurer, M. (2022). Rolling-time-dummy house price indexes: Window length, linking and options for dealing with low transaction volume. *Journal of Official Statistics*, 38(1), 127–151.
- Lyons, R. C., Shertzer, A., Gray, R., & Agorastos, D. (2026). The price of housing in the United States, 1890–2006. *The Quarterly Journal of Economics*.
<https://doi.org/10.1093/qje/qjaf047>
- Marcoot, J. L. (1985). Revision of consumer price index is now under way. *Monthly Lab. Rev.*, 108, 27.
- Marcoot, J. L., & Bahr, R. C. (1986). The revised consumer price index: Changes in definitions and availability. *Monthly Lab. Rev.*, 109, 15.
- Mason, C., & Butler, C. (1987). New basket of goods and services being priced in revised cpi. *Monthly Lab. Rev.*, 110, 3.
- Ozimek, A. (2013). *Sticky rents and the CPI for owner-occupied housing*. Temple University.

- Poole, R., Ptacek, F., & Verbrugge, R. (2005). Treatment of owner-occupied housing in the cpi. *Federal Economic Statistics Advisory Committee (FESAC) on December, 9, 2005*.
- Randolph, W. C. (1988). Housing depreciation and aging bias in the consumer price index. *Journal of Business & Economic Statistics*, 6(3), 359–371.
- Rees, A., & Jacobs, D. P. (1961). The cost-of-living index. In *Real wages in manufacturing, 1890-1914* (pp. 74–119). Princeton University Press.
- Rippy, D. (2014). The first hundred years of the consumer price index: A methodological and political history. *Monthly Labor Review*. <https://doi.org/10.21916/mlr.2014.13>
- Rivers, J. D., & Sommers, J. P. (1983). Vacancy imputation methodology for rents in the CPI. *Proceedings of the ASA Economics and Business Section*, 201–205.
- Silver, M. (2016). *How to better measure hedonic residential property price indexes*. International Monetary Fund.
- Triplett, J. E. (2001). Should the cost-of-living index provide the conceptual framework for a consumer price index? *The Economic Journal*, 111(472), 311–334.
- U.S. Bureau of Labor Statistics. (1983). *CPI Detailed Report: Changing the homeownership component of the Consumer Price Index to rental equivalence* (tech. rep.). Bureau of Labor Statistics.
- Vandell, K. D. (1995). *FHA restructuring proposals: Alternatives and implications*. Taylor & Francis.

A Data Appendix

Table A1: Real Median Price (2006 Dollars) vs Modal Size

City	Rent			
	Median Price		Modal Size	
	1914-1945	1946-2006	1914-1945	1946-2006
Atlanta	608	708	6	2
Baltimore	474	578	3	3
Boston	467	1,037	5	2
Charleston	392	545	5	2
Chicago	753	907	4	2
Cincinnati	522	582	4	1
Cleveland	568	625	5	2
Dallas	471	787	5	2
Detroit	629	546	5	2
Houston	479	697	5	2
Las Vegas	.	778	.	2
Los Angeles	562	1,057	4	1
Louisville	463	501	4	2
Memphis	517	443	4	2
Miami	464	895	4	2
Minneapolis	442	735	5	1
Nashville	471	609	5	2
New Orleans	466	598	5	2
New York City	992	1,263	3	1
Philadelphia	514	698	6	1
Phoenix	353	703	3	2
Pittsburgh	590	625	5	1
Portland	407	559	5	2
Salt Lake City	374	580	3	2
San Diego	364	667	5	1
San Francisco	487	1,124	3	2
Seattle	436	703	3	1
St. Louis	503	488	4	2
Tampa	356	559	5	2
Washington D.C.	683	925	2	1

Note: Modal size is measured in total rooms 1914-1945 and in bedrooms 1946-2006.


Table A2: Newspapers and Access Sites

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

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 to find a newspaper for Houston for 2006.


B Supplemental Results

Figure A1: Year-over-Year Percent Changes for HHP Rental Index and BLS Rent of Primary Residence




Note: This figure shows the year-over-year changes in the HHP rental index and the BLS Rent of Primary Residence series.

Figure A2: BLS New-Tenant and All-Tenant Rent Indexes



Note: This figure plots the New-Tenant Rent Index (R-CPI-NTR) in the quarter and the All-Tenant-Rent-Index (R-CPI-ATR) with a lead of three quarters. These series were taken from <https://www.bls.gov/cpi/research-series/r-cpi-ntr.htm> and we used the 2024q4 vintage.

Figure A3: CBRE Same Store Rental Index and HHP Rental Index by City



Note: This figure plots the city-level HHP rental index against the CBRE “same store” rental index.