Mimeo to:
"Financial Intermediaries as Suppliers of Housing Quality"*

Michael Reher†

April 2019

Abstract

This document is a companion to the paper "Financial Intermediaries as Suppliers of Housing Quality". I describe a methodology for computing quality-adjusted rent that does not require detailed property hedonic information. The index is rooted in a discrete choice model and borrows from quality-adjustment techniques used in the macroeconomics and trade literatures. The results suggest that quality accounts for 86% of rent growth over 2010-16. I discuss implications of this quality adjustment for inequality in housing costs, official rent indices, and asset pricing in real estate.

Keywords: Housing Quality, Housing Rents

JEL Classification: R30, R31

*I thank Trepp for providing me with academic pricing.

†University of California San Diego, Rady School of Management. Email: mreher@g.harvard.edu

1 Introduction

The paper "Financial Intermediaries as Suppliers of Housing Quality" (Reher 2019), henceforth FI, describes a structural approach to computing quality-adjusted rent as an alternative to the hedonic methodology used in that paper. I now describe such an approach and discuss the associated results. In Section 2, I derive the index and discuss results related to aggregate rent growth. I turn to cross-sectional results and the implications of this adjustment for official rent indices in Section 3. Appendices A, B, C have econometric, mathematical, and additional details.

2 Structural Rent Index

There are three advantages to introducing a preference structure. First, a structural methodology can be applied to datasets which lack detailed information on hedonic characteristics but still enable housing units to be ranked by quality (e.g. via property inspection ratings). Second, a structural approach enables inferences about time variation in absolute quality. For example, it can account for how dishwashers installed in 2016 may be better (e.g. less noisy) than those installed in 2010. This is more difficult to do with a hedonic approach because of practical limitations on the set of observables. Third, a structural approach allows rent growth in different quality segments to have a different impact on the resulting index, with weights based on households' revealed preference. For example, if supply-driven improvements make high quality units less expensive, the index weights top-tier rent growth according to households' implied willingness to move up the quality ladder.

In terms of underlying logic, a hedonic index holds the distribution of quality fixed, whereas the structural index holds the average household's utility fixed and asks how she must be compensated for changes in rent across the quality ladder. In particular, the notion of effective rent is the expenditure required to obtain a unit of housing utility (i.e. compensating variation). For reasons of space, I defer details on setup and implementation to Appendix 2. I compute the structural index using the Trepp data, which is described in FI.

2.1 Setup

Let $i \in \mathcal{I}$ index properties. Define housing quality h_i as a Cobb-Douglas aggregator of space s_i and amenities a_i such that $\log(h_i) = \mu \log(s_i) + \log(a_i)$. As discussed shortly, $\mu > 0$ will govern preferences for quality. Next, let H denote the highest quality in the market, $H \equiv \sup_{i \in \mathcal{I}} \{h_i\}$, which I will refer to as absolute quality. Finally, define the unit's quality segment as $\hat{h}_i \equiv \frac{h_i}{H} \in \{0, ..., 1\}$. For example, segment $\hat{h}_i = 1$ corresponds to units in Class A properties.

Households, indexed by j, are endowed with income y_j. They have additive random preferences over their choice of shelter with flow utility

$$u_{i,j} = \log(h_i) + \epsilon_{i,j}, \tag{1}$$

where $\epsilon_{i,j}$ is a taste shock. Incorporating consumption does not materially change the analysis.[1] Online Appendix B shows how these preferences give rise to a discrete choice problem where households choose a shelter to maximize a geometric average of quality h_i, personal appeal $\epsilon_{i,j}$, and inverse rent.

The structural index aims to track the utilitarian welfare associated with (1) or, equivalently, the dollar cost required to maintain this welfare at a fixed level. This dollar cost, which I call "welfare-relevant rent", is the sum of areas to the left of the Hicksian market demand curve for each segment $\hat{h}$, and its functional form depends on the distribution of $\epsilon_{i,j}$.[2] For the baseline exercise, $\epsilon_{i,j}$ follows a type 1 extreme value, or Gumbel distribution, which implies that the demand curves in each segment have a constant elasticity of substitution (CES) form.[3] The parameter $\sigma \equiv \mu + 1$ governs the shape of the market demand curve. When σ is high, individual households have less preference for quality and view units in different segments as substitutable.

The following proposition shows how, given σ, one can compute the growth in welfare-relevant rent using observed the rent and market share of each segment.[4] This growth is the social cost (i.e. compensating variation) associated with a change in the distribution of rent

[1] See Online Appendix B for details. All proofs are in Online Appendix B.

[2] Focusing on the market demand curve is standard technique for studying price and welfare in durable goods markets, such as automobiles. See Anderson, de Palma and Thisse (1992) or Eaton and Kortum (2002) for examples with CES market demand. Berry, Levinsohn and Pakes (2004) study automobile consumers with preferences similar to (1) when taste shocks do not imply CES market demand.

[3] See Online Appendix F. This is a well-known result due to Anderson, de Palma and Thisse (1992).

[4] As discussed below, there are a number of ways to partition the market into quality segments $\hat{h}$. I use official property inspection scores in my baseline analysis.

across the quality ladder.[5]

Proposition 2.1 (Structural Rent Index) *The compensating variation associated with a change in rent from t_0 to t is*

$$\pi_t^S = \exp\left[\sum_{\hat{h}\in\mathcal{H}} w_{\hat{h},t} \log\left(\frac{Rent_{\hat{h},t}}{Rent_{\hat{h},t_0}}\right)\right] \times \left[\left(\frac{Rent_{1,t}}{Rent_{1_0,t}}\right)^{\sigma} \frac{Share_{1,t}}{Share_{1_0,t}}\right]^{-\frac{1}{\sigma-1}} \equiv DQ_t \times GQ_t^{-\frac{1}{\sigma-1}}, \quad (2)$$

where $\mathcal{H} \subseteq [0,1]$ is the set of quality segments, $Rent_{\hat{h},t}$ and $Share_{\hat{h},t}$ are the rent and share of total units in segment $\hat{h}$ and year t; the Sato-Vartia weights $w_{\hat{h},t}$ are a function of $Rent_{\hat{h},t}$ and $Share_{\hat{h},t}$; and segment 1_0 contains units that were in segment 1 in year t_0.

The term DQ_t in (2) depends on the distribution (hence "D") of rent across the quality ladder.[6] It reflects how rent growth affects the average household differently depending on whether growth is at the top ($\hat{h}$ large) or bottom ($\hat{h}$ small) of the ladder, where each segment's weight $w_{\hat{h},t}$ encodes households' willingness to move to that segment. Next, GQ_t is growth (hence "G") in absolute quality.[7] When rent on top tier units is higher than units that were top tier in t_0 (i.e. $\text{Rent}_{1,t}$ large), it reflects an increase in absolute quality (e.g. less noisy dishwashers), and especially so if households have less preference for quality (i.e. σ large). However, if the relative share of top tier units is low (i.e. $\text{Share}_{1,t}$ small), then then this rent premium does not reflect absolute quality, but rather a scarcity of newly renovated units. Finally, growth in absolute quality dampens effective rent because $\sigma > 1$, but this effect is (exponentially) discounted by $\frac{1}{\sigma-1}$: when σ is large, households attach less value to quality, and thus its impact on effective rent growth is weak.

2.2 Implementation

I compute the structural rent index π_t^S in (2) using the Trepp data, which cover multifamily properties over 2010-2016. These data have detailed property improvement records that help me

[5] Strictly speaking, "compensating variation" is the difference in welfare-relevant rent (i.e. the market's minimized cost function) following a change in the observed distribution of rent. I will misuse the term slightly and refer to the growth in welfare-relevant rent as "compensating variation".

[6] The term is standard in price indices with CES market demand and has its origins in Diewert (1976), Sato (1976), and Vartia (1976). See Feenstra (1994) or Broda and Weinstein (2006) for additional discussion.

[7] This term relies on a similar insight as Redding and Weinstein (2018), who also exploit the CES first order condition to obtain an expression for the change in quality. However, Redding and Weinstein (2018) do not impose a hierarchy of quality and do not allow quality to grow on average.

identify the index's key parameter, σ, which I estimate using three strategies: (1) a property-level strategy utilizing idiosyncratic variation in payment timing; (2) a zip code level version of the property-level strategy; and (3) a zip code level strategy based on a widely-used GMM estimator proposed by Feenstra (1994). Due to space constraints, I defer details on these strategies to Appendix A.1. The average estimate for σ is 6.5, as shown in Table A1. Figure A1 performs an introspective exercise with photographs to interpret this magnitude.

The second piece of information needed to compute π_t^S is a sorting variable to partition units into quality segments $\hat{h}$. I partition the market using official property inspection ratings conducted by the Mortgage Bankers Association and Commercial Real Estate Finance Council (MBA/CREFC), although the results are similar when sorting properties by effective age.[8]

2.3 Results

Figure 1 summarizes growth in the structural rent index π_t^S and various related indices over 2010-2016. The baseline CES index saw real growth of 0.2% compared to 1.4% growth in average rent, which implies that improving quality can account for 86% of 2010-2016 real rent growth.[9] As before, the effect is stronger when benchmarking to an age adjusted index similar to that used by statistical agencies. Discounted growth in absolute quality, $GQ_t^{-\frac{1}{\sigma-1}}$, accounts for 73% of the wedge between age adjusted and structural rent growth.[10] This is consistent with descriptions of an "amenities arms race" in the apartment industry.[11]

In Appendix C, I show how to use a quasi hedonic methodology to infer growth in absolute quality, which can be applied to an arbitrary price index formula. This approach uses the same intuition from Proposition 2.1 that growth in absolute quality can be inferred from the premium of top tier units over units that were top tier in the base period. These non-CES indices all have real growth rates between 0.1% and 0.2%.

[8]This rating captures a property's quality relative to the top of its market, and it is regularly collected as part of the multifamily mortgage servicing protocol with the intent of minimizing agency frictions. Markets are defined as a geographic zone of competition and are between a county and an MSA in size. See FI for more details, including evidence that the MBA/CREFC rating does not suffer misreporting bias.

[9]The Trepp data are at the property-level, and so I weight properties by the number of units.

[10]Explicitly, fixing $GQ_t^{-\frac{1}{\sigma-1}} = 1$ leads to growth of 1.3% compared to 0.2% when using its estimated value. Growth in absolute quality therefore accounts for $\frac{1.3-0.2}{1.7-1.3} = 73\%$ of the wedge between age adjusted and structural rent growth.

[11]See the Washington Post article "An amenities arms race heats up in the apartment industry" (Orton, 2017).

Figure 1: Summary of Structural Rent Index

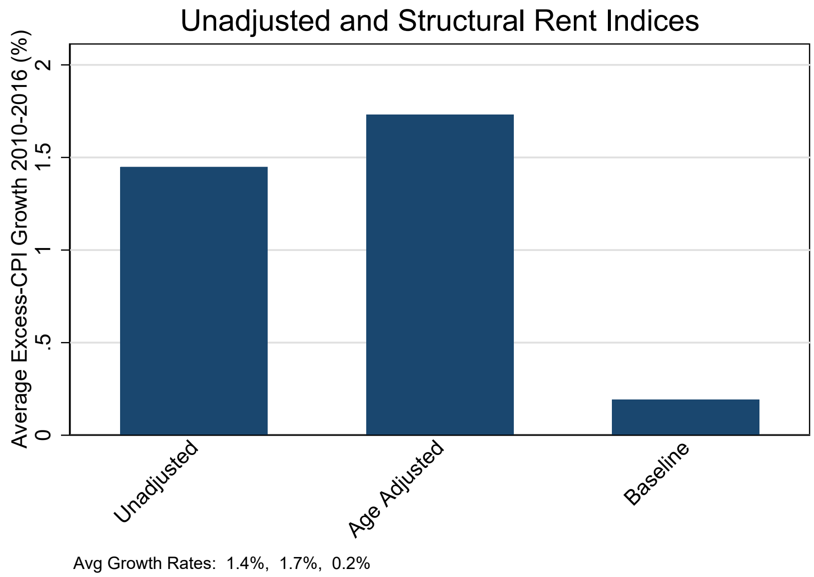

Note: This figure plots average annual growth in real rent over 2010-2016 for various rent indices. Unadjusted denotes average rent. Age Adjusted rent growth performs an age adjustment similar to that used by statistical agencies and is described in Appendix 3.3. Baseline denotes the structural index from (2). Data are from Trepp.

3 Implications of Accounting for Quality Improvements

I discuss implications of this quality adjustment for inequality in housing costs, official rent indices, and asset pricing in real estate.

3.1 Implications for Inequality in Housing Consumption

This extension asks how effective rent has varied by income. My focus is on heterogeneity across geographic markets and submarkets.[12] The structural index is the more natural tool for this exercise because it better accommodates heterogeneous valuations of quality, which would arise due to, say, non-homothetic preferences. I use the Trepp data because of its detailed information on property location.

For part of this analysis, I allow the preference parameter σ to vary in the cross-section.[13] Following Jaravel (2018), I partition the sample into brackets by zip code real income and then

[12] Markets are typically between a county and MSA in size, and submarkets are between a zip code and a county, per the MBA/CREFC property inspection guidelines.

[13] Appendix B.1 provides a microfoundation based on the notion that space s_i, unlike amenities a_i, is a necessity.

reestimate σ for each bracket.[14] The results in Table A2 show that the highest income zip codes have the lowest value of σ (4.9) and thus the highest willingness to pay for quality. Next, I partition the set of zip codes into an above and below median cohort according to average income over 2010-2016. Then I recompute π_t^S for the two cohorts using each zip code z's estimated preference parameter σ_z. One should interpret π_t^S as welfare-relevant rent for the average household in a given cohort.

Figure 2 plots real growth in real unadjusted rent and the structural index π_t^S over 2010-2016 by income cohort. Beginning with the left column, unadjusted real rent growth was, coincidentally, 1.3% for both cohorts. The middle column accounts for differences in quality, but constrains preferences to be the same. Whereas improving quality can explain all of real rent growth for the high income cohort, quality actually fell somewhat in low income markets. The right column relaxes the constraint on preferences, after which welfare-relevant rent growth falls by an additional 0.9 pps for the high income cohort. Altogether, household surplus from improving quality was 2.5 pps greater in high income markets, of which 64% (1.6 pps) was due to material changes in quality and 36% (0.9 pps) was due to greater preferences for it. The joint importance of quality and preferences is consistent with a model where investors make improvements where the equilibrium price of quality is highest.

Figure 3 obtains a similar finding when partitioning by within-MSA income, either averaged over 2010-2016 or as measured initially in 2010. The divergence in the gains from quality is strongest when partitioning by initial income, consistent with a view of "super gentrification" and the Guerrieri, Hartley and Hurst (2013) model of endogenous provision of amenities.

3.2 Implications for Asset Pricing

This extension studies how improvements vary by initial real estate valuations in a market. According to the logic of standard asset pricing, the cap rate (i.e. dividend-price ratio) should convey information about a unit's future rent (i.e. dividend) growth (Campbell and Shiller 1988).[15] To test this hypothesis, I sort zip codes − which I will call submarkets in this extension

[14] Income is measured using average adjusted gross income from the IRS. The real income brackets are 0-30%, 30-65%, and 65%-100%. I estimate σ with each of the three methodologies described in Appendix A and average across them.

[15] The cap rate equals the ratio of net operating income to appraised property value and is therefore similar to a dividend-price ratio.

Figure 2: Structural Rent Index by Income

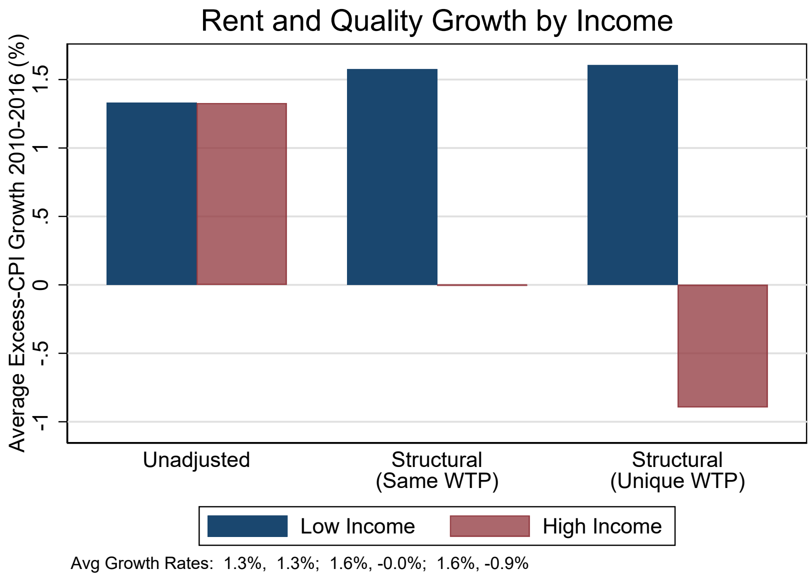

Note: This figure plots average 2010-2016 growth in unadjusted and structural rent indices for properties in zip codes with high or low income. High is defined as having average household income over 2010-2016 above the median across zip codes, and low is defined conversely. The leftmost column plots unadjusted average rent growth. The rightmost column plots growth in the baseline CES index in (2.1) using each zip code's estimated demand parameter σ_z. The middle column fixes σ_z at the average value for the low income cohort. Data are from Trepp.

− within each MSA according to the 2010 cap rate on multifamily properties and then compute quality-adjusted rent growth over 2010-2016 for zip codes with an above or below average value. The results in Figure 4 show how quality-adjusted rent growth was substantially lower in submarkets with a high initial cap rate (i.e. dividend-price ratio). This is consistent with rational expectations and the view that cheap properties necessitated substantial improvement to command their observed rent. By contrast, quality-adjusted rent growth was actually higher than observed growth in submarkets where the initial cap rate was lower. Panel (b) of Figure 4 sorts submarkets by house price decline during the 2006-2009 collapse and reveals a similar result: submarkets where property values fell by more during the crash saw subsequently greater improvements in housing quality.

Figure 3: Structural Rent Index by Income Partition

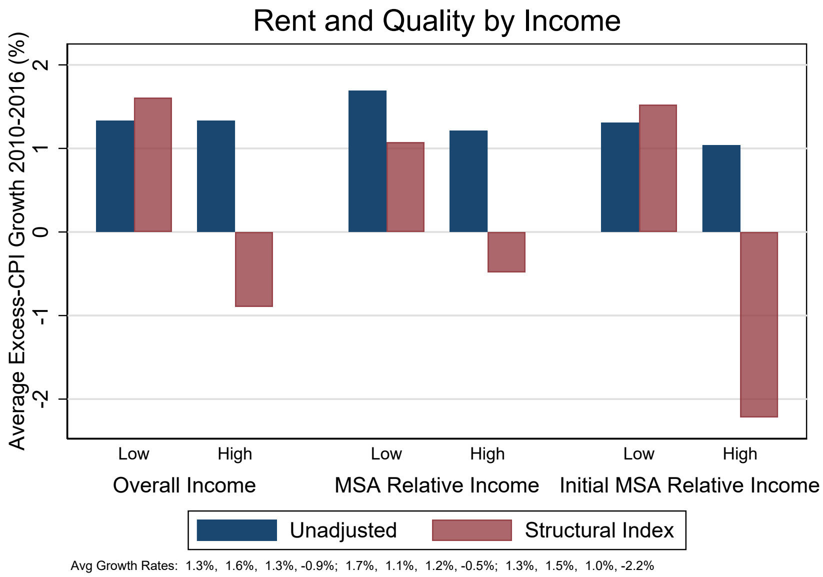

Note: This figure plots average 2010-2016 growth in unadjusted and structural rent indices for properties in zip codes with a high or low value of the indicated variable. High is defined as above median, and low is defined conversely. Overall Income denotes average household income over 2010-2016. MSA Relative Income denotes average household income over 2010-2016 after demeaning by the surrounding MSA. Initial MSA Relative Income denotes average household income in 2010 after demeaning by the surrounding MSA. The structural index is the baseline CES index in (2.1) using each zip code's income-based demand parameter σ_z. Data are from Trepp.

3.3 Relationship to Official Rent Indices

This section relates the results to what one would obtain from an age adjustment procedure similar to that used by statistical agencies.[16] Following Gallin and Verbrugge (2007), I define the age adjustment regression as

$$\log\left(\text{Rent}_{i,t}\right) = \gamma\left(\text{Age}_{i,t}; X_{i,t}\right) + u_{i,t}, \tag{3}$$

where $\text{Rent}_{i,t}$, $\text{Age}_{i,t}$, and $X_{i,t}$ are, respectively, a unit's rent, the age of the property, and a vector of structural features.[17] Then, one computes a unit's age adjusted rent as $\text{Rent}^A_{i,t} \equiv \text{Rent}_{i,t} e^{-\frac{\partial \gamma}{\partial \text{Age}_{i,t}}}$

[16] Age is the primary attribute the Bureau of Labor Statistics (BLS) corrects for when computing the Rent of Primary Residence (Ptacek 2013). The other corrections pertain to the changes in the inclusion of parking or utilities in rent, and the addition of a new room or central air conditioning.

[17] The function $\gamma\left(\text{Age}_{i,t}; X_{i,t}\right)$ approximates that used by the BLS as closely as possible given a different dataset. It includes age, its square, and its interaction with: the number of units in the property and an indicator for whether the property is over 85 years old. Since I do not observe a unit's location and thus neighborhood features in the AHS data, I estimate (5) as a panel regression and include a property fixed effect. When using

Figure 4: Forecasting Quality Growth with Submarket Indicators

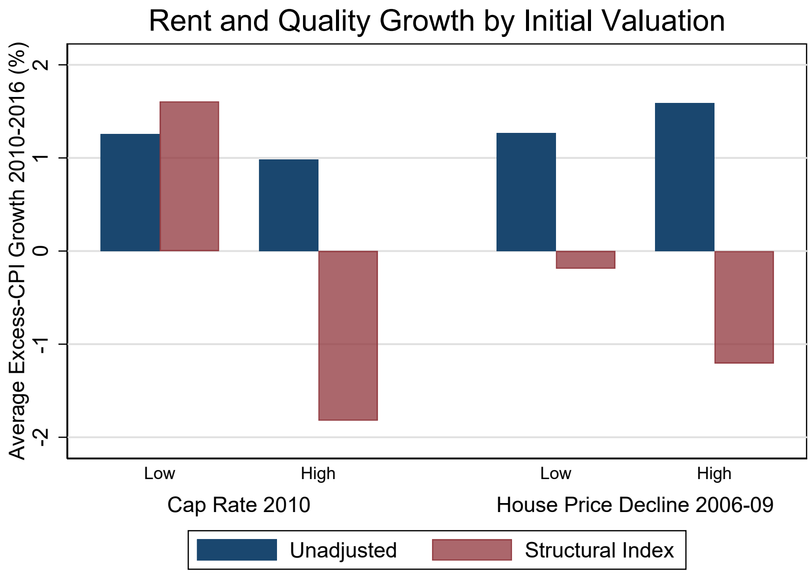

Note: This figure plots average 2010-2016 growth in adjusted and structural rent indices for properties in zip codes with a high or low value of the indicated variable. High is defined as above the average of the surrounding MSA, and low is defined conversely. House Price Decline 2006-09 denotes the change in the zip code level Zillow Home Value Index between 2006 and 2009. Cap Rate 2010 denotes the average ratio of net operating income to appraised value in the zip code in 2010. The structural index is the baseline CES index in (2.1). Data are from Trepp.

and aggregates $\text{Rent}^A_{i,t}$ across units to produce an average rent π^A_t that is benchmarked to the reference period,

$$\pi^A_t = \frac{\sum_{i \in \mathcal{I}} \text{Rent}^A_{i,t}}{\sum_{i \in \mathcal{I}} \text{Rent}_{i,t_0}}, \tag{4}$$

I next relate the results to what one would obtain from an age adjustment procedure similar to that used by statistical agencies.[18] I show age adjustments used to construct official indices can be biased upward and decompose the bias into two terms. The first term relates to the rate at which housing units depreciate, and the second relates to growth in absolute quality.

the Trepp data, I weight observations in (5) by number of units because the data are at the property-level.

[18]Age is the primary attribute the Bureau of Labor Statistics (BLS) corrects for when computing the Rent of Primary Residence (Ptacek 2013). The other corrections pertain to the changes in the inclusion of parking or utilities in rent, and the addition of a new room or central air conditioning.

First, following Gallin and Verbrugge (2007), I define the age adjustment regression as

$$\log\left(\text{Rent}_{i,t}\right) = \gamma\left(\text{Age}_{i,t}; X_{i,t}\right) + u_{i,t}, \tag{5}$$

where $\text{Rent}_{i,t}$, $\text{Age}_{i,t}$, and $X_{i,t}$ are, respectively, a unit's rent, the age of the property, and a vector of structural features.[19] Then, one computes a unit's age adjusted rent as $\text{Rent}^A_{i,t} \equiv \text{Rent}_{i,t} e^{-\frac{\partial \gamma}{\partial \text{Age}_{i,t}}}$ and aggregates $\text{Rent}^A_{i,t}$ across units to produce an average rent π^A_t that is benchmarked to the reference period, similarly to the expression for the hedonic index.

$$\pi^A_t = \frac{\sum_{i \in \mathcal{I}} \text{Rent}^A_{i,t}}{\sum_{i \in \mathcal{I}} \text{Rent}_{i,t_0}}, \tag{6}$$

The remainder of this section theoretically compares the structural and age adjusted rent indices π^S_t and π^A_t. Beginning with the setup described above, consider the following exercise. Suppose a unit's equilibrium rent is a function of its quality $h_{i,t}$.

$$\begin{aligned} \log\left(\text{Rent}_{i,t}\right) &= a_i + P \log\left(h_{i,t}\right) + u_{i,t} \\ &= a_i + P\left[\log\left(\hat{h}_{i,t}\right) + \log\left(H_t\right)\right] + u_{i,t}, \end{aligned} \tag{7}$$

where the notation is the same as in the paper with the addition of time subscripts, and $u_{i,t}$ is an iid shock. In particular, H_t is the highest quality in the market at t, which I call absolute quality, and $\hat{h}_{i,t} \equiv \frac{h_{i,t}}{H_t}$ is the relative quality of unit i. The parameter P is the equilibrium slope of the quality ladder, or price of quality. A unit's relative quality is also a function of its age,

$$\log\left(\hat{h}_{i,t}\right) = -\delta \text{Age}_{i,t} + v_{i,t}, \tag{8}$$

where δ is the rate of natural depreciation, and $v_{i,t}$ is not necessarily iid. The intuition is similar for more complicated depreciation schedules than (8). For the sake of argument, suppose all parameters in (7)-(8) are known, but quality $h_{i,t}$ is not observed. Let $\tilde{\mathbb{E}}$ denote a cross-sectional

[19] The function $\gamma\left(\text{Age}_{i,t}; X_{i,t}\right)$ approximates that used by the BLS as closely as possible given a different dataset. It includes age, its square, and its interaction with: the number of units in the property and an indicator for whether the property is over 85 years old. Since I do not observe a unit's location and thus neighborhood features in the AHS data, I estimate (5) as a panel regression and include a property fixed effect. When using the Trepp data, I weight observations in (5) by number of units because the data are at the property-level.

expectations operator: $\tilde{\mathbb{E}}[z_{i,t}] = \mathbb{E}[z_{i,t}|t]$. Then the age adjusted index (6) can be rewritten

$$\pi_t^A = e^{\delta P} \frac{\tilde{\mathbb{E}}[\text{Rent}_{i,t}]}{\tilde{\mathbb{E}}[\text{Rent}_{i,t-1}]}. \tag{9}$$

The relationship between π_t^S and π_t^A is described by the following proposition.

Proposition 3.1 (Bias in Age Adjusted Rent) *Suppose the total number of housing units is held fixed. Then age adjusted rent growth π_t^A is biased upward compared to the structural rent index π_t^S according to*

$$\frac{\pi_t^A}{\pi_t^S} = \left(\frac{H_t}{H_{t_0}}\right)^P \times \frac{\tilde{\mathbb{E}}[Rent_{i,t}]}{\tilde{\mathbb{E}}\left[Rent_{i,t} \times e^{-P\Delta v_{i,t}}\right]}. \tag{10}$$

The result in (10) decomposes bias in an age adjustment into two terms.[20] The intuition for the first term in (10) is that an age adjustment accounts for relative quality, not absolute quality. A top-tier unit 2010 may have lost no or very little quality by 2014, but it will still rent at a discount compared to a unit renovated to top-tier standards in 2014 if there is growth in absolute quality, $H_{2014} > H_{2010}$. The intuition for the second term is that age is an imperfect proxy for quality in the presence of improvement activity. To see this, first note from the depreciation process (8) that an improvement would generate a large disturbance term $\Delta\nu_{i,t} > 0$. Furthermore, because improvements move a unit up the quality ladder, there is positive covariance between $\Delta\nu_{i,t}$ and $\text{Rent}_{i,t}$. Together, these two features would make the second term in (10) greater than 1, leading to upward bias.

[20] The proof of Proposition 3.1 is as follows:

Proof Let $\bar{R}_{t_0}$ and $\bar{U}$ denote aggregate rent expenditure and welfare at t_0. Let $\bar{R}_t$ denote the minimized cost function in t from Lemma B.2 when target utility is $\bar{U}$. By definition, $\bar{R}_t$ is the aggregate expenditure required to maintain welfare at $\bar{U}$. Since the number of units is unchanged, $\bar{R}_t$ equals aggregate expenditure at t when holding each unit's quality fixed at its t_0 level,

$$\begin{aligned}\bar{R}_t &= \tilde{\mathbb{E}}\left[R_{i,t} e^{P\left[\Delta \log\left(\hat{h}_{i,t}\right) + \Delta \log(H_t)\right]}\right] \\ &= \tilde{\mathbb{E}}\left[R_{i,t} e^{P(\delta - \Delta v_{i,t} - \Delta \log(H_t))}\right],\end{aligned} \tag{11}$$

where the second line uses (8). Define $\pi_t \equiv \frac{\bar{R}_t}{\bar{R}_{t_0}}$. From the proof of Proposition 1 from the paper, $\pi_t^S = \pi_t$. Dividing each side of (11) by $\bar{R}_{t_0}$ gives expression in (10),

$$\frac{\pi_t^A}{\pi_t^S} = \left(\frac{H_t}{H_{t_0}}\right)^P \times \frac{\tilde{\mathbb{E}}[\text{Rent}_{i,t}]}{\tilde{\mathbb{E}}\left[\text{Rent}_{i,t} \times e^{-P\Delta v_{i,t}}\right]}$$

Recall that this exercise assumes knowledge of all parameters. If δP is unknown but estimated through OLS, then classical measurement error from the fact that age is an imperfect measure for relative quality would bias the estimated δP toward zero. This attenuation bias would partially offset the upward bias in (10).

Proposition 3.1 lays out two avenues for addressing upward bias in conventional age adjustments. First, collecting and incorporating information on improvement activity could reduce the measurement error from proxying quality with age. This could be accomplished by interviewing landlords, as opposed to the current practice of interviewing tenants. Ambrose, Coulson and Yoshida (2018) argue that reporting lags are an additional rationale for interviewing landlords. Second, addressing the bias from growth in absolute quality requires additional methodological tools. Several papers in the price adjustment literature (e.g. Redding and Weinstein 2018) are making progress on this margin.

References

Ambrose, B. W., Coulson, N. E. and Yoshida, J.: 2018, Housing Rents and Inflation Rates.

Anderson, S. P., de Palma, A. and Thisse, J.-F.: 1992, *Discrete Choice Theory of Product Differentiation*, The MIT Press.

Berry, S., Levinsohn, J. and Pakes, A.: 2004, Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market, *Journal of Political Economy* .

Broda, C. and Weinstein, D. E.: 2006, Globalization and the Gains from Variety, *The Quarterly Journal of Economics* .

Campbell, J. Y. and Shiller, R. J.: 1988, The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors, *Review of Financial Studies* .

Diewert, E.: 1976, Exact and Superlative Index Numbers, *Journal of Econometrics* .

Eaton, J. and Kortum, S.: 2002, Technology, Geography, and Trade, *Econometrica* .

Feenstra, R. C.: 1994, New Product Varieties and the Measurement of International Prices, *American Economic Review* .

Gallin, J. and Verbrugge, R.: 2007, Improving the CPI's Age-Bias Adjustment: Leverage, Disaggregation, and Model Averaging, *BLS Working Paper* .

Guerrieri, V., Hartley, D. and Hurst, E.: 2013, Endogenous Gentrication and Housing Price Dynamics, *Journal of Public Economics* .

Jaravel, X.: 2018, The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector.

Ptacek, F.: 2013, Updating the Rent Sample for the CPI Housing Survey, *Monthly Labor Review* .

Redding, S. J. and Weinstein, D. E.: 2018, Measuring Aggregate Price Indexes with Demand Shocks: Theory and Evidence for CES Preferences, *NBER Working Paper* .

Reher, M.: 2019, Financial Intermediaries as Suppliers of Housing Quality.

Sato, K.: 1976, The Ideal Log-Change Index Number, *Review of Economics and Statistics* .

Vartia, Y. O.: 1976, Ideal Log-Change Index Numbers, *Scandinavian Journal of Statistics* .

A Econometric Details

This appendix has econometric details related to the structural rent index.

A.1 Calculation of Structural Rent Index

Calculating the rent index π_t^S in (2) requires two pieces of information: (1) a ranking variable to partition the sample into segments $\{\hat{h}\}$, and, relatedly, an identifier for which units are top-end in period t and thus have quality H_t; and (2) the preference parameter σ. I obtain these objects from the Trepp dataset and compute π_t^S over 2010-2016. As described in Reher (2019), this dataset covers a geographically representative sample of multifamily properties, and it is particularly appealing for this exercise due to its detailed information on property upgrades, inspection ratings, and renovations which are collected as part of the multifamily mortgage servicing process.

First, to rank properties I use the MBA/CREFC property inspection rating. This rating captures a property's quality relative to a newly built unit and is obtained as part of the standard multifamily mortgage servicing protocol. The resulting index is robust to alternative partitioning variables such as the property's effective age. Importantly, because $\hat{h}$ does not directly appear in (2), this rating only needs to identify a property's rank and does not need to be accurate in the cardinal sense.

To classify top-end properties I use a combination of renovation and inspection data. I classify a unit as having quality H_t in period t if it is in a property that was newly built or renovated and first on the market in year t, and if it also was ranked in the top MBA/CREFC quality segment after construction or renovation.[21] Under the assumption that top-end units in year $t-1$ retain their absolute quality H_{t-1} through at least year t, one can use the number of units with quality H_t (i.e. in segment 1 at t) and H_{t-1} (i.e. in segment 1_0 at t) and their respective revenue shares in year t to compute GQ_t according to the expression in (2).[22] To account for the possibility that renovations and construction only occur in certain areas in year

[21] There is typically a 1-year lag between completion of construction and renovation and being rent-ready.

[22] Specifically, I compute ΔGQ_s for all $s = 1, ..., t$ and then take their product to obtain GQ_t. That is, I chain growth in absolute quality. When taking (2) to the data, I will use the theoretically equivalent expression for growth in absolute quality, $GQ_t = \left(\frac{\text{Expend}_{1,t}}{\text{Expend}_{1_0,t}}\right)^{\sigma}\left(\frac{\text{Share}_{1,t}}{\text{Share}_{1_0,t}}\right)^{1-\sigma}$, where $\text{Expend}_{\hat{h},t}$ is the share of aggregate rent expenditure on segment $\hat{h}$. Doing so reduces measurement error from unit level rent. Feenstra (1994) also relies on expenditure shares when possible, since they are subject to less measurement error.

t, I also compute GQ_t within each zip code-year bin and then average across zip codes that year, which yields very similar results.[23]

Second, I must estimate σ, which I do using three methodologies: a property-level credit supply shock using idiosyncratic variation in payment timing, a zip code level version of the property-level shock, and the Feenstra (1994) GMM estimator. I now describe these three methodologies. Before doing so, Table A1 summarizes the estimated σ from each of them. The average estimate is 6.5, and Figure A1 helps interpret this magnitude by performing an introspective exercise.

Table A1: Estimated Preferences for Quality

	Estimated σ		
Specification:	Property IV	Zip Code IV	GMM
Estimate	6.5	6.9	6.1
Confidence Interval	[2.8, 18.2]	[1.2, 19.6]	[4.1, 9.4]

Note: This table shows the estimated elasticity of substitution σ for various methodologies. Property IV, Zip Code IV, and GMM are discussed in Sections A.1.1, A.1.2, and A.1.3 below. Bootstrapped 95% confidence intervals are shown in brackets. Data are from Trepp.

Table A2 summarizes the estimates after partitioning the sample into income cohorts as described in Appendix 3.1. The highest income zip codes have the lowest value of σ (4.9) and thus the highest willingness to pay for quality.[24]

[23]Growth in the rent index π_t^S is 1.5% as opposed to 1.7% under the baseline.

[24]There is an apparent non-monotonicity in σ with respect to income bracket. This likely reflects the fact that σ is both the inverse willingness to pay for quality and the market level substitutability across segments. The estimated σ in very low income zip codes may be small if the data reveal limited movement across segments. By analogy, with CRRA preferences there is not a distinction between the coefficient of relative risk aversion and the inverse intertemporal elasticity of substitution.

Table A2: Estimated Preferences for Quality by Income

	Estimated σ		
Specification:	Property IV	Zip Code IV	GMM
Full Sample	6.5	6.9	6.1
By Real Income 2010-2016:			
Bottom 30%	5.7	14.9	8.6
Middle 35%	17.9	25.0	5.0
Upper 35%	4.7	3.7	6.2

Note: This table shows the estimated elasticity of substitution σ for various methodologies and income groups. Property IV, Zip Code IV, and GMM are discussed in Sections A.1.1, A.1.2, and A.1.3 below. Data are from Trepp.

Figure A1: Example Willingness to Pay for Quality

(a) Top Tier: $1,600

(b) Above Average General Market: $1,448

(c) Below Average General Market: $971

(d) Bottom Tier: $625

Note: This figure shows the willingness to pay for quality assuming the rent of a unit in the bottom quality segment is $625 per month, the elasticity of substitution across quality segments is $\sigma = 6.5$, and the properties only differ in their structural quality. The photographs are from the website of a large real estate investor in the Dallas, TX and San Antonio, TX markets. Panels (a)-(d) show properties corresponding to the 4 bins of the MBA/CREFC rating used in this paper, respectively. Moving from the top segment to the bottom segment entails consecutive reductions of 0.6, 2.6, and 2.9 log points of relative quality, respectively. Given the average estimated σ of 6.5, this implies a willingness to pay of 10% for "highest current market standards" versus "above average" (i.e. panel (a) vs. panel (b)); 40% for "minimal" versus "general" wear and tear (i.e. panel (b) vs. panel (c)); and 44% for "no" to "some" or "multiple" life safety violations (i.e. panel (c) vs. panel (d)).

A.1.1 Estimating σ: Property-Level Credit Supply Shock

This methodology estimates σ insofar as it is the inverse marginal willingness to pay for quality. Using Lemma B.2, $\frac{1}{\sigma}$ is the elasticity of rent with respect to quality, holding the distribution of units across quality segments constant, $\frac{1}{\sigma} = \frac{\partial \log(\text{Rent}_{\hat{h},t})}{\partial \log(\hat{h})}$.

As described in FI, the structure of most multifamily mortgage contracts generates spikes in improvement activity. Combining this institutional feature with the effectively exogenous variation in their due date established in FI, I construct an instrument for the change in log relative quality $\log(\hat{h})$. Suppose now that the quality of units in property i evolves according to

$$\log\left(h_{i,t}\right) = \log\left(h_{i,t-1}\right) + \log\left(\text{Improvements}_{i,t}\right) - \delta_{i,t}, \tag{A1}$$

where $\delta_{i,t}$ is a depreciation shock. As discussed in Reher (2019), having an impending loan due reduces the probability of making a quality improvement, lowering $\text{Improvements}_{i,t}$ and thus $\Delta \log\left(h_{i,t}\right)$ in (A1). This is because most multifamily mortgages are balloon loans which require renewal at the end of every loan term, with a modal term of 10 years. Moreover, refinancing is generally not an option and must be done through a process of defeasance.[25] Because of the possibility of cheaper borrowing costs after renewal, one would expect that having an impending loan due covaries negatively with a unit's change in quality.

Mapping to a regression equation, I estimate the system

$$\log\left(\text{Rent}_{i,z,t}\right) = \frac{1}{\sigma}\Delta \log\left(\text{Quality}_{i,z,t}\right) + \beta_0 \log\left(\text{Quality}_{i,z,t-1}\right) + a_i + a_{z,t} + u_{i,z,t} \tag{A2}$$

$$\Delta \log\left(\text{Quality}_{i,z,t}\right) = \tilde{\beta}_0 \text{Impending}_{i,z,t} + \beta_1 \log\left(\text{Quality}_{i,z,t-1}\right) + \tilde{a}_i + \tilde{a}_{z,t} + \tilde{u}_{i,z,t}, \tag{A3}$$

where i, z, and t index property, zip code, and year, and $\text{Impending}_{i,z,t}$ indicates if the property's loan is due in t or $t+1$.[26] Relative quality $\hat{h}_{i,z,t}$ is denoted $\text{Quality}_{i,z,t}$ and measured using the MBA/CREFC rating. The second-stage equation is (A2), and its first stage equation is (A3). However, I do not include information about the lender because my interest is on average improvement activity, not its dependence on whether the lender was affected by HVCRE regulation.

[25] Defeasance is a fairly complicated process in which the borrower must exchange the loan for another security of equal maturity, such as a Treasury.

[26] I weight observations in (A2)-(A3) by number of units because the Trepp data are at the property-level.

Table A3: Substitutability Across Quality Rungs with Property-Level IV

	First Stage	Second Stage
Outcome:	$\Delta \log \left(\text{Quality}_{i,t}\right)$	$\log\left(\text{Rent}_{i,t}\right)$
	(1)	(2)
$\text{Impending}_{i,t}$	-0.208**	
	(0.035)	
$\Delta \log \left(\text{Quality}_{i,t}\right)$		0.153**
		(0.072)
$\log \left(\text{Quality}_{i,t-1}\right)$	-0.805**	0.132**
	(0.008)	(0.058)
Estimator	OLS	2SLS
Property FE	Yes	Yes
Zip Code-Year FE	Yes	Yes
First Stage F		34.661
Number of Observations	67210	67210

Note: Subscripts i and t denote property and year. Columns 1 and 2 estimate (A3) and (A2), respectively. $\text{Quality}_{i,t}$ is relative quality based on the MBA/CREFC property inspection rating. $\text{Impending}_{i,t}$ indicates if the investor has a mortgage due in year t or year $t+1$. The estimator in column 2 is 2SLS, and the instrument for $\Delta \log \left(\text{Quality}_{i,t}\right)$ is $\text{Impending}_{i,t}$. Observations are property-years weighted by number of units. The sample period is 2010-2016. Standard errors are in parentheses. Data are from Trepp.

The zip code-year fixed effect $a_{z,t}$ absorbs local demand effects that would otherwise affect rent. Thus, any violation of the exclusion restriction due to expectations of future growth would need to require sub-zip code variation in demand. Also, FI show that having an impending loan due is uncorrelated with interest rate spreads or other measures of credit risk. The property fixed effects a_i absorb amenities.

Column 1 of Table A3 has the results of the first stage regression (A3). Having an impending loan leads to a deterioration in relative quality. The second stage in column 2 implies $\sigma = 6.5$, based on the point estimate of 0.15 on $\Delta \log \left(\text{Rel Quality}_{i,t}\right)$.[27]

A.1.2 Estimating σ: Zip Code Level Credit Supply Shock

If CES market demand is a poor approximation, one might be concerned that the previous strategy does not identify the appropriate parameter because it relies on a highly misspecified functional form. To address concerns about functional form, I propose a second strategy which

[27] When computing the bootstrapped standard errors in Table A1 and the heterogeneous preference parameters in Table A2, I estimate (A2) through a constrained optimization such that the implied value of σ lies in the interval [1, 25]. This remark also applies when using the zip code level credit supply instrument and the Feenstra (1994) GMM estimator.

estimates σ insofar as it is the aggregate elasticity of substitution across quality segments, and which thus obtains identification through a different functional form that is nonetheless consistent with the CES market demand structure. The source of variation is similar to the previous strategy, and the instrument used is a zip code level share of property owners with an impending loan due.

Using Lemma B.2, the CES market demand curve can be written

$$\log\left(\frac{\text{Expend}_{\hat{h}}}{\text{Expend}_{\hat{h}_0}}\right) = \left(1 - \frac{1}{\sigma}\right)\log\left(\frac{\text{Share}_{\hat{h}}}{\text{Share}_{\hat{h}_0}}\right) + \frac{1}{\sigma}\log\left(\frac{\hat{h}}{\hat{h}_0}\right), \tag{A4}$$

where, using the notation introduced in Appendix B, $\text{Expend}_{\hat{h}}$ is the aggregate share of rent expenditure on segment $\hat{h}$. I then estimate the following system through 2SLS,

$$\log\left(\frac{\text{Expend}_{\hat{h},z,t}}{\text{Expend}_{\hat{h}_0,z,t}}\right) = \left(1 - \frac{1}{\sigma}\right)\log\left(\frac{\text{Share}_{\hat{h},z,t}}{\text{Share}_{\hat{h}_0,z,t}}\right) + \gamma X_{z,t} + a_{\hat{h}} + a_{m,t} + u_{\hat{h},z,t} \tag{A5}$$

$$\log\left(\frac{\text{Share}_{\hat{h},z,t}}{\text{Share}_{\hat{h}_0,z,t}}\right) = \tilde{\beta}_0 \text{Impending}_{z,t} + \gamma X_{z,t} + \tilde{a}_{\hat{h}} + \tilde{a}_{m,t} + \tilde{u}_{\hat{h},z,t} \tag{A6}$$

where z indexes zip codes, t indexes years, m indexes MSAs $\hat{h}$ indexes quality segments according to the MBA/CREFC score; $\hat{h}_0$ is the reference segment, which I set at the lowest quality segment on the MBA/CREFC rating scale; the segment fixed effects $\alpha_{\hat{h}}$ and $\tilde{\alpha}_{\hat{h}}$ absorb the second term in (A4); and $\text{Impending}_{z,t}$ is the fraction of units in zip code z whose owner has a loan due in t or t+1, and it is a zip code level average of the instrument from the property-level system (A2)-(A3).[28] The equation of interest is the second stage (A5), and its first stage is (A6).

The MSA-year fixed effect $a_{m,t}$ and its first-stage counterpart in (A6) restrict variation within MSA m in which z is located and year t. The zip code controls include measures of investors' financial condition and local demand.[29] Since all variation comes from within MSA-year bins, one can think of (A5)-(A6) as comparing the timing of when most investors in a zip code took out their loan. Thus, the instruments are predetermined as of time t and do not

[28] Specifically, $\hat{h}_0$ corresponds to a raw MBA/CREFC rating of 3, 4, or 5. I group these three segments together because they are the lowest collective segment observed in all zip codes and years. See Appendix B of FI for full details on the interpretation of MBA/CREFC ratings.

[29] Financial controls are the average interest rate spread, securitization rate, and log term for loans on units in zip code z and year t. Demand controls are the log average income, log population, and fraction of households with social security benefits, capital gains, dividend income, and children in z and t, all based on IRS tax returns. See the footnote to Table A4 for more details on how these variables are proxied using the IRS data.

Table A4: Local Demand, Financial Condition, and the Timing of Loan Renewal

Outcome:	$\text{Impending}_{z,t}$
$\log\left(\text{Income}_{z,t}\right)$	0.002
	(0.005)
$\log\left(\text{Population}_{z,t}\right)$	-0.002
	(0.002)
$\text{Stock Ownership}_{z,t}$	-0.228
	(0.166)
$\text{Family Households}_{z,t}$	0.110
	(0.088)
$\text{Social Security Benefits}_{z,t}$	-0.060
	(0.060)
$\text{Capital Gains Income}_{z,t}$	0.209
	(0.190)
$\text{Rate Spread}_{z,t}$	0.117
	(0.145)
$\text{Securitized}_{z,t}$	0.000
	(0.004)
$\log\left(\text{Term}_{z,t}\right)$	-0.037**
	(0.003)
MSA-Year FE	Yes
R-squared	0.210
Number of Observations	9282

Note: Subscripts z and t denote zip code and year. $\text{Impending}_{z,t}$ is the fraction of units whose investor has a mortgage due in year t or t+1. $\text{Income}_{z,t}$ is average income per tax return. $\text{Population}_{z,t}$ is number of tax returns. $\text{Stock Ownership}_{z,t}$ is the fraction of households with dividend income. $\text{Family Households}_{z,t}$ is the fraction of returns with a child tax credit. $\text{Social Security Benefits}_{z,t}$ is the fraction of returns with social security income. $\text{Capital Gains Income}_{z,t}$ is the fraction of returns with capital gains. $\text{Rate Spread}_{z,t}$ is the average difference between the loan's current interest rate and the average loan interest rate the year of origination or renewal. $\text{Securitized}_{z,t}$ is the fraction of units whose loan was securitized within 3 months of origination. $\text{Term}_{z,t}$ is the average loan term in months. The sample period is 2010-2016. Standard errors are in parentheses.

capture contemporaneous demand shocks. However, it is plausible that the timing of borrowing decisions and the resulting interest rate reflected expectations about future demand in a given zip code with an MSA. These expectations would be reflected ex post in measures of local demand, or ex ante in the loan's rate spread or initial securitization status. To investigate this possibility, I project the instrument $\text{Impending}_{z,t}$ onto the control vector $X_{z,t}$. The results in Table A4 show that the only significant partial correlation with $\text{Impending}_{z,t}$ is the mechanical effect of having a shorter term. This finding suggests that, within the same MSA-year bin, $\text{Impending}_{z,t}$ does not reflect expectations about local demand.

Table A5 has the results of (A5)-(A6). Column 1 has the estimates from the first stage (A6).

Table A5: Substitutability Across Quality Rungs with Zip Code Level IV

	First Stage	Second Stage		
Outcome:	$\log\left(\frac{\text{Share}_{\hat{h},z,t}}{\text{Share}_{\hat{h}_0,z,t}}\right)$	$\log\left(\frac{\text{Expend}_{\hat{h},z,t}}{\text{Expend}_{\hat{h}_0,z,t}}\right)$		
	(1)	(2)	(3)	(4)
$\text{Impending}_{z,t}$	-0.538**			
	(0.173)			
$\log\left(\frac{\text{Share}_{\hat{h},z,t}}{\text{Share}_{\hat{h}_0,z,t}}\right)$		0.884**	0.857**	0.856**
		(0.143)	(0.138)	(0.146)
MSA-Year FE	Yes	Yes	Yes	Yes
Segment FE	Yes	Yes	Yes	Yes
Credit Controls	Yes	No	Yes	Yes
Demand Controls	Yes	No	No	Yes
First Stage F		9.596	10.677	9.683
Number of Observations	11574	11574	11574	11574

Note: Subscripts $\hat{h}$, z, and t denote quality segment, zip code, and year. Column 1 estimates (A6) and columns 2-4 estimate (A5). Quality segments are based on relative quality from the MBA/CREFC property inspection rating. Segment $\hat{h}_0$ is the lowest available in all zip codes and years. $\text{Impending}_{z,t}$ is the fraction of units whose property has a loan due in year t or $t+1$. $\text{Expend}_{\hat{h},z,t}$ and $\text{Share}_{\hat{h},z,t}$ are the aggregate share of rent expenditure and number of units in segment $\hat{h}$ within a given zip code-year. The estimator in columns 2-4 is 2SLS and the instrument for $\log\left(\frac{\text{Share}_{\hat{h},z,t}}{\text{Share}_{\hat{h}_0,z,t}}\right)$ is $\text{Impending}_{z,t}$. Credit and demand controls are those from Table A4. The sample period is 2010-2016. Standard errors are in parentheses. Data are from Trepp.

Consistent with the property-level specification, zip codes where more property owners have an impending loan due see fewer units in segments $\hat{h} > \hat{h}_0$, recalling that the reference segment $\hat{h}_0$ is the lowest on the MBA/CREFC rating scale. In these zip codes, there is a compositional shift toward lower quality units. The second stage results in columns 2-4 show how this compositional shift affected relative expenditure shares under the CES market demand curve (A4). The point estimate of 0.86 in column 4 implies σ of around 6.9, similar to the result of the property-level specification.[30]

A.1.3 Estimating σ: Feenstra GMM

Feenstra (1994) proposes an estimator for σ in the context of time-varying quality. It exploits the panel structure of the data to provide an identification condition. This approach is potentially problematic in my dataset because it requires a large number of time periods to produce consistent estimates. That said, I estimate σ using this method as well.

[30] Explicitly, $\sigma = \frac{1}{1-0.86}$.

To summarize the methodology briefly, one begins with the market demand curve implied by Lemma B.2 and previously expressed in (A4),

$$\log\left(\frac{\text{Expend}_{\hat{h},t}}{\text{Expend}_{\hat{h}_0,t}}\right) = \frac{1}{\sigma}\log\left(\frac{\hat{h}}{\hat{h}_0}\right) + \left(1-\frac{1}{\sigma}\right)\log\left(\frac{\text{Share}_{\hat{h},t}}{\text{Share}_{\hat{h}_0,t}}\right) + \nu_{\hat{h},t} \tag{A7}$$

where, using the notation introduced in Appendix B, $\text{Expend}_{\hat{h}}$ is the aggregate rent expenditure share on segment $\hat{h}$ and $\nu_{\hat{h},t}$ is a demand shifter.[31] Like with the zip code credit supply methodology, $\hat{h}_0$ is the reference segment, which I set at the lowest quality segment on the MBA/CREFC rating scale.

Then, one specifies the following isoelastic supply curve for a representative property owner deciding how many units in quality segment $\hat{h}$ to provide. This representative property owner aggregates the improvement decisions of individual property owners, giving rise to a supply curve which I express in terms of revenue, $\text{Expend}_{\hat{h},t}$,

$$\log\left(\frac{\text{Expend}_{\hat{h},t}}{\text{Expend}_{\hat{h}_0,t}}\right) = \alpha_0 + \alpha\log\left(\frac{\text{Share}_{\hat{h},t}}{\text{Share}_{\hat{h}_0,t}}\right) + \alpha_{\hat{h},t}. \tag{A8}$$

One takes differences of the demand and supply curves (A7) and (A8) to obtained the differenced shocks $\Delta\nu_{\hat{h},t}, \Delta\alpha_{\hat{h},t}$. These shocks give the moment condition

$$\mathbb{E}\left[\Delta\nu_{\hat{h},t}\Delta\alpha_{\hat{h},t}\right] = 0. \tag{A9}$$

Note that (A9) must apply to each zip code z. Therefore, rearranging (A9) gives the regression equation

$$\begin{aligned}\left[\Delta\log\left(\frac{\text{Expend}_{\hat{h},z,t}}{\text{Expend}_{\hat{h}_0,z,t}}\right)\right]^2 &= \theta_1\left[\Delta\log\left(\frac{\text{Expend}_{\hat{h},z,t}}{\text{Expend}_{\hat{h}_0,z,t}}\right)\Delta\log\left(\frac{\text{Share}_{\hat{h},z,t}}{\text{Share}_{\hat{h}_0,z,t}}\right)\right] + ... \\ &... + \theta_2\left[\Delta\log\left(\frac{\text{Share}_{\hat{h},z,t}}{\text{Share}_{\hat{h}_0,z,t}}\right)\right]^2 + u_{\hat{h},z,t},\end{aligned} \tag{A10}$$

[31] The estimator I derive is slightly different than the original proposed by Feenstra (1994) because I reason on quantity (i.e. number of units) rather than price (i.e. rent). Reasoning on quantity is more appropriate in my setting because the share of units in each segment must sum to 1, and doing so reduces measurement error from the fact that I approximate rent as revenue per occupied unit. However, the setup is effectively the same after replacing "goods" with "quality segments".

where the notation is the same as in previous specifications. Intuitively, equation (A10) expresses the relationship among the second moments of expenditure and unit shares, $\text{Expend}_{\hat{h},z,t}$ and $\text{Share}_{\hat{h},z,t}$. The coefficients θ_1 and θ_2 encode the elasticity of substitution σ and the supply elasticity α. In particular,

$$\sigma = \frac{1}{1+\theta_2\alpha}, \quad \alpha = \frac{\theta_1}{2} + \frac{1}{2}\sqrt{\theta_1^2 + 4\theta_2}. \tag{A11}$$

Since $u_{\hat{h},z,t}$ is a function of the differenced demand and supply shocks $\Delta\nu_{\hat{h},z,t}, \Delta\alpha_{\hat{h},z,t}$, one cannot estimate (A10) consistently. However, one can obtain consistent estimates by taking the average of (A10) across time periods, and estimating the resulting regression equation by weighted least squares. That is, $\text{plim}_{T\to\infty}\sum_{t=0}^{T} u_{\hat{h},z,t} = 0$. The resulting estimates of θ_1 and θ_2 imply $\sigma = 6.1$, shown in Table A1.

B Mathematical Details

This appendix has mathematical details related to the structural rent index.

B.1 Household Preferences: Structural Rent Index

This extension describes households' problem in greater detail. Reviewing the setup from Section 2, household j selects a unit i and derives additive random utility from the unit's quality h_i according to the preferences in (1),

$$u_{i,j} = \log(h_i) + \epsilon_{i,j}.$$

In the baseline case, $\epsilon_{i,j}$ follows a Gumbel, or type 1 extreme value distribution.[32] Moreover, quality is itself a composite of a unit's space s_i (e.g. square feet) and other amenities a_i (e.g. granite countertops) according to $\log(h_i) = \mu \log(s_i) + \log(a_i)$. For simplicity I assume that all units i in the same quality segment $\hat{h}$ have the same space $s_{\hat{h}}$. The next lemma describes how this preference structure gives rise to the discrete choice problem verbally articulated in Appendix 2.

Lemma B.1 (Discrete Choice) *A household with preferences (1) chooses her shelter according to*

$$\max_{i \in \mathcal{I}} \left\{ -\log(Rent_i) + \frac{1}{\mu}\log(h_i) + \frac{1}{\mu}\epsilon_{i,j} \right\}. \tag{B1}$$

Note that while households do not consume the numeraire, it is straightforward to allow for non-housing consumption. To do so, I follow Anderson, de Palma and Thisse (1992) and modify the baseline preferences (1) as follows

$$u_{i,j} = \kappa \log(c_j) + \log(h_i) + \epsilon_{i,j}, \tag{B2}$$

where c_j is household j's consumption of the numeraire. Maximizing (B2) under the budget

[32] The cumulative distribution function is $\Pr[\epsilon_{i,j} \leq \epsilon] = \exp\left[-\exp\left[-\left(\frac{\epsilon}{\tilde{\mu}} + \gamma^e\right)\right]\right]$, where $\gamma^e = 0.58$ is Euler's constant and $\tilde{\mu}$ is a scaling parameter. In particular, $\mathbb{E}[\epsilon_{i,j}] = 0$ and $\mathrm{Var}[\epsilon_{i,j}] = \tilde{\mu}^2 \frac{\pi^2}{6}$.

constraint $y_j = y_j^h + c_j$, where y_j^h is income available for housing expenditure,

$$c_j = \frac{\kappa}{\kappa + 1} y_j.$$

Thus, the analysis is effectively the same after replacing income y_j with income net of non-housing consumption, $y_j^h = \frac{1}{\kappa+1} y_j$.

The next lemma provides a useful aggregation result which enables Proposition 2.1.

Lemma B.2 (Anderson, De Palma, and Thisse 1992) *If $\epsilon_{i,j}$ follows a type 1 extreme value distribution with scaling parameter 1, then the distribution of number units across segments $\{Share_{\hat{h}}\}$ behaves according to*

$$\min_{\{Share_{\hat{h}}\}} \sum_{\hat{h} \in \mathcal{H}} Share_{\hat{h}} \times Rent_{\hat{h}} \quad s.t. \quad \bar{U} = \left[H^{\frac{1}{\sigma}} \sum_{\hat{h} \in \mathcal{H}} \hat{h}^{\frac{1}{\sigma}} Share_{\hat{h}}^{1-\frac{1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \tag{B3}$$

where $\sigma = \mu + 1$ and $\mathcal{H} \subseteq [0, 1]$ is the set of quality segments. Moreover, $\frac{\partial \log\left(Rent_{\hat{h}}\right)}{\partial \log\left(\hat{h}\right)} = \frac{1}{\sigma}$.

If households treat space as a necessity, one can suppose that μ is a decreasing function of income y_j. Thus, as y_j rises, households' relative preference for space versus amenities falls, so that high income households have low values of $\sigma = \mu + 1$, matching the estimates from Table A2.

B.2 Proofs

Proof of Lemma B.1

The preferences in (1) are preserved when multiplying by $\tilde{\mu} \equiv \frac{1}{\mu}$. Therefore, since rent Rent_i is denominated in numeraire per housing unit, a household with y_j available to spend on housing has the following utility, based on (1),

$$u_{i,j} = \log\left(\frac{y_j}{\text{Rent}_i / s_i}\right) + \tilde{\mu} \log(a_i) + \tilde{\epsilon}_{i,j} = \log(y_j) - \log(\text{Rent}_i) + \tilde{\mu} \log(h_i) + \tilde{\epsilon}_{i,j}, \tag{B4}$$

where the second equality uses $\log(h_i) = \mu \log(s_i) + \log(a_i)$, and $\tilde{\epsilon}_{i,j} \equiv \tilde{\mu}\epsilon_{i,j}$. From (B4) it follows that household j selects unit $i \in \mathcal{I}$ as the solution to

$$\max_{i \in \mathcal{I}} \left\{ -\log(\text{Rent}_i) + \frac{1}{\mu}\log(h_i) + \frac{1}{\mu}\epsilon_{i,j} \right\},$$

as in (B1).

Proof of Lemma B.2

First note that the problem in (B10) is equivalent to its corresponding primal problem,

$$\max_{\left\{\text{Share}_{\hat{h}}\right\}} \left[\sum_{\hat{h} \in \mathcal{H}} h^{\frac{1}{\sigma}} \text{Share}_{\hat{h}}^{1-\frac{1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad s.t. \quad \bar{y} = \sum_{\hat{h} \in \mathcal{H}} \text{Share}_{\hat{h}} \times \text{Rent}_{\hat{h}}, \tag{B5}$$

with

$$\bar{U} = \bar{y} \left(\left[\sum_{\hat{h} \in \mathcal{H}} h \times \text{Rent}_{\hat{h}}^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \right)^{-1} \equiv \bar{y}\bar{R}^{-1}, \tag{B6}$$

where $\bar{R}$ is the minimized cost function associated with the problem in (B10), which was referred to as "welfare relevant rent" in the text. In particular, using $h = \hat{h}H$, the solution implies that the share of aggregate expenditure on segment $\hat{h}$ is

$$\text{Expend}_{\hat{h}} = \frac{\hat{h}\text{Rent}_{\hat{h}}^{1-\sigma}}{\sum_{\hat{h} \in \mathcal{H}} \hat{h}\text{Rent}_{\hat{h}}^{1-\sigma}}. \tag{B7}$$

Continuing, it suffices to show that the aggregate demand generated by individual households solving (B1) behaves according to (B5) for some σ. I work with the normalized preferences in (B4) from the proof of Lemma B.1, where $\tilde{\epsilon}_{i,j} \equiv \tilde{\mu}\epsilon_{i,j}$ follows a Gumbel distribution with scaling parameter $\tilde{\mu} \equiv \frac{1}{\mu}$. Then, use equation (3.51) from Anderson, de Palma and Thisse (1992) for the case when there is the additional quality term $\tilde{\mu}\log(h)$, to write the probability a household chooses a unit in segment $h = \hat{h}H$ as

$$\varrho_{\hat{h}} = \frac{\hat{h} \times \text{Rent}_{\hat{h}}^{-\mu}}{\sum_{\hat{h} \in \mathcal{H}} \hat{h} \times \text{Rent}_{\hat{h}}^{-\mu}}. \tag{B8}$$

It follows that aggregate demand across households for units in segment $\hat{h}$ is

$$\text{Share}_{\hat{h}} = \frac{\bar{y}}{\text{Rent}_{\hat{h}}/s_{\hat{h}}} \cdot \varrho_{\hat{h}} \cdot \frac{1}{s_{\hat{h}}}, \tag{B9}$$

where $s_{\hat{h}}$ is the space afforded by units in segment $\hat{h}$. In particular, the three terms in (B9) are, respectively: (i) the space demanded by the average household, (ii) the share of households selecting a housing unit in segment $\hat{h}$, and (iii) the inverse space per housing unit.

Finally, using Proposition 3.8 from Anderson, de Palma and Thisse (1992), the aggregate demand system (B9) equals that of the representative household (B7) if and only if $\sigma = \mu + 1$. That is, the distribution of units across quality segments $\{\text{Share}_{\hat{h}}\}$ behaves according to the solution to the problem (B10). In addition, as pointed out by Anderson, de Palma and Thisse (1992), the value function associated with (B5) is a utilitarian welfare function.

To obtain the marginal willingness to pay, use the demand curve (B7) to write

$$\frac{\text{Rent}_{\hat{h}}}{\text{Rent}_{\hat{h}_0}} = \left(\frac{h}{h_0}\right)^{\frac{1}{\sigma}} \left(\frac{\text{Share}_{\hat{h}}}{\text{Share}_{\hat{h}_0}}\right)^{-\frac{1}{\sigma}}, \tag{B10}$$

for some reference segment $\hat{h}_0$, which gives $\frac{\partial \log(\text{Rent})}{\partial \log(h)} = \frac{\partial \log(\text{Rent})}{\partial \log(\hat{h})} = \frac{1}{\sigma}$ using $h = \hat{h}H$. This completes what needed to be shown.

Proof of Proposition 2.1

By definition, π_t^S is the growth in the unit cost function $\bar{R}_t$ from t_0 to t, explicitly $\pi_t^S \equiv \frac{\bar{R}_t}{\bar{R}_{t_0}}$. Using (B6) from the proof of Lemma B.2,

$$\bar{R} = \left[\sum_{\hat{h} \in \mathcal{H}} \hat{h}H_t \times \text{Rent}_{\hat{h},t}{}^{1-\sigma}\right]^{\frac{1}{1-\sigma}}, \tag{B11}$$

where $h_t = \hat{h} H_t$ is the absolute quality of segment $\hat{h}$ at t. Therefore, write π_t^S as

$$\pi_t^S = \left[\sum_{\hat{h} \in \mathcal{H}} \hat{h} \times \mathrm{Rent}_{\hat{h}}^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \times \left(\frac{H_t}{H_{t_0}} \right)^{-\frac{1}{\sigma-1}} \equiv DQ_t \times GQ_t^{-\frac{1}{\sigma-1}}. \tag{B12}$$

Next, following Feenstra (1994), use the results of Diewert (1976), Sato (1976), and Vartia (1976) to rewrite DQ_t as

$$DQ_t = \exp \left[\sum_{\hat{h} \in \mathcal{H}} w_{\hat{h},t} \log \left(\frac{\mathrm{Rent}_{\hat{h},t}}{\mathrm{Rent}_{\hat{h},t_0}} \right) \right], \tag{B13}$$

where the Sato-Vartia weights are

$$w_{\hat{h},t} = \frac{\frac{\mathrm{Expend}_{\hat{h},t} - \mathrm{Expend}_{\hat{h},t_0}}{\log\left(\mathrm{Expend}_{\hat{h},t}\right) - \log\left(\mathrm{Expend}_{\hat{h},0}\right)}}{\sum_{\hat{h} \in \mathcal{H}} \frac{\mathrm{Expend}_{\hat{h},t} - \mathrm{Expend}_{\hat{h},t_0}}{\log\left(\mathrm{Expend}_{\hat{h},t}\right) - \log\left(\mathrm{Expend}_{\hat{h},t_0}\right)}}, \tag{B14}$$

and, as in (B7) from the proof of Lemma B.2, $\mathrm{Expend}_{\hat{h},t}$ is the share of aggregate expenditure on segment $\hat{h}$ in t,

$$\mathrm{Expend}_{\hat{h}} = \frac{\mathrm{Share}_{\hat{h}} \times \mathrm{Rent}_{\hat{h}}}{\sum_{\hat{h} \in \mathcal{H}} \mathrm{Share}_{\hat{h}} \times \mathrm{Rent}_{\hat{h}}}. \tag{B15}$$

Finally, Lemma B.2 implies that the market demand curve has a CES structure, and using (B10),

$$\frac{\mathrm{Rent}_{1,t}}{\mathrm{Rent}_{1_0,t}} = \left(\frac{H_t}{H_{t_0}} \right)^{\frac{1}{\sigma}} \left(\frac{\mathrm{Share}_{1,t}}{\mathrm{Share}_{1_0,t}} \right)^{-\frac{1}{\sigma}}, \tag{B16}$$

where H_{t_0} is absolute quality in t_0 and segment $1_0 \equiv \frac{H_{t_0}}{H_t}$ contains units that were in segment 1 in year t_0.[33] Rearranging (B16) gives growth in absolute quality,

$$GQ_t = \left(\frac{\mathrm{Rent}_{1,t}}{\mathrm{Rent}_{1_0,t}} \right)^{\sigma} \frac{\mathrm{Share}_{1,t}}{\mathrm{Share}_{1_0,t}}. \tag{B17}$$

[33]While the setup does not feature depreciation, when taking π_t^S to the data I require that units in segment 1_0 retained their absolute quality through t, as discussed in Appendix A.1.

Combining (B13) and (B17) gives the expression in (2),

$$\pi_t^S = \exp\left[\sum_{\hat{h}\in\mathcal{H}} w_{\hat{h},t} \log\left(\frac{\text{Rent}_{\hat{h},t}}{\text{Rent}_{\hat{h},t_0}}\right)\right] \times \left[\left(\frac{\text{Rent}_{1,t}}{\text{Rent}_{1_0,t}}\right)^{\sigma} \frac{\text{Share}_{1,t}}{\text{Share}_{1_0,t}}\right]^{-\frac{1}{\sigma-1}} \quad \equiv DQ_t \times GQ_t^{-\frac{1}{\sigma-1}}.$$

C Non-CES Rent Indices

This section uses detailed data on property upgrade activity to infer time variation in quality with minimal structural assumptions. While the CES aggregator is one of the most commonly used in economics, one might be concerned that the expenditure function from Lemma B.2 is highly misspecified. To address this concern, this section performs a quasi-hedonic quality adjustment to correct each segment's rent for time-varying quality. The corrected rent can then be used in any non-CES price index formula (e.g. Tornqvist, Paasche, Laspeyres).

Write the rent on unit i in year t as

$$\log\left(\text{Rent}_{i,t}\right) = a_i + a_t + P\log\left(h_{i,t}\right) + u_{i,t} \tag{C1}$$

$$= a_i + a_t + P\left[\log\left(\hat{h}_{i,t}\right) + H_t\right] + u_{i,t}, \tag{C2}$$

where P is the equilibrium slope of the quality ladder, or price of quality. Let $\text{New}_{i,t}$ indicate if i is first on the market in t after renovation and is in the top quality segment. I maintain the assumption from the structural index in Appendix C of the paper that such units retain their absolute quality for at least one year. Then (C1) implies that rent growth conditional on being new ($\text{New}_{i,t} = 1$) or almost-new ($\text{New}_{i,t-1} = 1$) is

$$\Delta\log\left(\text{Rent}_{i,t}\right) = \Delta a_t + \underbrace{\beta_0}_{P\times\log(GQ_t)} \text{New}_{i,t} + \underbrace{\beta_1}_{-P}\left(\text{New}_{i,t}\times\log\left(\hat{h}_{i,t-1}\right)\right) + \Delta u_{i,t}, \tag{C3}$$

where $GQ_t = \frac{H_t}{H_{t-1}}$ is growth in absolute quality, using the notation from Proposition 1. In words, the rent growth differential between new and almost-new units reflects new unit quality growth, after controlling for the previous quality of new units. To measure relative quality, I again use the MBA/CREFC property inspection rating.[34]

Using the Trepp dataset, I estimate (C3) year-by-year on units such that

$$\max\left\{\text{New}_{i,t}, \text{New}_{i,t-1}\right\} = 1$$

and extract the coefficients $\{\beta_{0,t}\}$.[35] These point estimates give a sequence of quality growth rates

[34] Since inspections are rare during the year of renovation, I proxy for $\hat{h}_{i,t-1}$ using the most recent rating prior to renovation. Similarly, I proxy for $\Delta\log\left(\text{Rent}_{i,t}\right)$ using the most recent available rent data and annualizing.

[35] I weight observations in (C3) by number of units because the Trepp data are at the property-level.

in units of log rent, $\{P \times (GQ_t)\}$. The identifying assumption is that the renovation decision $\text{New}_{i,t}$ is orthogonal to unobserved changes in the rental market as contained in $\Delta u_{i,t}$. This assumption would be violated if, for example, renovations only occur in high-growth areas. To account for this difficulty, I estimate (C3) with MSA fixed effects and controls for property size. Thus, the comparison is strictly between units in new and almost-new properties of the same size and in the same MSA and year. While reducing bias, this approach substantially limits the available variation to estimate (C3) given the inclusion of so many covariates. Therefore, I use the James-Stein estimator, which optimally biases the point estimate toward 0.[36]

Given the estimated sequence of annualized growth rates $\{P \times \log(GQ_t)\}$, I correct each unit's rent according to

$$\text{Rent}_{i,t}^{GQ} = \text{Rent}_{i,t} \times \exp\left[-\sum_{\tau=t_0}^{t} P \times \log(GQ_\tau)\right]. \tag{C4}$$

Using (C4), one obtains the corrected rent in each quality segment $\{\text{Rent}_{\hat{h},t}^{GQ}\}$. Along with the appropriate data on aggregate expenditure shares $\{\text{Expend}_{\hat{h},t}\}$, one can compute effective rent using any price index formula. Some common formulae used in this paper are

$$\pi_t^{\text{Tornqvist}} = \exp\left[\sum_{\hat{h}\in\mathcal{H}} \frac{\text{Expend}_{\hat{h},t} + \text{Expend}_{\hat{h},t_0}}{2} \log\left(\frac{\text{Rent}_{\hat{h},t}^{GQ}}{\text{Rent}_{\hat{h},t_0}^{GQ}}\right)\right],$$

$$\pi_t^{\text{Paasche}} = \exp\left[\sum_{\hat{h}\in\mathcal{H}} \text{Expend}_{\hat{h},t} \log\left(\frac{\text{Rent}_{\hat{h},t}^{GQ}}{\text{Rent}_{\hat{h},t_0}^{GQ}}\right)\right],$$

$$\pi_t^{\text{Laspeyres}} = \exp\left[\sum_{\hat{h}\in\mathcal{H}} \text{Expend}_{\hat{h},t_0} \log\left(\frac{\text{Rent}_{\hat{h},t}^{GQ}}{\text{Rent}_{\hat{h},t_0}^{GQ}}\right)\right].$$

Excess-CPI growth in $\pi_t^{\text{Tornqvist}}, \pi_t^{\text{Paasche}}$, and $\pi_t^{\text{Laspeyres}}$ was between 0.1% and 0.2% over 2010-2016.

[36]The James-Stein estimator is $\hat{\beta}^{JS} = \max\left\{1 - \frac{c}{\text{F-statistic}}, 0\right\} \cdot \hat{\beta}^{OLS}$, where $\hat{\beta}^{OLS}$ is the OLS estimator. For $0 < c < \bar{c} < 2$ and at least three predictor variables, $\hat{\beta}^{JS}$ dominates $\hat{\beta}^{OLS}$ under the L^2 norm.