working paper

16 30

Panel Data Estimates of Age–Rent Profiles for Rental Housing

Joshua Gallin and Randal J. Verbrugge

FEDERAL RESERVE BANK OF CLEVELAND

Working Paper 16-30 December 2016

Panel Data Estimates of Age–Rent Profiles for Rental Housing

Joshua Gallin and Randal J. Verbrugge

This paper provides estimates of the net depreciation rate for rental housing using a unique confidential data set from the Bureau of Labor Statistics that covers over 30,000 rental units from 1998 to 2009. Our data and econometric approach allow us to add to the literature in three main ways. First, we can control for unobserved quality (including cohort effects) by allowing for unit-specific fixed effects. Our results suggest that estimates of the depreciation rate for rental housing that ignore unobserved heterogeneity suffer from omitted-variable bias and potentially from selection bias, and that these biases can be large. Second, we use a dummy-variable approach to estimate aging effects, thereby avoiding ad hoc assumptions about functional form. We find that rent for a typical housing unit at first falls rapidly with age, flattens, and then begins to rise with age for older units. This nonmonotonic pattern is a feature of many other studies of both age–rent and age–price profiles for housing, and it seems likely that the upward-sloped portion of such profiles is the result of unobserved improvements and changes in style. We show that the upward slopes of our estimated profiles are only partially eliminated when we attempt to control for major improvements by excluding housing units that see very large jumps in rent. Third, we present estimates of age–rent profiles separately for different types of structures and for different regions of the country.

Keywords: Housing depreciation, rental housing, housing economics.

JEL classification: R31,E31,E21.

Suggested citation: Gallin, Joshua, and Randal J. Verbrugge, "Panel Data Estimates of Age–Rent Profiles for Rental Housing," Federal Reserve Bank of Cleveland, Working Paper no. 16-30.

Randal J. Verbrugge is at the Federal Reserve Bank of Cleveland (randal.verbrugge@clev.frb.org). Joshua Gallin is at the Board of Governors of the Federal Reserve System (joshua.gallin@frb.gov). A portion of this work was completed while Randal Verbrugge was at the Bureau of Labor Statistics.

Introduction

Accurate estimates of depreciation rates for housing are necessary for a wide range of issues. For example, such estimates are needed for empirical studies of residential investment, and as an input into macro-housing models. Estimates of housing depreciation are required in studies of urban decline and renewal (Rosenthal, 2008), affordable housing (Quigley and Raphael, 2004), neighborhood filtering (Margolis, 1982), and rent control and tax policy (Malpezzi, Ozanne, and Thibodeau, 1987). Furthermore, housing depreciation is an ingredient for measuring house-price changes (Harding, Rosenthal, and Sirmans, 2007), consumer price inflation (Randolph, 1988; Gallin and Verbrugge, 2007), and national income and wealth (Leigh, 1980).

This paper provides new estimates of the net depreciation rate for rental housing using a unique and confidential data set from the Bureau of Labor Statistics (BLS) that covers over 30,000 thousand rental units from 1998 to 2009. The data are collected by the BLS to construct measures of rent inflation and, ultimately, overall inflation. For our purposes, we define net depreciation as the change in rent as age increases, holding constant other basic characteristics such as the number of rooms and the type of structure. Our data and econometric approach allow us to extend the literature in three main ways.

First, we can control for unobserved quality by allowing for individual fixed effects. While the importance of unobserved heterogeneity for estimating depreciation rates is clearly not a new idea (for example, see Randolph (1988)), studies of depreciation for rental housing have essentially bypassed the issue.[1] We show that estimates of the rental housing depreciation rate that ignore unobserved heterogeneity appear to suffer from omitted-variable bias and selection bias, and that these biases can be large. In particular, we find that estimates from a pooled regression, such as those in Randolph (1988), imply a net depreciation rate of about 0.25 percent per year during the first 99 years of a rental unit's life. Conversely, our model with unit-specific fixed effects finds only a roughly 0.1 percent rate of net depreciation per year. We argue that this discrepancy in estimates is consistent with higher-quality units being disproportionately removed from the stock of rental housing over a unit's life span through demolition or

[1] There a large literature on the depreciation in the context of repeat-sales house-price indexes (Harding, Rosenthal, and Sirmans, 2007).

conversion to other uses (such as owner-occupied housing). Removal of such units leaves relatively more unimproved or otherwise lower-quality housing units in the housing stock, which—in a pooled regression—biases the estimate of net depreciation upwards.

Second, by using age-specific fixed effects to estimate the age profile, we avoid reliance on ad hoc assumptions about the functional form of the relationship between age and rent common in the earlier literature.[2] We find that rent initially declines with a unit's age, but then rises with age for older units. Such non-monotonicity is a common finding in the literature (Coulson and McMillen, 2008; Clapp and Giocottto, 1998; Chinloy, 1978; and Gillingham, 1975). We present evidence that suggests that a large portion of this observed non-monotonicity stems from unobserved improvements to rental units during the sample period. Unit-specific fixed effects cannot absorb the effects of such improvements, but we argue that we can partially eliminate this bias by removing from our sample those units that experienced a very large rent change during the sample period. We think that very large rent increases are associated with unobserved renovations.

Third, we present estimates of age-rent profiles separately for different types of structures and for different regions of the country. By structure, we find a fairly clear ordering through the first 60 or so years of a structure's life: Single-family detached homes depreciate relatively little during the first three decades of life, and then experience essentially no subsequent net depreciation; single-family attached homes depreciate somewhat more than do detached homes, and then experience some net appreciation in later years; homes in apartment buildings with elevators depreciate at a rate similar to that of single-family attached homes, but then appreciate significantly in later years; and homes in apartments in buildings without elevators depreciate the most until about age 40 and then show some modest appreciation.

By region of the country, we find all regions show a similar qualitative pattern—an early period of rapid net depreciation, a middle period of roughly zero net depreciation, and a later period of net appreciation. We also find that, during the first 40

[2] Chinloy (1978) is an early paper that recognized that age-value profiles are not always monotone decreasing, which suggests that net "depreciation" for rental housing can actually be negative over some age ranges. Gallin and Verbrugge (2007) demonstrated that imposing typical functional-form constraints can cause misleading inferences.

years of a unit's life, rent declines least with age in the Midwest, followed by the Northeast, the South, and the West. We also examine the interaction of region and structure effects. We find that age-rent profiles by structure cannot be easily "explained" by the region profiles and that the region profiles cannot be explained in terms of the structure profiles: It appears that neither structure type nor location is a more "elemental" determinant of net depreciation than the other.

Econometric Considerations and Previous Research

Most studies of age-rent of age-price profiles can be represented as a version of

$$r_{it} = \lambda_t + \beta x_{it} + F(a_{it}) + q_{it} + u_{it} \tag{1}$$

where r_{it} is log rent, λ_t are aggregate year effects, x_{it} is a vector of observable characteristics such as the number of rooms and the type of structure, and a_{it} is the age of the housing unit; i indexes housing units and t indexes time.[3] We define q_i as unit-specific unobserved quality, which can include any unmeasured fixed characteristic that affects rents, including location.[4] The year effects term λ_t captures any common time-varying influences on rent, such as overall inflation, business-cycle influences, and changes in the relative price of housing. The "error-term" u_{it} captures any idiosyncratic time-varying characteristics that affect rent. Note that this term can include unobserved maintenance, changes to the structure, and unobserved changes to the neighborhood. The differences among various researchers' econometric approaches can largely be summarized in terms of their modeling choices for age effects, $F(\cdot)$, and unobserved quality, q_i.

Our paper follows in the steps of Chinloy (1979) in that we define net depreciation as the change in rent as age increases, holding constant other basic characteristics such as the number and types of rooms.[5] As such, our definition is an amalgam of physical depreciation, unobserved maintenance, and minor improvements.

[3] For the remainder of the paper we refer to all profiles as age-rent profiles when the meaning is clear.
[4] Note that the equation does not include an explicit "cohort" effect. Cohort effects are subsumed in the individual effects, and cannot be separately identified regardless, because our specification allows for age and time effects.
[5] While gross depreciation is also of interest, few data sets are rich enough to allow for their estimation.

Modeling Age Effects

Several early papers on the hedonics of rental housing imposed that depreciation is constant or that age effects enter as a quadratic, so that the net depreciation rate (defined as the partial derivative with respect to age) is a linear function of age (see Malpezzi, Ozanne, and Thibodeau, 1987 for a review of the early literature). A handful of exceptions exist, and these often come to different conclusions: age-rent and age-price profiles need not be quadratic, and indeed many not even be monotonic (Chinloy, 1977 and 1978; Goodman and Thibodeau, 1995; Malpezzi, Ozanne, and Thibodeau, 1987; Shilling, Sirmans, and Dombrow, 1991; Clapp and Giaccotto, 1998). Chinloy (1978), Clapp and Giaccotto (1998) and Lee, Chung, and Kim (2005) all provided theoretical explanations for why net depreciation could be negative over some age ranges. Chinloy (1978) argued that adverse selection could cause "perverse shapes for estimated depreciation functions". Clapp and Giaccotto (1998) argued that age coefficients include both a depreciation component and a demand-side component that changes over time. Lee, Chung, and Kim (2005) argued that older units are more likely to be redeveloped and that for older units, this redevelopment effect can outweigh normal net depreciation.

Although flexible functional forms are a useful way to summarize the shape of the age-rent profile, a highly flexible parametric approach may lead to false impressions about the linearity and monotonicity of the age-rent relationship. For example, a high-order polynomial has the flexibility to fit a highly nonlinear function. However, the estimated function can often exhibit large errors over ranges where there are few data points in order to fit well over ranges where there are a lot of data. In addition, estimates from parametric specifications can be sensitive to observations that have a lot of leverage. For example, Gallin and Verbrugge (2007) found that some outliers had significant leverage in a cross-sectional study of the BLS rent data.

Our approach is instead to estimate the age-rent profile "non-parametrically" as a set of age dummy variables. That is, for the specification in equation (1),

$$F\left(a_{it}\right)=\sum_{j=1}^{A}\delta_j D_{it}(j) \tag{1}$$

where δ_j is an age effect for each age group $j = 1,\ldots, A$ and $D_{it}(j)$ is a dummy variable equal to 1 if $\text{age} = j$. This approach is extremely flexible. However, it is costly in terms of degrees of freedom, and that cost goes up with the number of age groups, A. In addition, estimates based on a large number of age groups could be too noisy to interpret. One of the advantages of our approach is that we have a large panel data set, and can therefore afford to be somewhat profligate with our degrees of freedom.

Modeling Unobserved Quality and Selection

If the unobserved quality terms, q_i and u_{it}, are correlated with age, then estimates of the age effects that ignore unobserved heterogeneity will be biased (Chinloy, 1977; Randolph, 1988). It is worth describing the specific ways that this well-known problem of omitted-variables bias can be a problem in the context of rental housing.

Unobserved quality can be correlated with age because design philosophies and construction materials and methods change with time. In other words, cohorts of older rental units may differ from cohorts of younger units in ways that are unobservable in the data. If so, then even if all homes in all cohorts follow identical age-rent profiles, cross-section estimates of this profile will be biased. The simplest case is to imagine that older homes are of lower quality because, say, rooms are too small by current standards. In this case, the conditional age-rent profile estimated using a cross-section will be biased toward finding a negative slope.

Sample selection and sample attrition can also pose a problem to a cross-section approach if structures are torn down or converted to other uses in a way that varies systematically with the quality of the unit. To understand this effect, begin by noting that the cross-section sample available in 2000 is composed only of units that have survived until 2000. But not all rental homes built, say, in the 1920s survived until 2000; some were torn down and others were converted into other uses. Rent level is surely correlated with quality, and sample attrition of this nature is also likely correlated with quality. These correlations can bias conventional cross-section estimates.

For example, suppose that annual deterioration is randomly distributed around some constant and that physical depreciation is unavoidable and irreversible. If lower-quality homes are more likely to be torn down, units which experienced below-average physical deterioration will be more likely remain in the sample. In other words, surviving structures would differ from other units in their cohorts—their torn-down "siblings"—in a systematic way. In this case, a cross-section estimate of net depreciation would understate the true average rate of depreciation. Conversely, if higher-quality rental units are more likely to be converted to owner-occupied housing or some other use, a cross-section estimate of net depreciation would overstate the true average rate of depreciation because rental units that randomly experienced better—that is, smaller—depreciation shocks will tend to leave the sample, leaving behind only the "unlucky" units that depreciated more. These processes are likely to occur over the course of decades.

This sample-selection issue is often addressed in the cross-section literature by making extreme assumptions about the price of units which drop out of the sample (Hulten and Wykoff, 1981). Randolph (1988a) argued that cohort and individual effects could be effectively proxied by detailed observable neighborhood characteristics, which amounts to an assumption that omitted variable bias is not a problem once these control variables are included.

A Panel Approach

We address the problem of unobserved quality by using fixed-effects panel estimation from the 12-year panel. To our knowledge, this study is the first to apply a large geographically diverse panel approach to the age-rent profile issue. The panel approach addresses the problem that unobserved quality is likely correlated with age by allowing intercepts to vary across housing units. Putting this differently, q_i is estimated, and other coefficient estimates are conditional on this estimate. Our approach does *not* implicitly dismiss the likelihood that older homes have undergone significant renovations since they were built, because the assumption is merely that q_i is fixed *over the 12-year sample period*. Thus, q_i for a given home could be very different today from what it was 20 or 40 years ago. In addition, our fixed-effect approach addresses the selection problem by allowing for surviving structures to differ from torn-down or converted "siblings" in systematic ways.

Of course, some rental units no doubt experienced a major renovation during the sample period. To the extent that such renovation affected the unit's unobserved quality, our results will be distorted. Although we cannot identify the majority of renovations, we suspect that large rent changes are associated with large changes in unobserved quality. Thus, we address this issue by estimating the age-rent profiles on the full sample and on subsamples which exclude all observations for rental units that had a very large rent change in *any* year of the sample.

Profile Specifications

Our main regression specification is a fixed-effect model,

$$r_{it} = \lambda_t + q_i + \beta x_{it} + \sum_{j=1}^{A} \delta_j^f D_{it}(j) + u_{it} \tag{2}$$

where the variables are defined as in equation (1), except that the unit-specific unobserved quality q_i is estimated using fixed effects. We estimate the equation using standard panel methods for the entire country, by type of structure, and by Census region. The age groups j = 1, …, 34 are defined by three-year age bins:

$$D_{it}(j) = \begin{cases} 1 & \text{if } 3(j-1) < age_{it} \leq 3j \\ 0 & \text{otherwise} \\ & \forall j - 1, 2, \ldots, 33 \end{cases}$$

$$D_{it}(34) = \begin{cases} 1 \text{ if } age_{it} > 100 \\ 0 \quad \text{otherwise} \end{cases}$$

We also estimate an alternative specification that follows Randolph's (1988a) neighborhood approach.

$$r_{it} = \lambda_t + \theta z_n + \beta x_{it} + \sum_{j=1}^{A} \delta_a^p D_{it}(j) + v_{it} \tag{3}$$

Note that equation (4) does not have individual fixed effects, but does include a vector of time-invariant neighborhood characteristics z_n . We agree that using multiple proxy measures for unobserved variables is generally better than doing nothing (for example, see Bollinger and Minier, 2015); a comparison of the results from equations (3) and (4) will display the extent to which, in the present context, such proxies address the omitted-variable bias and selection issues that are relevant here. We refer to the estimated age effects from equation (3) as the fixed-effects age-rent profile and the estimated age effects from equation (4) as the pooled age-rent profile.

Data

The data were collected on a confidential basis by the BLS as part of their program to measure rent inflation for the CPI. For each of the metropolitan areas in the sample, the BLS randomly selected a geographically-diverse set of rental housing units via a geographic stratification procedure. The BLS then collected information about each unit, such as its age, structural characteristics, number of bedrooms and bathrooms, and utilities (including whether utilities are included in the rent). The housing sample is divided into six panels. Rent data on all the units in a particular panel are collected in the same month twice a year; in particular, panel 1 is “priced” in January and July, panel 2 in February and August, and so on. A typical unit remains in the sample for many years. In addition, changes in the structural characteristics of the unit over the sample period—such as the addition of new rooms—are recorded. See Ptacek and Baskin (1996) for more details. The BLS data we use are from 1998 to 2009. We merged the BLS data with Census tract-level data to provide fixed neighborhood effects as in Randolph (1988).

Figure 1 provides greater detail about the age distribution of rental housing. Most rental housing in our sample is between about 5 and 40 years old, and the distribution is skewed to the right. Figure 2 displays the distribution of structure types. Most rental housing can be found in smaller multifamily (apartment) buildings that do not have elevators, and in single-family detached homes. Single-family attached homes and multifamily buildings with elevators are less common kinds of rental housing. Figure 3 shows the regional distribution of housing units in the sample.

Figure 1

Age Distribution of Rental Units, 1998-2009

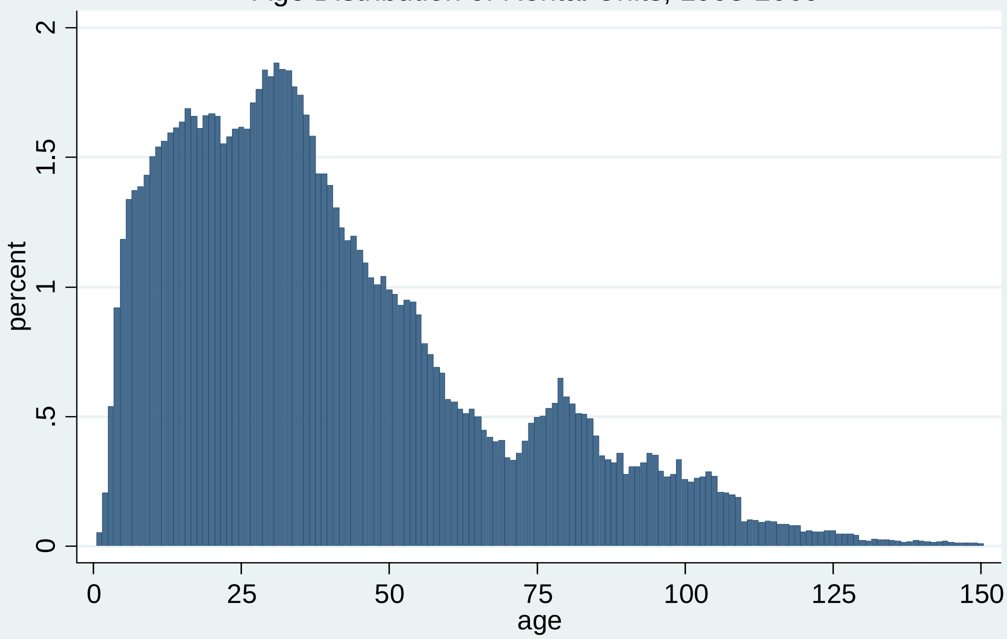

Figure 2

Structure Distribution of Rental Units, 1998-2009

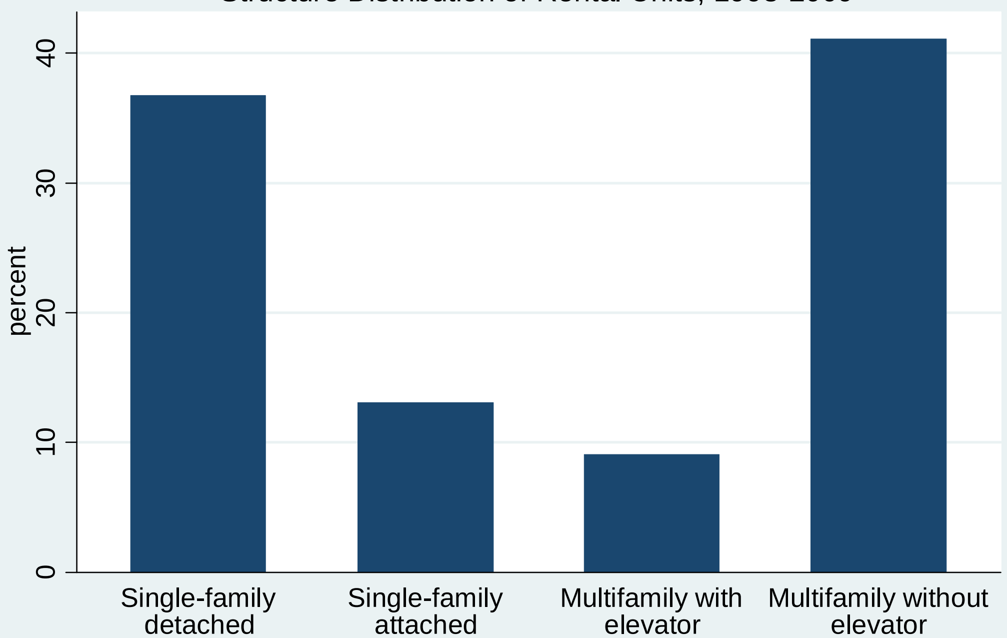

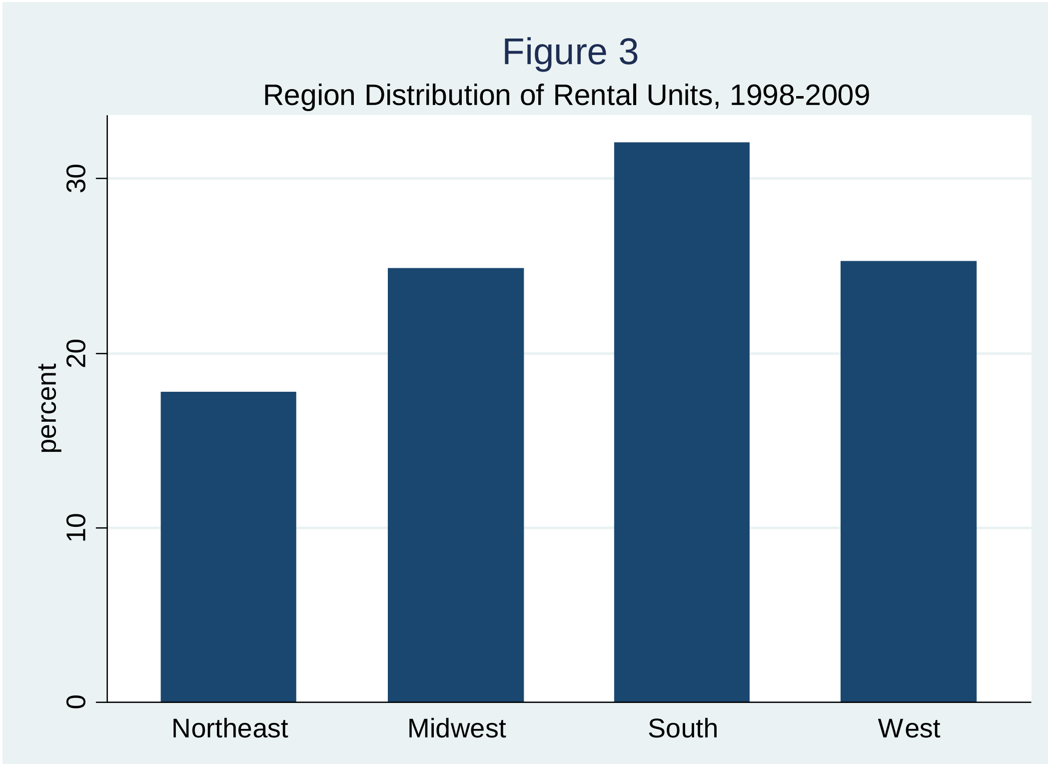

Results

National-level

Our main results for the estimated age effects are presented in Figure 4 and table 1.[6] The blue line presents the estimated age effects, $\hat{\delta}_a^p$, from the pooled regression in equation (4) in which, as discussed above, we use neighborhood characteristics rather than individual fixed effects to control for unobserved quality (as in Randolph, 1988). The black line presents the estimated age effects, $\hat{\delta}_a^f$, from the fixed-effect age-rent profile in equation (3). For both profiles, the area between the relevant dashed lines represents 95 percent confidence regions.

The figure presents the estimated age effects for the three-year age groups; the age dummy for 100 years and older was omitted, so we do not identify the *level* of the age effects. Note that Figure 4 displays cumulative net depreciation by age in log points

[6] The estimated age effects are the focus of this paper. The controls in x_{it} include items such as the number and type of rooms and the type of heating and cooling. The controls in z_n include items such as percent of population in the neighborhood that fall into demographic groups based on age, race, education, and poverty. See Appendix Table C.1 for estimated β coefficients from equations (3) and (4).

relative to a base of zero, which implies that the net depreciation rate is the *slope* of the line (see Table 1). For example, if depreciation were exponential ($R_a = R_0 e^{-\delta a}$), then the age-rent profile would be a straight line with slope $-\delta$, and intercept would be normalized to zero: $\ln R_a - \ln R_0 = -\delta a$.

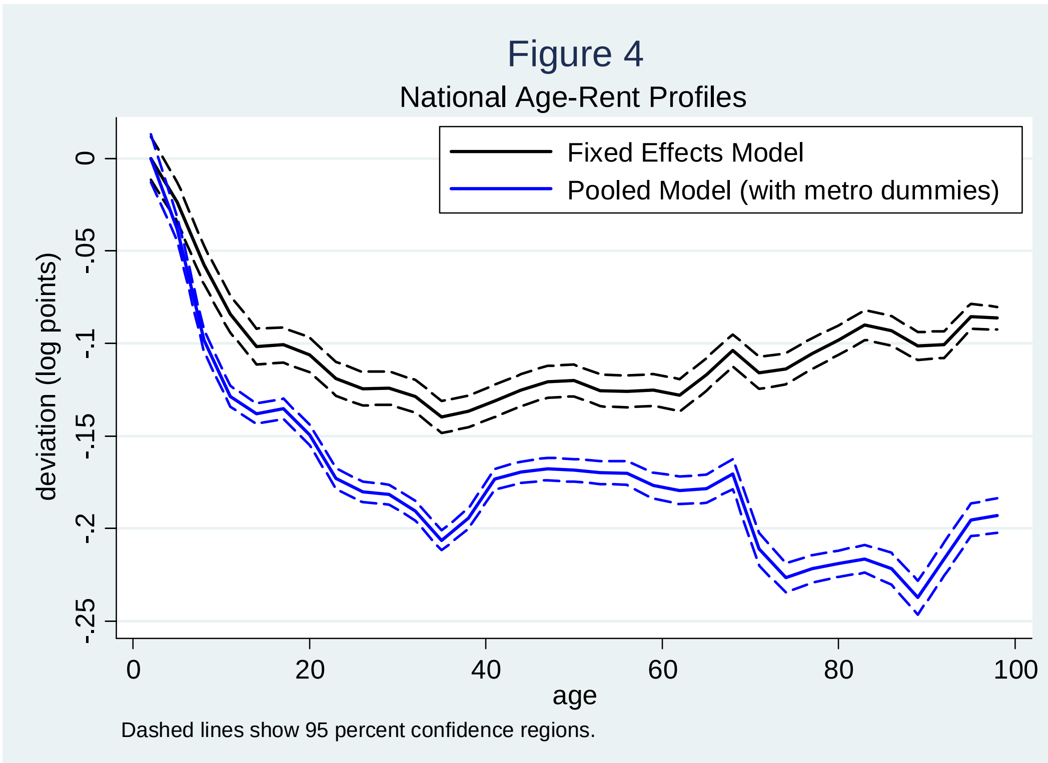

Figure 4 has two salient features, deriving from two distinct but related sources. The first feature is that the fixed-effect estimates of the age-rent profile show a shallower slope for the first 40 years of a structure's life than does the profile from the pooled regression. The differences are statistically significant and quantitatively large, as we discuss below. The divergence in profiles implies that the cross-sectional regression estimate is biased because unit quality—above and beyond that measured in x_i—is unobserved by us and hence omitted. To illustrate, consider a simplified model of rent in which r_{it} is determined according to $r_{it} = \alpha + \beta_1 age_{it} + \beta_2 x_i + \beta_3 q_i + \varepsilon_{it}$, where q_i is unit-specific quality that is unchanged over time for as long as the unit exists as a rental. We emphasize that aging of a unit exerts an influence on rent that is distinct from the influence of unit-specific quality. If q_i is omitted from the regression specification, the expected value of β_1 is given by $E\beta_1 = \beta_1 + \beta_3 \frac{\text{Cov}(age_{it}, q_i)}{\text{Var}(age_{it})}$. From this expression, we can see that the bias could be upwards or downwards, depending upon the correlation of age_{it} and q_{it} in the sample. In the present case, the pooled depreciation rate estimate is biased upwards (too much depreciation, i.e. the estimated effect of aging, β_1, is biased downwards). Thus our results indicate that the correlation between age and unit quality is negative. The perhaps surprising implication is that higher-aged units tend to be of lower unit-specific quality, independently of the aging process itself. There are two potential explanations of this correlation. The two explanations need not be mutually exclusive, and cannot be distinguished in our data. The first possibility is a cohort vintage effect: over time, unit-specific quality has been rising. The second possibility is that units with higher unobservable quality are more likely to be removed from the rental universe. This could happen either through conversions—higher quality units being converted to owned dwellings, possibly after extensive (unobserved) renovations—or demolitions—units in neighborhoods that are beginning to get trendy are more likely to undergo a "scrape-and-build." Note that the biases could have gone the other way. For instance, if lower quality rental housing tended to be removed from the rental stock, the pooled age-rent profile would have been elevated relative to the fixed-effect profiles, so that the associated estimate of the net depreciation rate would have been biased downwards. These results point to the importance of unobserved quality in estimating age-rent profiles: Whether one controls for unobserved quality of housing units appears to make an important difference for how one should characterize net depreciation of the rental housing stock. The panel approach gracefully includes unobserved unit quality, regardless of the source of this unchanged unit quality, in the fixed effect.

Table 1
Annualized Net Depreciation Rates: Nation
(percent)

	Pooled	Fixed Effect	
Age	Full sample	Full sample	Restricted Sample
0 to 9	1.09	0.64	0.57
10 to 18	0.07	0.19	0.06
19 to 27	0.34	0.20	0.11
28 to 45	-0.06	0.00	-0.05
46 to 63	0.05	0.01	0.03
64 to 99	0.04	-0.12	-0.05

Note: The restricted sample excludes all housing that experienced any rent increases during the sample period of greater than .3 log points.

Focusing now only on the fixed-effect results, the second salient feature of Figure 4 is that while rents appear to depreciate quite rapidly for young housing units (the "Fixed Effect/Full Sample" column of Table 1), depreciation essentially ceases, on average, for rental units between 28 and 63 years, and that rents actually *appreciate*—that is, the estimated depreciation is negative—for older units. Our finding that the age-rent profile is non-monotonic is consistent with the results of Chinloy (1978), Clapp and Giaccotto (1998), and Lee, Chung, and Kim (2005). However, our results are based in part on the assumption that unobserved quality is fixed during the twelve-year period covered by our sample. As mentioned previously, a renovation could easily violate this assumption. Some renovations are observable in principle. For example, an addition of rooms would be identifiable in the data. However, a home that was fully gutted and refurbished but did not experience a change in observable characteristics in the data such as an increase in rooms or the addition of central air would be observationally equivalent in our data to its previous, less desirable self, since this type of renovation would not be recorded by the BLS. With this in mind, we tested the sensitivity of our results to this assumption by examining the impact of excluding all the observations of any rental unit that had rent a change exceeding various thresholds in any period. Gallin and Verbrugge (2013) show that average rent changes for most rental units that experience a rent change are below 3 percent. This suggests that large rent changes are associated with increases in unobserved quality. We cannot be sure that such units experienced a change in their unobserved quality. However, we think it highly likely that units that experienced outsized rent increases were in fact renovated.

Figure 5 presents the estimated fixed-effect age-rent profiles for the full sample (in black, repeated from Figure 4) and a restricted subsample that excludes all observations for units that experienced *any* rent increases during the sample period of greater than .3 log points, or about 35 percent (in red).[7] The exclusion of units that likely experienced renovations clearly has a large effect on the estimated profile: While younger units appear to experience about the same amount of net depreciation in the three samples, units appear to experience much less net appreciation at higher ages (see also the “Fixed Effects/Restricted Sample column in Table 1). Thus, the result of “perverse shapes for estimated depreciation functions” (Chinloy, 1978) appears to be, at least in part, caused by renovations of older units that affect unobserved quality. As noted above, this study aims to provide estimates of the net depreciation rate, rather than the gross depreciation rate.

[7] The restricted fixed-effect regression had about 6 percent fewer observations that did the full-sample regression.

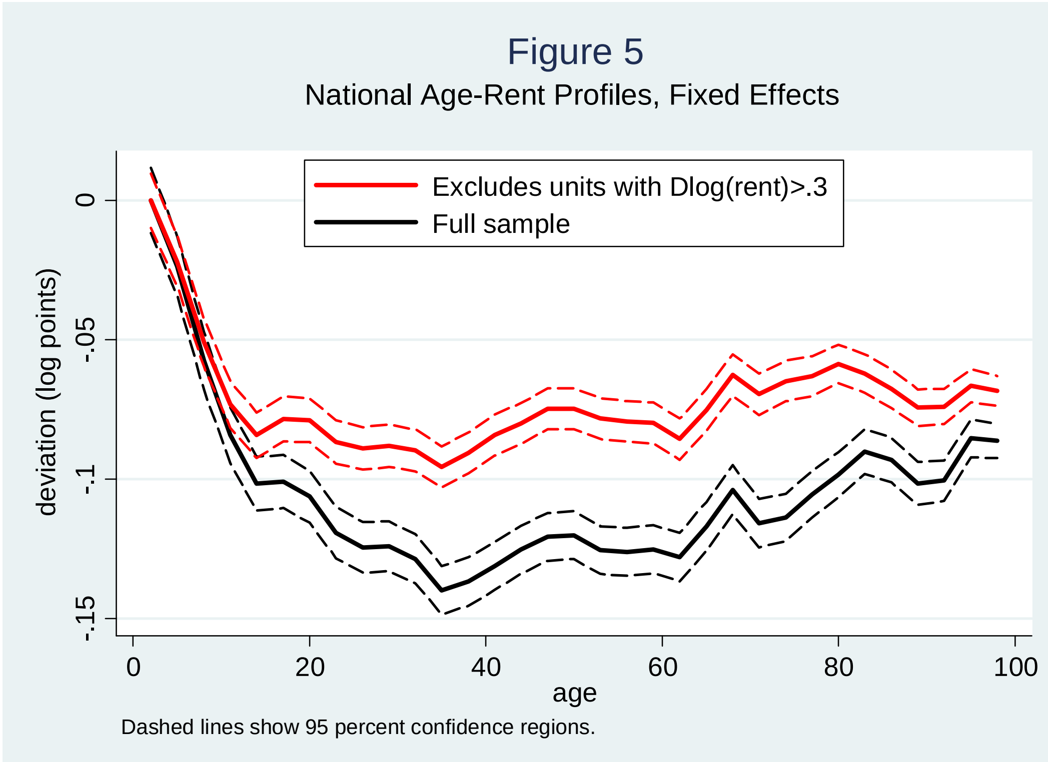

Figure 6 and Table 2 display age-rent profiles separately for four types of structures: single-family detached, single-family attached, apartment buildings with an elevator, and apartment buildings without an elevator; see the appendix for figures which include standard errors. The most striking feature of the profiles is that different structures have such different profiles, with fairly clear ordering through the first 60 or so years of a structure's life: single-family detached homes depreciate relatively little during the first three decades of life, and then experience essentially no net depreciation after that (the black line in Figure 6 and column 1 of Table 2). This shallow depreciation profile is consistent with the fact that land is a larger component of the rent for this structure type. Single-family attached homes depreciate somewhat more, and then experience some net appreciation in later years (the blue line in the figure and column 3 in the table). Homes in apartment buildings with elevators depreciate at a rate similar to that of single-family attached homes, experience no net depreciation from age 40 to 60, and then appreciate significantly in later years (the red line and column 5). Finally, homes in apartments in buildings without elevators depreciate the most till about age 40 and then show some modest appreciation (the green line and column 7).

Table 2
Annualized Net Depreciation Rates: by Structure
(percent)

Age	Single-family detached: Full sample	Single-family detached: Restricted Sample	Single-family attached: Full sample	Single-family attached: Restricted Sample	Apartment with elevator: Full sample	Apartment with elevator: Restricted Sample	Apartment without elevator: Full sample	Apartment without elevator: Restricted Sample
0 to 9	0.2	0.2	0.7	0.6	0.2	0.3	1.1	0.9
10 to 18	0.1	0.0	0.1	-0.1	0.2	0.0	0.6	0.3
19 to 27	0.1	0.1	0.1	0.0	0.9	0.7	0.3	0.1
28 to 45	0.0	0.0	0.1	0.0	-0.1	-0.1	0.2	0.0
46 to 63	-0.1	0.0	-0.1	0.0	-0.4	-0.3	-0.1	-0.2
64 to 99	0.0	0.0	-0.6	-0.3	0.0	0.0	-0.1	0.0

Note: The restricted sample excludes all housing that experienced any rent increases during the sample period of greater than .3 log points.

Figure 6
Age-Rent Profiles by Structure Type
(fixed effects)

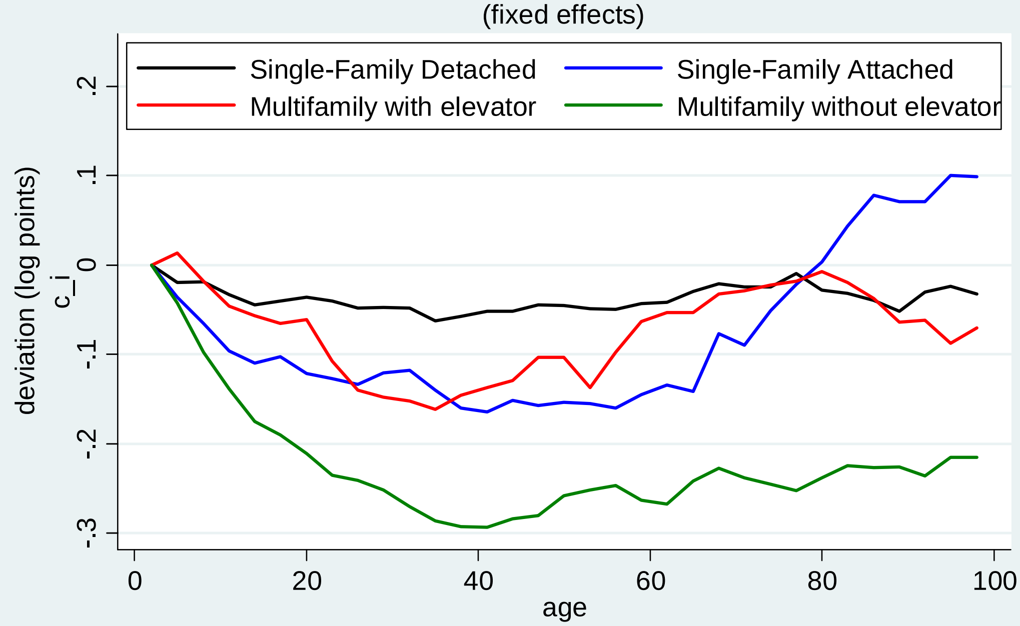

Figure 7 and Table 3 display the age-rent profiles for the four major regions of the country: the Northeast (in black), the Midwest (in blue), the South (in red), and the West (in green); see the appendix for region definitions and standard errors. All regions show a similar qualitative pattern—an early period of rapid net depreciation, a middle period of roughly zero net depreciation, and a later period of net appreciation. However, the pattern varies by region. In particular, there is an economically significant variation in the rate of net depreciation for young housing units, with rental units the South and West showing the most net depreciation and those in the Midwest showing the least. For older units, the pattern is more jumbled, with the West experiencing the most net appreciation and the Northeast the least.

Table 3
Annualized Net Depreciation Rates: by Region
(percent)

	Northeast		Midwest		South		West	
Age	Full sample	Restricted Sample	Full sample	Restricted Sample	Full sample	Restricted Sample	Full sample	Restricted Sample
0 to 9	0.5	0.3	0.1	0.4	0.7	0.6	0.8	0.7
10 to 18	-0.1	-0.1	0.2	0.1	0.3	0.1	0.1	0.0
19 to 27	0.3	0.4	0.2	0.1	0.2	0.1	0.2	0.1
28 to 45	0.1	0.0	0.0	-0.1	-0.1	-0.1	0.0	0.0
46 to 63	0.1	0.1	0.1	0.0	-0.1	0.0	-0.1	0.0
64 to 99	-0.1	0.0	-0.1	0.0	0.0	-0.1	-0.2	-0.1

Note: The restricted sample excludes all housing that experienced any rent increases during the sample period of greater than .3 log points.

Figure 7
Age-Rent Profiles by Region
(fixed effects)

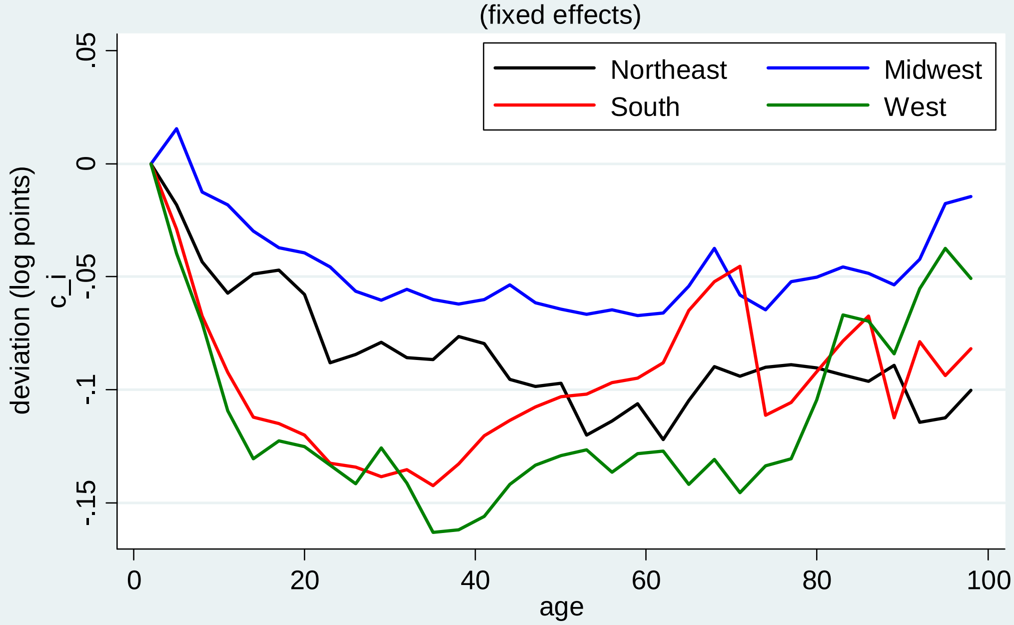

A full explanation of the different estimated rates of depreciation by structure and region would require a description of the interactions of climate, engineering, tenant behavior, and landlords' maintenance choices. Such an explanation is beyond the scope of this paper. However, we can use our estimated profiles and the distribution of units by structure and location to examine how structure type and location interact. For example, if we assume that age-rent profiles vary only by structure, we can ask if differences in regional profiles can be explained in terms of structure profiles and regional differences in the distribution of structures. Similarly, if we assume that age-rent profiles vary only by region, we can ask if structure profiles be explained in terms of regional profiles and variation in the location of different structure types.

To answer these questions, let $\hat{\delta}_{r,a}$ be the estimated fixed-effect age-rent profile for region r, let $N_{s,a}$ be the number of housing units of structure type s, and $N_{s,r,a}$ be the number of housing units of structure type s that are found in region r, where $N_{s,a} = \sum_r N_{s,r,a}$. Then one can construct

$$\tilde{\delta}_{s,a} = \sum_r \hat{\delta}_{r,a} \frac{N_{s,r,a}}{N_{s,a}} \tag{4}$$

as an estimate of the age-rent profile for structure type s based solely on age-rent profiles by region and the region distribution by structures. Similarly, one can construct

$$\tilde{\delta}_{r,a} = \sum_s \hat{\delta}_{s,a} \frac{N_{r,s,a}}{N_{r,a}} \tag{5}$$

as an estimate of the age-rent profile for region r based solely on age-rent profiles by structure and the structure distribution by regions.

Figures 8 present the structure estimates from equation (4) and Figure 9 presents the region results from equation (5).[8] The figures show that the age-rent profiles by structure are not well "explained" by the region profiles and the region profiles are not well explained by the structure profiles. To put it another way, the age-rent profile in, for example, the South does not look the way it does because of the underlying nature of depreciation rates for different types of structures interacted with the types of structures found in the south. Similarly, the age-rent profile for single-family detached homes, for example, does not look as it does because of the underlying nature of depreciation in different parts of the country interacted with where such homes are typically found. Thus, in a rough sense at least, it appears that neither structure type nor location is particularly "elemental" as a determinant of net depreciation.

[8] The oldest multifamily units in the south in our sample are 93 years old.

Figure 8

Age-rent profiles by region: The explanatory power of structure profiles

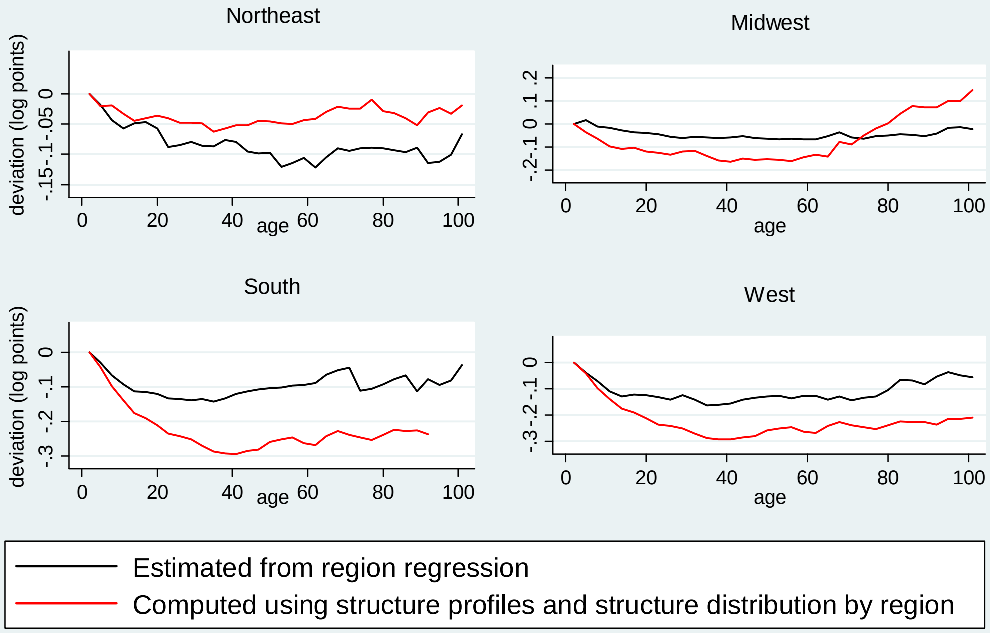

Figure 9

Age-rent profiles by structure type: The explanatory power of region profiles

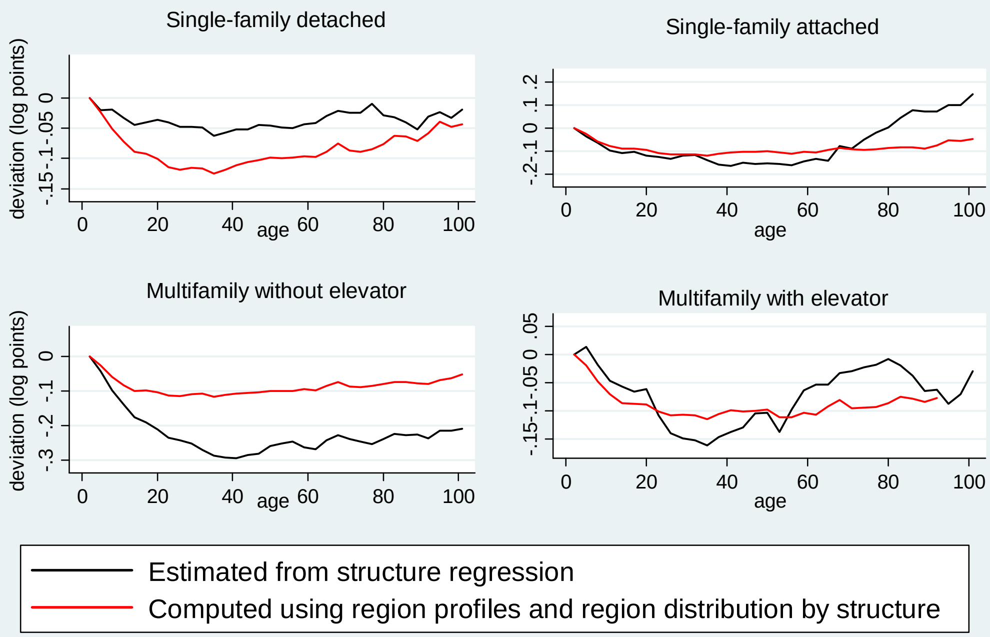

Conclusion

In this paper we use a large panel of rental housing units to estimate age-rent profiles for such properties. There is a long literature on this subject, and researchers have grappled with two main issues. The first is unobserved quality: Two units might have the same age, floor plan, location, etc., and thus be observationally equivalent to the econometrician, but have completely different unobserved quality because of differences in say, kitchen cabinets, flooring, and general upkeep. The second is sample selection: Rental homes can be torn down, redeveloped, or converted to owner-occupied housing. Both unobserved quality and sample selection are likely correlated with age, and thus—if ignored—will contaminate estimates of the age-rent profiles.

We control for unobserved quality and selection effects by including in our regressions unit-specific fixed effects and by examining the effect of excluding from our regressions units with large rent changes. Our regressions with unit-specific fixed effects suggest that older, higher quality rental units are more likely to be removed from the sample—presumably through demolition or conversion to other uses (such as to owner-occupied status)—than they would be if randomly selected.

These contributions to the literature are possible because of the quality and nature of the BLS rental sample. In particular, the dataset is large, follows units over time, and includes a significant number of structure characteristics. Data on unit characteristics allow us to control for observable improvements such as the addition of rooms. The panel aspect of the data allows us to control for a significant portion of unobserved quality by allowing for unit-specific fixed effects. The size of the dataset allow us to exclude units that showed large rent gains for any period, which eliminates units that had large renovations that greatly enhanced the rental's unobservable (to us) quality. Our results show that estimates of the depreciation rate for rental housing that ignore unobserved heterogeneity and selection suffer from significant biases.

Despite our controls for unobserved heterogeneity and sample selection, the puzzle of the "perverse shapes for estimated depreciation functions" (Chinloy, 1978), though lessened, remains. One interpretation of the data and results is that we are unable to completely control for unobserved improvements to older units *during the sample period*. In particular, that a fair bit of the non-monotonicity in age-rent profiles disappears when units with large rent increases are eliminated suggests that unit-specific effects are not completely fixed. Thus, although our results show the importance of controlling for unobserved heterogeneity and selection effects when estimating net depreciation of the rental housing stock, a finer identification of renovations or improvements is still desirable.

References

[1] Bollinger, Christopher R. and Minier Jenny. 2015. "On the Robustness of Coefficient Estimates to the Inclusion of Proxy Variables." *Journal of Econometric Methods*, 2015, 4(1).

[2] Chinloy, Peter T. "Hedonic Price and Depreciation Indexes for Residential Housing: A Longitudinal Approach." *Journal of Urban Economics*, October 1977, v. 4, iss. 4, pp. 469-82.

[3] Chinloy, Peter T. 1978. "Depreciation, Adverse Selection and Housing Markets." *Journal of Urban Economics*, April 1978, v. 5, iss. 2, pp. 172-87.

[4] Chinloy, Peter T. 1979. "Hedonic Price and Depreciation Indexes for Residential Housing Again." *Journal of Urban Economics*, April 1979, v. 6, iss. 2, pp. 272.

[5] Chinloy, Peter T. 1979. "The Estimation of Net Depreciation Rates on Housing." *Journal of Urban Economics*, October 1979, v. 6, iss. 4, pp. 432-43.

[6] Clapp, John M.; Giaccotto, Carmelo. 1998. "Residential Hedonic Models: A Rational Expectations Approach to Age Effects." *Journal of Urban Economics*, November 1998, v. 44, iss. 3, pp. 415-37.

[7] Coulson, N. Edward; McMillen, Daniel P. 2008. "Estimating Time, Age and Vintage Effects in Housing Prices." *Journal of Housing Economics*, June 2008, v. 17, iss. 2, pp. 138-51.

[8] Gallin Joshua and Verbrugge Randal. 2013. "Sticky Rents." Working Paper, 2013.

[9] Gallin, Joshua; Verbrugge, RandaI. 2007. "Improving the CPI's Age-Bias Adjustment: Leverage, Disaggregation and Model Averaging." U.S. Bureau of Labor Statistics, Working Papers: 410.

[10] Gillingham, Robert. 1975. "Place to Place Rent Comparisons." *Annals of Economic and Social Measurement*, Winter 1975, v. 4, iss. 1, pp. 153-73.

[11] Goodman, Allen C. and Thibodeau, Thomas G. 1995. "Age-Related Heteroskedasticity in Hedonic House Price Equations." *Journal of Housing Research*, 1995, v. 6, iss. 1, pp. 25-42.

[12] Harding, John P.; Rosenthal, Stuart S.; and Sirmans, C. F. 1988. "Depreciation of Housing Capital, Maintenance, and House Price Inflation: Estimates from a Repeat Sales Model." *Journal of Urban Economics*, March 2007, v. 61, iss. 2, pp. 193-217.

[13] Hulten, Charles R.; Wykoff, Frank C. 1981. "The Estimation of Economic Depreciation Using Vintage Asset Prices: An Application of the Box-Cox Power Transformation." *Journal of Econometrics*, April 1981, v. 15, iss. 3, pp. 367-96.

[14] Lane, Walter F.; Randolph, William C.; Berenson, Stephen A. "Adjusting the CPI Shelter Index to Compensate for Effect of Depreciation." *Monthly Labor Review*, October 1988, v. 111, iss. 10, pp. 34-37.

[15] Lee, Bun Song; Chung, Eui-Chul; Kim, Yong Hyun. 2005. "Dwelling Age, Redevelopment, and Housing Prices: The Case of Apartment Complexes in Seoul." *Journal of Real Estate Finance and Economics*, February 2005, v. 30, iss. 1, pp. 55-80.

[16] Leigh, Wilhelmina A. 1980. "Economic Depreciation of the Residential Housing Stock of the United States, 1950-1970." *The Review of Economics and Statistics*, Vol. 62, No. 2. May, 1980, pp. 200-206.

[17] Malpezzi, Stephen; Ozanne, Larry; and Thibodeau, Thomas G. 1987. "Microeconomic Estimates of Housing Depreciation." *Land Economics*, November 1987, v. 63, iss. 4, pp. 372-85.

[18] Margolis, Stephen E. 1982 "Depreciation of Housing: An Empirical Consideration of the Filtering Hypothesis." *Review of Economics and Statistics*, February 1982, v. 64, iss. 1, pp. 90-96

[19] Ptacek, Frank; Baskin, Robert M. 1996. "Revision of the CPI Housing Sample and Estimators." *Monthly Labor Review*, December 1996, v. 119, iss. 12, pp. 31-39.

[20] Quigley, John M. and Raphael, Steven. 2004. "Is Housing Unaffordable? Why Isn't It More Affordable?" *Journal of Economic Perspectives*, Winter 2004, v. 18, iss. 1, pp. 191-214

[21] Randolph, William C. 1988. "Estimation of Housing Depreciation: Short-term Quality Change and Long-term Vintage Effects." *Journal of Urban Economics*, March 1988, v. 23, iss. 2, pp. 162-78.

[22] Shilling, James D.; Sirmans, C. F.; Dombrow, Jonathan F. 1991. "Measuring Depreciation in Single-Family Rental and Owner-Occupied Housing." *Journal of Housing Economics*, December 1991, v. 1, iss. 4, pp. 368-83.

Appendix A: Charts

Figure A.1
Age-Rent Profiles by Structure Type
(fixed effects)

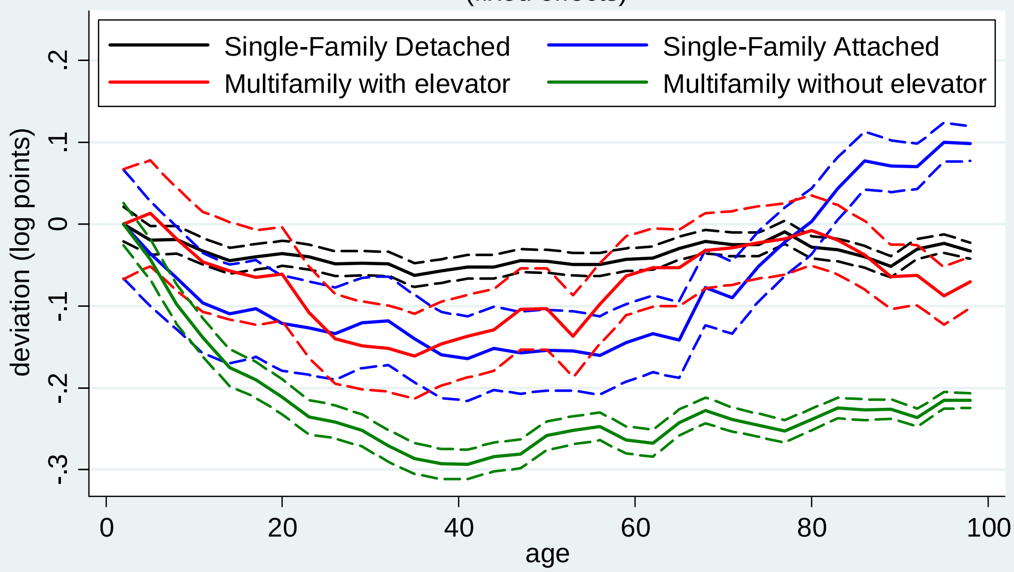

Dashed lines show 95 percent confidence regions.

Figure A.2
Age-Rent Profiles by Region
(fixed effects)

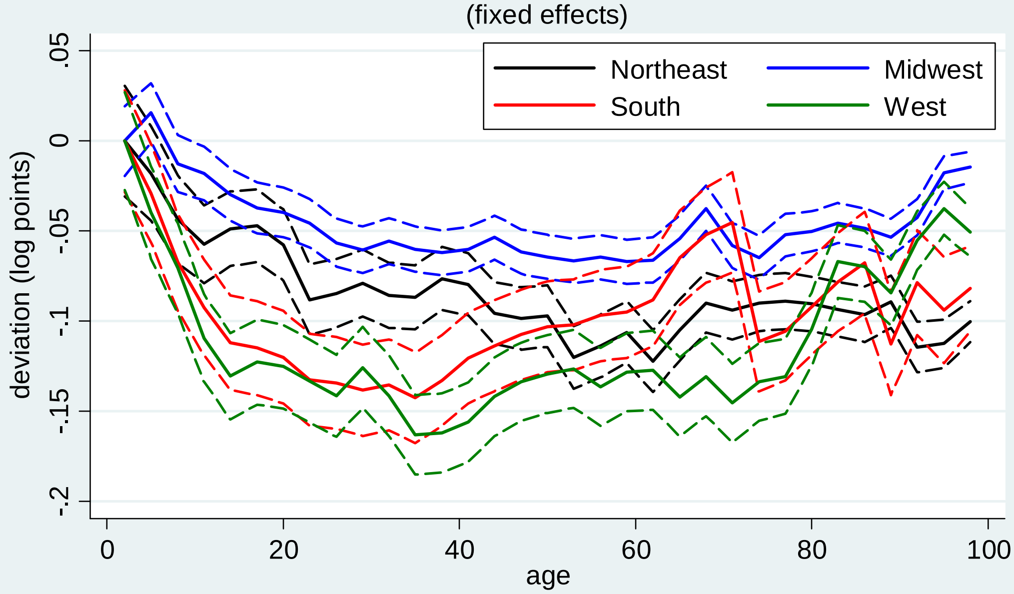

Dashed lines show 95 percent confidence regions.

Appendix B: Census Region Definitions

Northeast: Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, Vermont, New Jersey, New York, Pennsylvania.

Midwest: Illinois, Indiana, Michigan, Ohio, Wisconsin, Iowa, Kansas, Minnesota, Missouri, Nebraska, North Dakota, South Dakota.

South: Delaware, District of Columbia, Florida, Georgia, Maryland, North Carolina, South Carolina, Virginia, West Virginia, Alabama, Kentucky, Mississippi, Tennessee, Arkansas, Louisiana, Oklahoma, Texas.

West: Arizona, Colorado, Idaho, Montana, Nevada, New Mexico, Utah, Wyoming, Alaska, California, Hawaii, Oregon, Washington.

Appendix C: Estimates of hedonic controls

Table C.1
Hedonic Regressions
(1998 to 2009)

	Fixed Effects (equation 3)	Pooled (equation 4)
Electric heating	0.0018	0.0152
	(0.0012)	(0.0011)
Natural gas heating	0.0230	-0.1107
	(0.0039)	(0.0056)
Other heating	-0.0070	-0.1349
	(0.005)	(0.0052)
Central air	0.0311	0.1423
	(0.0021)	(0.0015)
Other air	-0.0043	0.0718
	(0.002)	(0.0017)
Window air	-0.0128	-0.0084
	(0.0014)	(0.0015)
Heat included	0.0057	0.0597
	(0.0015)	(0.0015)
Electricity included	0.0243	0.0859
	(0.0019)	(0.002)
Free parking	0.0105	0.0063
	(0.0017)	(0.0016)
Number of bathrooms	0.0622	0.1812
	(0.0076)	(0.0036)
Number of bathrooms squared	-0.0062	-0.0181
	(0.002)	(0.001)
Number of bedrooms	0.0561	0.1644
	(0.0037)	(0.0016)
Number of bedrooms squared	0.0014	-0.0034
	(0.0007)	(0.0003)
Number of other rooms	0.0361	0.0659
	(0.0031)	(0.0019)
Number of other rooms squared.	-0.0051	-0.0040
	(0.0006)	(0.0004)
R^2	0.945	0.687
Observations	448,307	416,361

Note: The standard error is displayed below the coefficient estimate. Estimated time dummy variables, capturing rent inflation over time, are nearly identical and not reported.

Pooled regressions include the following time-invariant controls: type of structure; building size; percent of population with some college, under poverty level, of school age, 65 or older, white, living in mobile home, and living in rental housing. Coefficient estimates are available from the authors upon request. Seven percent of the observations in the pooled regression were dropped due to missing data.