Federal Reserve Bank of Cleveland Working Paper Series

Disentangling Rent Index Differences: Data, Methods, and Scope

Brian Adams, Lara P. Loewenstein, Hugh Montag, and Randal J. Verbrugge

Working Paper No. 22-38

December 2022

Suggested citation: Adams, Brian, Lara P. Loewenstein, Hugh Montag, and Randal J. Verbrugge. 2022. "Disentangling Rent Index Differences: Data, Methods, and Scope." Working Paper No. 22-38. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202238.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Disentangling Rent Index Differences: Data, Methods, and Scope*

Brian Adams[1], Lara Loewenstein[2], Hugh Montag[1], and Randal Verbrugge[2]

[1]US Bureau of Labor Statistics
[2]Federal Reserve Bank of Cleveland

December 5, 2022

Abstract

Prominent rent growth indices often give strikingly different measurements of rent inflation. We create new indices from Bureau of Labor Statistics (BLS) rent microdata using a repeat-rent index methodology and show that this discrepancy is almost entirely explained by differences in rent growth for new tenants relative to the average rent growth for all tenants. Rent inflation for new tenants leads the official BLS rent inflation by four quarters. As rent is the largest component of the consumer price index, this has implications for our understanding of aggregate inflation dynamics and guiding monetary policy.

Keywords: house prices, rent growth, inflation measurement, monetary policy, forecasting

JEL Classification: E31, E37, E27, H31

*The views expressed in this paper are solely those of the authors and do not necessarily reflect the opinions of the Federal Reserve Bank of Cleveland or the Federal Reserve System. Emails: adams.brian@bls.gov, lara.loewenstein@clev.frb.org, montag.hugh@bls.gov, randal.verbrugge@clev.frb.org

1 Introduction

Shelter comprises 32 percent of the consumer price index (CPI).[1] Accurate inflation measurement therefore depends critically on accurate rent inflation measurement, which drives both CPI rent[2] and owners' equivalent rent (OER). It is therefore concerning that rent indices differ so greatly. For example, in 2022q1 inflation rates in the Zillow Observed Rent Index (ZORI)[3] and the marginal rent index (ACY MRI)[4] reached an annualized 15 percent and 12 percent, respectively, while the official CPI for rent read 5.5 percent (see Figure 1a).

If the Zillow reading were to replace the official CPI rent measure, then the 12-month headline May 2022 CPI reading of 8.6 percent would have been over 3 percentage points higher. These are consequential discrepancies, larger than any of the historical CPI biases noted by the Boskin commission (Boskin et al. (1997)) and much greater than any of the current biases noted in Lebow and Rudd (2003) and Moulton (2018). Differences of this magnitude have consequences for housing economics, monetary policy, contract escalation, and GDP and welfare measurement (Ambrose et al. 2018; Hill et al. 2020; Ambrose et al. 2022).[5]

The divergence in rent inflation readings might stem from differences in sample representativity, index construction methods, scope of the underlying rental data, or quality adjustments. The BLS Housing Survey, which underlies CPI rent, is a random sample designed for measuring rent growth; it is fully representative of the rental housing stock in US cities. In contrast, both ZORI and the CoreLogic Single-Family Rent Index (SFRI),[6] are based off samples of mainly higher-tier detached rental units that advertise in the Multiple Listing Service (MLS), and the ACY MRI covers larger apartment complexes in a restricted number of cities (Ambrose et al. 2018). CPI rent and OER indices use 6-month changes in average rent growth over a fixed sample of rental units; in contrast, ZORI and the SFRI are both repeat-rent indices, and the ACY MRI is the product of two aggregate indices, a price and an expected cap rate. The CPI measures rent growth facing all occupants, while the alternative measures track rent growth facing new tenants. Only the CPI adjusts for aging, structural changes, and changes in utilities provision.

[1]This is the sum of the CPI aggregation weights of rent and owners' equivalent rent. By comparison, food accounts for approximately 14 percent.

[2]CPI rent is based on the rent paid by tenants for their primary residence.

[3]See Clark (2022).

[4]See Ambrose et al. (2022).

[5]Furthermore, other rent inflation measures are also more cyclical and less sticky than inflation based on CPI rent, leading to important implications for macroeconomic modeling. For instance, Appendix D shows that Phillips curve parameter estimates (following Ashley and Verbrugge (2020)) and estimated impulse response functions of New Keynesian models (following Gelain and Manganelli (2020)) are very sensitive to the rent index used.

[6]See Boesel et al. (2021) and Nothaft (2018).

We use the microdata underlying official measures of rent inflation in the CPI to assess the differences between the CPI rent index and other measures of rent growth. We create several weighted repeat-rent indices in the style of Case and Shiller (1989), such as a weighted new-tenant repeat-rent (NTRR) index (using only leases of tenants who recently moved in), and a weighted all-tenant repeat-rent (ATRR) index (using all tenants, whether they recently moved in or not). These indices allow us to determine whether the differences between CPI rent inflation and alternative measures are due to differences in the representativity of the sample, the scope of the sample, or the methodology employed.

We find that most of the discrepancy between CPI rent and other measures is due to differences in rent increases for *all* tenants versus *new* tenants. In 2022q2, our ATRR index was recording 5.94 percent year-over-year inflation, while the NTRR inflation rate was at 11.88 percent. CPI rent inflation was at 5.14 percent.

The CoreLogic SFRI, despite its non-representative nature, is a fair approximation to our NTRR index. Some authors (Ambrose et al. (2015, 2022)) have criticized the all-tenant approach, arguing that the CPI should track new-tenant rents to better capture current market conditions. Conversely, the price statistics literature generally favors the use of an all-tenant index, to capture changes in the purchasing power of typical renters. We contribute to this debate by clarifying the difference that the use of a new-tenant index would make, using the same data source that underlies the CPI, and noting practical challenges that would accompany such usage.[7]

More generally, a price index measure should be chosen based upon its intended purpose. For studying generalized changes in living standards or for Social Security benefit escalation, an all-tenant index is preferred to reflect what is happening to the typical household. In such contexts, data revisions — an inherent feature of repeat-rent measures — are somewhat problematic. Regarding inflationary-pressure signals, new-tenant indices more quickly reflect market condition shifts (our NTRR leads the CPI by about four quarters,[8] while the ATRR leads it by about one quarter).

For monetary policy considerations, it is not immediately clear which rent index is preferable. Different models yield different conclusions. In some models, the central bank should distinguish between price developments in different sectors, and (for instance) target more persistent prices (for example, La'O and Tahbaz-Salehi (2022)). But these models tend to take existing component indices as given, merely consider optimal weighting. Further-

[7]A related argument against using a new-tenant rent index in inflation measurement is that if the fraction of households that move varies significantly over time, we could be capturing price changes for an ever-changing portion of the market. We show that, consistent with other research (Ganong and Shoag 2017; Molloy et al. 2011), the share of units with new tenants has slowed in the BLS housing sample over time and shows a seasonal pattern.

[8]SFRI and ZORI are roughly coincident with NTRR, while ACY MRI lags it a tad, perhaps reflecting a property-owner's expectations lag.

more, relatively few models studying optimal policy include housing. As far as we know, no monetary policy model has addressed the topic of optimal rental inflation measurement per se.

2 The BLS Housing Survey

We construct new repeat-rent indices using the BLS Housing Survey data, the same data that the BLS uses to compute the CPI rent index. The survey follows a sample of renter-occupied housing units, surveying the same rental units every six months. The sample is evenly split across six panels, each with a different set of survey months (for example, January and July, or February and August). Observations include the contract rent, the utilities and services included with the rent, the tenant's move-in date, and many other unit-level characteristics. Units are not necessarily surveyed at the start of leases; many observations are of continuing renters in the middle of a lease or after a lease renewal. There is nothing in the data to directly indicate if a lease was renewed.

The BLS Housing Survey uses a multistage sampling design meant to draw a sample representative of rental expenditure.[9] In Table 1 we show that summary statistics from the BLS Housing Survey are very similar to those from the American Housing Survey (AHS), which is another carefully crafted survey of housing units.[10] The first stage of the BLS Housing Survey sampling selects large geographic areas, called "primary sampling units" (PSUs), to represent all metropolitan and micropolitan areas in the United States.[11] Each PSU is divided into segments, which become the fundamental units for sampling and weighting. Segments consist of one or more contiguous Census blocks, often Census block groups. Segments are selected using a probability-proportional-to-size (PPS) method, where "size" is an estimate of total shelter expenditure within the segment. Finally, the BLS randomly samples enough rental units to yield at least five responding units per segment. The sample size is typically around 40,000 units across the six panels.

The BLS selected a new sample in 1999. Subsequently, the survey lost units to demolition, to conversion to other uses, or to respondent non-cooperation.[12] The survey periodically added new units sampled from construction permit data. More recently, the BLS

[9]For more details on the design of the Housing Survey sample see "Updating the rent sample for the CPI Housing Survey" available at https://www.bls.gov/opub/mlr/2013/article/updating-the-rent-sample-for-the-cpi-housing-survey.htm.

[10]The construction of the AHS is described in more detail in Section **??** in the appendix.

[11]The BLS redesigned its geographic sample in 2018; now, PSU definitions match core-based statistical areas. Previously, PSUs had been modified metropolitan statistical areas and groups of counties with smaller towns (Paben et al. 2016).

[12]Gallin and Verbrugge (2016) suggest that sample attrition was concentrated in higher-quality units; such attrition influences aging bias estimates, among other things.

implemented a rolling sample replacement design, with a new sample drawn starting in 2012. Since 2016, units remain in the sample for only six years; one-sixth of the sample is replaced annually.

Repeat-rent indices require paired observations of the same unit. Because units are surveyed every six months, an all-units repeat-rent index can be calculated from nearly the beginning of the data in 1999. A repeat-rent index based on observations of new tenants can only add a unit to its calculation after the second observed move-in. Figure 2a shows that between 13 and 25 percent of units have new tenants in a month. Several years are needed for a new-tenant repeat-rent index to achieve a steady sample size, so we begin our NTRR series in 2005q1, which corresponds with the time frame covered by other rent growth measures. The number of observations (pairs of new tenant observations counted as of the date of the second transaction) used to construct our NTRR is plotted in Figure 2b.

The official CPI dates observations to their survey collection month. A rent change in a unit may happen in the months in between surveys, and therefore several months before the survey collection period. We seek to identify the month in which the rent change happens. Accordingly, we date observations either to their recorded move-in date or the completion of the most recent six-month interval since move-in.[13] This is the most likely date of the rent change, because most rental contracts in the US are annual, and six-month contracts are also common (Crone et al. 2010). Because we identify, and use, the date of the rent change —which typically occurs prior to the collection period —our indices will reflect rent changes sooner than will official indices.

CPI rent and our repeat-rent indices use a CPI-constructed rent measure called "economic rent." Economic rent is the result of adjusting the contract rent to account for services rendered in lieu of rent, changes in utilities bundled with rent, and for the aging of units (Crone et al. 2010; Gallin and Verbrugge 2007). The CPI rent index further includes vacancy adjustments and adjustments for structural changes to the housing units. Instead of estimating the value of structural changes, we exclude observations for which the number of rooms, bathrooms or half-bathrooms changes, or where a field note includes the words "remodel," "renovate," or "refurbish." We exclude imputed rents for vacant units while addressing the vacancy bias described in Sommers and Rivers (1983) by other means. In particular, in our NTRR, the last unit observation used is the final observed rent change. In our ATRR, we exclude all observations after the last date at which a new tenant moves in.[14] None of the alternative rent growth indices discussed in this paper make adjustments

[13]For example, consider a tenant who moved into a housing unit in February 2011, and the housing unit is sampled on an April-October cycle. If the BLS microdata show that the rent changed from October 2011 to April 2012, we assume that the month it changed was February 2012, a 6-month multiple of the move-in date.

[14]Vacancy bias arises from two factors: unit attrition from the sample, and unit-level rent-setting upon tenant turnover versus upon contract renewal. Rent changes are often small (or zero) when leases are

for structural change, aging, or changes in the provision of utilities and services.

CPI rent is calculated using the average six-month change in that month's sample. The index converts it into a monthly change by taking its sixth root. Let $\text{rent}_i^*(t)$ denote economic rent. Then the rent index at time t for a particular geographic region is constructed as

$$I^R(t) = \left(\frac{\Sigma_i w_i \text{rent}_i^*(t)}{\Sigma_i w_i e^{F_{i,t}} \text{rent}_i^*(t-6)} \right)^{1/6} I^R(t-1) \tag{1}$$

where w_i is the unit-specific weight,[15] and F_i, t is an age-bias factor that lowers the rent level in period $t-6$ to account for the fact that the observed change in rent will understate the constant-quality change in rent.[16]

A drawback of the BLS Housing Survey is its modest size. As depicted in Figure 2a, every quarter we have observations on between 11 and 17 thousand housing units. The share of these observations that are new tenants varies over time, fluctuating between 14 and 23 percent. This share has a clear seasonal pattern, reflecting the seasonal pattern of relocations, which are higher in summer and fall. The share of housing units with new tenants fell during the Great Recession, but has remained stable since 2012.

For an observations to be included in our NTRR, we need to observe a prior new-tenant observation for that housing unit. Consequently, the sample size underlying our NTRR is smaller than the total sample of new-tenant observations, and varies more over time. The sample size is plotted in Figure 2b. For most of the 2000s, we had between 1500 and 2000 observations every quarter. The sample size of new tenancies subsequently fell to about 750 in the quarters before the COVID pandemic. The recent decline in repeat observations for new tenants is driven by a feature unique to BLS microdata. Starting in 2012, the BLS sample began converting to a six-year rotation whereby each rental unit is included in the sample for only six years. Prior to 2012, a unit would typically remain in the sample for much longer. To include a unit in our NTRR index, we need to observe two separate new tenant move-ins. Given the limitation on the length of time a unit remains in the sample, we are much less likely to observe two new tenant observations for a given unit with the rotating sample. The change also results in a change in the distribution of property types in our NTRR index. As apartments tend to see higher turnover in tenants than do single-family rentals, we see an increase in the share of observations attributed to apartment building

renewed, but are much larger when new tenants move in. If most sample attrition occurs during tenant turnover, then the CPI for rent would have a downward bias without some sort of adjustment. The CPI practice is to impute a "final" rent that reflects a typical tenant turnover rent increase, before dropping the unit out of the sample. Our indices have no vacancy bias, because the last data from a unit coincide with a turnover rent.

[15]For each unit, there is a segment-level rent weight (corresponding to the segment that the unit is in), and a second weight that adjusts for nonresponse; see Section 4.

[16]For more details on the construction of CPI rent see Verbrugge and Poole (2010) or the BLS Handbook of Methods.

rents, although this share does not increase until 2016.

3 Other Rent Data Sources and Indices

3.1 CoreLogic SFRI

The CoreLogic SFRI employs an arithmetic repeat-rent methodology using rental listings of single-family properties in the Multiple Listing Service (MLS). CoreLogic collects these data from participating realtor boards. By 2020, CoreLogic had on average 10 years of history for these boards, and it had more than 20 years of data in some markets. CoreLogic creates rent indices for CBSAs for which it has sufficient data. The national SFRI is then a weighted average of the available CBSAs, where the weight is based on the value of the rental housing stock in each CBSA (Boesel et al. 2021).

3.2 Zillow Observed Rent Index

ZORI is a repeat-rent index that begins in 2014. It is based on Zillow's proprietary rental data from rental listings on its website and from MLS listing data. Its estimation methodology proceeds in three stages. First, Zillow estimates equation 2 unweighted. In the second stage, Zillow regresses the squared residuals from the first stage on weights created by comparing the distribution of structure type and age of rental properties in Zillow's data to that in the American Community Survey (ACS) in each respective year. The predicted values from this second stage are used in a weighted least squares regression of equation 2; this index forms the ZORI. Once constructed, the index is smoothed using a three-month exponentially weighted moving average.

Both SFRI and ZORI are based solely on the rents paid by new tenants, not tenants renewing a lease. The MLS data set underlying both is not representative of the general rental market. The Census's 2018 Rental Housing Finance Survey estimates that only 11 percent of single-unit rental properties are listed using a real estate agent (and thus listed in the MLS).[17] On average, rental listings in the MLS are more expensive, larger, and newer than newly occupied rental units in the AHS (see Table 1).

3.3 The ACY Marginal Rent Index

The ACY MRI of Ambrose et al. (2022) is based on the product of two series for large multifamily properties from Real Capital Analytics (RCA): the commercial property price index (CPPI), which is a repeat-transaction index, and the monthly average multifamily

[17] See Choi and Young (2020) for the differential advertising strategies of landlords.

capitalization rate for transacting properties (or income yield). The product of these two series produces a baseline net rent index that is then re-scaled to match a former index created by the same authors from Experian RentBureau data, called the repeat-rent index (RRI),[18] which was discontinued in 2010.

The CPPI and the multifamily capitalization rate are based on RCA's database of commercial properties. The capitalization rate is based on the last month's net operating income for each property (not a historical average) and is therefore forward looking. However, while the database maintained by RCA is comprehensive, it is limited to properties worth at least $2.5 million or more. The data underlying the ACY MRI are therefore very different from those underlying either the SFRI or the CPI rent index.

4 Constructing a Repeat-Rent Index

The repeat sales method of Bailey et al. (1963) measures price growth, controlling for the time-invariant components of unobserved quality by using observed housing unit-level changes in price. Repeat-transaction indices have been used for house prices (Case and Shiller 1989) and rents (Ambrose et al. 2015; Boesel et al. 2021; Clark 2022).

Suppose our data set of rental prices has observations sampled from periods $\{1, \ldots, N\}$. Let a unit i have rent observations in period s and period $t > s$. This observation pair enters a regression as

$$\ln P_{it} - \ln P_{is} = \gamma_1 D_{i1} + ... + \gamma_N D_{iN} + u_{it}, \tag{2}$$

where $D_{ij} = 1$ if the second observation in the pair took place in period j, and $D_{ij} = -1$ if the first observation in the pair took place in period j, and for each other period $D_{ij} = 0$. For our example observation, $D_t = 1$ and $D_s = -1$. By using log prices, the parameters γ approximate percentage differences in prices from the base year; the base period index value is normalized to 1.

Using the BLS housing survey data, we construct two repeat-rent indices, which differ only in their scope. The first is the new-tenant repeat-rent (NTRR) index; like the SFRI and its peers, it uses only observations with a new tenant. Observation pairs thus bookend the tenure of a renter within a housing unit: the first date records when the renter moved in, and the second date when the next renter moved in. Tenure lengths average about three years but vary substantially. The NTRR reflects prices that a new renter would face if she changed her housing unit every period.

The second repeat-rent index is an all-tenant repeat-rent (ATRR) index. Its scope is

[18] See Ambrose et al. (2015).

broader: its sample includes all housing units and dates. In this case, each observation pair is based on the two consecutive occasions on which a housing unit is surveyed as part of the BLS rental data set. The ATRR represents the prices paid by an average renter (new and continuing). Thus, it differs in scope from the NTRR, SFRI, and several other rent indices. Comparisons of the ATRR with the NTRR will isolate the effects of changing the scope, since both draw from the same sample and share the repeat-rent methodology. Comparisons of the ATRR and the CPI rent index will isolate the effects of changing methodology.

We drop any observations in the top and bottom 1 percent of annualized rent changes. To further mitigate volatility, both NTRR and ATRR are quarterly (instead of monthly) indices. While CPI rent includes an adjustment for structural changes in properties, we instead drop any properties with a change in the number of bedrooms, number of other rooms, number of half-bathrooms, number of full bathrooms, type of AC equipment, type of heating equipment, or that indicate that they are newly remodeled.

Our repeat-rent indices are estimated using weights from the CPI rent index. First, we estimate equation 2 to create PSU-specific rent indices. The ATRR indices use segment weights (which is the inverse of the probability of selecting that segment within a given PSU) to estimate equation 2 for each PSU. We adjust segment weights by the number of observations that we omitted due to remodelling, vacancy, and outlier adjustments as well as survey non-responses to maintain representivity within the PSU. CPI segment weights are calculated for each collection period. Each of our all-tenant repeat-rent observations is weighted using the segment weight from the more recent collection date. The NTRR PSU-specific indices use uniform lower weighting, as the segment weights do not correspond to the universe of new tenant units. Second, we use a second set of "upper" weights, which represent the share of national rent expenditure accounted for by that PSU, to aggregate the PSU-specific indices and form a national index. Upper weights are updated every two years for the CPI.[19]

Differences in tenure lengths may make the error term in equation 2 heteroskedastic, particularly for our NTRR index. Case and Shiller (1989) and Goetzmann (1992) propose a three-stage procedure to address it. First, we estimate equation 2 and obtain the residuals. Second, we regress the residuals squared on a constant and the time between observations. Third, we use the resulting predicted values to estimate a GLS version of equation 2. Like Clark (2022), we found that this heteroskedasticity correction had a negligible effect on our final index.

We calculate rent inflation by taking the annual log difference in index values. Standard

[19] The geographic areas included in the BLS Housing Survey changed in 2016; some non-self-representative PSUs were dropped and others were added. Baltimore and Washington, DC were split starting in 2016. Other areas, like New York City, the New Jersey suburbs, and the New York suburbs were combined. We create a static set of PSUs to construct our ATRR and NTRR.

errors for our inflation estimates are created using a bootstrap method described in Appendix ??.

5 Disentangling Rent Index Differences

We start by comparing our two new repeat-rent indices to each other and to CPI rent. Although CPI rent and ATRR differ slightly, they generally track each other. By contrast, NTRR is more volatile and shows much higher recent rent growth. We then compare our indices to measures of rent inflation based on alternative data sources.

5.1 ATRR and CPI Rent

Our ATRR and CPI rent are both measures of rent growth for *all* tenants, whether new or continuing an existing lease. In this sense they should be very similar. However, besides their construction methodology, they may differ for a number of reasons:

1. CPI rent includes a vacancy adjustment with missing rents, while for the ATRR, units with missing rent are dropped.

2. The BLS applies a quality adjustment for large structural changes to CPI rent, while we drop observations with structural changes from the ATRR.

3. CPI rent is constructed using all rental observations in the housing survey, while we drop outliers as described in Section 4.

4. CPI rent is based on rents collected that month, while observations in the ATRR are based on dates relative to move-in dates.

5. The two indices are weighted slightly differently, as described in Section 4.

Figure 1b contains a comparison of the rent inflation implied by the ATRR (in blue) and that based on CPI rent (in purple). The two track each other with some notable differences. The ATRR is lower than CPI rent on average, especially early in the sample, although starting in 2021 and most of 2022 ATRR is above CPI rent inflation.

It appears from Figure 1b that ATRR leads BLS CPI rent. We confirm this in Figure 3a, which shows cross-correlations between the indices at various lag lengths. The ATRR leads the CPI by about one quarter, most likely reflecting the delay between when rents are changed and when that change is reflected in the BLS Housing Survey.

5.2 NTRR versus ATRR

We next compare our NTRR and ATRR; see Figure 1b. The two indices differ mostly in the scope of their underlying data. The NTRR is limited to new tenants, while the ATRR is based on both new and continuing rental leases.

The NTRR leads the ATRR. This is evident both in a visual comparison of their time series (in Figure 1b) and in their intertemporal cross-correlations (Figure 3a). The NTRR's trough in rent change after the housing crisis is deepest in 2009q4, whereas ATRR reaches its trough later in 2010q1and stays negative throughout 2010. Likewise, the current price spike begins in NTRR well before it begins in ATRR. The correlation in quarterly changes is only .79, but the correlation of the quarterly changes in ATRR with a three-quarter lag of NTRR is .85. As Figure 3a shows, the NTRR leads the ATRR by about three quarters, while it leads the CPI for rent by about four quarters.

The NTRR is much more volatile. The large rent decline in 2009-2010 and the inflation spike in 2021-2022 are more extreme in the NTRR. NTRR inflation also has many more smaller fluctuations over our entire sample period. This noise is also reflected in the standard errors. Standard errors for the ATRR are small, averaging 0.03 percentage points. In contrast, the NTRR has bigger standard errors, averaging 0.47 percentage points. The chief reason for this difference is that the NTRR is calculated from a subset of the observations in the ATRR — those with newly moved-in tenants. But a second reason is that continuing renters in the ATRR, even when signing new leases, tend to have sticky rents (Gallin and Verbrugge 2017).

5.3 NTRR and Alternate Measures of Rent Inflation

We compare our NTRR measure of rent inflation to the CoreLogic SFRI, ZORI, and the ACY MRI. Our NTRR and the CoreLogic SFRI have similar time series. The differences in their inflation rates are often not statistically significant. In Figure 1a, the SFRI generally appears to be a smoothed version of the NTRR. Before 2022, their quarterly year-on-year inflation rates never differ by more than 1.73 percentage points and have an average absolute difference of 0.73 percentage points. Quarterly changes in the two indices are highly correlated ($\rho =$ .93), and neither series leads the other. Figure 3a depicts the intertemporal cross-correlations of the SFRI, ACY MRI, ATRR and the CPI rent inflation with the NTRR.

The similarity of the NTRR and the SFRI suggests that the lack of representativeness of SFRI data is not driving the divergence between the SFRI and the official BLS rent index. Instead, differences mainly stem from methodology (repeat-rent versus the CPI methododology described in Section 2), scope (new tenant versus all tenants), and/or rent adjustments (the BLS performs quality adjustments that the SFRI does not). The SFRI data approximate rent change in the BLS data, which might support their use in macroeconomic and housing studies. Because the SFRI and the NTRR are similar, comparing the NTRR and the ATRR — which changes scope, holding methodology constant — shows that most of the difference between the SFRI and CPI rent inflation is due to scope.[20]

ZORI begins in 2014, providing a shorter comparison period with our NTRR. The ZORI quarterly year-on-year rent inflation rate is often similar to the NTRR, but not always. Indeed, in recent quarters, the ZORI inflation rate exceeds not only our NTRR, but all other rental inflation rates. Nevertheless, ZORI and the NTRR are highly correlated ($\rho = 0.97$).

The ACY MRI displays much higher volatility than the other rent inflation measures (see Figure 1a) and lags the NTRR by about one quarter (see Figure 3a), perhaps reflecting an information lag in the expectations of property sellers. Given that the MRI is already created using a scaling factor (see Section 3.3), it is possible that a different scaling factor could reduce its deviations from the NTRR.

We estimate a new scaling factor by minimizing the year-on-year percentage change in the ACY MRI's mean squared error from the year-on-year percentage change of NTRR. The minimization problem is

$$\min_{r,a} \sum_t \left(r \left(\pi^{y/y}_{t,rMRI} \right) + a - \pi^{y/y}_{t,NTRR} \right)^2 \tag{3}$$

where $\pi^{y/y}_{t,rMRI}$ is the year-over-year inflation in the ACY MRI, $\pi^{y/y}_{t,NTRR}$ is the year-over-year inflation in the NTRR, r is the scaling factor, and a is a constant. The mean squared error-minimizing value of r is 0.450 and a is 1.43. The resulting (rescaled) index, which is depicted in Figure 3b, matches the dynamics of the NTRR and the SFRI fairly well; it generally lies within the error bounds of the SFRI except briefly in 2011 and 2012.

6 Implications for Measuring Rent Inflation

The NTRR leads the ATRR and the CPI rent Index, and so in some sense the NTRR is a timelier index. This is one reason Ambrose et al. (2022) suggested that the BLS should use a rent inflation measure based on new tenants. Yet, instability at the end of the sample hinders its use as a precise estimate of real-time rent inflation. Repeat transaction indices are subject to revisions, because new observations inform estimates for previous periods. The

[20] The fact that our NTRR inflation and the SFRI rent inflation track each other so closely implies that the data underlying the SFRI may be useful to researchers to measure rent growth. In Section **??** in the Appendix, we create CBSA-level new-tenant repeat-rent indices using both the MLS microdata and the BLS Housing Survey. Even at the local level, the two data sets provide similar estimates of rent growth. This is despite the fact that even within cities, the coverage of the MLS data is starkly different from that of the BLS Housing Survey (see Section **??** in the Appendix for details).

new observation provides an estimate of a unit's rent increase since the previous observation for that same unit. It therefore influences the index over the entire time spanned by those repeat observations. The latest period of a repeat transaction index is especially volatile because the sample size available to estimate the index at time s is smallest at time s, and then grows for time $t > s$ as new repeat observations span that period. Thus, the index estimate for period s gradually improves, as more rent observations accumulate in later periods.

The NTRR based on BLS Housing Survey data is especially susceptible to end-of-sample revisions because housing units are surveyed every six months. As a result, the full sample of new tenants for a given quarter is unavailable until six months after the end of that quarter. A new tenant rent measure would also be a complex addition to the CPI, which currently limits ex-post revisions to seasonal adjustment and serious errors.

We provide an example of inflation-estimate revision, graphed in Figure 3c. The historical inflation rate that would have to be estimated using only the data available through 2015q4 (represented by the blue dashed line in 3c) increasingly gyrates around the historical inflation rate estimated from all of the data (represented by the black line). The deviation represents the influence of additional sample. In the 2015q4 estimates, confidence intervals are much wider for quarters 2015q1–q4. As more data became available, point estimates received revisions as large as 0.41 percentage points. Repeat-rent indices are inherently prone to this behavior, although the effect will be exacerbated if the sample size is small to begin with.[21]

7 Dynamic Relationships and CPI Rent Forecasting Implications

To explore the dynamic relationships between the various rent index inflation rates, as well as to assess potential forecast gains for CPI rent using SFRI, we estimate vector error-correction models (VECM) on pairwise sets of series. These highlight both the long-term relationship, and their shorter-run dynamics. The VECM are specified as

$$\Delta y_t = \alpha(\beta' y_{t-1}) + \upsilon + \sum_{i=1}^{3} \Lambda \Delta y_{t-i} + \epsilon_t \tag{4}$$

where $y_t = (y_{1,t}, y_{2,t})'$ is a vector of two (12-month growth-rate) rent indices (for example, $y_{1,t}$ = CPI rent, $y_{2,t}$ = SFRI); $\Delta y_t = (y_t - y_{t-1})$; $\upsilon = (\upsilon_1, \upsilon_2)'$ is a constant; Λ is a matrix of coefficients on lag terms; the vector $\beta = (\beta_1, \beta_2)'$ describes the long-term cointegration

[21]Different rent measures also have implications for PCE inflation, macroeconomic models, and Phillips curve estimation. Appendix ?? discusses this in more detail.

relationship between y_1 and y_2 such that $\beta' y_{t-1}$ is stationary; and $\alpha = (\alpha_1, \alpha_2)'$ determines the speed at which each variable adjusts back toward this cointegrating relationship. We normalize β_1 to 1. We estimate these relationships on pre-pandemic data, to avoid overfitting based on one extreme episode. Table 2 reports the values of α and β along with some standard errors.

Long-term relationship estimates suggest that long-term averages of all four of the series investigated will coincide. In responding to deviations from their long-term relationships (with either CoreLogic or with CPI rent), the NTRR and the ATRR do most (or all) of the adjusting to eliminate said deviations; the NTRR also does most of the adjusting toward the ATRR and toward the CPI rent. Conversely, and somewhat surprisingly, in CPI rent-CoreLogic relationship, neither variable strongly moves to eliminate the gap.

We explore the predictive content of the SFRI for CPI rent using the Bayesian information criterion (BIC). Dropping all SFRI terms from equation 4 — that is, using a univariate model in ΔCPI rent$_t$ — the BIC is -0.127. Inclusion of lags of the SFRI causes the BIC to fall to -0.491, indicating substantial predictive content for these terms. However, if one then adds the cointegration term ($\beta' y_{t-1}$), the BIC *rises* to -0.456. These results indicate that the SFRI has predictive content for CPI rent, but that the cointegration relationship is not a useful predictor.

8 Conclusion

We show that the main differences between alternative rent growth measures from CoreLogic and Zillow and the CPI rent index are due to differences in the scope of the underlying data sources. The CPI rent index is based on rent of *all* renters, while the CoreLogic SFRI and ZORI are based on the rent of new tenants. We create a repeat-rent index from a sample of the BLS Housing Survey that is limited to new tenants. The resulting rent index is much closer to both alternative indices.

Rent inflation based on the CoreLogic SFRI has a surprisingly close relationship to the inflation based on our NTRR. This is despite the fact that the data underlying the SFRI are not representative: they pertain only to larger and more expensive single-family units, and are not fully geographically representative. This has implications for researchers who use MLS microdata to measure rent growth.

References

Ambrose, Brent W, N Edward Coulson, and Jiro Yoshida (2015). “The repeat rent index.” *Review of Economics and Statistics*, 97(5), pp. 939–950. doi:10.1162/REST_a_00500.

Ambrose, Brent W, N Edward Coulson, and Jiro Yoshida (2018). “Reassessing Taylor rules using improved housing rent data.” *Journal of Macroeconomics*, 56, pp. 243–257. doi:10.1016/j.jmacro.2018.03.001.

Ambrose, Brent W, N Edward Coulson, and Jiro Yoshida (2022). “Housing rents and inflation rates.” *Journal of Money, Credit and Banking*. doi:10.1111/jmcb.12971.

Ashley, Richard and Randal Verbrugge (2020). “Finding a stable Phillips curve relationship: A persistence-dependent regression model.” Working Paper 19-09r, Federal Reserve Bank of Cleveland. doi:10.26509/frbc-wp-201909r.

Bailey, Martin J, Richard F Muth, and Hugh O Nourse (1963). “A regression method for real estate price index construction.” *Journal of the American Statistical Association*, 58(304), pp. 933–942. doi:10.1080/01621459.1963.10480679.

Boesel, Molly, Shu Chen, and Frank E. Nothaft (2021). “Housing preferences during the pandemic: effect on home price, rent, and inflation measurement.” *Business Economics*, 56, pp. 200 – 211. doi:10.1057/s11369-021-00241-4.

Boskin, Michael J, Ellen R Dulberger, Robert J Gordon, Zvi Griliches, and Dale W Jorgenson (1997). “The CPI commission: Findings and recommendations.” *The American Economic Review*, 87(2), pp. 78–83. URL https://www.jstor.org/stable/2950888.

Case, Karl E and Robert J Shiller (1989). “The efficiency of the market for single-family homes.” *The American Economic Review*, 79(1), p. 125. URL https://www.jstor.org/stable/1804778.

Choi, Jung Hyun and Caitlin Young (2020). “Owners and renters of 6.2 million units in small buildings are particularly vulnerable during the pandemic.” *Urban Wire*. URL https://www.urban.org/urban-wire/owners-and-renters-62-million-units-small-buildings-are-particularly-vulnerable-during-pandemic.

Clark, Joshua (2022). “Methodology: Zillow Observed rent Index (ZORI).” https://www.zillow.com/research/methodology-zori-repeat-rent-27092/. Accessed: 2022-11-11.

Crone, Theodore M, Leonard I Nakamura, and Richard Voith (2010). “Rents have been rising, not falling, in the postwar period.” *The Review of Economics and Statistics*, 92(3), pp. 628–642. doi:10.1162/REST_a_00015.

Gallin, Joshua and Randal J. Verbrugge (2007). "Improving the CPI's age-bias adjustment: Leverage, disaggregation and model averaging." Working Paper 411, US Bureau of Labor Statistics. URL https://www.bls.gov/osmr/research-papers/2007/pdf/ec070100.pdf.

Gallin, Joshua and Randal J. Verbrugge (2016). "Panel data estimates of age-rent profiles for rental housing." Working Paper 16-30, Federal Reserve Bank of Cleveland. URL https://www.clevelandfed.org/en/newsroom-and-events/publications/working-papers/2016-working-papers/wp-1630-age-rent-profiles.aspx.

Gallin, Joshua and Randal J. Verbrugge (2017). "Sticky rents." Unpublished Manuscript.

Ganong, Peter and Daniel Shoag (2017). "Why has regional income convergence in the US declined?" *Journal of Urban Economics*, 102, pp. 76–90. doi:10.1016/j.jue.2017.07.002.

Gelain, Paolo and Simone Manganelli (2020). "Monetary policy with judgment." Working Paper 2404, European Central Bank. URL https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2404~bce81d28a5.en.pdf.

Goetzmann, William Nelson (1992). "The accuracy of real estate indices: Repeat sale estimators." *The Journal of Real Estate Finance and Economics*, 5(1), pp. 5–53. doi:10.1007/BF00153997.

Hill, Robert J., Miriam Steurer, and Sofie R. Waltl (2020). "Owner occupied housing, inflation and monetary policy." Working Paper 2020-18, University of Graz. URL http://www100.uni-graz.at/vwlwww/forschung/RePEc/wpaper/2020-18.pdf.

La'O, Jennifer and Alireza Tahbaz-Salehi (2022). "Optimal monetary policy in production networks." *Econometrica*, 90(3), pp. 1295–1336. doi:10.3982/ECTA18627.

Lebow, David E and Jeremy B Rudd (2003). "Measurement error in the Consumer Price Index: Where do we stand?" *Journal of Economic Literature*, 41(1), pp. 159–201. doi:10.1257/002205103321544729.

Molloy, Raven, Christopher L Smith, and Abigail Wozniak (2011). "Internal migration in the United States." *Journal of Economic Perspectives*, 25(3), pp. 173–96. doi:10.1257/jep.25.3.173.

Moulton, Brent R (2018). "The measurement of output, prices, and productivity." Technical report, Hutchins Center on Fiscal & Monetary Policy at Brookings, Washington, DC. URL https://www.brookings.edu/research/the-measurement-of-output-prices-and-productivity/.

Nothaft, Frank (2018). "Single-family rental: Measuring rent growth and return." *Realtor University: The Journal of the Center for Real Estate Studies*, 6(1), pp. 5–14.

Paben, Steven P., William H. Johnson, and John F. Schilp (2016). "The 2018 revision of the consumer price index geographic sample." *Monthly Labor Review*. URL https://www.bls.gov/opub/mlr/2016/article/the-2018-revision-of-the-CPI-geographic-sample.htm.

Sommers, J.P. and J.D. Rivers (1983). "Vacancy imputation methodology for rents in the CPI." In *Proceedings of the American Statistical Association, Business and Economic Statistics Section*. American Statistical Association, Alexandria, VA.

Verbrugge, Randal and Robert Poole (2010). "Explaining the rent–OER inflation divergence, 1999–2007." *Real Estate Economics*, 38(4), pp. 633–657. doi:10.1111/j.1540-6229.2010.00278.x.

	All rental units		Units with new tenants		
	AHS	BLS	AHS	BLS	MLS
Rent (2015 $)	1,043	1,046	1,017	1,043	1,739
Years Between Obs.	2.8	0.7	4.1	1.4	2.5
Characteristics					
Year Built	1967	1973	1970	1977	1981
Rooms (#)	5	4	4	4	6
Bedrooms (#)	2	2	2	2	3
Bathrooms (#)	2	1	2	2	2
Air Conditioning					
Central (%)	59.3	57.1	64.9	64.0	76.2
Other or None (%)	40.7	42.9	35.1	36.0	18.7
Property Type					
Detached (%)	28.1	21.9	25.5	15.0	46.5
Semidetached (%)	8.8	15.1	8.6	14.0	28.0

Table 1. Summary Statistics. *Note:* Values from the AHS are from the 2015, 2017, and 2019 surveys. Values from MLS are from 2015 onwards. Values for units in the BLS may occasionally be missing due to non-response of the residents, so the time between observations is greater than 0.5 years. *Source:* BLS Housing Survey, AHS, MLS.

	vecout				
	NTRR-CoreLogic	ATRR-CoreLogic	NTRR-CPI Rent	CPI Rent-CoreLogic	NTRR-ATRR
Cointegrating Equation					
β_1	1.000	1.000	1.000	1.000	1.000
β_2	-0.974	-1.007	-0.741	-1.284	-0.995
(Std Error)	(0.063)	(0.078)	(0.056)	(0.115)	(0.049)
Speed of Adjustment					
α_1	-0.410	-0.140	-0.208	0.054	-0.360
(Std Error)	(0.020)	(0.003)	(0.014)	(0.000)	(0.027)
α_2	-0.059	-0.008	0.078	0.037	0.103
(Std Error)	(0.009)	(0.007)	(0.001)	(0.003)	(0.003)

Table 2. Pairwise Vector Error Correction Results

Figure 1

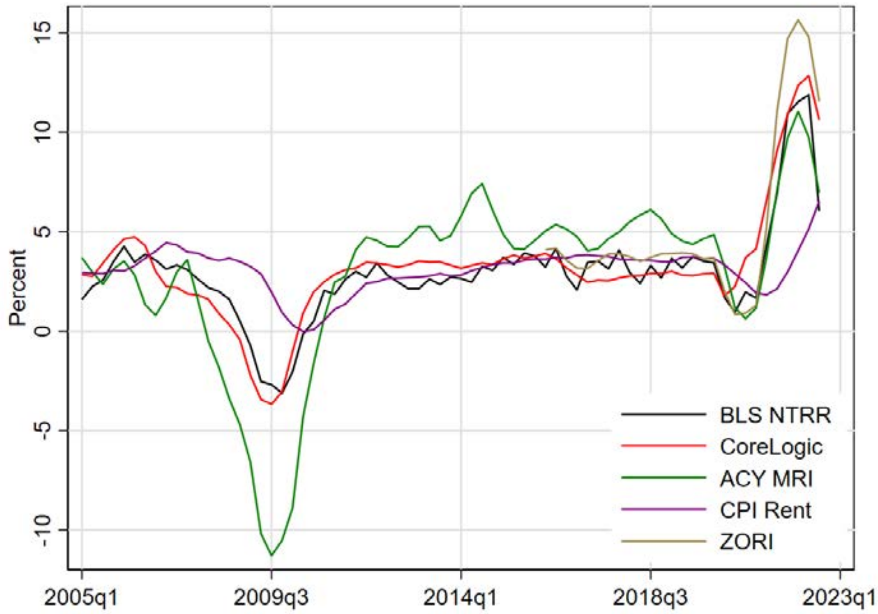

(a) VARIOUS MEASURES OF RENT INFLATION. *Note:* CPI rent is rent of primary residence. The construction of NTRR is described in Section 4. *Source:* BLS Housing Survey, Corelogic SFRI, Zillow (ZORI) and Ambrose et al. (2022) for the ACY MRI.

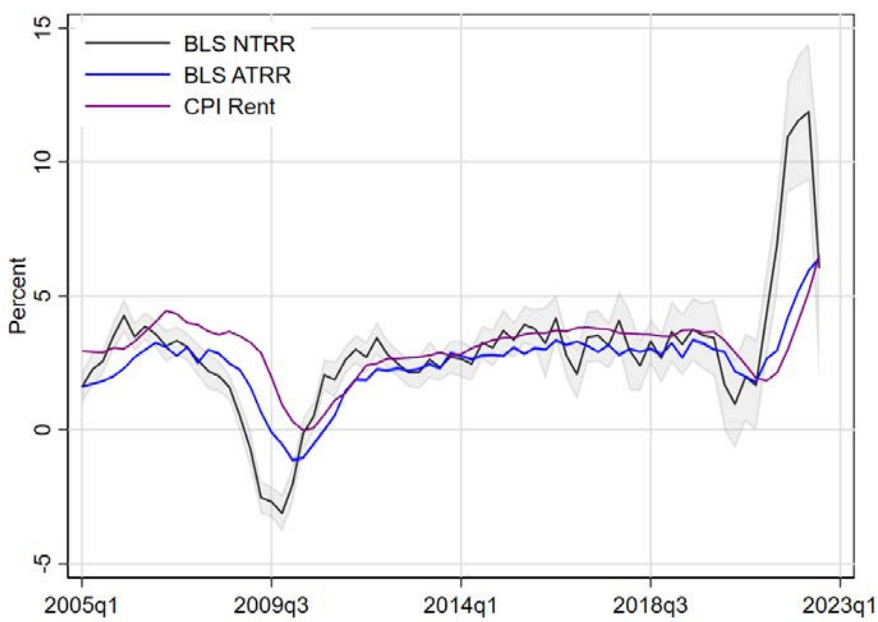

(b) NTRR VERUS ATRR. *Note:* CPI rent is rent of primary residence. The construction of ATRR and NTRR are described in Section 4. *Source:* BLS Housing Survey.

Figure 2

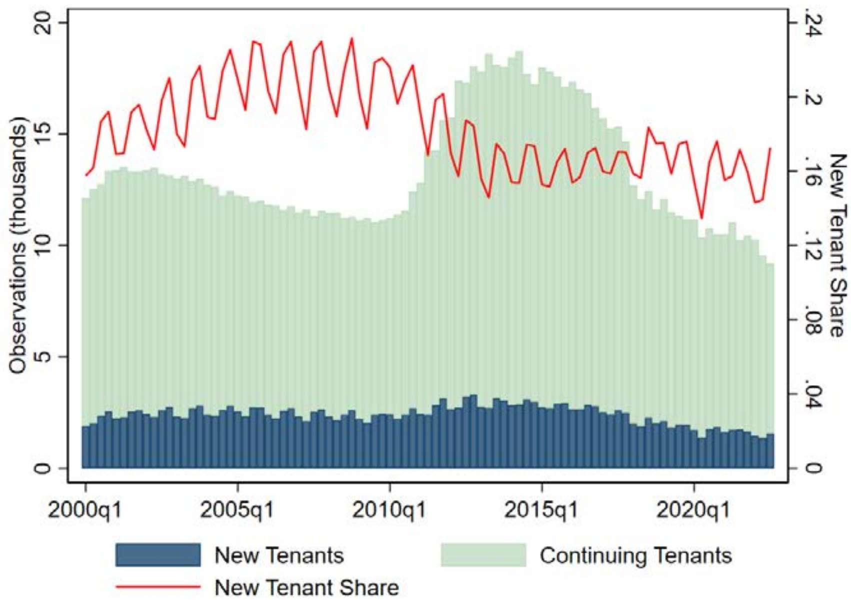

(a) Observations in the BLS Housing Survey. *Note:* The bars show the total number of housing-unit observations with non-missing rent in each quarter, broken down by those for new tenants (in blue) and continuing tenants (in green). The red line plots the share attributed to new tenants over time. *Source:* BLS Housing Survey.

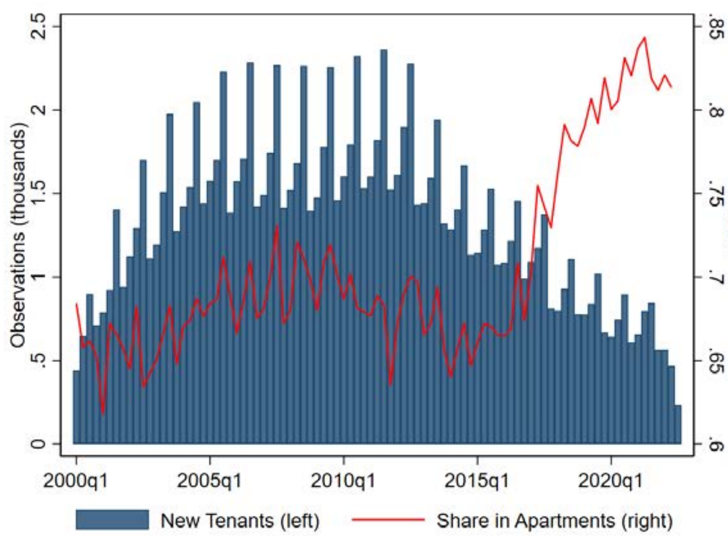

(b) Number of Observations Used in Construction of New Tenant Repeat-Rent Index. *Note:* Observations in the NTRR require two new-tenant observations on a housing unit. The x-axis is the date of the second new-tenant observation in each repeat pair. *Source:* BLS Housing Survey.

Figure 3

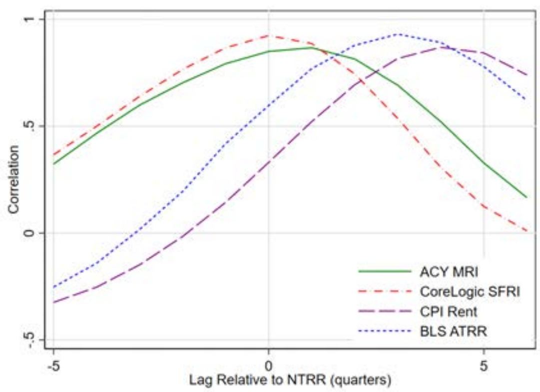

(a) Lagged Correlation with the NTRR. *Note:* Values are cross-correlations between indices at various lag lengths. *Source:* Authors' calculations on data from BLS, Corelogic, Zillow and Ambrose et al. (2022).

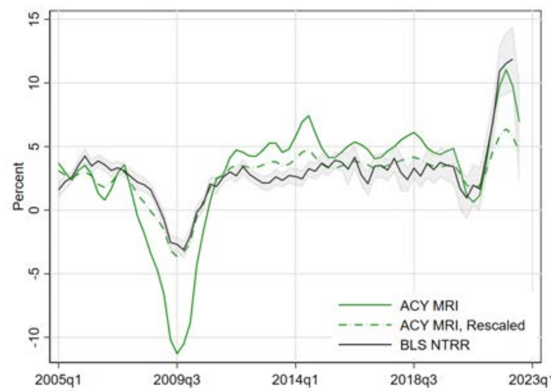

(b) Rescaled Version of Marginal Rent Index from Ambrose et al. (2022). *Note:* The NTRR is plotted along with 95 percent confidence intervals. *Source:* BLS and Ambrose et al. (2022).

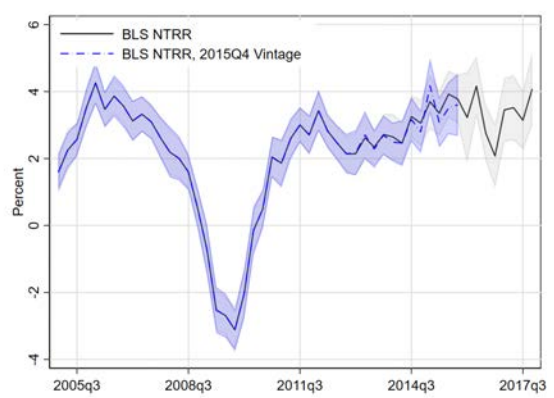

(c) Comparing Repeat-Rent Indices Using Different Data Vintages. *Note:* Both series are plotted with their respective 95 percent confidence intervals. *Source:* BLS.