

Discussion of “Nonresponse Imputations and Related Measurement Issues in the CPI for Shelter” (Loewenstein, Montag, and Verbrugge)

Boaz Abramson

Columbia Business School

The paper addresses an important and under-appreciated problem in the measurement of U.S. rent inflation: the growing prevalence of survey nonresponse in the Bureau of Labor Statistics (BLS) Housing Survey, which underlies the CPI shelter index. The paper makes several valuable contributions. It documents that nonresponse rates in the BLS Housing Survey have been rising steadily and now account for approximately 40 percent of all observations. It demonstrates that nonresponse is not random—both nonresponse rates and rent inflation co-vary with observable characteristics of housing units—which implies that imputing rents for nonresponses is necessary. The paper then details the current BLS imputation procedure and its limitations. Finally, it proposes and evaluates alternative imputation strategies.

This discussion is organized around three main comments. The first situates the nonresponse problem in a broader context. The second evaluates the alternative imputation methods proposed by the authors. The third offers suggestions for extending the analysis.

1. Survey Nonresponse: A Broad and Growing Challenge

The paper begins with the observation that nonresponse in the BLS Housing Survey has increased substantially over recent decades and now accounts for roughly 40 percent of the data. This is an important problem that extends far beyond the BLS Housing Survey. Figure 1, which is taken from

the paper, illustrates response rates across a range of BLS surveys—including the Consumer Expenditure Survey (CE-Diary and CE-Interview), the American Time Use Survey (ATUS), the BLS Housing Survey (CPI-Housing), and the Current Population Survey (CPS)—from April 2015 through April 2025. The figure documents that low and declining response rates are a general feature of BLS administered surveys, not an idiosyncrasy of the Housing Survey.

Household survey response rates, April 2015–April 2025

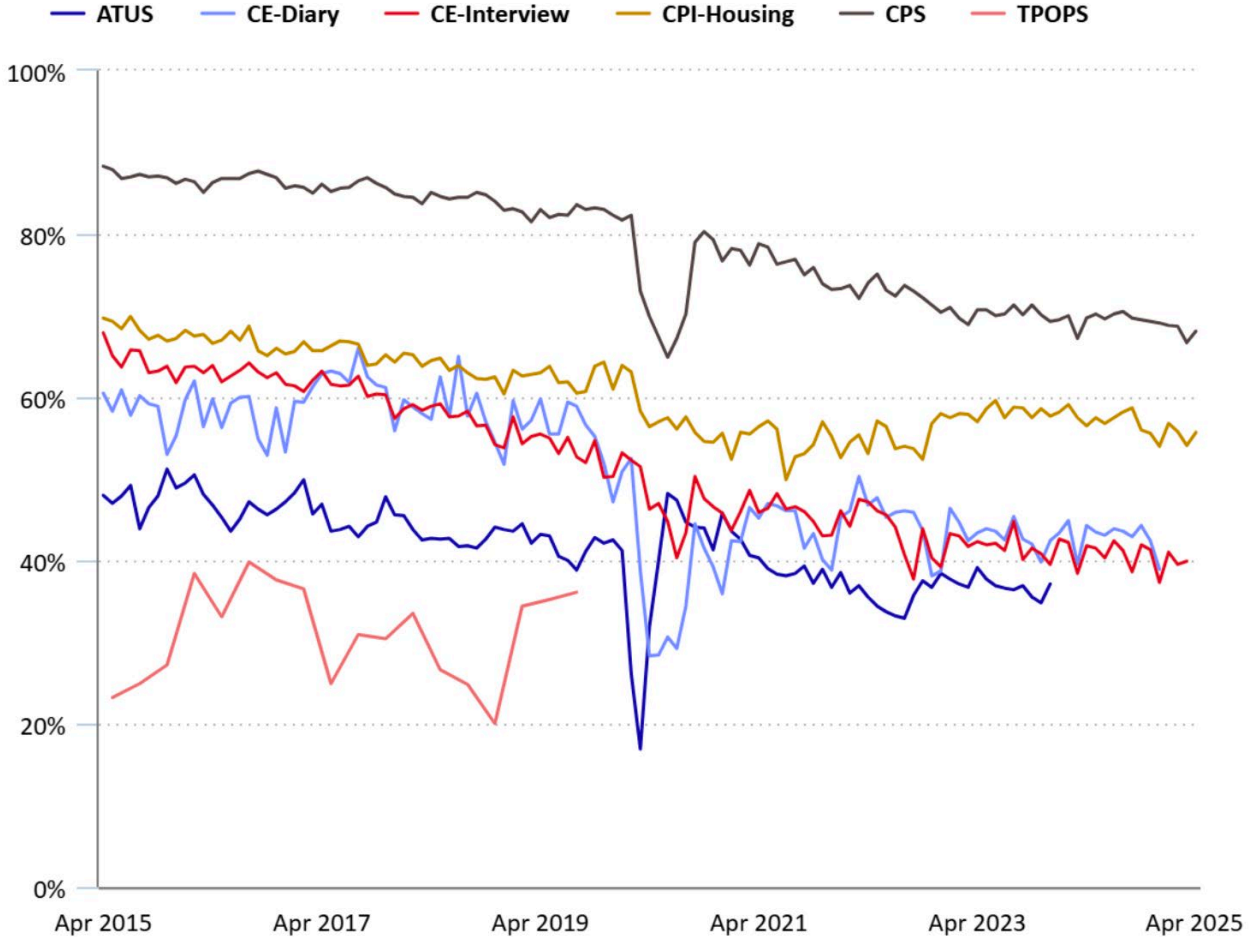


Figure 1. Response Rates for BLS Administered Surveys (reproduced from Loewenstein, Montag, and Verbrugge).

This pattern extends to surveys administered by other federal agencies. Figure 2 displays response rates for three additional surveys: the American Community Survey (ACS), administered by the Census Bureau; the National Health Interview Survey (NHIS), administered by the CDC; and the

CPS Annual Social and Economic Supplement (CPS-ASEC), administered jointly by the Census Bureau and the BLS. Across all surveys, response rates have declined markedly, particularly since 2019. The ACS, which had maintained response rates above 90 percent for most of the period shown, experienced a sharp drop during the COVID-19 pandemic and has not fully recovered. Response rates in the NHIS and CPS-ASEC have fallen to around 50 percent and 60 percent, respectively.

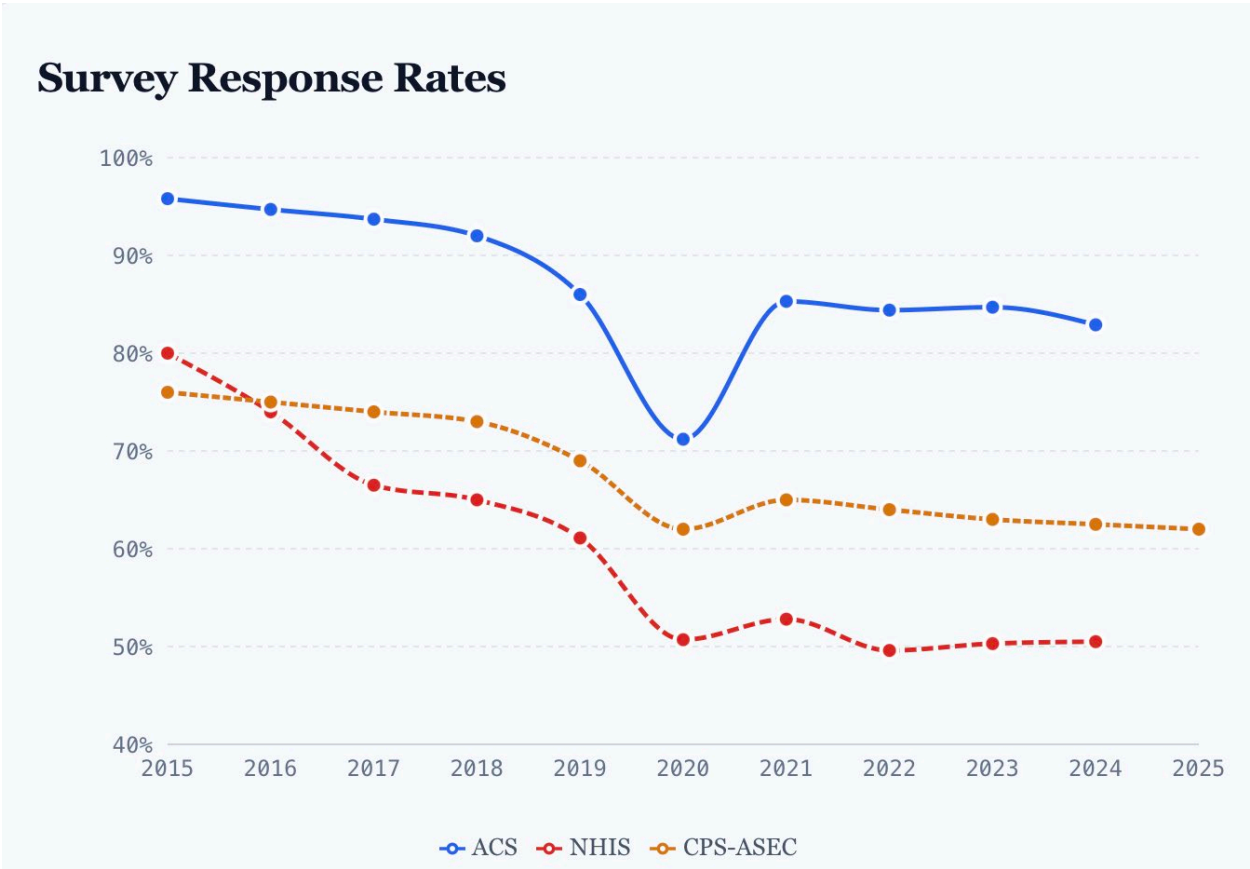


Figure 2. Response Rates for Surveys Administered by Other Agencies.

The important takeaway is that a large and increasing fraction of the data underlying U.S. employment, consumption, poverty, health, and housing statistics is currently being imputed. This observation should be alarming to any researcher using these data sources. The overwhelming magnitude of the non-response problem also means that the potential contribution of this paper

extends well beyond improving the measurement of rent inflation in the BLS Housing Survey. If the authors are able to develop accurate imputation methods that can be adapted and applied to other surveys, this will benefit the entire profession.

2. Alternative Imputation Methods and Their Performance

The current BLS imputation procedure is a class-means approach based on city, month, and lagged rent tercile. For each nonresponding housing unit, the BLS imputes rent based on the unit's lagged rent and the average rent change observed for responding units in the same city and month that are in the same lagged rent-tercile. The authors consider alternative imputation schemes that are similar to the one utilized by the BLS. In particular, they consider a similar class-means approach but differ based on the observable variables they use to partition the data. They consider an unconditional imputation based on all responding units in the same city and month; an imputation based on all responding units in the same city and month that have the same structure type; and an imputation based on all responding units in the same city and month whose tenants have the same tenancy length. For the latter, since tenancy length is unobserved for nonrespondents, the authors consider several approaches to impute tenure, including a carryforward of the last observed move-in date, upper- and lower-bound scenarios, and a probabilistic method that uses ACS mobility rates.

How do the resulting indexes compare to the current CPI-shelter index published by the BLS? Figure 3 below, which corresponds to Figure 7 in the paper, compares the alternative rent indexes in terms of year-on-year changes (Panel (a)) and levels (Panel (b)). The main takeaway is that all four indexes—the current BLS index, the index based on unconditional imputation, the index

based on structure-type imputation, and the index based on lagged (or “raw”) tenancy imputation—track each other closely over the entire sample period.

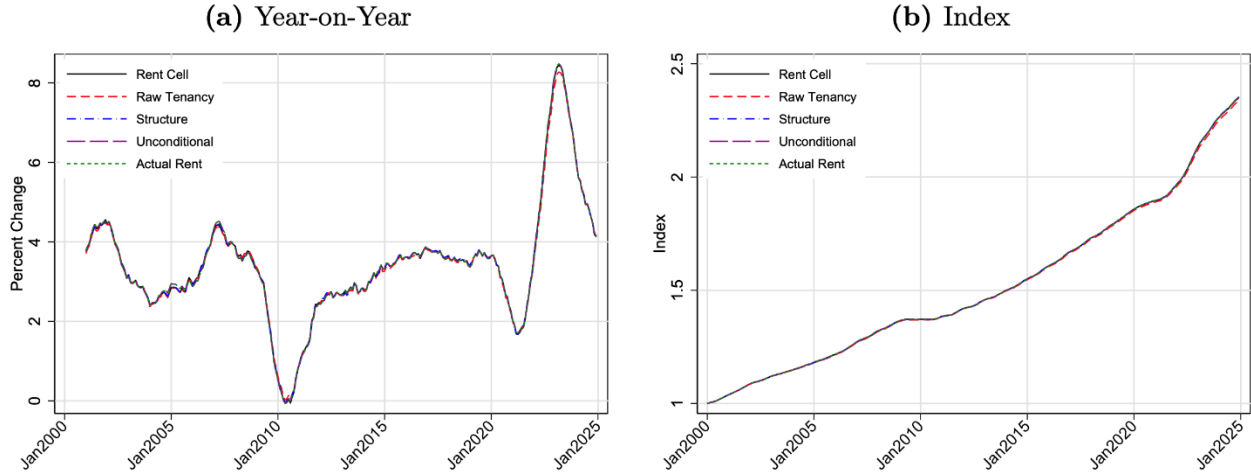


Figure 3. *Alternative Rent Indexes: Year-on-Year Change (left) and Index Level (right). Reproduced from Loewenstein, Montag, and Verbrugge, Figure 7.*

The authors’ preferred index is the one based on ACS-imputed tenancy. Figure 4 zooms in on the comparison between the current BLS index and the index based on ACS-imputed tenancy. The ACS-tenancy index tracks the BLS index closely until approximately 2020, after which it diverges somewhat, ending approximately two percentage points below the BLS index by early 2025. Over a span of 15 years, this is arguably a minor cumulative difference. Moreover, as illustrated in Panel (a) of Figure 8 in the paper, the differences in terms of year-on-year rent inflation seem negligible.

(b) Index

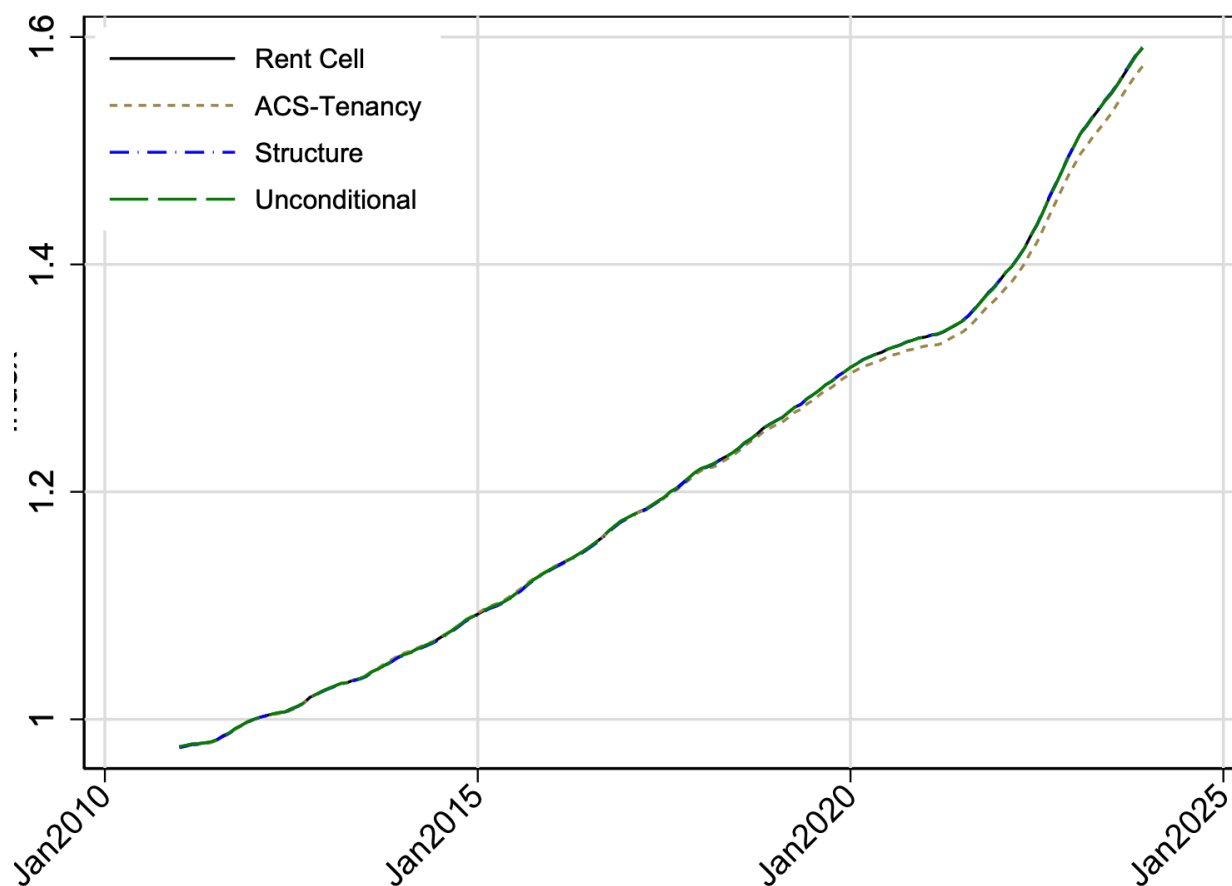


Figure 4. Alternative Rent Indexes: Reproduced from Loewenstein, Montag, and Verbrugge, Figure 8(b).

Do the alternative rent indexes that the authors construct reduce the imputation bias relative to the current BLS rent index? The authors examine this using a jackknife resampling analysis. Figure 5 presents histograms of the jackknife imputation errors for the current BLS index (Panel (a)) and the preferred ACS-imputed-tenancy index (Panel (b)). The distributions of errors are very similar across the two methods. The mean imputation error is 0.20 percent for the current BLS index rent and 0.18 percent for the ACS-imputed-tenancy—an arguably minor difference.




Figure 5. Jackknife Imputation Errors. Reproduced from Loewenstein, Montag, and Verbrugge, Figure 9.

The overall takeaway is that the alternative imputation methods the authors consider yield indexes that are very similar to the current BLS index and that do not substantially reduce the imputation bias. This raises the question of whether a class-means procedure is the right imputation method to consider, and whether different imputation methods can more substantially improve upon the imputation bias in the BLS Housing Survey.

3. Suggestions for Future Research

In this section, I propose alternative imputation methods that may reduce the imputation bias relative to the current BLS index. To begin, note that any imputation can be formalized as a function F that satisfies:

$$Rent_{i,t} = F(X_{i,s<t}, X_{i',s<t}, X_{i,t}, t),$$

where $Rent_{i,t}$ is the imputed rent for nonresponding unit i at time t , $X_{i,s<t}$ denotes lagged characteristics of unit i (such as its lagged rent), $X_{i',s<t}$ denotes lagged characteristics of all responding units (such as their lagged rent and tenant tenure status), and $X_{i,t}$ denotes

contemporaneous characteristics of all responding units (such as their current rent and tenant tenure status). In principle, the parameters of the imputation function can depend on time t .

The econometrician's goal is to find the function F that best imputes rents, or more formally that minimizes some measure of imputation bias—for example, the sum of squared jackknife imputation errors. Note that the authors impose relatively strong restrictions on the function F : they consider only class-means imputation methods, they consider only a prespecified set of observable characteristics the imputation can be based on, they require the imputation parameters to be time-invariant, and they generally allow the imputation to utilize information only from one-period lags.

There is no obvious reason to impose these restrictions if the objective is to minimize the imputation error. Fortunately, computer scientists have developed imputation techniques that impose little to no restrictions on the imputation function. Specifically, Machine Learning techniques such as Gradient Boosting, Random Forests, and Deep Neural Networks allow imputing rents without pre-specifying the characteristics of the imputation function. These methods do not require pre-specifying the observable variables that can be used for imputation, they are not constrained to class-means, they allow the imputation function to vary over time through periodic algorithm retraining as new data become available, and are well-suited to capturing the nonlinear and time-varying relationships that appear to characterize nonresponse and rent inflation in the BLS data. Given that the nature of nonresponse appears to exhibit substantial temporal variation, this flexibility seems particularly well-suited to the problem at hand.

An additional advantage of these machine learning approaches is that they circumvent the need to separately impute tenancy length for nonrespondents. This is a challenging intermediate step in

the authors' framework which requires auxiliary assumptions about unobserved tenant turnover. Machine learning methods sidestep this complication by directly mapping observed characteristics to imputed rents. Given that the tenancy-based methods in the paper require substantial assumptions, a more general approach seems attractive.

Finally, the paper would benefit from a case study that showcases a difference between the indexes the authors impute and the original BLS index. As discussed above, the various indexes seem to track each other closely over most of the sample period. As a result, a reader may reasonably wonder whether the choice of imputation method matters for practical applications. Providing an example where the alternative indexes lead to meaningfully different conclusions would help make the case for the importance of these alternatives.

One natural candidate involves the relationship between monetary policy and rents. Since rent is the largest component of the CPI, the ability of monetary policy to control inflation depends in large on how it affects rents. Dias and Duarte (2019) and Abramson, De Llanos and Han (2025) find that monetary tightening increases rents by shifting household demand from the owner-occupied to the rental market, a finding with significant implications for monetary policymakers. It would be informative to examine whether this estimated relationship depends on which rent index is used. If the impulse response of rents to a monetary policy shock differs materially across the BLS and the new indexes the authors propose, that would constitute strong evidence that the choice of imputation method has important consequences.

An alternative exercise would compare the Taylor rule prescription for the federal funds rate under the official CPI, which uses the current BLS index for measuring shelter inflation, and an alternative CPI, which would instead use the authors' indexes for measuring the shelter inflation

component. Ambrose, Coulson, and Yoshida (2018) document that the Taylor rule is sensitive to rent inflation measurement choices. If the implied Taylor rule prescriptions differ substantially across indexes, that would showcase that the imputation method matters.

Conclusion

This paper addresses an important and growing problem. Survey nonresponse has become pervasive across government household surveys, and the accuracy of official statistics—for employment, consumption, poverty, health, and housing increasingly depends on the quality of imputation methods. The paper makes a valuable contribution by documenting the nature of nonresponse in the BLS Housing Survey, characterizing how nonresponse and rent inflation covary with observable characteristics, and proposing imputation methods that are based on structure type and tenancy length.

At the same time, the empirical results suggest that class-means imputations do not materially reduce the imputation bias relative to the current BLS index. Extending the analysis to more flexible imputation methods—particularly machine learning techniques—might allow the authors to improve their rent index and would strengthen the paper's contribution.

References

Abramson, Boaz, Pablo De Llanos, and Lu Han (2025). "Monetary Policy and Rents." Working paper, Columbia Business School.

Ambrose, Brent W., N. Edward Coulson, and Jiro Yoshida. "Reassessing Taylor rules using improved housing rent data." *Journal of Macroeconomics* 56 (2018): 243-257.

Dias, Daniel A, and João B Duarte. 2019. "Monetary policy, housing rents, and inflation

dynamics.” *Journal of Applied Econometrics*, 34(5): 673–687.