Valuing Housing Services in the Era of Big Data: A User Cost Approach Leveraging Zillow Microdata

Marina Gindelsky[a+]
U.S. Bureau of Economic Analysis

Jeremy G. Moulton[b]
University of North Carolina – Chapel Hill

Scott A. Wentland[c+]
U.S. Bureau of Economic Analysis

March 1, 2019

Abstract

Historically, residential housing services or "space rent" for owner-occupied housing has made up a substantial portion (approximately 10%) of U.S. GDP final expenditures. The current methods and imputations for this estimate employed by the Bureau of Economic Analysis (BEA) rely primarily on designed survey data from the Census Bureau. In this study, we develop new, proof-of-concept estimates valuing housing services based on a user cost approach, utilizing detailed microdata from Zillow (ZTRAX), a "big data" set that contains detailed information on hundreds of millions of market transactions. Methodologically, this kind of data allows us to incorporate actual market prices into the estimates more directly for property-level hedonic imputations, providing an example for statistical agencies to consider as they improve the national accounts by incorporating additional big data sources. Further, we are able to include other property-level information into the estimates, reducing potential measurement error associated with aggregation of markets that vary extensively by region and locality. Finally, we compare our estimates to the corresponding series of BEA statistics, which are based on a rental-equivalence method. Because the user-cost approach depends more directly on the market prices of homes, we find that since 2001 our initial results track aggregate home price indices more closely than the current estimates.

Keywords: residential housing, Big Data, housing services, owner-occupied, space rent, home prices
JEL Classifications: E01, C80, R00

Acknowledgements: We would like to thank the organizing committee and participants of the 2018 NBER-CRIW Pre-Conference, Erwin Diewert, Dennis Fixler, Kyle Hood, Kurt Kunze, Raven Molloy, Brent Moulton, Mick Silver, Dylan Rassier, Brian Smith, and Randal Verbrugge for their helpful comments. All errors are our own.

+Disclaimer: Any views expressed here are those of the authors and not necessarily those of the Bureau of Economic Analysis or the U.S. Department of Commerce. Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the author(s) and do not reflect the position of Zillow Group.

[a] Office of the Chief Economist 4600 Silver Hill Rd, Suitland, MD 20746; marina.gindelsky@bea.gov
[b] Department of Public Policy, Abernethy Hall, CB 3435, Chapel Hill, NC 27599; moulton@email.unc.edu
[c] *Contact Author*. Office of the Chief Economist 4600 Silver Hill Rd, Suitland, MD 20746; scott.wentland@bea.gov

1. Introduction

Housing is an important part of the economy and the national economic accounts. As part of its tabulation of Personal Consumption Expenditures (PCE) within Gross Domestic Product (GDP), the Bureau of Economic Analysis (BEA) estimates aggregate expenditure on housing, measuring what households in the United States spend on housing services. Because a house is generally a long-lasting asset and the flow of its services is not consumed in its entirety in a single year, housing is not measured like many other consumption expenditures as simply the aggregate of home prices and quantities. The flow of housing services in GDP are, as a result, measured as conceptually most similar to rent for these services in a given period. For renters (tenant-occupied housing), this tabulation is straightforward, both intuitively and from an economic measurement standpoint, as it amounts to the aggregate sum of rents paid for all residential units over a given period. But, for conceptual consistency due to the fact that homeowners do not pay rent explicitly, the analogous calculation imputes market rents (also called "space rent") for the owner-occupied housing stock as if owners "rent" to themselves. The 2008 System of National Accounts (SNA) recommends this imputation for owner-occupied housing so the estimate of housing services is not arbitrarily distorted based on the decision to rent versus own a home.[1] Historically, these aggregate housing estimates for both tenant and owner-occupied housing have accounted for a substantial proportion of overall consumer expenditures and the economy more generally (approximately 16% of PCE, or about 10% of GDP final expenditures), and have been relatively stable over recent decades as shown in Figure 1 below.

[1] Specifically, the 2008 SNA states: "The production of housing services for their own final consumption by owner occupiers has always been included within the production boundary in national accounts, although it constitutes an exception to the general exclusion of own-account service production. The ratio of owner-occupied to rented dwellings can vary significantly between countries, between regions of a country and even over short periods of time within a single country or region, so that both international and inter-temporal comparisons of the production and consumption of housing services could be distorted if no imputation were made for the value of own-account housing services." (SNA 2008, 6.35, p. 99).

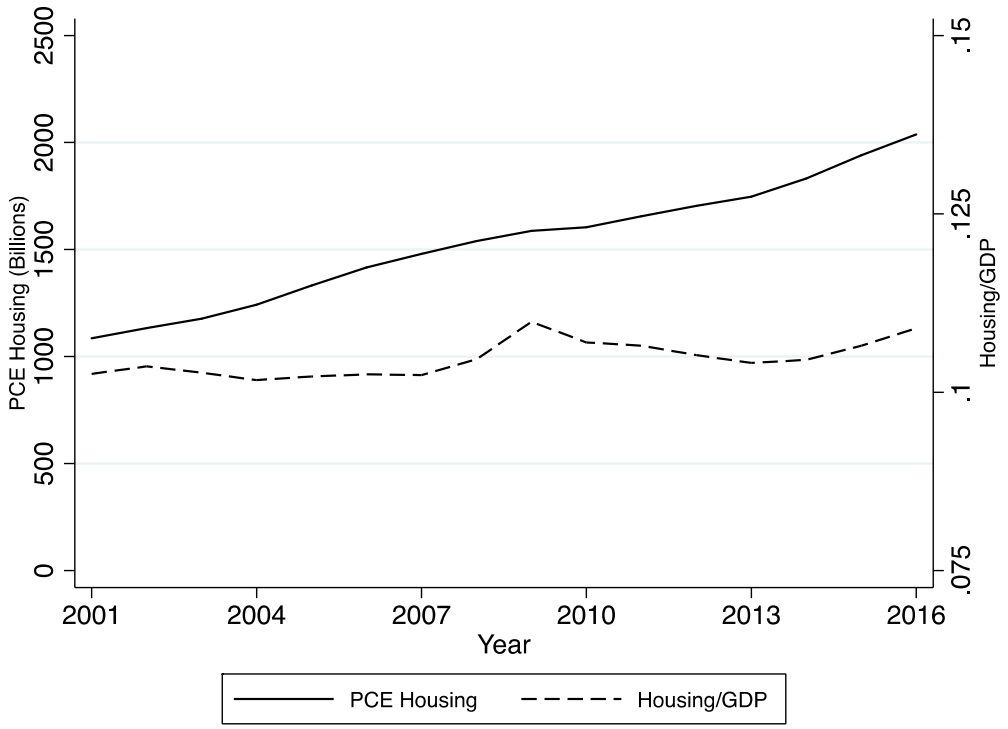


Figure 1: Nominal PCE Housing and PCE Housing/GDP
Source: U.S. Bureau of Economic Analysis, "Table 2.5.5: Personal Consumption Expenditures (PCE) by Function," bea.gov.

While existing price indices of the national housing market, such as either the FHFA or Case-Shiller Price Index, do not exactly measure the same construct, they show considerably more variation over time compared to housing services in PCE. To illustrate this contrast over the same time period, both the FHFA *All Transaction Price Index* and the Case-Shiller *U.S. National Home Price Index* are depicted below in Figure 2. A critical part of this difference is how housing services are measured and the corresponding underlying data. While indices like Case-Shiller are based on home *prices*, the BEA's current imputations of owner-occupied housing services primarily rely on designed survey data from the Census Bureau and a rental-equivalence method that bases its imputations on market *rents* of tenant occupied-homes. Hence, the purpose of this paper is to explore a method that relies more directly on market prices of the homes themselves, a user-cost approach, which utilizes "big data" from Zillow to provide a proof-of-concept alternative to the

current rental-equivalence method used by BEA. However, we should state at the outset that this is not a paper about constructing an official account or arguing explicitly for a particular method; rather, we simply take the necessary first step of exploring its feasibility with a new "big data" source and provide initial estimates. Further, this also allows us to evaluate the extent to which the user cost method reflects broader price trends as compared to other data series.

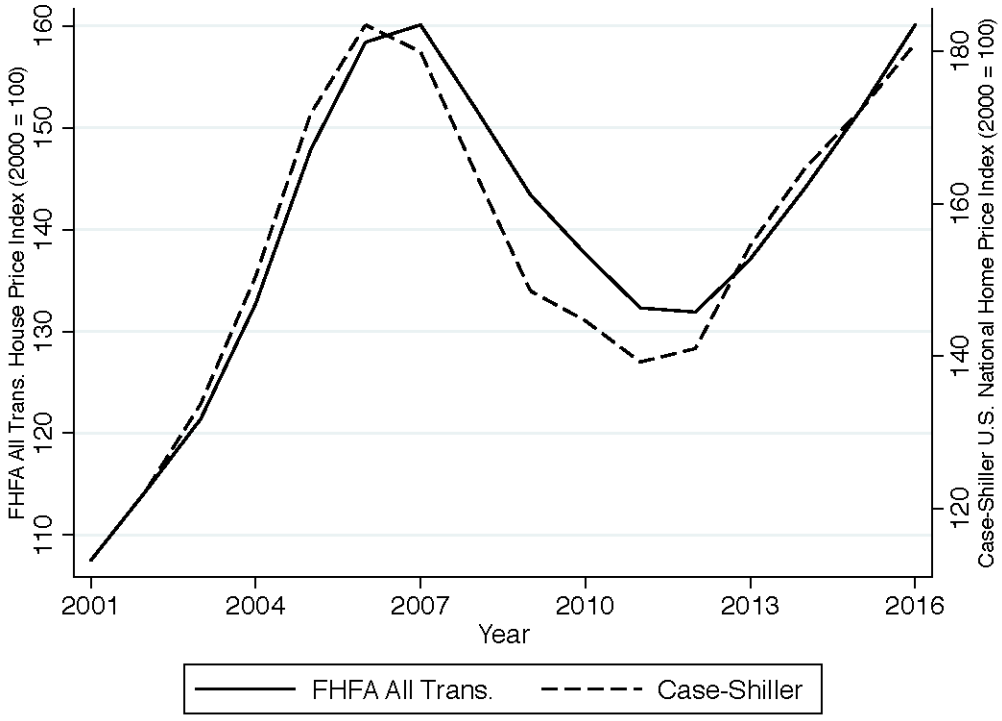


Figure 2: FHFA and Case-Shiller U.S. National Home Price Index
Source: https://fred.stlouisfed.org/series/CSUSHPINSA; /USSTHPI

The BEA's current approach based on rental data is the most common method used by national statistical agencies around the world (Katz 2017), in part due to the fact that countries collect high quality data on rents from nationally representative, specifically designed surveys of tenants and other sources. In contrast, data on home sales and corresponding home characteristics are primarily recorded by local municipalities, and thus reporting often differs by locale, making a national effort to collect this data quite costly. Indeed, only in recent decades have most localities

digitized these records, making rental survey data the most practical data source prior to the era of "big data." But, in the modern era companies like Zillow have privately undertaken a laudable effort to collect, compile, and organize a massive database of public data from local tax assessors' offices across the U.S. for the purposes of providing this information to users of their website. Zillow has recently provided much of their microdata to researchers free of charge, including those at BEA, which makes a user cost method based on fine-level price and home characteristic data more tractable, at least as a proof-of-concept effort to show how estimates built from national microdata stack up against current methods. A new investigation using "big data" is important given how prior studies (for example, Verbrugge (2008), Garner and Verbrugge (2009), Aten (2018), and others) have found persistent and sizable differences between rental-equivalence and user cost methods using data from the Census Bureau and other aggregate data sources.

2. Background – Rental-equivalence vs. User Cost Approach

A central problem for statistical agencies is finding the right data; and, this is particularly true for imputing owner-occupied housing (OOH) statistics where the challenge emanates from accounting for transactions that are not *directly* measurable or observable. Statistical agencies like the BEA measure the value of housing services for OOH indirectly by using data that should closely approximate market rent that homeowners would expend. The two approaches briefly discussed above are the two approaches recommended by the 2008 SNA statistical framework: rental-equivalence and user cost.[2] Conceptually, absent transactions costs and other market frictions,

[2] More generally, the OECD Manual, Measuring Capital, summarizes the broader concept for the user cost of capital as follows: "Suppose the owner of an asset wants to determine the minimum price (before adding on costs of associated labour and overheads) at which he is willing to rent the asset during one period of time. In the simplest case, three main cost elements have to be considered: *(i)* the cost of financing or the opportunity cost of the financial capital tied up through the purchase of the asset; *(ii)* depreciation, *i.e.* the value loss due to ageing; *(iii)* revaluation, *i.e.* the expected price change of the class of assets under consideration." (OECD 2009, p.65)

basic economic principles predict that market rents should approximately equal average cost (in the long run) if markets are competitive. The underpinning theory of this (approximate) equality can be derived from capital theory, which is based on Jorgenson's (1963, 1967) theory of capital and investment, where the rental cost of capital will equal its *ex ante* user cost (Katz 2009).[3] For example, if rent for an identical home was much higher than its user cost incurred by a homeowner, then more people would buy this preferred capital asset and fewer would rent, bidding down rents and bidding up home prices to the point where rents and costs are approximately equal.[4]

BEA's current method used in the national accounts follows a rental-equivalence approach using data from the Census's Residential Finance Survey (RFS) to benchmark rent-to-value ratios for different value classes of properties, which is then used to impute average contract rent for owner-occupied properties across similar dimensions.[5] This weighted rental imputation constitutes what is often referred to as "space rent," which is then multiplied by corresponding aggregate housing unit counts from the Census's American Housing Survey (AHS) to obtain the aggregate estimate of the total imputed rent of owner-occupied housing. For a more detailed discussion of the BEA's current method, refer to Mayerhauser and McBride (2007) and Katz (2017).

The rental-equivalence method is often cited as a preferred method for this imputation because most countries have relatively thick rental markets with substantial data on market rents.

[3] As a thought experiment, one can think of user cost in this context as measuring the net expenditure associated with purchasing a home at the beginning of a period, incurring cost during the period, and selling the home at the end of the period, abstracting away from transaction costs and other market frictions. According to Jorgensonian capital theory, the rental rate for this home set at the beginning of the period would equal this expected cost, ex ante. See also McFadyen and Hobart (1978) for an instructive cross-walk from Jorgenson (1967) to a user cost for housing.

[4] Of course, this abstracts from risk, market imperfections, and transactions costs, which is particularly significant in housing (Bian, Waller, Wentland 2016). Thus, some gap might persist, but generally rents and user costs should move together over longer periods of time. In fact, recent empirical work by Goeyvaerts and Buyst (2019) has found a "strong correspondence" between rents and user costs using detailed micro-data.

[5] The BEA had last benchmarked these rent-to-value ratios using the 2001 RFS, the last time the data was available. Since then, the BEA has made quality and price adjustments primarily based on data from the BLS, which also relies on a rental-equivalence method for the CPI.

In fact, more than one-third of all housing units in the U.S. are rented to tenants. However, while the U.S. has a large number of tenant-occupied housing units, the distribution of rental units is not the same as owner-occupied units (Glaeser and Gyourko 2009), owner-occupied units have disproportionally more detached single-family residences (SFRs) and the distribution is skewed toward higher value homes. For additional discussion of this point and an illustration of these differences using recent Census data, see Aten (2018).

A key weakness of the rental equivalence method is that rents must then be interpolated to represent the full housing market, despite the stark differences between the distribution of characteristics of tenant versus owner-occupied homes and issues emanating from the thinness of certain rental markets (Glaeser and Gyourko 2009). When rental markets are thin, the SNA recommends "other means of estimating the value of housing services," (SNA 2008, p. 109) which has led researchers and statistical agencies to explore alternative methods like a user cost approach. This approach utilizes data on the cost to the user of owning a home (e.g., interest, taxes, maintenance/depreciation, etc., which varies directly with the price of a home) rather than rents of different, possibly unrepresentative tenant-occupied homes. For an instructive review of this voluminous literature and novel examples of developing user cost estimates, see Diewert (2008a, 2008b), Katz (2004), Verbrugge (2008), Davis, Lehnert, and Martin (2008), Haffner and Heylen (2011), Hill and Syed (2016), Aten (2018) and numerous other papers on this topic.

An important advantage of the user cost approach is coverage of directly observable data across more of the housing market. Because neither approach relies on a universe of observable data (i.e., rents are not observed for owner-occupied homes, nor are transaction prices observed in every period for all homes), estimates are derived from samples of survey data or market transactions data in a given period. Yet, transaction prices, the backbone of the user cost method,

are readily available for virtually all strata of the housing market, both tenant-occupied and owner-occupied homes in particular. While Gillingham (1983), Verbrugge (2008) and Diewert, Nakamura, and Nakamura (2009) and others have noted that the user cost approach has a number of weaknesses (e.g., greater volatility, sensitivity to interest rates, and conceptual issues with ex ante and ex post measurement), these would need to be weighed against weaknesses with the rental-equivalence approach (or any other approach, for that matter) to make the ultimate determination of which method to pursue. Nonetheless, weighing in on this debate falls outside the scope of this paper, as two necessary prerequisites for even considering a new approach are assessing whether it is feasible and conducting an initial evaluation of how the new estimates compare to the current approach, which is the aim of this paper.

3. Data

The novelty of this paper primarily resides with usage of new data. As we alluded to in the introduction, we use residential housing microdata from Zillow's ZTRAX data set. It contains transaction data as well as a large set of individual property characteristics for sales recorded from local tax assessor's data.[6] The data coverage is generally representative of the United States' national housing market, initially comprising 374 million detailed records of transactions across more than 2,750 counties.[7] This includes information regarding each home's sale price, sale date, mortgage information, foreclosure status, and other information commonly disclosed by a local

[6] Data are provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the authors and do not reflect the position of Zillow Group. Nonproprietary code used to generate the results for this paper is available upon request of the authors.

[7] Because some states do not require mandatory disclosure of the sale price, we currently do not have price data for the following states: Idaho, Indiana, Kansas, Mississippi, Missouri, Montana, New Mexico, North Dakota, South Dakota, Texas, Utah, and Wyoming. In addition, Maine has substantial missing data in our current sample, and we omit this state as well. Our method aggregates to the Census Division level by using housing unit counts from the ACS at the regional level. As a result, we assume that the states with data within a Census Division are reasonably representative of a state left out, which is an assumption we hope to explore in further research with supplemental data.

tax assessor's office. We link each transaction to each home's property characteristics into a single dataset. The assessment data typically includes an array of characteristics one would find on Zillow's website or a local tax assessor's office describing the home, namely the size of the home (in square feet), number of bedrooms and bathrooms, year built, and a variety of other characteristics.[8] We received all of this data in a somewhat raw form, requiring additional cleaning for research purposes.

We carefully scrutinized missing data and extreme values as part of our initial culling of outliers and general cleaning. The initial data set from Zillow contains sales of empty plots of land, some commercial property transactions, agricultural sales, and a host of types of properties that are outside the scope of the housing services estimates we aim to measure. Therefore, we limit the sample to single family homes, townhouses, rowhomes, apartments, condos, and properties that are most closely associated with the current estimates. We winsorize acreage at five acres (limiting the influence of large farms) and outlier homes that are on the upper tail of the distribution (i.e. are larger than 10,000 square feet or have more than five bedrooms, more than three bathrooms).[9] We also drop homes that sold for less than 30,000 dollars, the bulk of which are not arms-length transactions or they may represent teardowns. We cull homes that were built prior to 1865 or report a negative age of home (i.e. sale year – year built). While the Zillow data set contains a vast number of property characteristics, in our initial analysis we primarily rely on the variables above that have the most coverage nationally to limit how much data we would effectively discard.[10] We limit the

[8] Zillow's ZTRAX data contains separate transaction files by state, where all transactions need to be linked to corresponding assessment records. With guidance from Zillow, we were able to merge the bulk of the data, but not without some data loss (which figures into the size of our final sample).

[9] We also create indicator variables equal to one if the property had missing characteristic values or reported a lot size of zero or there are missing bedrooms or bathrooms.

[10] In untabulated regressions, we conducted a sensitivity analysis for subsets of the sample that employed more property characteristics to determine whether the results are sensitive to omitted variables for which we can control. Our results were generally robust to omitting variables that have more limited coverage.

sample years to 2001 through 2016, as the data is most complete for the vast majority of the states in our sample.

To ensure the quality of the final sample, we compared our cleaned Zillow sample to the U.S. Census American Community Survey (ACS) to ensure that this administrative data aligned with carefully collected (albeit more limited) survey data provided by the Census. Generally, there are only a limited set of characteristics of homes that were in both the ZTRAX data and the ACS (e.g., number of bedrooms, year built, number of rooms, tax amount, and an indicator for whether the property has more than 10 acres). When we compare them in aggregate, we find that they are quite similar in terms of their summary statistics. In untabulated results, we found that these shared variables across data sets had median and mean values that fell within a few percentage points of one another.

4. Methodology – An Idiosyncratic User Cost Approach

4.A. Overview

Generally, our approach using this microdata is motivated by constructing estimates from the bottom-up, as we estimate a user cost for *each individual property* in our data set and then aggregate upward to produce a weighted national-level estimate. We begin by estimating a simplified user cost of housing services for each home in the data set based on the formula:

$$U_{it} = P_{it}(i_t + \delta_i + \tau_{it} + \gamma_i - \mathrm{E}[\pi_i])$$

where for a given property (i) in quarter (t) P is the price of an individual home, i is the owner's discount rate or financial opportunity cost for a long-term asset like a home (as a proxy we use the nominal interest rate on a 10-year Treasury rate in quarter t),[11] δ is a constant representing

[11] While the data set includes individual interest rates for transacted properties, the coverage is not as universal as other variables. However, it is customary for user cost estimates to use a single market interest rate to reflect the financial opportunity cost of the long-term asset (e.g., see Aten 2018 for a recent example, among numerous others). Conceptually, if

depreciation and maintenance costs of 3.5%,[12] τ is the individual property's effective tax rate, γ is a constant 2% risk premium associated with owning relative to renting,[13] and $E[\pi]$ is expected appreciation (revaluation) for a given year as 2%, which assumes homeowners have a very long-term view of home prices appreciating approximately the same as overall inflation in the economy.[14] We vary the latter assumption in a second user cost calculation we discuss later in the paper, where price expectations are based on recent home price appreciation/depreciation in one's local area. Overall, our primary contribution to the literature is estimating national property-level user costs using idiosyncratic price and property tax data, which we describe in more detail below. While we simplify this method using some constants, we return to a discussion of these simplifications and ways of creating a more idiosyncratic and possibly more precise estimate later in section 6.

a homeowner purchased a home when rates were at 4%, but rates have since risen to 7%, the latter rate more closely represents the opportunity cost in that time period, as the homeowner could alternatively be earning a return on that equity of a similar long-term asset. The time series dynamics are similar if we use average 30-year mortgage rates, which we show later in the paper for robustness.

[12] A depreciation rate of 1.5% is common to the literature (e.g., Aten (2018) and Verbrugge (2008)), and Gill and Haurin (1991) use a constant of (1.5% + 2% =) 3.5% for the combined maintenance and depreciation term. Conceptually, there is wear and tear on a home that would be similar to what a renter would incur in the analogous tenant-occupied counterfactual, but primarily this is structural depreciation due to the property itself aging. Because these costs (on average) would be priced into a tenant's rent, it is logical to factor this into the imputation for owner-occupied properties. However, given that homes depreciate at different rates depending on age and other maintenance costs may vary by region and home type, we acknowledge that a constant rate here is likely an oversimplification, which we wish to vary in a later draft of this paper.

[13] This risk premium was used by Himmelberg, Mayer, and Sinai (2005) "to compensate homeowners for the higher risk of owning versus renting" (p.75). The constant 2% value they use originates from Flavin and Yamashita (2002), but Himmelberg et al. (2005) point out that a more sophisticated model would allow this risk premium to vary over time as the risk of owning relative to renting changes over time.

[14] Verbrugge (2008) rigorously considered a variety of measures of $E[\pi]$ using different forecast techniques, concluding that, "a very long horizon appreciation forecast (such as a long moving average), or an inflation forecast, should be used in the user cost formula" (p. 694). Preference for an *ex ante* long-horizon measure is consistent with Diewert's (2006) argument that, "it is unlikely that landlords use econometric forecasts of housing price appreciation one year away and adjust rents for their tenants every year based on these forecasts. Tenants do not like tremendous volatility in their rents and any landlord that attempted to set such volatile rents would soon have very high vacancy rates on his or her properties. It is, however, possible that landlords may have some idea of the long run average rate of property inflation for the type of property that they manage and this long run average annual rate of price appreciation could be inserted into the user cost formula." During the period we study, the Federal Reserve had maintained either an explicit or implicit target of 2% inflation over the long run (e.g., see their policy statements on their website regarding 2%: https://www.federalreserve.gov/faqs/money_12848.htm). Ex post, inflation, particularly in the housing market, departed from this target; but, use as an *ex ante* measure may not be unreasonable. For robustness, we consider a method where $E[\pi]$ is determined by recent experience in one's local area.

4.B Idiosyncratic P – Actual and Predicted

While Zillow already constructs property-level valuation estimates (Zestimates) using their propriety automated valuation model (AVM), for transparency we rely on a combination of actual transaction prices and, for homes that did not transact during our sample period, our own hedonic valuations based on their microdata. Because we have fine, transaction-level price data, we are able to first use actual market prices for *P* (when available). For example, if property *i* was purchased in the first quarter of 2010, then for that quarter the *actual* price was used for the transacted property (P in the formula above).[15] Turnover varies considerably by state and locality; approximately one-third of properties in our dataset sold at least once within the window we study (from 2001-2016). For the value of the home in the following quarter we posit that the price is simply the transacted price plus the average price appreciation/depreciation of the housing stock of the county (which we estimate using the same hedonic model we use for our price imputations discussed below). We use the same logic for the quarters proceeding that sale until there is a new sale of that property.[16] Broadly, using more direct price data conforms most closely to the principles of valuation laid out by the System of National Accounts (SNA), where market prices are "the basic reference for valuation in the SNA" (SNA 2008, p. 22),[17] and thus much of our aggregate calculation flows directly from millions of observed market prices underlying the housing stock.

[15] The ACS has home price data with reasonably good coverage; however, this data comes from asking survey respondents to place a value on their own home. An advantage of the Zillow data is that we have actual market transactions and predictions based on market data. With linked microdata, eventually we would like to explore the differences between these data sets for use in the national accounts.

[16] This method would likely be altered if it were implemented in national accounts over a longer time-series, as a single transaction price adjusted for inflation may be less predictive of the actual price in other years as the time series becomes much longer. For example, we may limit interpolations to a single five or ten year window; but, because our time series here only covers fifteen years we take the simplified approach.

[17] More specifically, the SNA recommends that statistical agencies use market prices when market prices are available, but "in the absence of market transactions, valuation is made according to costs incurred (for example, non-market services produced by government) or by reference to market prices for analogous goods or services (for example, services of owner-occupied dwellings)" (SNA 2008, p. 22).

For homes that did not sell during our sample period, we impute their prices based on transactions of similar homes that sold in each quarter using a hedonic model.[18] Conceptually, most of a home's value can be explained by its physical characteristics, location, and time; hence, our hedonic model uses sale prices of similar homes along these dimensions to estimate an imputed market valuation for each home in our data set. While this approach is somewhat simplified compared to more complex machine learning techniques as used by Zillow's proprietary AVM, an advantage of this hedonic approach is transparency, an important pillar of national accounting methods, where the model can be fully described to the public or users of the accounts if an approach like this were to be formally adopted. Therefore, we impute $\hat{P}$ based on the following hedonic model for each quarter separately:

$$Sale\ Price_{ij} = \alpha + \sum \beta\, X_i + \gamma LOCATION_j + \sum \delta\, sq.ft._i * LOCATION_j + \sum \varphi\, acreage_i * LOCATION_j + \varepsilon$$

where X is a set of physical characteristics (bedrooms, bathrooms, age of the structure, living area measured by square feet, lot size measured by acreage, whether the home was a single story, whether it had a pool, whether the home had a basement, whether it had a porch, and whether the home was new construction), location fixed effects, and interaction of location fixed effects with square footage and acreage, respectively.[19] For practicality in estimation, we initially use Census tract fixed effects, although we obtain similar estimates using finer-level geographic fixed effects

[18] Within-quarter hedonic regressions avoid issues of controlling for macro-level relevant time-varying factors that could bias predictions if not properly accounted for in the model.

[19] While the Zillow ZTRAX data contains a lot more information about individual properties that would help with valuation, we chose the variables with extensive coverage across all states in the data set. When compared to a fuller model that includes many more home characteristics, the marginal gain in precision was small compared to the potential loss in observations due to missing data in states/localities that do not regularly report certain variables. When one of the key characteristics (e.g. bedrooms) was missing, we bottom coded it and included a missing indicator in the regression rather than drop the sale entirely. We also included an indicator in the regression for whether the home had extreme values for any of these characteristics to account for non-linearities, as opposed to just dropping these observations as well.

like Census block groups or blocks.[20] To avoid making predictions with thin cells, we specify that a given tract have at least ten sales in the quarter of estimation. If this condition is not met within a given tract in a given quarter, we then estimate the same model only for observations that do not meet this threshold using county (FIPS) level fixed effects. While intensive for processing, allowing square footage and acreage to vary by location encapsulates the idea that valuation of these attributes varies widely across areas. For example, an additional 500 square feet in a home in New York City, will be valued much differently than the same addition upstate in Syracuse.[21] For non-single family residencies (non-SFRs), which we estimate separately from detached single family residencies (SFRs), we omit acreage and other SFR-specific characteristics from the hedonic model.

Despite this relatively simple hedonic model construction, for most states and most quarters, the model fit (R^2) fell within 0.8 and 0.9, producing errors that stack up quite reasonably compared to more sophisticated techniques. In order to assess the accuracy of our model's price predictions, we constructed a measure of error for each record for which we have an observed price as follows:

$$Average\ Percent\ Error\ (APE) = \frac{Predicted\ Price - Actual\ Price}{Actual\ Price}$$

Then, to obtain an aggregate error, the median of all APEs in a state in a given time quarter is multiplied by the share of the observations in that state in the total observations. This measure is plotted for each quarter for the aggregated SFRs in 36 states in Figure 3 below. Our model

[20] Smaller geographic units like block groups and blocks have fewer sales, which we found to be less ideal for quarterly predictions. In a previous draft, we had similar (albeit somewhat less precise) results to tracts using zip code fixed effects. We have also explored a variety of other specifications to improve model fit and predictions, including a semi-log specification, where sale price is logged, but these models produced similar results overall.

[21] This approach is used commonly in the hedonic valuation literature for housing and land. See, for example, Kuminoff and Pope (2013).

appears to slightly overestimate sales price for the majority of the period (2000-2016) and underestimate during the Great Recession and its immediate aftermath.

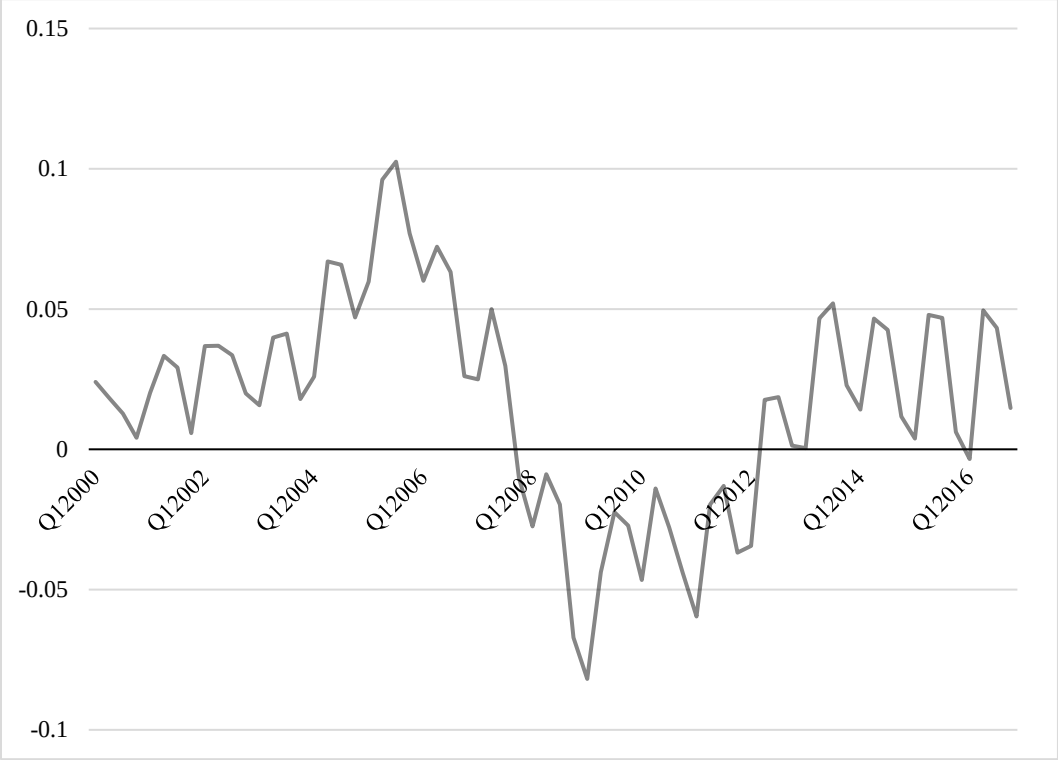


Figure 3: Median Percentage Error (Predicted-Actual), Weighted by State

4.C. Property Taxes

Property taxes vary widely across states and municipalities. As of 2017, the highest property tax state was New Jersey with an average effective tax rate of 2.31%, whereas Hawaii and Alabama have average rates of 0.32% and 0.48%, respectively.[22] Even within states there is considerable variation. Hence, for accurate estimates of user cost we attempt to account for the idiosyncratic nature of a property's taxes. Because the Zillow data is collected primarily from local tax assessor office databases, the coverage of property taxes is quite good. We use individual tax data to

[22] Variation in property taxes across states gained attention during the national coverage of the Tax Cuts and Jobs Act of 2017. For example, USA Today ran a story comparing effective property tax rates across the U.S.: https://www.usatoday.com/story/money/personalfinance/2017/04/16/comparing-average-property-taxes-all-50-states-and-dc/100314754/

determine a property's effective tax rate based on a denominator of *P* (actual or predicted price) rather than the corresponding assessment value associated with each property in the data.[23]

We made this choice for a couple reasons. First, regarding the denominator, the assessment value is often much lower than the market value, so if we apply the rate based on the assessed value to the market value of *P* in the user cost calculation we would overestimate the amount homeowners pay in our calculation. The degree of mis-assessment of value varies considerably by locale, and in some cases it is by design of local policies for states like California to have assessments tied to historical values for longer tenured homeowners. Second, this approach better reflects the average effective tax rate, because like other elements of the tax code, homeowners do not all pay the same posted rate due to local property tax relief exemptions and relief for special groups (Moulton, Waller, and Wentland 2018). Finally, in the present study we are unable to accurately determine the *net* tax bill for each homeowner or precisely consider the full range of offsetting tax benefits that come with homeownership (namely, mortgage interest deductions and state/local tax deductions); but, if we are able to successfully link this data to administrative data, then constructing more credible estimate of these benefits would be possible for future work.[24]

4.D. Quantity, Housing Counts, and Aggregation

Once we obtain user-cost estimates for millions of individual properties across the United States, we then aggregate to a weighted national estimate of housing services based on the

[23] We currently have one year of tax amount data from Zillow, but updating this data more often (preferably annually) may be required if this method is to be used for national accounts measurement. In rare cases where our computed tax rate estimates far exceeded the average tax rate of the state (by a factor of 3), with winsorize these observations to the state average. When they were much smaller (by a factor of 1/3), we also replaced them with the state average.

[24] Other studies like Himmelberg, Mayer and Sinai (2005) use a constant income tax rate for homeowners to back out an estimate for the mortgage interest and property tax deduction benefit, but our ambition is to use linked administrative data to back out a more precise estimate. Linkages to Census administrative data records, for example, would also allow us to better estimate maintenance and other costs for households (or, at least regionally – where wear and tear from climate and other factors may contribute to households reporting systematically different levels of maintenance expenditures) and to better understand housing market dynamics of populations of homeowners vs. renters. We return to this point in section 6.

corresponding quantities of the housing stock by location/region, type of home (single family residence (SFR) vs. non-SFR), and number of bedrooms. We use the weighted unit counts of the housing stock from Census's American Community Survey (ACS), which provide a yearly count of the aggregate number of residential housing units by Census Division, depicted in Figure 4. For illustrative purposes, refer to Table 1, where we show the calculation of our national estimate for Q4 of 2016. For each Census Division or region of the U.S., we multiply the average user cost for each type of home (SFR vs. non-SFR) for each bedroom category.[25]

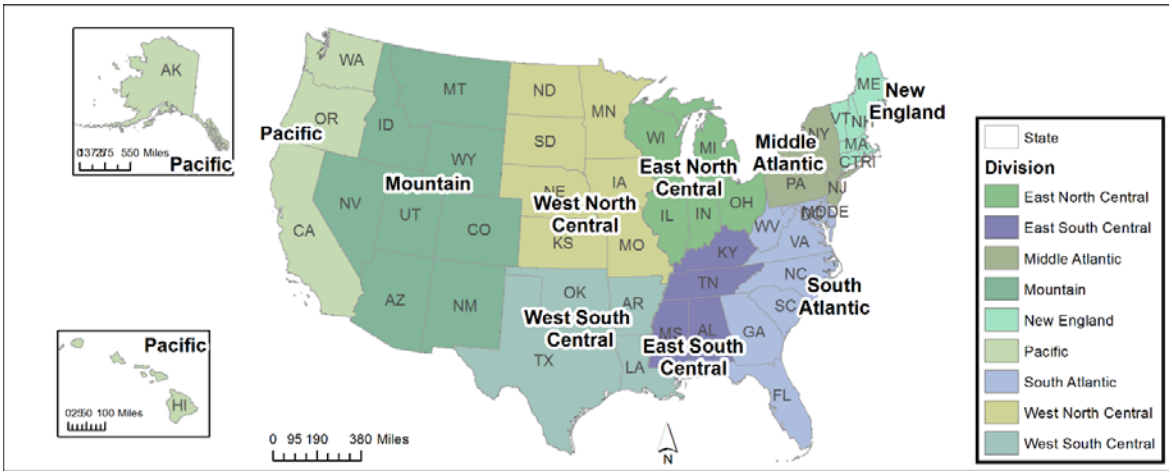


Figure 4: Census Divisions

Source: https://www.census.gov/geo/reference/webatlas/divisions.html

This method of aggregation assumes that the non-missing data is reasonably representative of the missing data. For example, Indiana's sale prices are missing from the ZTRAX data set, as it is among the non-disclosure states that does not ordinarily record sale prices in public use tax assessor data. Hence, our final aggregate estimates must assume that the average user costs imputed from sales in its Census region (Illinois, Michigan, Ohio, and Wisconsin) reflect the

[25] We use bedrooms as a proxy for size of the home to create categorical differences that more accurately reflect the weighted total. The bins are numbered 1 through 5+ in Table 1. However, for states that did not have good coverage of the number of bedrooms, we assumed that the distribution of user cost approximately aligned with the distribution of bedrooms and assigned homes to corresponding bins of bedrooms. In future work, we will explore using county-level quantity counts, as finer location averages could be more relevant that averages by physical characteristics.

Table 1: User Cost Aggregation – Summary Calculation for 2016q4

		SFR			Non-SFR		
Division	Bedrooms	Ave. User Cost ($)	Q	P*Q ($B)	Ave. User Cost ($)	Q	P*Q ($B)
1	1	12,067	78,841	1	13,914	767,907	10
	2	14,063	490,837	7	17,077	1,008,062	17
	3	18,411	1,605,454	30	21,276	530,671	11
	4	27,695	838,200	23			
	5	41,531	206,494	9			
2	1	8,094	141,935	1	10,972	2,591,742	28
	2	10,352	1,031,546	10	18,475	2,622,026	48
	3	15,287	3,609,221	55	34,374	2,175,227	75
	4	25,385	2,237,124	57			
	5	44,348	583,540	25			
3	1	4,975	218,903	1	6,136	1,755,527	11
	2	6,166	1,938,344	12	10,396	2,484,126	26
	3	9,972	6,567,881	66	16,363	940,174	15
	4	17,602	2,990,114	53			
	5	28,715	670,174	19			
4	1	6,271	146,868	1	9,327	767,992	7
	2	7,954	1,043,513	8	11,423	955,087	11
	3	10,649	2,680,432	28	14,673	349,577	5
	4	13,907	1,526,842	21			
	5	19,126	474,931	9			
5	1	6,232	196,669	1	8,986	2,046,952	18
	2	7,768	1,919,499	15	12,665	3,258,395	41
	3	12,311	7,543,817	93	21,940	1,877,190	41
	4	22,078	3,747,649	83			
	5	37,967	1,105,643	42			
6	1	4,864	93,315	0	6,833	446,856	3
	2	4,543	734,721	3	8,860	693,547	6
	3	7,099	2,895,815	21	13,902	247,577	3
	4	13,538	1,058,912	14			
	5	21,924	246,716	5			
7	1	2,627	192,651	5	6,147	1,382,770	9
	2	5,032	1,171,105	6	10,334	1,346,042	14
	3	8,400	4,647,022	40	16,195	450,554	7
	4	13,830	2,158,298	30			
	5	26,371	415,247	10			
8	1	10,246	129,086	1	10,511	786,290	8
	2	12,877	762,322	10	12,341	1,061,530	13
	3	14,641	2,602,678	38	14,585	432,452	6
	4	23,784	1,598,170	38			
	5	35,644	626,676	22			
9	1	17,596	316,702	6	19,425	2,534,134	49
	2	22,625	1,578,474	36	25,308	2,880,140	73
	3	26,428	5,078,692	134	34,368	1,132,054	39
	4	36,299	2,935,190	107			
	5	49,116	758,066	37			
			SFR Total =	**1,230**		**Non-SFR Total =**	**598**
						Total User Cost: 1,230 + 598 =	**1,828**

Indiana market.[26] Missing data itself is not a prohibitive limitation for constructing national accounts, as statistical agencies always have limited data; but, the issue is more a matter of the extent of the representativeness of the data we do have. While many of these states are reasonably represented by their neighboring states' housing markets, as the Indiana case may be, one exception might be Texas (the largest state for which we have missing price data) where the current method may be the most problematic, simply because of the variability of the housing markets within the state. If this method, or some variation of it using similar data, were to be adopted by the BEA, supplemental data would be required to verify these assumptions or to re-weight the estimates to better represent the missing states' housing markets. The scope of this study, however, is to explore how far this particular "big data" set can go toward this end.[27]

4.E Varying Ex Ante Expected Price Appreciation/Depreciation

Finally, we vary the $E[\pi]$ term of ex ante expected price appreciation for robustness. Our default specification assumes a very long-run view of home price inflation of a constant 2% per year, despite the fact that homeowners during this period may very well have perceived price appreciation quite differently, particularly for some regions that experienced steep price fluctuations. To test how the results differ if homeowners had drastically different expectations than we assume in our default specification, establishing a lower bound of sorts, we assume the opposite end of the spectrum for our alternative specification. That is, if our default is that homeowners take a *constant long-run, national* view of price expectations, then the opposite might be a *variable short-run, local* view of price expectations. Thus, our alternative specification

[26] Recall that one of the limitations of this data set is that there is no price data from the following states: Idaho, Indiana, Kansas, Mississippi, Missouri, Montana, New Mexico, North Dakota, South Dakota, Texas, Utah, and Wyoming. Maine is also excluded due to limited data in a number of quarters of our sample period.

[27] The American Housing Survey (AHS) also has high quality data on the unit counts of the housing stock, but the survey is only available every other year. While the counts are not always identical across surveys, the differences are relatively small. In future work, we plan to use linked Census data to construct our own unit weights with our data.

assumes that homeowners expect ex ante price appreciation to be their local (county-level) average price inflation from the prior year (quarter $t - 4$). This is calculated by taking the percent change of the median predicted price by county over the previous four quarters from our hedonic model estimates discussed above.[28] While this is somewhat simplistic, our goal is to provide a sense of a reasonable range of possible estimates, as a more moderate moving average as in Verbrugge (2008) may produce an estimate somewhere in between this range of results, albeit closer to the long-run default specification.[29]

5. Results

Our full set of results for all years and quarters in our sample appear in Table 2, which shows both the total and average user cost estimates of housing services as well as the corresponding estimates by housing type (SFR vs. non-SFR) by quarter. A better visual of this data is shown in Figures 5 and 6. Specifically, Figure 5 illustrates the default specification graphically over time, broken out by housing type using the default user cost specification, showing similar time series dynamics and that the user cost of detached SFRs are consistently higher than non-SFRs, as one would expect.

The key figure of the paper is Figure 6, where we compare our average yearly user cost measure of housing services with the BEA's yearly estimate of housing services from PCE. Note that we compare the full estimates of aggregate housing services because we are estimating user

[28] Note that this is not seasonally adjusted. Some of the volatility in prices will be from purely seasonal factors. This can be augmented by applying a standard seasonal adjustment. For now we are reporting the raw, unadjusted results.

[29] Generally, countries that employ a user cost method for housing omit the $E[\pi]$ term entirely, simplifying the calculation (Diewert and Nakamura 2009). One way of thinking about this simplification involves referring back to the reason why the $E[\pi]$ term is factored in the calculation in the first place. As a thought experiment, the user cost method is often pitched as calculating the cost of an owner who purchases a home at the beginning of a period and sells it at the end (assuming away transactions costs). The $E[\pi]$ term in that case would simply be the capital gain/loss during a given period; but, if the next period begins with repurchasing the same home at the price from the end of the last period, then the capital gain/loss is essentially erased immediately. For now, we remain somewhat agnostic to the different approaches by offering results for multiple ways of incorporating $E[\pi]$ into user cost; our default specification comes at the suggestion of feedback we received from the NBER-CRIW Pre-Conference in 2018.

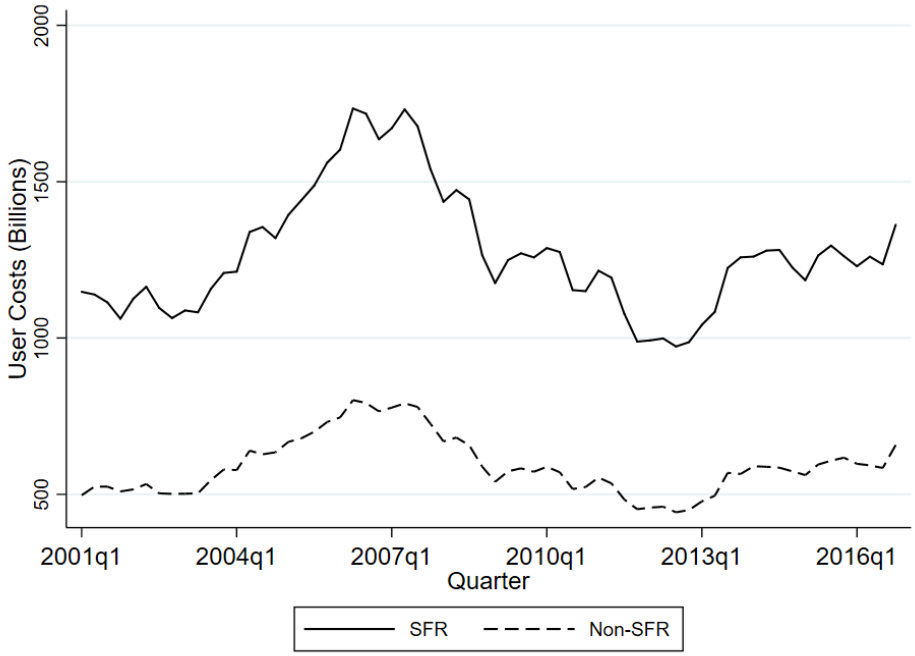


Figure 5: Total Quarterly User Costs by SFR/Non-SFR (Default Specification)

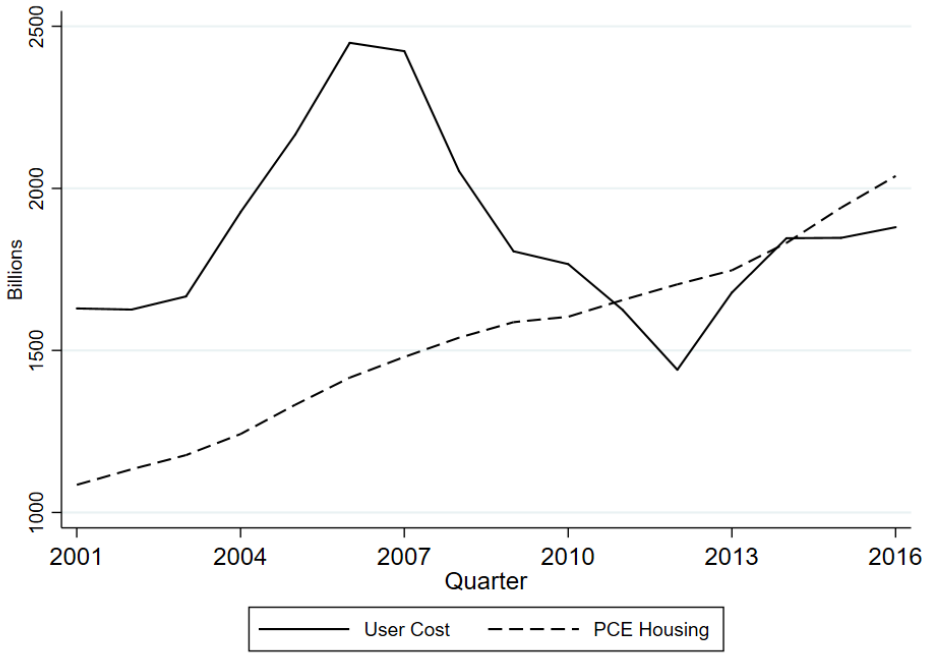


Figure 6: Total Yearly User Cost (Default) Compared to PCE Housing Estimates

Table 2: Housing User Costs by Quarter from 2001 through 2016

	Full Sample		SFR		Non-SFR	
	Total User Cost ($B)	Ave. User Cost ($)	Total User Cost ($B)	Ave. User Cost ($)	Total User Cost ($B)	Ave. User Cost ($)
2001q1	1,645	16,751	1,148	17,262	497	13,503
2001q2	1,664	16,944	1,139	17,534	525	14,199
2001q3	1,639	16,683	1,114	17,267	525	14,587
2001q4	1,571	15,997	1,062	16,496	509	13,961
2002q1	1,642	16,562	1,126	17,150	516	14,216
2002q2	1,697	17,119	1,164	17,811	533	14,678
2002q3	1,600	16,138	1,096	16,942	503	13,860
2002q4	1,566	15,792	1,064	16,310	502	13,840
2003q1	1,590	15,863	1,088	16,475	502	13,668
2003q2	1,586	15,816	1,082	16,339	503	13,767
2003q3	1,703	16,983	1,157	17,375	545	14,670
2003q4	1,788	17,835	1,208	18,124	580	15,741
2004q1	1,790	17,594	1,212	17,818	578	15,742
2004q2	1,979	19,453	1,339	19,673	640	17,320
2004q3	1,983	19,495	1,355	19,825	628	17,096
2004q4	1,954	19,211	1,320	19,355	635	17,160
2005q1	2,062	20,035	1,394	20,165	668	17,587
2005q2	2,120	20,598	1,441	20,767	679	17,913
2005q3	2,188	21,263	1,488	21,512	700	18,406
2005q4	2,292	22,276	1,561	22,540	732	19,147
2006q1	2,349	22,717	1,603	22,833	746	19,595
2006q2	2,536	24,525	1,734	24,689	801	21,041
2006q3	2,510	24,278	1,718	24,438	792	20,752
2006q4	2,401	23,226	1,636	23,271	766	20,052
2007q1	2,449	23,459	1,671	23,489	778	20,278
2007q2	2,522	24,164	1,731	24,375	791	20,603
2007q3	2,457	23,542	1,678	23,596	780	20,338
2007q4	2,265	21,704	1,540	21,762	726	19,012
2008q1	2,105	20,025	1,436	20,193	669	17,298
2008q2	2,155	20,497	1,473	20,680	682	17,801
2008q3	2,100	19,976	1,443	20,355	657	17,084
2008q4	1,854	17,633	1,265	17,882	589	15,296
2009q1	1,716	16,237	1,176	16,619	540	13,859
2009q2	1,823	17,251	1,249	17,601	574	14,692
2009q3	1,854	17,544	1,271	17,946	583	15,010
2009q4	1,830	17,318	1,258	17,750	572	14,782
2010q1	1,876	17,594	1,288	18,008	588	15,002
2010q2	1,846	17,314	1,275	17,849	571	14,580
2010q3	1,671	15,667	1,153	16,172	517	13,208
2010q4	1,674	15,696	1,150	16,135	524	13,357
2011q1	1,769	16,520	1,215	17,006	554	13,942
2011q2	1,728	16,138	1,193	16,674	535	13,493
2011q3	1,561	14,576	1,077	15,057	484	12,128
2011q4	1,441	13,458	989	13,865	453	11,416
2012q1	1,450	13,420	993	13,770	458	11,427
2012q2	1,459	13,504	999	13,841	461	11,617
2012q3	1,416	13,100	973	13,517	443	11,023
2012q4	1,437	13,296	987	13,704	450	11,199
2013q1	1,520	14,010	1,042	14,389	478	11,795
2013q2	1,579	14,559	1,084	14,995	496	12,228
2013q3	1,793	16,531	1,224	16,926	569	13,971
2013q4	1,824	16,809	1,258	17,384	565	13,995
2014q1	1,850	16,903	1,261	17,265	590	14,303
2014q2	1,868	17,064	1,280	17,535	588	14,274
2014q3	1,867	17,056	1,282	17,604	585	14,201
2014q4	1,799	16,435	1,226	16,788	574	13,966
2015q1	1,747	15,821	1,185	16,075	562	13,398
2015q2	1,860	16,844	1,264	17,185	596	14,250
2015q3	1,904	17,241	1,296	17,636	608	14,519
2015q4	1,879	17,020	1,262	17,178	617	14,773
2016q1	1,828	16,447	1,230	16,539	598	14,219
2016q2	1,853	16,677	1,260	16,986	593	14,050
2016q3	1,820	16,381	1,236	16,634	585	13,851
2016q4	2,020	18,182	1,362	18,167	658	15,706

cost for all residential homes in our sample, applying the same method to all homes whether they are owner-occupied or not.[30] Our aggregate measure of housing was initially much higher than the BEA's estimate in 2001, but this gap widened precisely when home prices throughout much of the U.S. appreciated considerably during the run up to the financial crisis and Great Recession.

The more pronounced path of the user cost-based estimate from 2001 through 2010, during the infamous bubble-bust years, bears a striking resemblance to national house price indices like Case-Shiller's or FHFA's, rising just under $1 trillion from 2001 to the peak in 2006-7, with a similarly precipitous fall in the several years that followed. Broadly, this result is consistent with other recent work like Braga and Lerman (2019) who assess the divergence in consumer price index (CPI) measures using a user cost vs. rental-equivalence approach, also finding a stark contrast between the two approaches over these years. However, beginning around 2010, the user cost-based estimate of housing services using Zillow data has tracked much more closely to the housing estimate based on the BEA's current rental-equivalence method.

One alternative specification of the user cost method, factoring in very recent (one year), very local (county-level) price expectations, depicts a more pronounced bubble and bust in its measurement of housing services of the same time period. Figure 7 shows a user cost closer to the rental equivalence estimates early in the 2000s, but also shows price expectations producing a much sharper peak and trough with the alternative specification, with the level in recent years being considerably smaller than current BEA estimates of housing. But, given that this specification is much more aggressive in its price expectations assumptions, this result should be

[30] Also note that aside from methodology, there are other small differences that remain. For example, we do not include the imputed rent for farm dwellings, as we cull properties zoned for agriculture and we do not have separate estimates for group homes, nor do we differentiate between vacant and occupied-dwellings. But, these estimates are small and relatively constant over time, so they would not account for much of the differences in price dynamics over time in Figure 4. With linked administrative data, future work could make vacancy rate adjustments to our user cost estimates.

seen as one of the more volatile series this data can produce with this approach, and therefore interpreted with more than a grain of salt, so-to-speak. Indeed, this is one reason why most countries that actually employ the user cost method for housing in their national accounts or price indices often simplify this method further by omitting the price appreciation term in the user cost calculation (Diewert and Nakamura 2009).

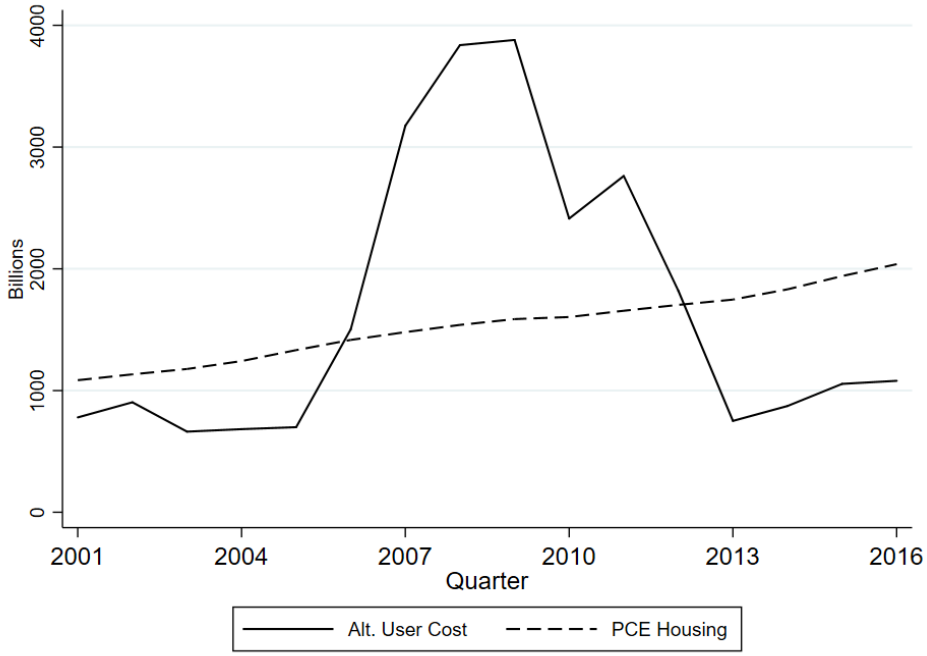


Figure 7: Total Alternative User Cost Compared to PCE Housing

For robustness, we vary some of the assumptions underlying the user cost formula, which we show in Figure 8. First, rather than incorporate a fixed homeownership risk premium of 2%, one alternative would be to use the average 30-year fixed mortgage rate as a stand-in for the 10-year treasury rate and this 2% constant. The 30-year mortgage rate generally tracks the time-series dynamics of other long-term interest rates like the 10-year Treasury, but it contains this additional risk premium that can vary slightly over time due to market conditions. Not surprisingly, this specification produces very similar results to our default specification, due to the stability of this premium over our sample period. Second, if we omit the $E[\pi]$ term entirely, a practice that some

countries have elected to do when implementing a user cost approach, this shifts the series upward, effectively reflecting more costly housing services across the entire time series. Finally, note that because our E[π] term and risk premium term are both constants, one can also think of our default specification as simply including offsetting terms; but, in future work we intend to explore additional ways to measure these more idiosyncratically as a further robustness test.



Figure 8: Comparing Different User Cost Methods with PCE

An important benefit to calculating user cost estimates with microdata is that there is greater scope for separating estimates geographically or by housing type. More generally, national statistical offices face increasing demands by users for finer partitions of the national accounts, which is a key advantage of "big data" over traditional designed survey data that suffers to a greater extent from a thin cell problem. As an example, Figures 9 and 10 show average user cost by region (Census Division) for single family residences (SFR) and non-SFR's respectively, although the data easily allows us to provide measures at the county or tract level (except, of course, for states

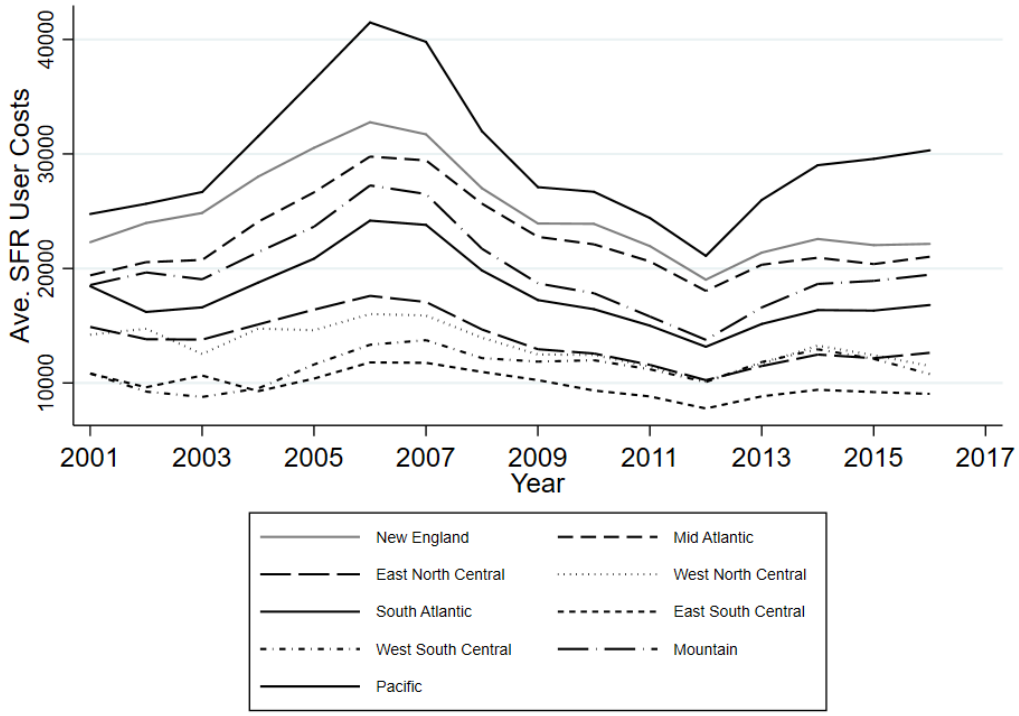


Figure 9: Average Yearly User Costs for SFR by Census Division

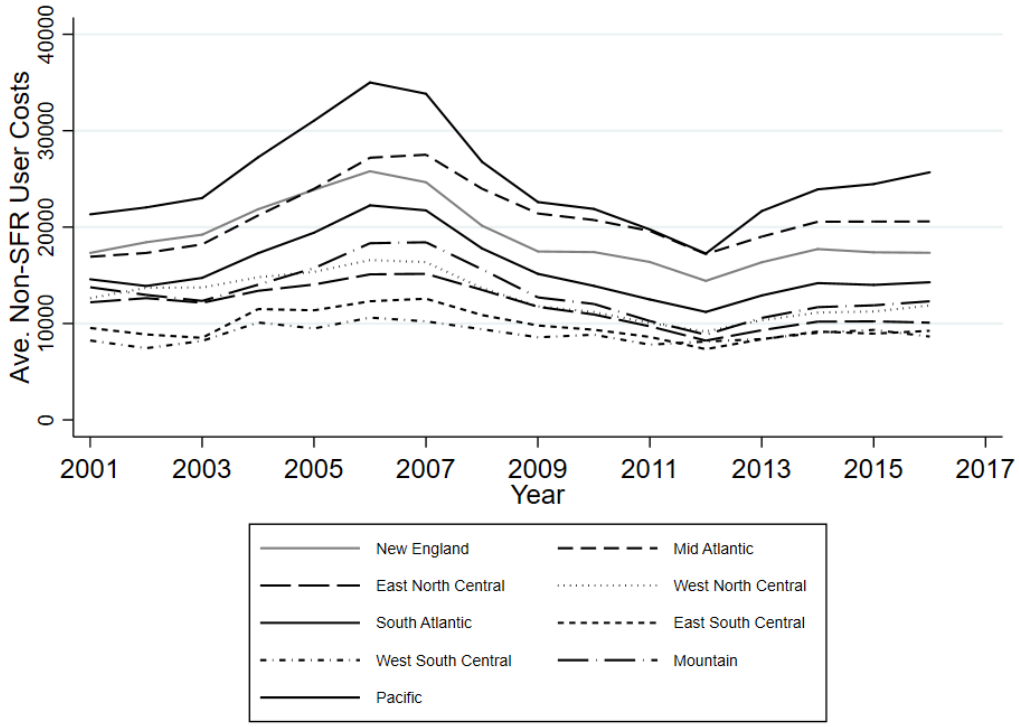


Figure 10: Average Yearly User Costs for Non-SFR by Census Division

with missing price data). As a reasonableness check, the estimates produce expected results that the Pacific region has the highest average user costs of housing, followed by New England, with several regions at the bottom experiencing mild, if any, bubble-bust market dynamics. This is consistent with numerous other regional metrics of the housing market over this same period.

Finally, while large aggregate estimates are often the focus of NIPA estimates, many users prefer per unit averages. Figure 11 depicts average user cost per residential unit for three different specifications and the corresponding BEA per unit space rent estimate. While the shape is nearly identical to Figure 8, the magnitudes may be helpful for assessing reasonability of the estimates.

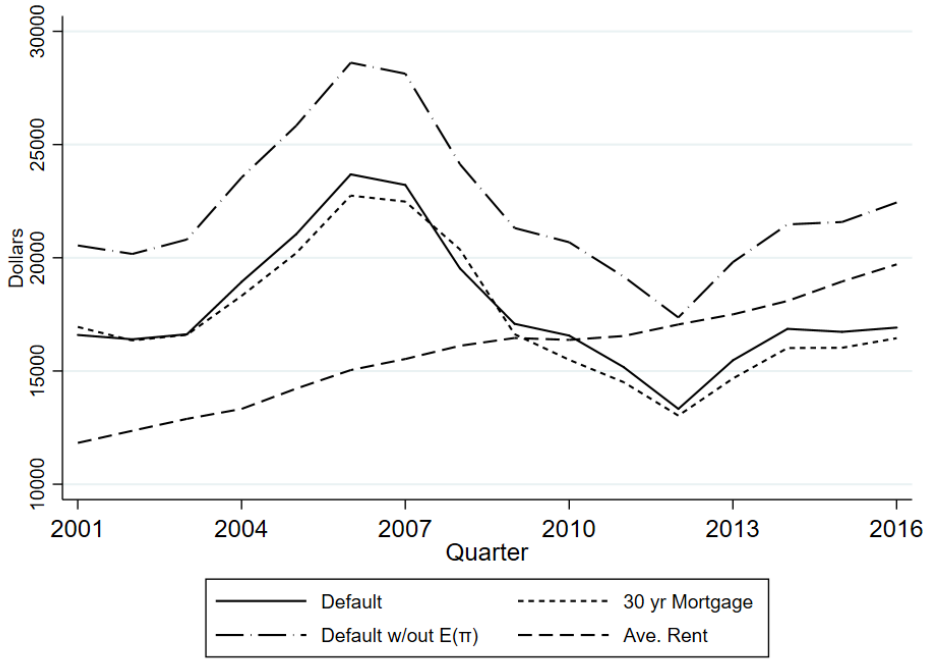


Figure 11: Average User Costs and PCE Average Rent

6. Discussion

We find that a user cost method using fine-microdata from Zillow can produce estimates of housing services comparable to the BEA's current method, at least for the most recent years we estimate. However, the departure from the rental-equivalence method during the first decade of

this century (and, extended periods prior to that, based on other studies using different data) shows that convergence of these estimates is far from guaranteed. And, if there are systematic divergences, particularly when the housing sector is experiencing a pronounced boom-bust cycle, a central question for national statistical offices will be: to what extent should housing estimates reflect underlying asset appreciation (that does not appear in rental data), which may or may not be temporary? And, which conception of aggregate housing is more relevant to users of the data and policymakers?[31]

We made a number of methodological simplifications and assumptions, which we intend to explore further in future iterations and refinements of these estimates. Additional precision gained from refining these estimates may help, at least in part, bridge the aforementioned gap. For example, other studies have accounted for mortgage interest and property tax deductibility in the user cost calculation for housing. This benefit was higher during the boom (lowering user costs) and lower during the bust, potentially accounting for some of the cyclical departure of user cost from the rental equivalence estimates. Or, insofar as maintenance and depreciation vary idiosyncratically or by region, a more sophisticated approach could exacerbate user cost differences if high price areas experienced relatively higher costs during the boom period. In either case, linked administrative data could help us answer these questions by creating idiosyncratic, property-specific estimates of the tax benefits, maintenance and depreciation costs, and a host of other refinements that could generate even more precise estimates. Further, there is additional scope for making refinements based on our current data, depreciation varying by the age of the home, which we will also explore. Finally, linked data administrative data may also help bridge

[31] There is evidence that the economic decisions of homeowners are, in fact, influenced by price appreciation/depreciation of their homes and housing wealth. See, for example, Mian and Sufi (2011), Mian, Rao, and Sufi (2013), Campbell and Cocco (2007), and Lowenstein (2018). Further, a related question would be: which conception of aggregate housing would be the most useful to monetary policymakers? We leave this, however, to future research.

the gap of our understanding of which user cost assumptions most directly compare to market rents, particularly for tenant-occupied homes for which we have rent data and user cost estimates based on Zillow data, as this comparison would show the most direct apples-to-apples comparison of the two methodological approaches.

While these and other fundamental questions remain, there is a great deal of potential upside to incorporating new data and exploring new methods into the national accounts, which is a key motivator of this study. Statistical agencies are continuously seeking ways to lower response burden for survey respondents, which is of increasing concern in an era of falling response rates more generally, and to find more cost-effective means for delivering statistics to users. If "big data" sources can substantially improve precision for regional and type stratification, for example, or even supplement parts of the current method where data may be thin, then a wholesale replacement of the current method may be a false dichotomy, as a hybrid or supplemental approach could be a valid consideration as well. We leave this, however, to future research.

References

Aten, Bettina. 2018. Valuing Owner-Occupied Housing: an empirical exercise using the American Community Survey (ACS) Housing files, BEA Working Paper, March 2018.

Bian, X., Waller, B.D. and Wentland, S.A., 2016. The Role of Transaction Costs in Impeding Market Exchange in Real Estate. *Journal of Housing Research*, 25(2), pp.115-135.

Braga, B. and Lerman, R.I., 2019. Accounting for homeownership in estimating real income growth. *Economics Letters*, 174, pp.9-12.

Bureau of Economic Analysis. 2018. Table 2.5.5 Personal Consumption Expenditures by Function. Accessed from https://apps.bea.gov/iTable/iTable.cfm?reqid=19&step=2#reqid=19&step=2&isuri=1&1921=survey on December 17, 2018.

Campbell, J. Y. and Cocco, J. F. (2007). How do house prices affect consumption? evidence from micro data. *Journal of Monetary Economics*, 54 (3), 591–621.

Davis, M.A., Lehnert, A. and Martin, R.F., 2008. The Rent-price ratio for the aggregate stock of owner-occupied housing. *Review of Income and Wealth*, 54(2), pp.279-284.

Diewert, W. Erwin, 2006. "Conclusions and Future Directions", Summary paper of the OECD-IMF Workshop on Real Estate Price Indices, Paris, available at http://www.oecd.org/dataoecd/32/21/37848333.pdf

Diewert, W. E. 2008a, The Treatment of Owner Occupied Housing and Other Durables in a Consumer Price Index, in W. E. Diewert, J. Greenlees and C. Hulten (eds.), Price Index Concepts and Measurement, NBER Studies in Income and Wealth, University of Chicago Press, 2008.

Diewert, W. E. 2008b, "Conclusions and Future Directions", The Paris OECD-IMF Workshop on Real Estate Price Indexes, Paris, November 6-7m 2006.

Diewert, W.E., Nakamura, A.O., 2009. Accounting for housing in a CPI. In: Diewert, W.E., Balk, B.M., Fixler, D., Fox, K.J., Nakamura, A.O., 2009. Price And Productivity Measurement, vol. 1 – Housing, Trafford Press, pp. 7–32 (Chapter 2). Available from: www.vancouvervolumes.com and <www.indexmeasures.com>

Diewert, W.E., Nakamura, A.O. and Nakamura, L.I., 2009. The housing bubble and a new approach to accounting for housing in a CPI. *Journal of Housing Economics*, 18(3), pp.156-171.

Flavin, Marjorie and Takashi Yamashita, 2002. "Owner-Occupied Housing and the Composition of the Household Portfolio." *American Economic Review.* March, 92:1, pp. 345–62.

Garner, T.I. and Verbrugge, R., 2009. Reconciling user costs and rental equivalence: Evidence from the US consumer expenditure survey. *Journal of Housing Economics*, 18(3), pp.172-192.

Gill, H.L. and Haurin, D.R., 1991. User cost and the demand for housing attributes. *Real Estate Economics*, 19(3), pp.383-396.

Gillingham, Robert. 1983. Measuring the Cost of Shelter for Homeowners: Theoretical and Empirical Considerations, *Review of Economics and Statistics*, 65, 254–65.

Glaeser, Edward L. and Joseph Gyourko, 2009. "Arbitrage in Housing Markets," in E. L. Glaeser and John M. Quigley editors, Housing markets and the economy: risk, regulation, and policy: essays in honor of Karl E. Case, Lincoln Institute of Land Policy, Cambridge, MA.

Goeyvaerts, G. and Buyst, E., 2019. Do market rents reflect user costs? *Journal of Housing Economics*. Forthcoming.

Haffner, M. and Heylen, K., 2011. User costs and housing expenses. Towards a more comprehensive approach to affordability. *Housing Studies*, 26(04), pp.593-614.

Hill, R.J. and Syed, I.A., 2016. Hedonic price–rent ratios, user cost, and departures from equilibrium in the housing market. *Regional Science and Urban Economics*, 56, pp.60-72.

Himmelberg, C., Mayer, C. and Sinai, T., 2005. Assessing high house prices: Bubbles, fundamentals and misperceptions. *Journal of Economic Perspectives*, *19*(4), pp.67-92.

Jorgenson, D.W., 1963. Capital theory and investment behavior. *American Economic Review* 53 (2), 247–259.

Jorgenson, D., 1967. The theory of investment behavior. In *Determinants of Investment Behavior* (pp. 129-175). NBER.

Katz, Arnold J. (2004), "Estimating Dwelling Services in the Candidate Countries: Theoretical and Practical Considerations in Developing Methodologies Based on the User Cost of Capital Measure," Bureau of Economic Analysis #BE-54.

Katz, A.J., 2009. Estimating dwelling services in the candidate countries: theoretical and practical considerations in developing methodologies based on a user cost of capital measure. In: Diewert, W.E., Balk, B.M., Fixler, D., Fox, K.J., Nakamura, A.O., 2009. Price and Productivity Measurement, vol. 1 – Housing, Trafford Press, pp. 33–50 (Chapter 3). Available from: <www.vancouvervolumes.com> and <www.indexmeasures.com>.

Katz, Arnold J. (2017), "Imputing Rents to Owner-Occupied Housing by Directly Modelling Their Distribution", (WP2017-7), BEA Working Paper, August 2017.

Kuminoff, N.V. and Pope, J.C., 2013. The value of residential land and structures during the great housing boom and bust. *Land Economics*, 89(1), pp.1-29.

Lowenstein, Lara. 2018. Consumption of Housing During the 2000s Boom: Evidence and Theory, Working Paper.

Mayerhauser, Nicole and Denise McBride. 2007. "Treatment of Housing in the National Income and Product Accounts," BEA Staff Study presented before the Society of Government Economists at the Annual Convention of the Allied Social Science Associations, December.

McFadyen, S. and Hobart, R., 1978. An alternative measurement of housing costs and the consumer price index. *The Canadian Journal of Economics*, *11*(1), pp.105-112.

Mian, Atif, and Amir Sufi. 2011. House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis. *American Economic Review*, 101 (5): 2132-56.

Mian, A., Rao, K. and Sufi, A. 2013. Household balance sheets, consumption, and the economic slump. *Quarterly Journal of Economics*, pp. 1–40.

Moulton, J.G., Waller, B.D. and Wentland, S.A., 2018. Who Benefits from Targeted Property Tax Relief? Evidence from Virginia Elections. *Journal of Policy Analysis and Management*, 37(2), pp.240-264.

United Nations, Commission of the European Communities, International Monetary Fund, Organisation for Economic Co-operation and Development, and World Bank, 2008 System of National Accounts, United Nations, NY, 2008. https://unstats.un.org/unsd/nationalaccount/sna2008.asp

Verbrugge, R., 2008. The puzzling divergence of rents and user costs, 1980–2004. *Review of Income and Wealth*, 54(4), pp.671-699.