Nonresponse Imputations and Related Measurement Issues in the CPI for Shelter*

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

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

March 13, 2025

Disclosure: Hugh Montag is a full-time employee of the US Bureau of Labor Statistics, which has reviewed a draft of this extended abstract and approved it for release. Lara Loewenstein and Randal Verbrugge are full-time employees of the Federal Reserve Bank of Cleveland, but have an intergovernmental contract with the BLS that gives them access to the data. They have no other potential conflicts to disclose.

1 EXTENDED ABSTRACT

The BLS Housing Survey provides the data for CPI tenant rent and owner-equivalent rent, which together comprise the majority of CPI shelter inflation. Shelter inflation makes up over 30 percent of overall CPI inflation and nearly 20 percent of PCE inflation, and even higher percentages of their ex-food-and-energy counterparts. It therefore behooves any economist interested in price stability to understand rental pricing and specifically the manner in which the shelter indexes are constructed, which are based BLS Housing Survey.

The ability of the US government to measure rental prices is central to its ability to create accurate price indices, which are central to measuring inflation. As such, the fact that almost 40 percent of current observations in the BLS Housing Survey are imputed, with the majority of these (over 30 percent of all observations) being imputed due to nonresponse, is

*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, the Federal Reserve System, or the BLS. Disclaimer: This paper provides a summary of research results. The information is being released for statistical purposes, to inform interested parties, and to encourage discussion of work in progress. The paper does not represent an existing, or a forthcoming new, official BLS statistical data product or production series.

worthy of concern. It is especially worthy of concern because the share of observations that are imputed has been increasing, and has effectively *doubled* since 2000 (see in Figure 1).[1]

If the imputed responses in the BLS indicated the same average inflation rate as non-imputed responses, then they would have no impact on the aggregate index. But imputed observations do not have the same average inflation rate as non-imputed observations. Furthermore, the inflation rate of imputed observations is much more volatile than the overall shelter inflation rate (see in Figure 2). This raises some important questions, including:

1. Is the current method of imputation optimal and/or unbiased?

2. How much uncertainty is introduced into shelter inflation measures given the rate and nature of the imputation?

In this paper we will address these questions directly. First, we will provide some basic statistics on the rate of imputation in the BLS Housing Survey with the chiefgoal of ascertaining if it is reasonable to assume that missing data in the BLS Housing Survey is random. In other words, are the missing data ignorable? Our initial analysis indicates that the answer is no. In particular, imputation rates vary dramatically by property type (see Figure 3a), city (see Figure 3b), and tenure length (see Figure 3c). We will use further statistical models, such as further pattern and logistic regression analysis to compare rates of missing data with observable covariates.

There are two reasons for missing data in the BLS Housing Survey: vacancy and nonresponse. Nonresponse is by far the more common reason, especially more recently. As such, we will largely focus on this category, with a more brief discussion of the vacancy imputation.

2 NONRESPONSE

The current nonresponse imputation methods used by the BLS makes assumptions about data being missing-at-random and about reasonable comparison groups for imputing missing data. The current method involves a hot-deck procedure similar to what is described for income imputations in the CPS described in (Hokayem et al. 2015) where missing rent changes are imputed from the rent changes of similar units.[2] If the desired pool of similar

[1]The BLS Housing Survey is not the only government survey to suffer from nonresponse and other forms of missing data. This project will contribute to a literature in economics on missing data and measurement official statistics that includes Hokayem et al. (2015), who show that income nonresponse in the March Supplement in the CPS leads to a 1 percent bias in official US poverty rates, and Heffetz and Reeves (2021) who find that measurement issues (including nonresponse) resulted in a potential 1.5 percent bias (in either direction) in the unemployment rate during the COVID19 pandemic. Other BLS surveys, such as the Consumer Expenditure Survey, have also implemented imputation methods to address missing data.

[2]The BLS's Commodities and Services Survey (C&S) performs similar hot-deck imputations. If an item is missing or if there is a non-comparable substitution, the C&S will often impute a price using the average

units is too small, then the criteria for similarity is relaxed. Currently, units are divided into thirds by their their within-city rent levels from the previous survey (low-rent, medium-rent, and high-rent). If the BLS has more than 5 respondents in the same rent tier and city as the nonrespondent, then the rent is imputed using the average rent change from this comparison pool. If the BLS has fewer than 5 respondents in this comparison pool, it increases the comparison group by relaxing the rent-tier criteria. For example, they may combine the low- and medium rent tiers. Finally, if this does not result in a sufficient comparison pool, the BLS combines all respondent for that time period in the city.

This method is only valid if the pool of similar units is representative of the nonrespondent, and there are many reasons to think that this may not be the case. For example, the BLS does not know at the time of nonresponse whether the nonrespondent is a new tenant, or the same tenant that previously occupied the unit. It therefore does not include tenancy in the list of criteria to create comparable units. But it is well-known that the size of rent changes depends critically on tenancy (Gallin et al. 2024). This will create inaccuracy in the imputation if the tenancy of the nonrespondent is not equal to the average of the tenancy of the comparison pool, which Figure 3c indicates is almost certainly not the case. This is one example, among others, of potential inaccuracies caused by the current imputation procedure.[3]

We will attempt to understand the determinants of nonresponse in order to better create comparison pools for imputation. Previous studies on nonresponse have been assessed using links to adminstrative data (Bollinger and Hirsch 2006; Heffetz et al. 2024). We have effectively no information on the tenant themselves, so are not able to link respondents (and nonrespondents) with any administrative data. The BLS Housing Survey has limited information on the method of contact used, number of times contact was attempted, and reasons for nonresponse. However, we can make use of other available information to understand determinants of nonresponse. As a panel survey, the BLS Housing Survey surveys the same housing unit multiple times, so imputed values can be compared to future reported values. We can also assess the importance of serial nonresponse—nonresponse of a given unit over multiple periods—in overall nonresponse rates. Second, the data include the exact location of the property, allowing us to correlate nonresponse rates with local area demographic including local area income, race, and share of homeowners from the American Community Survey. While backward revision are not permitted to the official CPI print, information in the BLS Housing Survey such as move-in dates is backfilled, allowing us to compare imputations to a small sample of actual responses. Last, we can answer questions such as whether

price relative for similar items in the same geographic area. See https://www.bls.gov/opub/hom/cpi/calculation.htm#imputation for further details.

[3]It is also important to note that, even if nonresponse is random, there can still be substantial biases in the imputation (Hirsch and Schumacher 2004; Bollinger and Hirsch 2006).

response rates vary by rent levels and previous respondent type.

We will follow-up by offering a some suggestions for improving the current imputation method given our analysis of the data. We will assess whether this method is ideal by considering counterfactual inflation measures using other imputation approaches. These could include both the refinement of comparison groups, or other methods such as inverse probability weighting (Rosenbaum 1987; Bang and Robins 2005; Wooldridge 2007).

Last, we will provide estimates of the uncertainty induced by nonresponse and the current imputation procedure. This can be done using a variety of methods including Monte Carlo simulations and boundary exercises that involve creating upper and lower bounds such as in Cochran (1977) and Manski (1989).

3 VACANCY

Vacancy is the other reason missing values are imputed. The BLS imputes the rent of vacant units to avoid a bias associated with sample exit. The imputation when a unit first becomes vacant is based upon the rent changes observed in units in the sample that, over the past six months, switched from one tenant to another tenant (i.e., the unit turned over). However, as the research of Gallin et al. (2024) demonstrates, the gap between a unit's actual rent and the rent that it would command were the unit to become vacant is systematically related to tenure. This suggests that the first vacancy imputation could be improved by making use of the tenure information in the sample, and perhaps by making use of recent movements in new-tenant rent indexes.

4 OTHER MEASUREMENT ISSUES IN SHELTER INFLATION

If time and space allows, we will also explore other forms of potential bias in the creation of shelter inflation, including the age bias adjustment and potential rotation bias due to the sample design.

Aging bias. The aging bias adjustment is intended to account for the *average* net depreciation of rental units, given the location and age of the unit: the quality of the rental unit is (on average) decreasing over time, so even if a rent is unchanged, a tenant is paying more on a quality-adjusted basis.

The current method for aging bias is elucidated in Houck (2023). In brief, in January of each year, linear regression is used to model the statistical relationship between rent on the one hand, and age-of-unit on the other hand, controlling for a host of other covariates, some of which derive from other data sources such as the American Community Survey. The coefficients on the age terms in the regression are mapped into aging factor adjustments, which then enter into the computation of shelter inflation. In particular, estimated inflation

typically increases as result of the aging factor, which controls for quality change in the unit—by taking into account the fact that (on average) the unit has deteriorated by a small amount over the previous six months.

What is the aging factor measuring? The BLS housing survey includes information on structural renovations, such as a change in the number or types of rooms in a unit, but does not include information on non-structural renovations. For example, the survey does not include kitchen renovations or other significant changes to the unit (such as new carpet, installation of a dishwasher, or a rodent infestation), all of which have a notable impact on rent.[4] In keeping with this, the underlying aging factor regression estimates rent as a function of age and location, but also any renovations that the BLS categorizes as structural. Thus, it can be interpreted as estimating the rent, conditional on *average* net spending on upkeep and the like. Thus the aging factor estimates *average net* deterioration.

The aging factor regression includes a number of covariates that are not directly related to unit quality, such as neighborhood income. These variables are correlated with rent, but do not directly measure deterioration. We will explore the extent to which the inclusion of said covariates alters the inferences about net aging.

Following Verbrugge and Gallin (2016), we will also study aging-bias estimation that exploits the panel dimension of the data. A panel approach allows one to control for unobserved characteristics that do not change, although this does nothing to address changes in rent due to unobserved changes in said characteristics. The previous findings demonstrate the importance of understanding sample exit and its role in aging bias estimates.

Rotation bias: Around 2012, the BLS switched to a 6-year rotating panel whereas they had previously maintained a more static panel. This rotation implies that at any given point, units have been in the sample for different lengths of time, exposing the BLS Housing Survey to potential rotation bias where responses depend on the length of time in the survey (Bailar 1975; Krueger et al. 2017). We will explore potential rotation bias in the BLS Housing Survey, along with any relationship between rotation bias, tenancy lengths and nonresponse. It is feasible, for example, that tenants are less likely to respond (or to provide accurate answers) the longer they have been in the sample, increasing nonresponse rates or accuracy for properties with low turnover and for properties at the end of the 6-year period in the sample.

[4]We will continue to work with the BLS CPI Housing staff to illuminate how renovations are categorized.

References

Bailar, Barbara A (1975). "The effects of rotation group bias on estimates from panel surveys." *Journal of the American Statistical Association*, 70(349), pp. 23–30.

Bang, Heejung and James M Robins (2005). "Doubly robust estimation in missing data and causal inference models." *Biometrics*, 61(4), pp. 962–973.

Bollinger, Christopher R and Barry T Hirsch (2006). "Match bias from earnings imputation in the current population survey: The case of imperfect matching." *Journal of Labor Economics*, 24(3), pp. 483–519.

Cochran, William G (1977). "Sampling techniques." *Johan Wiley & Sons Inc.*

Gallin, Joshua, Lara P Loewenstein, Hugh Montag, and Randal J Verbrugge (2024). "Sticky continuing-tenant rents." *Work in progress.*

Heffetz, Ori, Guy Lichtinger, and Daniel B Reeves (2024). "Assessing nonresponse bias in macro indicators by combining para-, administrative, and survey data." Technical report, National Bureau of Economic Research.

Heffetz, Ori and Daniel Reeves (2021). "Measuring unemployment in crisis: effects of covid-19 on potential biases in the cps." Technical report, National Bureau of Economic Research.

Hirsch, Barry T and Edward J Schumacher (2004). "Match bias in wage gap estimates due to earnings imputation." *Journal of labor economics*, 22(3), pp. 689–722.

Hokayem, Charles, Christopher Bollinger, and James P Ziliak (2015). "The role of cps nonresponse in the measurement of poverty." *Journal of the American Statistical Association*, 110(511), pp. 935–945.

Houck, Ben (2023). "A review of recent improvements to the cpi's housing age-bias adjustment." *Monthly Labor Review.*

Krueger, Alan B, Alexandre Mas, and Xiaotong Niu (2017). "The evolution of rotation group bias: Will the real unemployment rate please stand up?" *Review of Economics and Statistics*, 99(2), pp. 258–264.

Manski, Charles F (1989). "Anatomy of the selection problem." *Journal of Human resources*, pp. 343–360.

Rosenbaum, Paul R (1987). "Model-based direct adjustment." *Journal of the American statistical Association*, 82(398), pp. 387–394.

Verbrugge, Randal J and Joshua Gallin (2016). "Panel data estimates of age-rent profiles for rental housing." Technical report, Federal Reserve Bank of Cleveland.

Wooldridge, Jeffrey M (2007). "Inverse probability weighted estimation for general missing data problems." *Journal of econometrics*, 141(2), pp. 1281–1301.

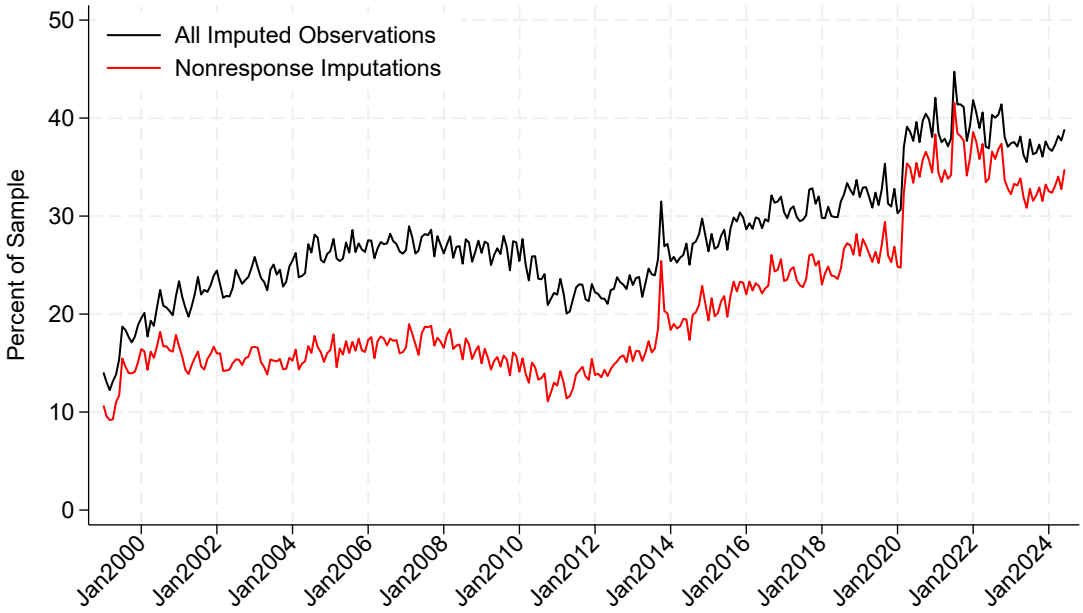


Figure 1. Share of Observations that are Imputed. *Note:* Observations are imputed due to nonresponse or vacancy. *Source:* Authors calculations using the BLS Housing Survey.

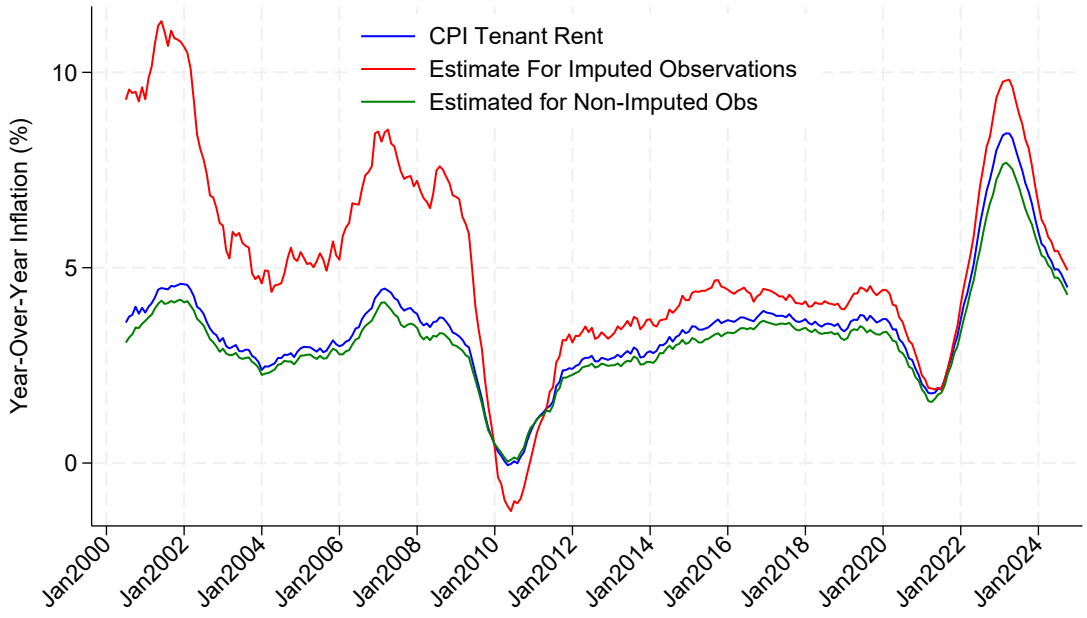
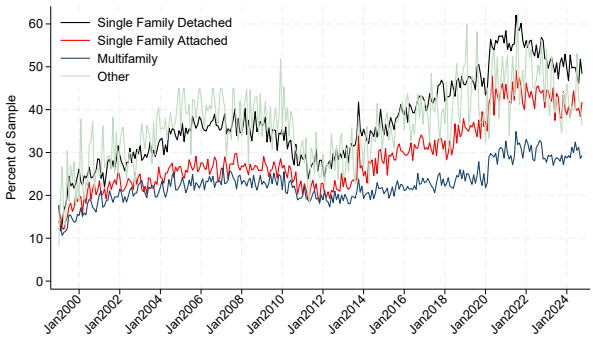
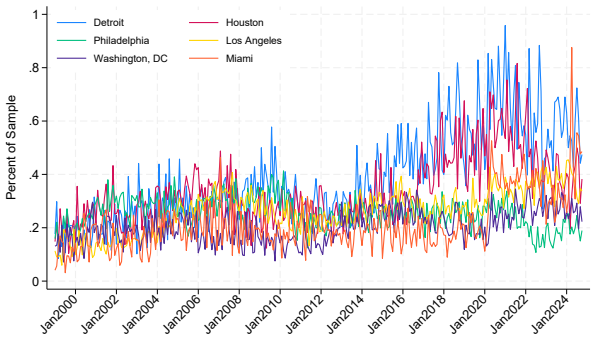


Figure 2. Inflation Rates for Imputed and Non-Imputed Observations. *Note:* Inflation is calculated using the method used for official CPI shelter inflation, but limiting the observations to imputed and non-imputed observations. *Source:* Authors' calculations using BLS Housing Survey.

(a) By Property Type



(b) By City



(c) By Tenure Length (Years in Unit)

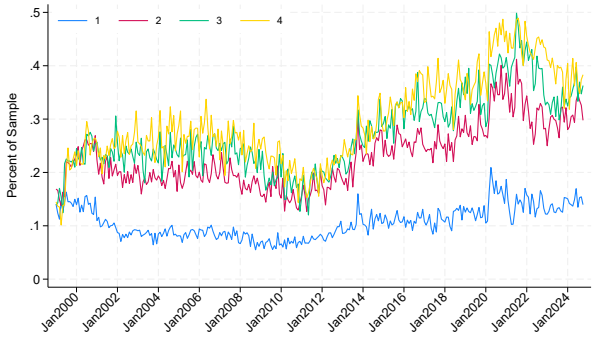


Figure 3. Imputation Rates by Category