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. 22-38R

September 2023

Suggested citation: Adams, Brian, Lara P. Loewenstein, Hugh Montag, and Randal J. Verbrugge. 2023. "Disentangling Rent Index Differences: Data, Methods, and Scope." Working Paper No. 22-38R. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202238r.

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

September 25, 2023

Rent measurement determines 32 percent of the CPI. Accurate rent measurement is therefore essential for accurate inflation measurement, but the CPI rent index often differs from alternative measures of rent inflation. Using repeat-rent inflation measures created from CPI microdata, we show that this discrepancy is largely explained by differences in rent growth for new tenants relative to all tenants. New-tenant rent inflation provides information about future all-tenant rent inflation, but the use of new-tenant rents is contraindicated in a cost-of-living index such as the CPI. Nevertheless, policymakers should integrate new-tenant inflation into inflation forecasts and monetary policy decisions.

Keywords: house prices, rent growth, inflation measurement, monetary policy, forecasting, marginal rent growth, average rent growth
JEL Classification: E31, E37, E27, H31

*We thank Rob Cage, Robert Poole, Thesia Garner, Jeff Hill, Paul Liegey, and Jeff Medlar for their help in understanding CPI weights; audiences at the 2022 North American Meeting of the Urban Economics Association and the 2023 Federal Economic Statistics Advisory Committee meeting; Kentaro Nakajima, Alan Detmeister, and Fiona Greig for their thoughtful discussions; and audiences at presentations at the BLS, the Federal Reserve Bank of Cleveland, and the Federal Reserve Board (with a special thanks for insightful comments at the latter from Ekaterina Peneva, Mark J. Bognanni, and Daniel Villar). Last, but very much not least, we extend a big thank you to Paolo Gelain for providing IRF estimates from his DSGE model. 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

The CPI is one of the most important aggregate economic indicators. It is the basis of Social Security cost-of-living adjustments, it is used in financial contracts, and its movements contribute to the personal consumption expenditures price index (PCE) that forms the inflation objective of the Federal Open Market Committee (FOMC). The biggest component in the CPI is shelter. Rental housing has 8 percent of the relative importance in the CPI. Owner-occupied housing has an additional 24 percent. Because it is measured using the owners' equivalent rent (OER) method, rents ultimately drive about 32 percent of the CPI.[1] Accurate measurement of inflation therefore depends critically on accurate measurement of rent inflation.

Given the importance of rent inflation, the large differences in measures of rent inflation are concerning. For example, in 2022q1 inflation rates in the Zillow Observed Rent Index (ZORI)[2] and the Marginal Rent Index (ACY MRI)[3] reached an annualized 15 percent and 12 percent, respectively, while the official CPI for rent was 5.5 percent (see Figure 1a). If the Zillow reading were to replace the official CPI rent and OER measures, then the May 2022 12-month all-items CPI reading of 8.6 percent would have been more than 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 decisions, contract escalation, and GDP and welfare measurement (Ambrose et al. 2018; Hill et al. 2020; Ambrose et al. 2022). For instance, had ZORI replaced the official CPI rent and OER measures, then real wage estimates would have fallen by nearly 5 percent between January 2021 and January 2022 (as opposed to 1.5 percent) and Social Security cost-of-living increases would have been $17 billion higher in 2022.

We study the possible sources of divergence between CPI rent measures and three prominent alternatives: the ZORI, the MRI, and the CoreLogic Single Family Rent Index (SFRI). These divergences might stem from differences in sample representativity (only the CPI is representative), index construction methods (ZORI and the SFRI are repeat-rent indices, the CPI is based upon six-month changes in rents, and the ACY MRI is the product of two aggregate indexes), quality adjustments (only the CPI does this), or the scope of the underlying rental data (the CPI measures rent growth facing all occupants, while the alternative measures track rent growth facing new tenants).

We use the microdata underlying official measures of CPI shelter inflation to create two weighted repeat-rent indices in the style of Case and Shiller (1989): the new-tenant repeat-

[1]By comparison, food accounts for approximately 14 percent of the CPI. In the PCE, shelter represents over 15 percent, and food about 13 percent.

[2]See Clark (2022).

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

rent (NTRR) index (using only leases of tenants who recently moved in), and the all-tenant repeat-rent (ATRR) index (using all tenants, whether they recently moved in or not). We find that most of the discrepancy between CPI rent and other measures is due to scope: to differences in rent increases for *all* tenants versus *new* tenants (See Figure 1b). In 2022q2, our ATRR index was recording 6.73 percent year-over-year inflation, while the NTRR inflation rate was at 11.95 percent. CPI rent inflation was at 5.14 percent. Perhaps surprisingly, the SFRI, despite its non-representative nature, is generally a fair approximation to our NTRR index over our sample period.

Should the CPI change the way it measures rents? Some economists, both in research papers and policy commentary,[4] have criticized the all-tenant approach, arguing that the CPI should track new-tenant rents to better capture current market conditions. However, the CPI is designed to be a cost-of-living index and is therefore meant to measure the change in prices needed to keep welfare, or utility, constant for a typical household. For that reason, the price statistics literature has long argued that one should use a rent measure capturing the rent movements facing a typical renter. We review this argument below, then contribute to this debate by clarifying the difference that the use of a new-tenant index would make by using the same data source that underlies the CPI and noting some practical impediments.

Last, we demonstrate the first-order importance of the method of rent measurement for macroeconomic modeling and discuss the implications for monetary policymakers. We show that new-tenant rent indexes provide an earlier read on changes in rental market conditions. We then construct an alternative core PCE index, replacing its shelter indexes with the SFRI, and demonstrate that this replacement has a first-order impact on the behavior of macroeconomic models and on the nature and strength of the Phillips curve. While our results are a first step, we hope future research will further explore the implication of policymakers targeting new-tenant rent inflation.

1 THE BLS HOUSING SURVEY

We use the BLS Housing Survey data, the same data used to compute the CPI rent index. The BLS Housing Survey follows a sample of around 40,000 renter-occupied housing units. It surveys the same units every six months, recording the rent, the utilities and services included with the rent, the tenant's move-in date, and other unit characteristics. The survey's multistage sampling design, described in Appendix Section A.1, aims to create a sample representative of rental expenditure. Table 1 shows 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.

[4]For example, see Paul Krugman's January 27, 2023, opinion column in the *New York Times*: "Wonking Out: Inflation and the Imputation Game," as well as Ambrose et al. (2015, 2022).

Limiting the BLS Housing Survey sample to new renters reduces its size. As depicted in Figure 2a, every quarter we have observations on between 11 and 17 thousand housing units. The share of these observations reflecting tenants that moved in over the last six months fluctuates between 13 and 25 percent. This share is higher in summer and fall, reflecting the seasonal pattern of moving. It fell during the Great Recession, but has remained stable since 2012.

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. With less time in the sample, fewer units give two new-tenant observations after sample rotation begins. The sample rotation also alters the distribution of property types. 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 rents, although this share does not increase until 2016.

2 CONSTRUCTING A BLS HOUSING SURVEY 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

$$\ln P_{it} - \ln P_{is} = \gamma_1 D_{i1} + ... + \gamma_N D_{iN} + u_{it}, \tag{1}$$

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 every other period $D_{ij} = 0$. For our example observation, $D_{it} = 1$ and $D_{is} = -1$. 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.

We construct two indices. The first is the NTRR; 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 the ATRR index. It 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.[5] 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.

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. 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 coincides with the coverage of other rent growth measures.

We seek to identify the month in which the rent change happens. The official CPI dates observations to their survey collection month. A rent change in a unit may happen in the months between surveys, and therefore several months before the survey collection period. Accordingly, when constructing the NTRR and ATRR, we date observations either to their recorded move-in date or to the completion of the most recent six-month interval since move-in.[6] 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.[7]

We use economic rent as our measure of rent. This is the measure used in the official CPI rent index. Economic rent adjusts the contract rent to account for services rendered in lieu of rent, for 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 adjusting for structural changes, we 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 for which a field note includes the words "remodel," "renovate," or "refurbish."

With regard to vacancy adjustments, as the last observation for every unit in our NTRR

[5]Another approach would limit observations to new lease signings to remove repeat observations within the same lease contract (and therefore with zero rent growth). However, our goal with the ATRR is to create a rent index as similar as possible to the CPI rent index to isolate any changes due to the index construction methodology, so we include all observations. The data also do not indicate when a new lease is signed.

[6]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 six-month multiple of the move-in date.

[7]Using imputed dates of rent changes is also necessary to create a single ATRR and NTRR index. If we used survey dates we would need to create six separate indices, one for each panel of units that is surveyed in a given month.

is the final observed rent change, it does not suffer from the vacancy bias described in Sommers and Rivers (1983). In our ATRR, instead of using the CPI's vacancy adjustment, we account for vacancy bias by excluding all observations after the last date in which a new tenant moves in unless the housing unit has been sampled within the previous six months.[8]

We also drop any observations in the top and bottom 1 percent of annualized rent changes. To further mitigate volatility, both the NTRR and the ATRR are quarterly (instead of monthly) indices.

We replicate the weighting used in the official CPI tenant rent index as closely as possible. First, we estimate equation 1 to create index-area-specific rent indices. As described in Section A.1, each index area is composed of a self-representing primary sampling unit (PSU) or multiple non-self-representing PSUs, which in turn are made up of segments. For the ATRR index, we use segment weights (the inverse of the probability of selecting that segment within a given index area) to estimate equation 1 for each index area. We adjust the official segment weights to account for the number of observations that we omit due to remodelling, vacancy, and outlier adjustments as well as survey non-responses to maintain representativity within the index area. CPI segment weights are available for each collection period, and we weight each repeat-rent observation using the segment weight from the more recent collection date. The NTRR index-area-specific indices use uniform lower weighting, as the segment weights do not correspond to the universe of new-tenant units.

Differences in tenure lengths may make the error term in equation 1 heteroskedastic, particularly for our NTRR index. Case and Shiller (1989) and Goetzmann (1992) propose a three-stage procedure to address it. When creating our index-area-level indices, we estimate equation 1 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 1. Like Clark (2022), we found that this heteroskedasticity correction had a negligible effect on our final index.

We combine the index-area-level quarter-on-quarter changes using a second set of "upper" weights, which represent the share of national rent expenditure accounted for by that index area. Upper weights are updated every two years for the CPI.[9] Once we have a national measure of quarter-on-quater inflation, we can back out a national rent index.

[8] 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 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. 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.

[9] The geographic areas included in the BLS Housing Survey changed in 2016; some non-self-representative index areas 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 index areas to construct our ATRR and NTRR.

We calculate rent inflation by taking the annual log difference in our national index values. Standard errors for our inflation estimates are created using a bootstrap method described in Appendix D.

3 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 the ATRR differ slightly, they generally track each other. By contrast, the 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.

3.1 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.[10] In this sense they should be similar. However, besides their construction methodology, they may differ for several reasons:

1. CPI rent includes a vacancy adjustment for 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 2.
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 2.

Figure 1b compares the annual rent inflation implied by the ATRR (in blue) and that based on CPI rent (in purple). The ATRR is lower than CPI rent on average, especially early in the sample, although starting in 2021 and most of 2022 the ATRR is above CPI rent inflation. The two generally track each other. Their differences are sizeable by the standards of precision for official statistics, but tiny in comparison to the difference between CPI rent and rent indices based on new tenants.

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

[10] Details on the construction of the CPI tenant rent index are included in Section A.1 in the Appendix.

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.[11]

3.2 NTRR versus ATRR

We next compare our NTRR and ATRR; see Figure 1b. 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 2009, whereas the ATRR reaches its trough later in 2010q1. Likewise, the current price spike begins in the NTRR well before it begins in the ATRR. The correlation in quarterly changes is only .65, but the correlation of the quarterly changes in the ATRR with a three-quarter lag of the NTRR is .93. As Figure 3a shows, the NTRR leads the ATRR by about three quarters, while it leads CPI rent by about four quarters.

The NTRR is much more volatile. The deflation 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 than the ATRR. This noise is 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 0.44 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).

3.3 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 1.80 percentage points and have an average absolute difference of 0.74 percentage points. Quarterly changes in the two indices are highly correlated ($\rho =$ .92), 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 methodology described in Section 1), scope (new tenant versus all tenants), and/or rent adjustments (the

[11] Robust estimates generally support the conclusions drawn from the classically estimated correlograms, though the correlations with the ATRR and the ACY MRI peak slightly earlier.

BLS performs quality adjustments that the SFRI does not). 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.[12] Comparing the ATRR and the official BLS rent index shows that the repeat-transaction methodology and adjustments have noticeable effects, but are not the main drivers of differences with other indices.

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 ($\rho = 0.92$).

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 A.4), 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 the NTRR. The minimization problem is

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

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 0.462 and a is 1.51. The resulting (rescaled) index, which is depicted in Figure 3b, matches the dynamics of the NTRR fairly well.

[12] 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 C 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 B.2 in the Appendix for details).

4 IMPLICATIONS FOR MEASURING RENT INFLATION IN THE CPI

As noted in the introduction, some authors have argued for the use of new-tenant rents in the CPI, since these rents better capture current market conditions. But the CPI is designed to be a cost-of-living index, capturing changes in the purchasing power of a dollar for a typical household. For this reason, the price statistics literature generally favors the use of an all-tenant index as it capture changes in the purchasing power of typical renters. Only households that move every time period would experience the rent changes reflected in a new-tenant index. Similarly, OER is based upon an opportunity cost concept: by deciding to live in my own home, what am I giving up in terms of forgone rental payments? But only owners finding new tenants every period would obtain a rental payment flow matching the rent changes in a new-tenant index.

Beyond the theoretical reasons for not using new-tenant rent inflation in the CPI, there are also practical difficulties with using a new-tenant rent index. We highlight two.

First, instability at the end of the sample hinders the use of new-tenant rent measures 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 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 observation pairs 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. 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 data. 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 with smaller samples.

The CPI currently limits ex-post revisions to seasonal adjustment, except when serious errors are detected. The CPI is used in Social Security benefit escalation, in the indexing of wage contracts, and in other uses in which revisions would be unwelcome. Not only is an all-tenant index preferred in this context to reflect what is happening to the typical household, but also, data revisions — an inherent feature of repeat-rent measures — are problematic.

The second reason is that, at present, the CPI rent sample is far too small to support accurate rent inflation measurement below the national level. And one could not simply use the SFRI in the CPI: while the SFRI has historically tracked the NTRR well, the series diverged notably at the start of the pandemic. Furthermore, since the SFRI sample is non-representative and dependent upon MLS listings, there is no guarantee that the SFRI will track the NTRR in the future.

Thus the NTRR is not suitable for real-time use in a CPI. That said, because new-tenant indices more quickly reflect shifts in market conditions, their signal is clearly useful for agents who must forecast future inflation, and in particular for monetary policymakers. In that regard, not only are new-tenant indices useful for predicting CPI shelter indices (see Appendix E), but new-tenant rents may also provide a better signal of the responsiveness of rents to policy changes. New-tenant rent indices are also useful for other purposes; for example, they are more appropriate for comparisons with homeowner marginal user costs. And they may prove useful in unofficial aggregate price indices as well: since the intended purposes of a price index should guide its construction, researchers may well decide that constructing a different price index may be useful for addressing the particular questions that they seek to answer.

5 IMPLICATIONS OF ALTERNATIVE RENT MEASURES FOR MACROECONOMIC MODELING

Should monetary policymakers target a price index that includes a measure such as the NTRR? To begin to answer this question, we next investigate the importance of using alternative rent measures in aggregate price measures for statistical inferences in Phillips curve estimation and in New Keynesian dynamic stochastic general equilibrium (DSGE) models.

For this purpose, we construct an alternative price index, the "core SFRI" PCE price index. The PCE price index, rather than the CPI, dominates monetary policy decision-making. The inflation target is specified in terms of the headline PCE index, and core PCE (that is, PCE-less-food-and-energy) plays a focal role in monetary policy deliberations, since it is thought to be a good signal of trend movements in the headline PCE index. Accordingly, Phillips curve models and DSGE models that are focused on monetary policy (especially those within the Federal Reserve System) are often specified in terms of core PCE.

The PCE shelter inflation indices (that is, rent, and owners' equivalent rent (OER)) are driven by the corresponding CPI shelter indices. The main difference between the CPI OER index and the CPI rent index is that different weights are used in equation (A.1); both are driven by rent movements in the CPI rental sample. However, use of a new-tenant rent index in the PCE price index would likely yield different dynamics. For illustrative purposes, we construct a price index in which all rent and OER categories in the core PCE are replaced by the CoreLogic SFRI rent.[13]

We display the monthly core SFRI growth rate in Figure F.3 in the Appendix, along with the growth rate of core PCE. While the core SFRI is more volatile, in most months the difference between the inflation rates is modest, and usually those differences are not persistent. The median difference is a mere 0.08 percentage points. But differences can be much larger; for example, from January to May 2009, and from April 2021 to April 2022, differences exceed a full percentage point. Despite what look to be modest differences between these alternative indices, Phillips curve parameter estimates and estimated DSGE impulse response functions turn out to be quite sensitive to these differences.

First, we specify a Phillips curve, following Ashley and Verbrugge (2023), who demonstrate that the Phillips curve is persistence-dependent (fluctuations in inflation respond differently to persistent fluctuations in unemployment, versus less-persistent fluctuations in unemployment), and that these relationships are asymmetric. The specification is stable in the 1985-2019 data, and performs well in out-of-sample forecast tests.[14] Furthermore, the specification resolves numerous inflation puzzles (such as "missing disinflation"), and explains why the Phillips curve is thought to have recently steepened.

Let π_t^{core} denote monthly inflation in the 12-month core PCE index, $\pi_t^{coreSFRI}$ denote monthly inflation in the 12-month core SFRI index, $u_t^{neg.lowgap}$ denote the negative portion of persistent fluctuations in the unemployment rate gap, and $u_t^{pos.medgap}$ denote the positive portion of the moderately persistent fluctuations in the unemployment rate gap. Below we display our specification. Immediately below the two Phillips curve coefficients ϕ^{low} and ϕ^{med}, we display the corresponding coefficient (and standard error) estimates.

$$\pi_t^{core} = \alpha + \sum_{k=1}^{3} \beta_k \pi_{t-k}^{core} + \underbrace{\phi^{low}}_{\substack{-0.06 \\ (0.02)}} u_{t-1}^{neg.lowgap} + \underbrace{\phi^{med}}_{\substack{-0.11 \\ (0.08)}} u_{t-2}^{pos.medgap} + \epsilon_t \tag{3}$$

$$\pi_t^{coreSFRI} = \alpha + \sum_{k=1}^{3} \beta_k \pi_{t-k}^{coreSFRI} + \underbrace{\phi^{low}}_{\substack{-0.02 \\ (0.02)}} u_{t-1}^{neg.lowgap} + \underbrace{\phi^{med}}_{\substack{-0.26 \\ (0.12)}} u_{t-2}^{pos.medgap} + \epsilon_t \tag{4}$$

[13]Montly SFRI growth rates are seasonally adjusted. We are not the first to investigate the importance of alternative rent measures for inflation measurement and its many consequences. For instance, Ambrose et al. (2022) investigate the use of the MRI as the measure of rent and OER in the CPI. Differences are quite stark. However, the CPI places a greater weight on housing consumption than does the PCE, so it is of interest to investigate the impact of alternative rent measures in the core PCE index.

[14]Indeed, Verbrugge and Zaman (2023) construct a quarterly modification of the specification and demonstrate that it predicts the post-COVID run-up in inflation.

Over the 2006-2019 estimation period,[15] core PCE inflation has a weak Phillips curve: it responds weakly to "overheating" forces (that is, when the highly persistent gap is below zero), and not at all to "recessionary" forces (that is, when the moderately persistent gap is above zero). Conversely, core SFRI inflation has a robust Phillips curve, with a different character: it responds strongly to recessionary forces, but is unresponsive to overheating.

Next, we investigate the extent to which DSGE model parameter estimates are sensitive to the rent measure used in the core PCE. In particular, we estimate the DSGE model in Gelain and Manganelli (2020) using both core PCE and core SFRI PCE and examine the resulting estimated impulse response function (IRF) error bands to the three structural shocks in the model.[16] These are plotted in Figure 4.

Three observations stand out. First, the return of the core SFRI to its steady state following a structural shock is noticeably more rapid. Second, the error bands around the core SFRI IRFs are generally tighter. And third, the inflation IRFs are statistically distinct; in response to all three structural shocks, there is disjointness of the error bands, either up to five months, or after five months. For instance, in response to a monetary policy shock, core SFRI returns much more rapidly to the steady state, with disjoint IRF error bands after three months. In keeping with this, the policy rate and real GDP growth return much more rapidly to trend, with disjointness of IRF error bands after three months and five months, respectively.

In short, using an alternative rental series in the core PCE leads to materially different inferences about the very existence of a Phillips curve and about the responses of inflation (and other variables) to structural shocks in a DSGE model.

The results above suggest that the question about the particular price index that monetary policymakers should target is an important one. But —leaving aside the practical issues noted in the previous section —definitively answering this question will be challenging. To our knowledge, research studying the optimal choice of a price index for policymakers has limited itself to asking about weights attributed to existing sub-indices (for example, La'O and Tahbaz-Salehi (2022)), not about choosing between alternative sub-indices that have very different properties. Moreover, a DSGE model studying the optimal price index construction for policymakers would require a model that features both renters and owners facing aggregate risk, such as in Sommer et al. (2013)), and recognizes that different agents in the model (policymakers versus households) would use different aggregate price indices. And last, it would need to recognize that household welfare losses would need to be computed using the price index most relevant for households.

[15] We estimate on pre-pandemic data to avoid overfitting based on one extreme episode.

[16] IRFs are nonlinear functions of the model parameter estimates. These plot the dynamic response of the model's endogenous variables to a shock in an exogenous variable.

6 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. This 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. However, new-tenant indices do not belong in the CPI, and further research is required to understand the implications of their use in an inflation target.

References

Adams, Brian and Randal Verbrugge (2021). "Location, location, structure type: Rent divergence within neighborhoods." Working Paper 21-03, Federal Reserve Bank of Cleveland. doi:10.26509/frbc-wp-202103.

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

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. Publisher: Elsevier.

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 J Verbrugge (2023). "The Intermittent Phillips Curve: Finding a Stable (But Persistence-Dependent) Phillips Curve Model Specification." Working Paper 19-09R2. doi:10.26509/frbc-wp-201909r2. Institution: Federal Reserve Bank of Cleveland.

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, publisher: American Economic Association.

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)." URL https://www.zillow.com/research/methodology-zori-repeat-rent-27092/.

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. Publisher: MIT Press.

Foote, Christopher L, Lara Loewenstein, and Paul S Willen (2020). "Cross-sectional patterns of mortgage debt during the housing boom: Evidence and implications." *The Review of Economic Studies*, 88(1), pp. 229–259. doi:10.1093/restud/rdaa034.

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. doi:10.26509/frbc-wp-201630. 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). "A theory of sticky rents: search and bargaining with incomplete information." Working Paper 17-05, Federal Reserve Bank of Cleveland. doi:10.26509/frbc-wp-201705.

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.

La'O, Jennifer and Alireza Tahbaz-Salehi (2022). "Optimal monetary policy in production networks." *Econometrica : journal of the Econometric Society*, 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.

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/.

Paben, Steven P., William H. Johnson, and John F. Schilp (2016). "The 2018 revision of the Consumer Price Index geographic sample." *Monthly Labor Review*. doi:10.21916/mlr.2016.47. URL https://www.bls.gov/opub/mlr/2016/article/the-2018-revision-of-the-CPI-geographic-sample.htm.

Ptacek, Frank (2013). "Updating the rent sample for the CPI Housing Survey." *Monthly Labor Review*. doi:10.21916/mlr.2013.25. URL https://www.bls.gov/opub/mlr/2013/article/updating-the-rent-sample-for-the-cpi-housing-survey.htm, institution: US Bureau of Labor Statistics.

Sommer, Kamila, Paul Sullivan, and Randal Verbrugge (2013). "The equilibrium effect of fundamentals on house prices and rents." *Journal of Monetary Economics*, 60(7), pp. 854–870. doi:10.1016/j.jmoneco.2013.04.017. Publisher: Elsevier.

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.

Verbrugge, Randal J and Saeed Zaman (2023). "The Hard Road to a Soft Landing: Evidence from a (Modestly) Nonlinear Structural Model." Working Paper 23-03. doi:10.26509/frbc-wp-202303. Institution: Federal Reserve Bank of Cleveland.

Wolter, Kirk M. (2007). *Introduction to Variance Estimation*. Springer, New York, 2 edition. doi:10.1007/978-0-387-35099-8.

Table 1. Summary Statistics

	All rental units		Units with new tenants		
	AHS	BLS	AHS	BLS	MLS
Rent (2015 $)	1,043	1,081	1,017	1,101	1,742
Years Between Obs.	2.8	0.6	4.1	1.2	2.6
Annual Rent Growth (%)	2.7	3.8	2.6	4.1	2.4
Characteristics					
Year Built	1967	1975	1970	1978	1981
Rooms (#)	5	4	4	4	6
Bedrooms (#)	2	2	2	2	3
Bathrooms (#)	2	1	2	2	2
Air Conditioning					
Central (%)	59.3	57.3	64.9	64.3	77.1
Other or None (%)	40.7	42.7	35.1	35.7	18.7
Property Type					
Detached (%)	28.1	23.7	25.5	16.2	51.2
Semidetached (%)	8.8	15.8	8.6	14.6	25.3

Values from the AHS are from the 2015, 2017, and 2019 surveys. Values from the MLS and the the BLS Housing Survey are from 2015–2019. 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.

Figure 1

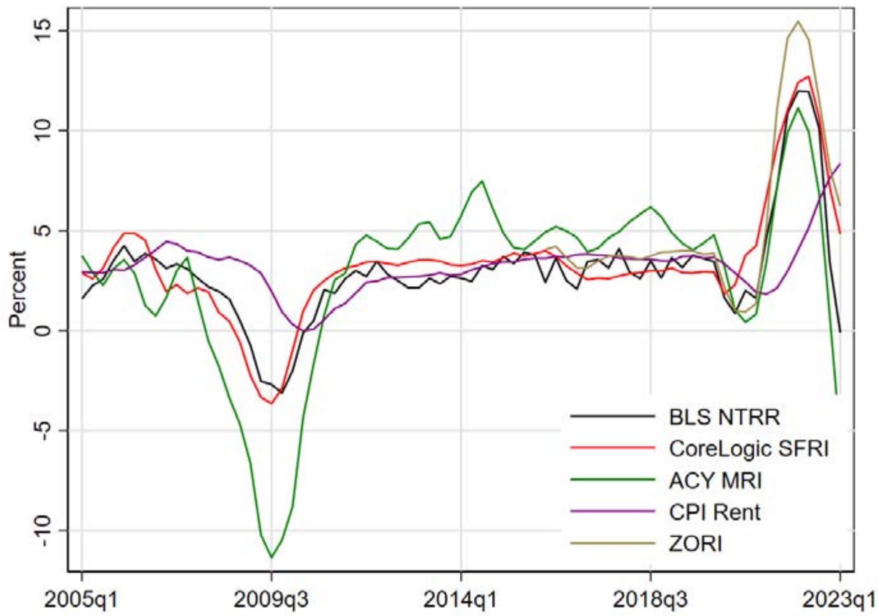

(a) Various Measures of Rent Inflation. *Note:* CPI rent is rent of primary residence. The construction of the NTRR is described in Section 2. *Source:* BLS Housing Survey, CoreLogic SFRI, Zillow (ZORI), and Ambrose et al. (2022) for the ACY MRI.

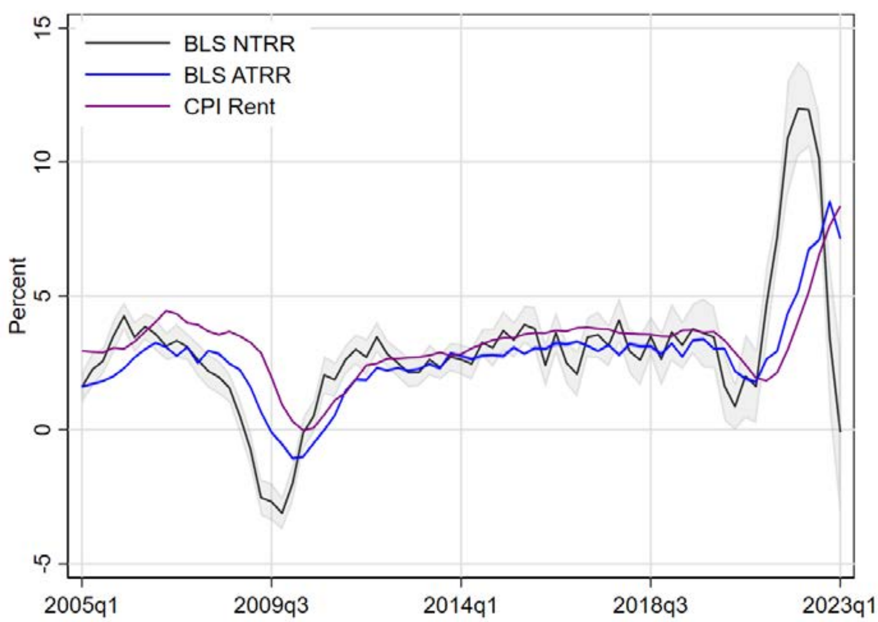

(b) NTRR verus ATRR. *Note:* CPI rent is rent of primary residence. The construction of the ATRR and the NTRR is described in Section 2. *Source:* BLS Housing Survey.

Figure 2

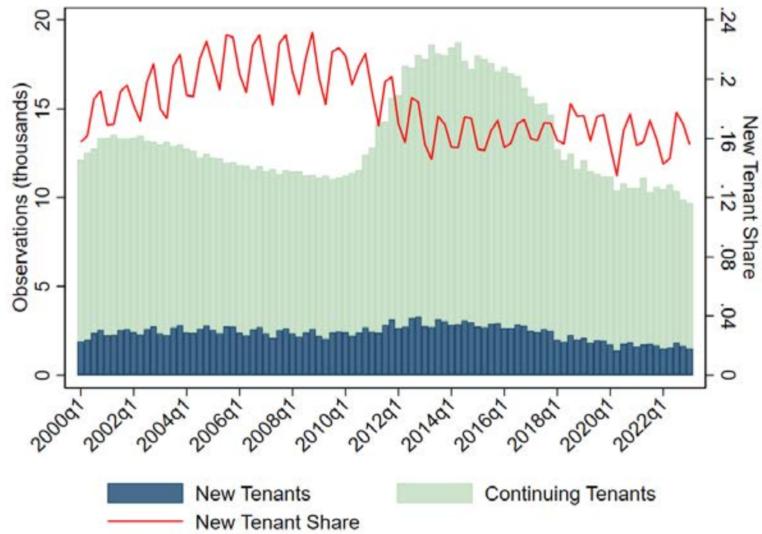

(a) Observations in the BLS Housing Survey. *Note:* The bars show the total number of housing-unit observations with non-missing rent in each quarter, broken down by those for new tenants (in blue) and continuing tenants (in green). The red line plots the share of observations in each quarter that got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

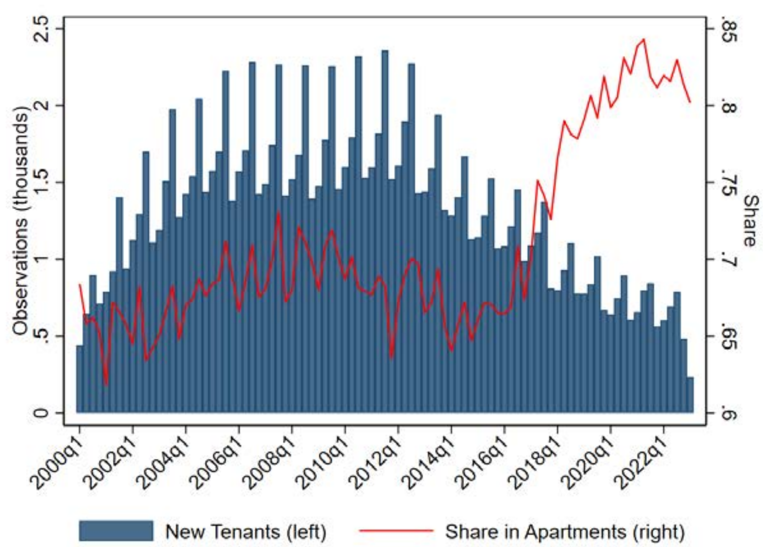

(b) Number of Observations Used in Construction of New-Tenant Repeat-Rent Index. *Note:* Observations in the NTRR require two new-tenant observations on a housing unit. The x-axis is the date of the second new-tenant observation in each repeat pair. *Source:* BLS Housing Survey.

Figure 3

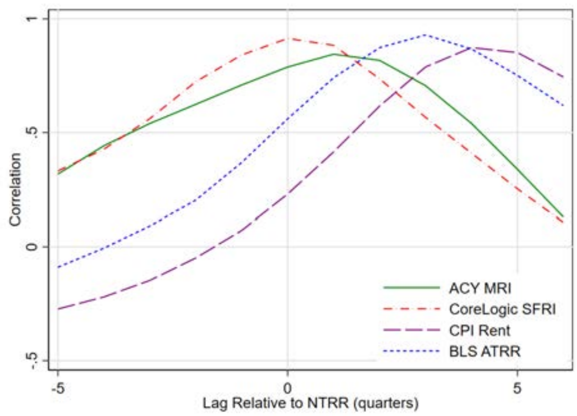

(a) Lagged Correlation with the NTRR. *Note:* Values are cross-correlations between indices at various lag lengths. *Source:* BLS Housing Survey, CoreLogic, and Ambrose et al. (2022).

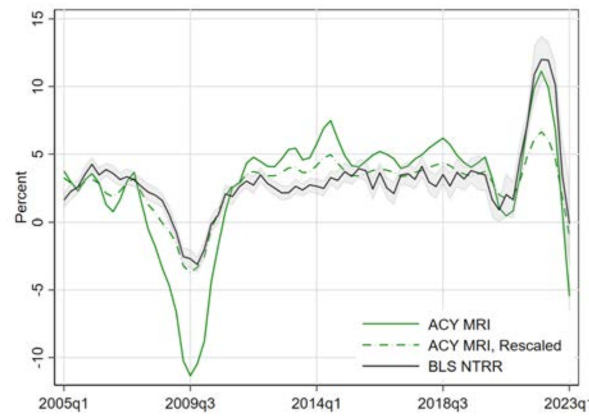

(b) Rescaled Version of Marginal Rent Index from Ambrose et al. (2022). *Note:* The NTRR has 95 percent confidence intervals. *Source:* BLS Housing Survey and Ambrose et al. (2022).

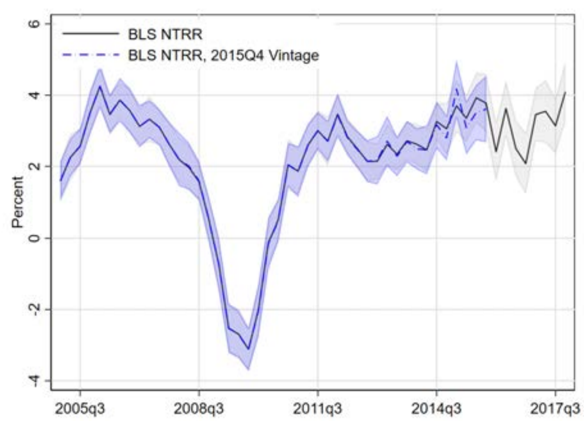

(c) Comparing Repeat-Rent Indices Using Different Data Vintages. *Note:* Both series are plotted with their respective 95 percent confidence intervals. *Source:* BLS.

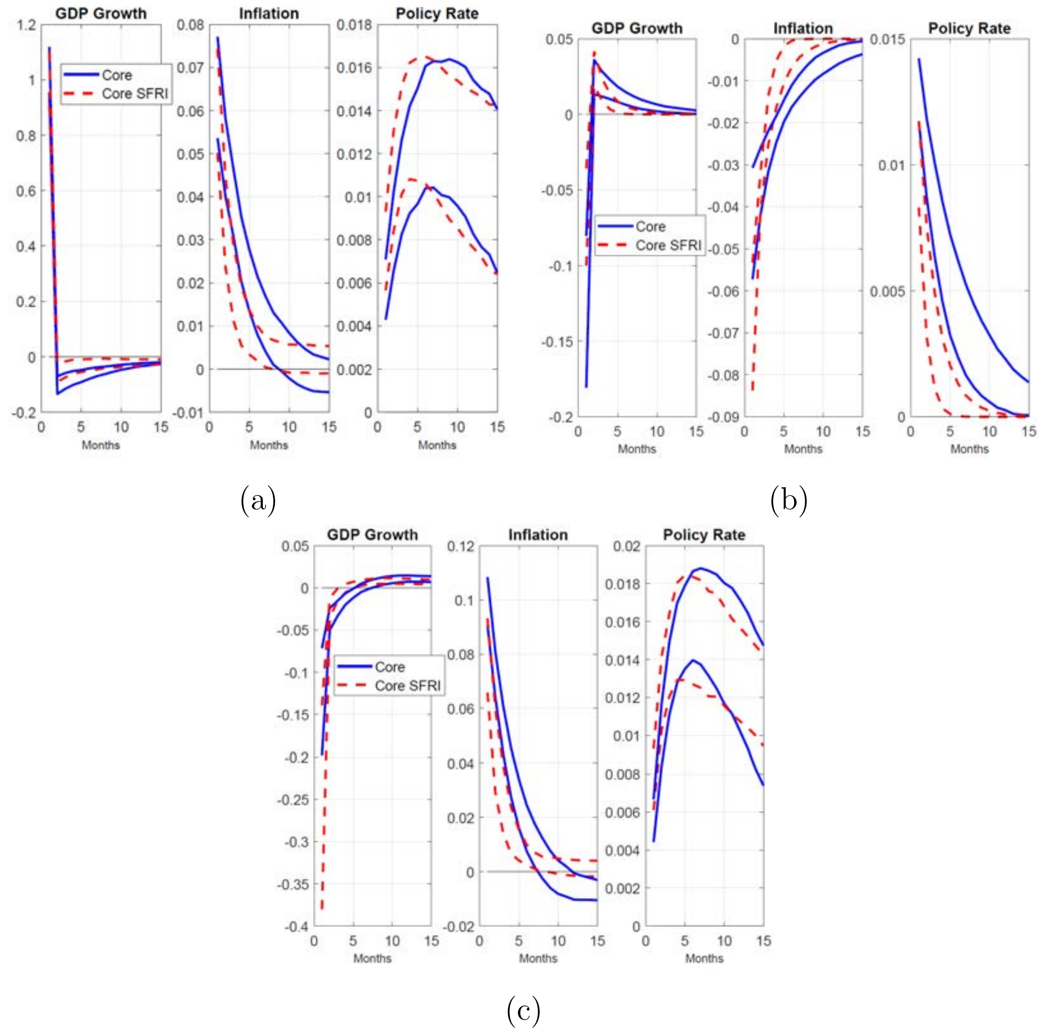

Figure 4. Impulse Response Functions in a New Keynesian Model. *Note:* IRFs display the dynamic response of endogenous variables to exogenous shocks in the model. They are nonlinear functions of model parameter estimates. We display only the IRF error bands, since our purpose is to assess whether IRFs for different inflation measures are statistically distinct. Panel (a) plots the IRF error bands to a government spending shock, panel (b) plots the IRF error bands to a monetary policy shock, and panel (c) plots the IRF error bands to a price markup shock. *Source:* Authors' calculations.

ONLINE APPENDIX

Disentangling Rent Index Differences: Data, Methods, and Scope

by Brian Adams, Lara Loewenstein, Hugh Montag, and Randal Verbrugge

A RENT DATA SOURCES AND INDICES

A.1 The BLS Housing Survey and the CPI Rent Index

The BLS Housing Survey uses a multistage sampling design meant to draw a sample representative of rental expenditure.[17] The first stage selects large geographic areas called "primary sampling units" (PSUs). PSU definitions now match metropolitan and micropolitan statistical areas. Before the BLS redesigned its geographic sample in 2018, PSUs had been modified metropolitan statistical areas and groups of counties with smaller towns (Paben et al. 2016). Each PSU is subdivided into segments, which become the fundamental units for sampling and weighting. Segments are 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 BLS selected a new sample in 1999. Subsequently, the survey lost units to demolition, to conversion to other uses, or to respondent non-cooperation.[18] The survey periodically added new units sampled from construction permit data. More recently, the BLS 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.

CPI rent is calculated using the average six-month change in economic rent in that month's sample, which is converted 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 index area is constructed as

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

where w_i is the weight for unit i,[19] 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.[20]

---

[17] For more details on the design of the Housing Survey sample see Ptacek (2013).

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

[19] A unit's weight in the rent index depends on the estimated aggregate rent payments from its segment and the response rate for the segment.

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

Indices are calculated for each index area, which is either a large PSU or the set of PSUs representing the smaller cities in a Census division. The national index derived from the average of changes in the index area indices weighted by rent expenditure in that index area. Until January 2023 the aggregation weights were updated every two years, so that the indices in year t are aggregated using expenditures from $t - 1$ or $t - 2$. Starting in January 2023, the aggregation weights are updated annually.

A.2 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).

A.3 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 1 in the main text 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 1; 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).[21] On average, rental listings in the MLS are more expensive, larger, and newer than newly occupied rental units in the AHS (see Table 1).

of Methods.

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

A.4 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 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),[22] 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.

B REPRESENTATIVITY: FURTHER DETAILS

In the main body, we briefly discussed the representativity of data underlying the SFRI, the ZORI, the MRI, and indices based on BLS data. In this section, we discuss the sample representativity of two other data sources, as well as other information pertinent to comparison studies like this one.

Why is sample representativity important? A non-representative sample is, effectively, a sample that has been conditioned on a variable, such as geography or structure type. (Equivalently, non-response bias is a chief concern in many contexts.) "Location-location-location" has been an aphorism in real estate since at least the 1920s, and rent growth can vary significantly within and across cities (Verbrugge and Poole 2010). Real estate markets are segmented by location, but also by structure type (Adams and Verbrugge 2021). Thus, rental market dynamics vary not only by location, but also by structure type (within a location). A data source that is restricted along one of these two dimensions will feature rent movements that may differ from the average.

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

B.1 The American Housing Survey

In Table 1 we compare the BLS Housing Survey to the American Housing Survey (AHS). We discuss the construction of the AHS here.

The AHS is a longitudinal housing unit survey conducted biennially by the US Department of Housing and Urban Development in odd-numbered years and designed (after weighting) to represent the US housing stock (and not US housing expenditure). Based upon 1980 Census data, the national sample underwent a redesign in 1985, with a base sample size of approximately 47,000 housing units (owned and rented). However, few homes remained in the panel over its entire length; over the 1985-2013 period, 100,000 different homes were included.[23] In 2005, the national sample was improved in two ways: first, mobile home coverage was adjusted by replacing the units currently in the sample with mobile homes selected from Census 2000, and, second, assisted living housing units selected from Census 2000 were introduced into the sample. A new representative national sample of approximately 85,000 housing units was drawn for the 2015 AHS using the master address file (MAF) as the sampling frame, with additional oversampling of selected metropolitan areas and HUD-assisted housing units. The total sample size beginning in 2015 is about 115,000 housing units.

The AHS collects information about units' physical characteristics (including the physical condition of homes), information on neighborhoods, information on the characteristics of people who live in the homes, vacancies, home improvements, and housing costs. In Table 1, we use the national sample in 2015; the AHS sample was redrawn at this date, so the sample is discontinuous there.[24] AHS data do not identify whether utilities are included in the contract rent.

B.2 Geographic representativity of MLS and BLS Housing Survey

BLS data are representative of expenditures across urban areas in the US, and AHS data (and ACS data, to a somewhat lesser extent) are representative of housing units across the entire US. To convey a sense of the coverage of MLS data versus BLS data, Figure B.1 maps what locations are most sampled in Los Angeles, where both data sources have many observations. The BLS sample is concentrated in its selected segments, but these segments are spread throughout the metropolitan area.

[23]An interesting aspect of the AHS is that a housing unit's transitions between owner-occupied and renter-occupied are observable; see, for example, Foote et al. (2020). Conversely, in the BLS data, transitions from owner-occupied to renter-occupied occur outside of the sample, and a transition from renter-occupied to owner-occupied will typically imply that the unit drops from the sample.

[24]We do not have access to Zillow microdata or to RCA CPPI microdata, so their corresponding summary statistics are not included in the table.

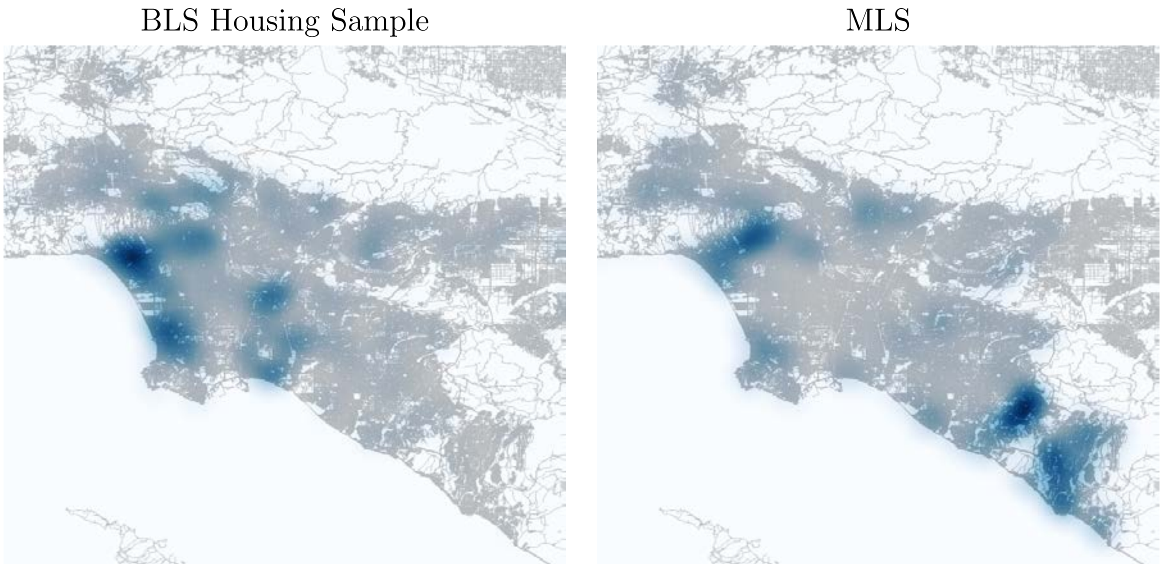

Figure B.1. Heatmap of Sample Locations in Los Angeles. *Note:* The left panel shows the geographic distribution of the sample in the BLS Housing Survey. The right panel shows the geographic distribution in the MLS data. *Source:* BLS Housing Survey and CoreLogic.

C MLS-BASED REPEAT-RENT INDEX

The CoreLogic SFRI is based on housing units listed for rent in the MLS. Because CoreLogic also provides access to the underlying data, these data are often used by researchers. While the rental rates and geographic dispersion of rental units in the MLS data are not representative (see Table 1 and Section B.2 above), our results suggest that over our sample period, the rent growth of properties listed in the MLS is representative. Most researchers use these data for a specific area. We therefore created a series of rent indices using the MLS data for areas that match the PSUs in the BLS Housing Survey. Our methodology is identical to that described in Section 2, including that we remove any properties that the listing indicates were recently renovated or remodeled. We then compared the resulting MLS and CPI-data-based indices, and found that they consistently gave similar results — although the CPI-based indices are more volatile, reflecting their smaller sample size. Figure C.2 contains two examples. Our findings should provide some confidence to researchers who wish to use the MLS data to measure local rent growth.

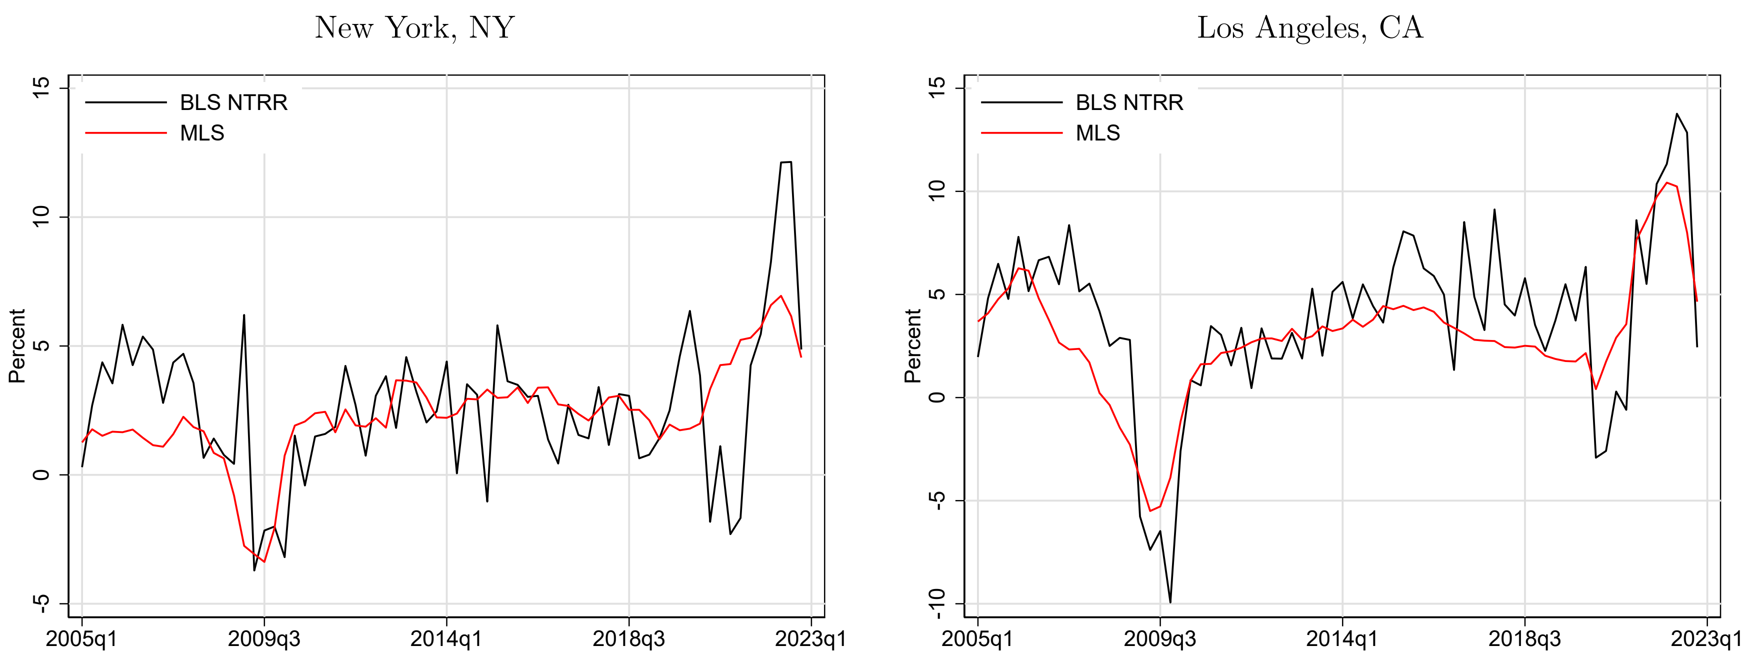

Figure C.2. PSU-LEVEL INDICES USING BLS NEW-TENANT DATA AND MLS RENTAL LISTINGS. *Note:* Areas are defined as the respective CBSA. *Source:* BLS Housing Survey and CoreLogic Multiple Listing Service data.

D CONSTRUCTING VARIANCE ESTIMATES FOR REPEAT-RENT INDICES

The BLS microdata are derived from a multistage sampling design. These data are then used to create a repeat-rent index, whose four-quarter growth rate is then computed. This is a nonlinear function of the data. In such cases, variances are unknown. To determine whether these indices are statistically indistinguishable, we estimate variances of the quarterly estimates using a bootstrap analysis, following Wolter (2007). These methods are even applicable to estimators deriving from complex sample survey designs.

The basic idea involves forming random groups by resampling housing units at random with replacement within each PSU in the BLS rental sample. We create k groups of housing units for each PSU, and then use the groups to create k PSU-specific repeat-rent indices. Next, we use the upper-level weights to aggregate the resampled PSU quarter-on-quarter changes to create k national repeat-rent quarter-on-quarter changes. The estimate of the variance $v\left(\hat{\theta}\right)$ in any given month is given by:

$$v\left(\hat{\theta}\right) = \frac{1}{k-1}\sum_{j=1}^{k}\left(\hat{\theta}_j - \hat{\theta}\right)^2, \tag{D.2}$$

where $\hat{\theta}$ is the average estimate across the k groups, and $\hat{\theta}_j$ is the estimate from group j, and we have suppressed the time subscript. Since k is not large in our applications, the confidence interval takes the form

$$\hat{\theta} \pm t_{k-1,\alpha/2}\sqrt{v\left(\hat{\theta}\right)}, \tag{D.3}$$

where $t_{k-1,\alpha/2}$ is the upper $\alpha/2$ percentage point of the t distribution.

To ensure that the random group estimator has acceptable statistical properties, the random groups must be formed so that each group has the same sampling design as the original sample. Thus, in the multistage sampling undertaken by the BLS for its rent sample, random groups must be formed by dividing the ultimate clusters, which are Census block groups, into k groups.

E 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 the 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

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

where $y_t = (y_{1,t}, y_{2,t})'$ is a vector of two rent indices (e.g., $y_{1,t} = \ln(\text{CPI rent}_t)$, $y_{2,t} = \ln(\text{SFRI}_t)$); $\Delta y_t = y_t - y_{t-1}$; γ and $\upsilon = (\upsilon_1, \upsilon_2)'$ are constants; Λ is a matrix of coefficients on lag terms; the parenthetical expression $(\gamma + \beta' y_{t-1})$, which is the (stationary) error-correction term, describes the long-term cointegration relationship between y_1 and y_2; 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 E.1 reports the values of α and β along with some standard errors.

In responding to deviations from their long-term relationships (with either CoreLogic or CPI rent), point estimates suggest that the NTRR and the ATRR do most of the adjusting to eliminate said deviations; the NTRR also does most of the adjusting toward the ATRR. In the 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). A stepwise model-selection search found the optimal model for quarterly CPI rent growth included lags 1 to 4 of quarterly CPI rent growth, lags 2 to 4 of quarterly CoreLogic growth, and the error-correction term. The BIC for this model is 1.761. Dropping the error-correction term from the model reduces the BIC to 1.740. Finally, dropping all CoreLogic terms from the model yields a BIC of 2.218. These results indicate that the SFRI has strong predictive content for CPI rent, but that the error-correction term is not a useful predictor.

Table E.1. Pairwise Vector Error Correction Results

	vecout				
	NTRR-CoreLogic	ATRR-CoreLogic	CPI Rent-NTRR	CPI Rent-CoreLogic	NTRR-ATRR
Cointegrating Equation					
γ	0.588	0.995	-1.201	-0.686	-0.168
(Std Error)	(0.044)	(0.070)	(0.166)	(0.168)	(0.080)
β_1	1.000	1.000	1.000	1.000	1.000
β_2	0.893	0.818	1.227	1.133	1.032
(Std Error)	(0.008)	(0.013)	(0.030)	(0.030)	(0.014)
Speed of Adjustment					
α_1	-31.925	13.190	-13.483	-7.776	-101.688
(Std Error)	(29.200)	(5.425)	(2.424)	(4.130)	(34.030)
α_2	-5.286	1.264	32.885	12.578	-0.388
(Std Error)	(14.806)	(11.331)	(20.386)	(4.461)	(10.976)

F ADDITIONAL FIGURES REFERENCED IN TEXT

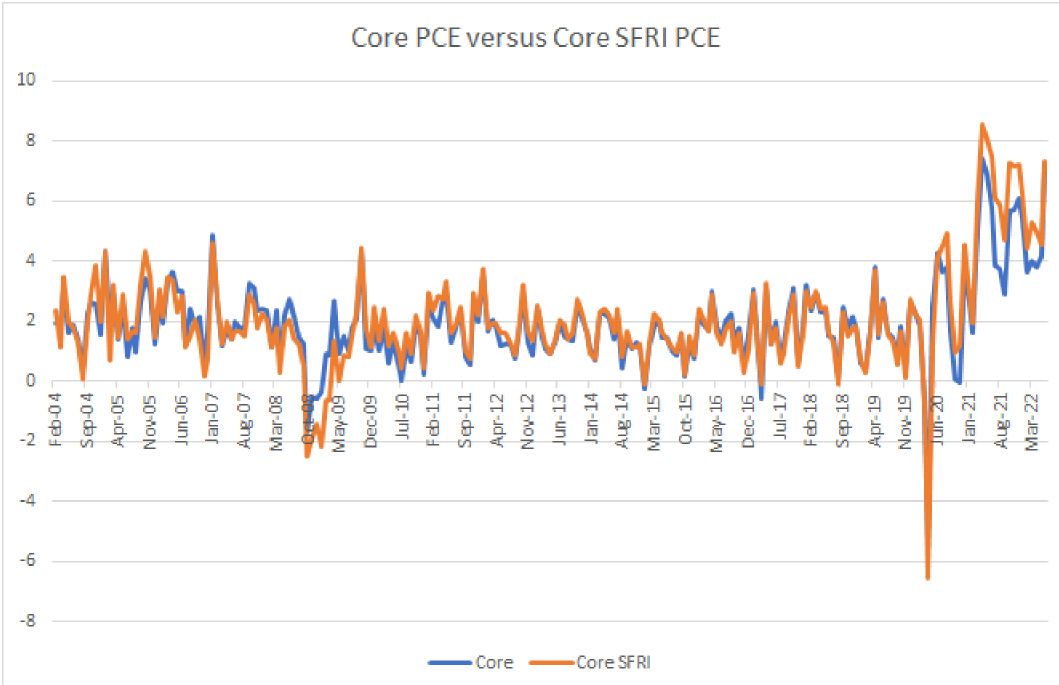

Figure F.3. Monthly Annualized Core PCE inflation versus Core-SFRI PCE inflation. *Note:* Core is standard core PCE inflation, which includes CPI rent inflation as a component. Core SFRI is an alternative inflation series that replaces CPI rent with rent inflation based on the SFRI . *Source:* Bureau of Economic Analysis, CoreLogic, and authors' calculations.