This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Big Data for Twenty-First-Century Economic Statistics

Volume Authors/Editors: Katharine G. Abraham, Ron S. Jarmin, Brian Moyer, and Matthew D. Shapiro, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-80125-4 (cloth), 978-0-226-80139-1 (electronic)

Volume URL: https://www.nber.org/books-and-chapters/big-data-twenty-first-century-economic-statistics

Conference Date: March 15-16, 2019

Publication Date: Februrary 2022

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

Chapter Author(s): Marina Gindelsky, Jeremy G. Moulton, Scott A. Wentland

Chapter URL: https://www.nber.org/books-and-chapters/big-data-twenty-first-century-economic-statistics/valuing-housing-services-era-big-data-user-cost-approach-leveraging-zillow-microdata

Chapter pages in book: p. 339 – 370

12

Valuing Housing Services in the Era of Big Data

A User Cost Approach Leveraging Zillow Microdata

Marina Gindelsky, Jeremy G. Moulton,
and Scott A. Wentland

12.1 Introduction

Housing is an important part of the economy and the national economic accounts. As part of the 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

Marina Gindelsky is a research economist in the Office of the Chief Economist at the Bureau of Economic Analysis.

Jeremy G. Moulton is an associate professor of public policy at the University of North Carolina at Chapel Hill.

Scott A. Wentland is a research economist in the Office of the Chief Economist at the Bureau of Economic Analysis.

We would like to thank the organizing committee and participants of the 2018 NBER-CRIW Pre-Conference and corresponding 2019 Conference on Big Data for 21st Century Economic Statistics, as well as the following individuals for their valuable input: Katharine Abraham, Erwin Diewert, Dennis Fixler, Kyle Hood, Kurt Kunze, Han Liu, Raven Molloy, Brent Moulton, Mick Silver, Dylan Rassier, Matthew Shapiro, Brian Smith, and Randal Verbrugge for their helpful comments. All errors are our own. Any views expressed here are those of the authors and not necessarily those of the Bureau of Economic Analysis or the US 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 authors and do not reflect the position of Zillow Group. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see https://www.nber.org/books-and-chapters/big-data-21st-century-economic-statistics/valuing-housing-services-era-big-data-user-cost-approach-leveraging-zillow-microdata.

quantities.[1] The flow of housing services in GDP is, 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 because it amounts to the aggregate sum of rents paid for all residential units over a given period. The analogous calculation for homeowners 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 that the estimate of housing services is not arbitrarily distorted based on the decision to rent versus own a home, which can vary substantially across time and space.[2] Historically, both tenant- and owner-occupied housing have accounted for a substantial proportion of overall consumer expenditures and the economy more generally (approximately 16 percent of PCE, or about 10 percent of GDP final expenditures), and have been relatively stable over recent decades, as shown in figure 12.1 below.

The PCE housing series has risen steadily over the last couple of decades, congruent with other official series like the Consumer Price Index (CPI) Rent Index and the CPI Owners' Equivalent Rent Index, both depicted in figure 12.2 below. A common element among these statistics is that they rely on reported rents from survey data, as the BEA's current method follows a rental-equivalence approach leveraging survey data. Moreover, the BEA's housing estimates were adjusted over this time period using the owner-occupied rent series directly (for reasons we discuss in more depth in the next section). Recently, however, the academic literature has begun to reexamine the rental market over this period using "Big Data" sources, finding that using alternative data and methods reveals a different picture. For example, when rents are measured using different data, as shown by the Ambrose-Coulson-Yoshida (ACY) Repeat Rent Index (also depicted in figure 12.2) using market transaction data from Experian RentBureau, a conflicting story emerges as rents flatten out earlier than the CPI series and even fall in absolute terms in 2008–2009.[3] This drop in rents, while less dra-

1. Housing is included in both consumption and investment expenditures in GDP statistics, where new construction is accounted for in Residential Fixed Investment. The focus of this paper is on Housing Services within Personal Consumption Expenditures.

2. 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" (United Nations et al. 2010, 99).

3. This index is derived from Ambrose, Coulson, and Yoshida's (2015) recent work constructing a rent index more similar to Case-Shiller's repeat sales method using Big Data, although the series only goes through 2010 at the time of this publication.

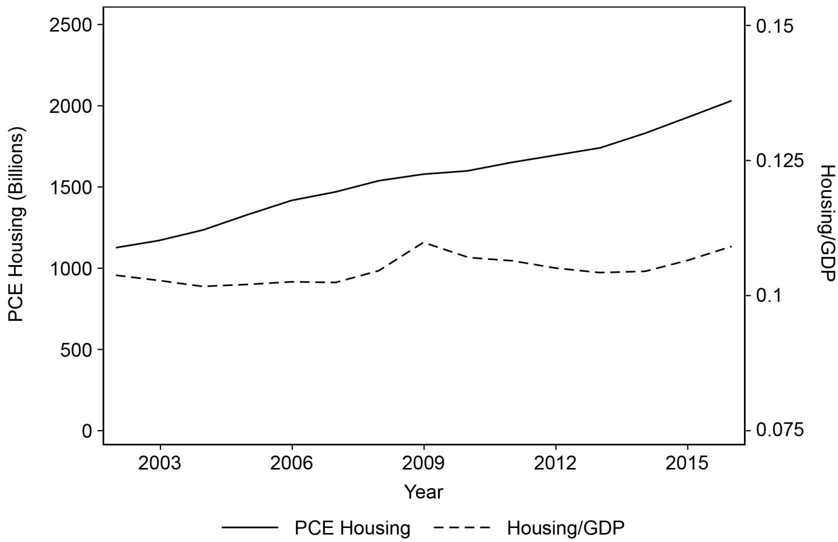

Fig. 12.1 Nominal PCE housing and PCE housing/GDP

Source: US Bureau of Economic Analysis, "Table 2.5.5: Personal Consumption Expenditures (PCE) by Function," bea.gov.

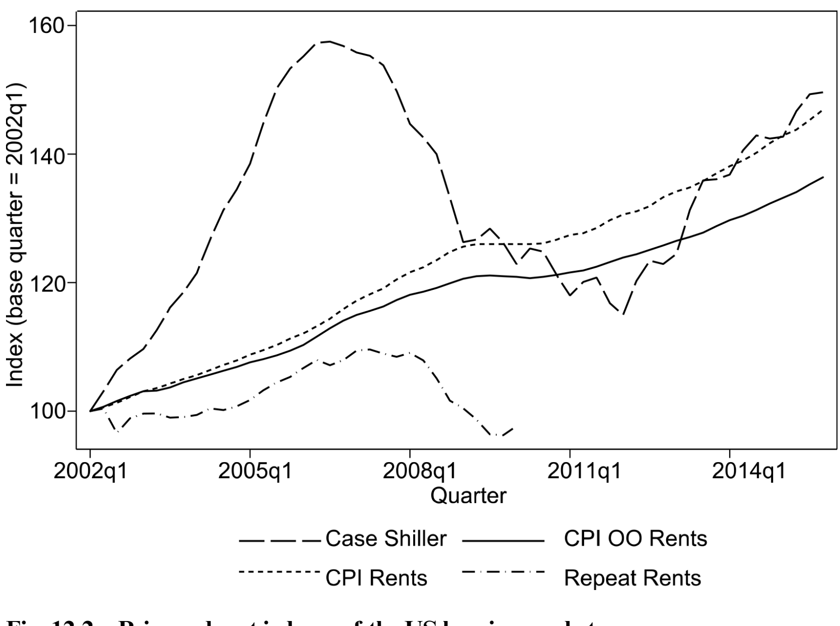

Fig. 12.2 Price and rent indexes of the US housing market

Sources: ACY; https://fred.stlouisfed.org/series/CSUSHPINSA; /CUUR0000SEHA; /CUSR0000SEHC.

matic in magnitude, was more consistent with the freefall in home prices as shown by the Case-Shiller National Home Price Index amid the (in)famous boom-bust-recovery in home prices over the broader period.

The divergence among these series stems from the underlying data and method.[4] Ambrose, Coulson, and Yoshida's (2015) finding, where market data and an alternative method paint a different picture of the rental market, motivates further research into other housing statistics and whether Big Data can find a similar pattern of divergence or whether this phenomenon is unique to the rental market they study.

The purpose of this paper is to explore the extent to which alternative data sources, namely Big Data from Zillow containing information on hundreds of millions of home transactions, can be used to construct an estimate of housing services. The data are suited to a user-cost approach, which we use to construct a time series and compare it to the BEA's current rental equivalence–based estimates since the early 2000s. The goal of this paper is not to construct an official account or argue for a particular method; rather, we investigate the implications of a new Big Data source and compare the results of associated methods to current nominal estimates.[5]

This paper also contributes to literature on user cost methods that are both well suited to Big Data sources and commonly used in academic literatures beyond national accounts. This is particularly true in cases where rental market data are inadequate (as in many countries).[6] For example, Himmelberg, Mayer, and Sinai (2005) employ a user cost approach to assess price fundamentals of the housing market, while others have used housing user costs in a number of applications from evaluating tax policy to interest deductions (e.g., Albouy and Hanson 2014; Poterba 1992; Poterba,

4. Critiques of the BLS's rental series, which fall outside the scope of our paper, are the subject of numerous papers, including Ambrose, Coulson, and Yoshida (2015). This topic is covered in an earlier review of this literature by Lebow and Rudd (2003). Ambrose, Coulson, and Yoshida (2015) argue that the CPI method and underlying data sources understated the extent to which rental market prices fell during the housing bust. See also Gordon and vanGoethem (2007), McCarthy and Peach (2010), and Ozimek (2014) for related critiques.

5. Constructing user cost estimates is also a prerequisite for a statistical agency to consider constructing a hybrid series that blends rental equivalence and user cost estimates like the opportunity cost approach proposed by Diewert (2009), as part of a comprehensive look at competing methods from the literature. A nominal series is also a necessary first step to take prior to constructing a real series based on these data, which we leave for future research.

6. A number of European and African countries have employed a user cost approach (or a variant thereof) for measuring housing services, often as a result of data limitations of thin unsubsidized rental markets (Katz 2009). A (nonexhaustive) list of such countries includes: Botswana, Central African Republic, Croatia, Estonia, Ghana, Hungary, Latvia, Lithuania, Malta, Montenegro, Nigeria, Poland, São Tomé, Serbia, Slovak Republic, Slovenia, Tunisia, Uganda, Zambia, Zimbabwe. According to Eurostat in 2016, nearly 70 percent of the population in EU28 countries own their own homes, with a sizable fraction of households living in subsidized or rent-free housing (e.g., over 80 percent in Lithuania, Malta, Bulgaria, and Croatia), limiting the representativeness of market rents in many countries (Komolafe 2018).

Weil, and Shiller 1991).[7] We provide a transparent method for constructing a nominal user cost-based series that can be built from the bottom up with similar microdata (e.g., data from vendors like CoreLogic) and could be replicated for a variety of uses in the literature.

12.2 Rental Equivalence versus User Cost

12.2.1 Background

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 mentioned above are those recommended by the 2008 SNA statistical framework: rental-equivalence and user cost.[8] The former estimates what market rent would be for a given owner-occupied home if it were rented, while the latter focuses on the cost to the homeowner.[9] Conceptually, absent transaction costs and other market frictions, 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).[10] For example, if

7. Poterba and Sinai (2008) note: "the neoclassical investment model, which focuses on the user cost of capital, is a standard tool for studying housing demand and for analyzing the equilibrium value of the imputed rental income accruing to homeowners under various tax regimes" (p. 86).

8. Specifically, the SNA states that, "This approach can take either a user-cost formulation that attempts to measure the changes in the cost to owner-occupiers of using the dwelling, or a rental-equivalence formulation based on how much owner-occupiers would have to pay to rent their dwellings. The latter method is more generally adopted for CPIs" (United Nations et al. 2010, §15.141). However, some countries have adopted a variant of the user cost approach for their CPI measurement, including Canada, Estonia, Iceland, Slovakia, and Sweden (Baldwin, Nakamura, and Prud'homme 2009).

9. 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, 65).

10. 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.

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.[11]

12.2.2 Current Approach of the BEA

The BEA's current approach, based on a rental-equivalence method, 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. Specifically, the BEA's current method uses the Residential Finance Survey (RFS, Census Bureau) to benchmark rent-to-value ratios for different value classes of properties, which are then used to impute average contract rent for owner-occupied properties across similar dimensions. This weighted rental imputation constitutes what is often referred to as "space rent," which is then multiplied by corresponding aggregate housing unit counts to obtain the aggregate estimate of the total imputed rent of owner-occupied housing.[12] During benchmark years, BEA used Decennial Census for quantity counts, while in nonbenchmark years either American Housing Survey (available biannually) or Current Population Survey data from the Census Bureau were used. The BEA last benchmarked the rent-to-value ratios used to derive space rent using the 2001 RFS, the last time the requisite data from this survey were available. Since then, the BEA has made quality adjustments and price adjustments, with the latter based on data from the BLS's CPI Owners' Equivalent Rent Index (which also relies on a rental-equivalence method).[13] This method is generally regarded as the

See also McFadyen and Hobart (1978) for an instructive cross-walk from Jorgenson (1967) to a user cost for housing.

11. Of course, this abstracts from risk, market imperfections, and transaction costs, which are particularly significant in housing (Bian, Waller, and 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 microdata.

12. For a more detailed discussion of the BEA's current method, refer to Mayerhauser and McBride (2007) and Katz (2017). To summarize, the 2001 method assumed OOH homes with comparable values as tenant-occupied homes also have comparable rent-to-value ratios, so the method takes weighted-average rent-to-value ratios by value class for tenant units from the RFS and applies the mid-point market value to owner-occupied units within the corresponding value classes reported in the American Housing Survey. This imputed total rental value is then weighted by the number of owner-occupied units reported in the American Housing Survey in each class to calculate an average annual rental value (AARV), which is then used to generate a total value of aggregate OONFP housing services by multiplying AARV by the number of owner-occupied housing units reported in the decennial Census.

13. BEA's weighting adjustment based on rent-to-value introduces a measure of home value into the imputation of owner-occupied space rent, as does the housing quality adjustment used since 2001. However, since 2001 the rent-to-value ratios have not changed due to the expiration of the underlying survey data, which is why the series has primarily moved with the CPI Owners' Equivalent Rent Index. Because the CPI for OOH is a constant-quality index, the purpose

preferred method for this imputation because most countries have relatively thick rental markets with substantial data on market rents. Indeed, more than one third of all housing units in the US are rented to tenants.

12.2.3 Methodology: A Comparison

The rental-equivalence approach, however, is not without its limitations due to the nature of the data. While a sizable fraction of homes are tenant occupied, rental data are not necessarily representative of the entire housing stock. Specifically, the distribution of rental units is not the same as owner-occupied units (Glaeser and Gyourko 2009); the share of detached single-family residences (SFRs) is higher for owner-occupied units as is the share of higher-value homes, as the market for rental units thins out and quality and home value increase.[14] Coulson and Li (2013) review the voluminous literature regarding these differences and provide additional evidence of homeowners taking better care of (and investing more in) their homes, resulting in difficult to measure qualitative differences between owner-occupied and tenant-occupied homes.[15] Also, because surveys record a snapshot of the market, rent surveys may overrepresent renewal rent for existing tenants and underrepresent new leases—a problem that may be exacerbated by business cycle fluctuations (Ambrose, Coulson, and Yoshida 2015). Verbrugge (2008) argues that this may oversmooth the series as someone surveyed in December may have signed their lease earlier in the year (in, say, February), reflecting lagged market conditions in the rental market.[16] While subject to its own limitations (as we discuss below), the user cost approach relies on different data than the rental-equivalence approach, which has led researchers and some statistical agencies to explore it as an alternative for estimating housing services. This approach instead utilizes data on the cost to the user of owning a home (e.g., interest, taxes, maintenance/depreciation),

of the additional quality adjustment is "to account for changes in the real value of housing per unit," which is the percent change in the "real dollar stocks of owner-occupied structures, of additions and alterations, and of major replacements" using values from BEA's fixed assets accounts divided by the number of owner-occupied units" (Mayerhauser and McBride 2007).

14. For additional discussion of this point and an illustration of these differences using recent Census data, see Aten (2018).

15. Crone, Nakamura, and Voith (2000) cite a number of reasons that complicate the BLS's attempts to compensate for the differences in owner-occupied versus tenant units by oversampling rental units that have characteristics like rentals: "First, these units are often temporary rentals that drop out of the sample in a short time, so that reporting is spotty. Second, the market for these units is very thin, so that the observed rents may not be good proxies for the implicit value of the unit's service flow if it were an owner-occupied unit. Third, rental units are subject to double-sided moral hazard, which leads to long-term contracts and price regulation. Fourth, rental units are professionally managed while owner-occupied units are not."

16. In addition, because the BEA has used the CPI Owners' Equivalent Rent Index to make adjustments to space rent, this also introduces potential measurement issues associated with the CPI. See Lebow and Rudd (2003) for a review of the literature on mismeasurement in the CPI, and Crone, Nakamura, and Voith (2000, 2009) for more on mismeasurement of CPI rents in particular.

which varies directly with the price of a home, rather than rents of different, possibly unrepresentative tenant-occupied homes.[17] Detailed microdata on home sales and corresponding home characteristics are primarily recorded by local municipalities; and because reporting often differs by locale, this has previously made a national effort to collect these 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 collected, compiled, and organized a massive database of public data from local tax assessors' offices across the US 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 it feasible to implement a user cost approach based on fine-level price and home characteristic data to compare with current methods.

One benefit of the approach we are assessing is that it relies on directly observable data that cover a significant share of the housing market. While rents are not directly observable for owner-occupied homes, transaction prices, the backbone of the user cost method,[18] are readily available for virtually all strata of the housing market, both tenant-occupied and owner-occupied homes. As a result, given the differences in rental and owner-occupied housing units documented in the literature discussed above, the user cost method does not suffer from the same selection issues as rent-based approaches. Indeed, when rental markets are thin, the SNA recommends "other means of estimating the value of housing services," (United Nations et al. 2010, 109) like a user cost approach that does not rely on rent data.

There is, however, a sizable literature noting potential weaknesses of a user cost approach or conceptual departures that fundamentally differ from rental equivalence. For example, Gillingham (1983), Verbrugge (2008) and Diewert, Nakamura, and Nakamura (2009) and others have noted that the user cost approach often has greater volatility, sensitivity to interest rates, and introduces deeper conceptual issues with the role of asset prices in this estimate with ex ante and ex post measurement. For instance, the degree of volatility of Verbrugge's (2008) user cost estimates largely hinged on how he estimated expected (*ex ante*) appreciation/depreciation, which can vary

17. For an instructive review of this voluminous literature and novel examples of developing user cost estimates, see Diewert (2003, 2008), Katz (2009), Verbrugge (2008), Davis, Lehnert, and Martin (2008), Haffner and Heylen (2011), Hill and Syed (2016), Aten (2018), and numerous other papers on this topic.

18. Despite transaction costs and substantial frictions in the housing market, a thick literature has documented that home prices respond relatively quickly to a host of different types of shocks to demand, whether they are very local, neighborhood level shocks (e.g., Anenberg and Kung 2014; Linden and Rockoff 2008; Wentland, Waller, and Brastow (2014) or aggregate-level or informational shocks (e.g., Bernstein, Gustafson, and Lewis 2019; Bui and Mayer 2003; Moulton and Wentland 2018).

substantially depending on the assumptions used to construct this component. This literature also voices disagreements on precisely what parameter values should be used in the computation, including which interest rate is most appropriate or whether to include expected appreciation/depreciation at all. Small changes to these parameter values can change the estimates substantially, as we document in more detail below in our discussion of figure 12.7 and the alternative user cost estimates we produce by varying these parameters. Finally, as a more general conceptual point, by tying estimates of housing services more closely to the asset value of a home and interest rates, the user cost approach begs the question: to what extent should a measure of housing services vary with interest rates and asset prices? We return to this point in the Discussion section below.

12.3 Data

The novelty of this paper primarily arises from usage of new data, specifically residential housing microdata from Zillow's ZTRAX dataset. It contains transaction data as well as a large set of individual property characteristics for sales recorded from local tax assessors' data.[19] The data coverage is generally representative of the United States' national housing market, comprising 374 million detailed records of transactions across more than 2,750 counties.[20] This includes information regarding each home's sale price, sale date, mortgage information, foreclosure status, and other information commonly disclosed by a local tax assessor's office. We link each transaction to each home's property characteristics into a single dataset. The assessment data include 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.[21] We received all these data in a somewhat raw form, requiring significant cleaning for research purposes.

19. 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.

20. 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 a substantial share of missing data in our current sample and is accordingly omitted. 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.

21. Zillow's ZTRAX dataset contains separate transaction and assessment files by state—that is, 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).

We carefully scrutinized missing data and extreme values as part of our initial culling of outliers and general cleaning. The initial dataset from Zillow contains sales of empty plots of land, some commercial property transactions, agricultural sales, and other 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, row-homes, apartments, condos, and properties that are most closely associated with the current estimates. While we estimate rural properties separately (properties with 1 to 100 acres), we drop homes that have greater than 100 acres (limiting the influence of large farms) and winsorize homes that are in 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). When we construct our final user cost estimates we also drop homes that sold for less than \$30,000 for SFRs (\$15,000 for non-SFRs), homes in the top percentile of predicted price, or that had a price 10 times higher than the county median.[22] We cull homes that report a negative age (i.e., sale year < year built). While the Zillow dataset 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.[23] We limit the results to the years from 2002 through 2015, when the data are most complete for the vast majority of the states in our sample.

To assess the quality of the final sample, we compared our cleaned Zillow sample to the ACS to ensure that these administrative data aligned with carefully collected (albeit more limited) survey data provided by the Census Bureau. Generally, there is only a limited set of home characteristics found 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 datasets had median and mean values that fell within a few percentage points of one another.

22. To limit the influence of outliers or measurement error on model coefficients in our regressions, we drop homes that sold for less than \$1,000 and extreme outliers at the top end (10 times the county median), and then the tails of the distribution for sale price at the 2.5 and 97.5 percentiles within each county within each quarter. This is a more restrictive culling at the regression stage because the main objective of the regressions is to obtain coefficients that provide the most reasonable price predictions, whereas when we construct the final user cost estimates we aim to exploit a somewhat less restrictive sample to maintain better representativeness (while still drawing a line to cull suspicious outliers).

23. 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.

12.4 Methodology—An Idiosyncratic User Cost Approach

12.4.1 Overview

Generally, our approach using the Zillow microdata is motivated by constructing estimates from the bottom up, as we estimate a user cost for *each individual property* in our dataset *for each quarter* 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 dataset based on the formula

$$U_{it} = P_{it}(r_{it}^{rf} + \delta_i + \tau_{it} + \gamma_i - \varphi^m(r_{it}^m) - \varphi^\tau(\tau_{it}) - \mathrm{E}[\pi_i]),$$

where for a given property (i) in quarter (t), P is the price of an individual home, r^{rf} is the owner's discount rate or financial opportunity cost for a long-term asset like a home (we use the nominal interest rate on a 10-year Treasury note for an appropriate risk-free rate in quarter t),[24] δ is a constant 3.5 percent representing depreciation and maintenance costs,[25] τ is the individual property's effective tax rate, and γ is a constant 2 percent risk premium associated with owning relative to renting.[26]

The latter three terms consist of potential offsetting benefits to homeownership, which are subtracted from the preceding costs such that user cost

24. While the dataset 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 a homeowner purchased a home when rates were at 4 percent, but rates have since risen to 7 percent, 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.

25. A depreciation rate of 1.5 percent 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. 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 is a simplification.

26. This risk premium was used by Himmelberg, Mayer, and Sinai (2005) "to compensate homeowners for the higher risk of owning versus renting" (p. 75). While a risk premium was used as early as Poterba (1992), the constant of 2 percent was used by Flavin and Yamashita (2002) and Poterba and Sinai (2008). The latter study argues that this accounts for the fact that, "homeowners bear both asset-class risk and idiosyncratic, house-specific risk" (p. 86). Himmelberg, Mayer, and Sinai (2005) also use a 2 percent constant but point out that a more sophisticated model would allow this premium to vary over time as the risk of owning relative to renting changes over time. This risk, however, is separate from rental risk which, as Sinai and Souleles (2005) point out, is hedged with homeownership. Sinai and Souleles (2005) find that this rental risk is directly capitalized in home prices.

represents the net cost to the homeowner. Mortgage interest and property taxes are tax deductible in the US (to a point), regardless of occupancy status. Himmelberg, Mayer, and Sinai (2005) use a constant average marginal tax rate (MTR) for all homes, which they multiply by the average 30-year mortgage rate (r_{it}^{m}) in period t. However, their approach assumes (1) all homeowners itemize their tax returns, (2) the interest is on the entire principal of the home, and (3) there is little variation in income across regions of the United States. Instead, we construct a multiplier, φ, to allow variation in our approximation of the average benefit to mortgage interest and property taxes, using the ACS to determine the average household income for homeowners by home type (SFR versus non-SFR) and home size (number of bedrooms) by each Census Division.[27] Based on average household income, we assign an MTR and a probability that the homeowner itemizes based on the percent of people who itemize in their income stratum. This allows a five-bedroom home in a high-income region like New England or the Pacific region to have a proportionately higher tax benefit than a two-bedroom home in a poorer region.

The φ^{τ} multiplier consists of this MTR and itemization probability, while the φ^{m} multiplier incorporates an additional product of the average loan-to-value (LTV) ratio by Census region to account for the fact that a homeowner can only write off interest on an outstanding loan amount (i.e., if the LTV ratio was zero for all homes, there would be no realized mortgage interest tax benefit).[28] Finally, $E[\pi]$ is expected appreciation (revaluation) for a given year. We set this to 2 percent, which assumes homeowners have a very long-term view of home prices appreciating approximately the same as overall inflation in the economy.[29] While approximately 2 percent is common

27. We use data from the IRS's Statistics of Income (Table 1.2) and the following adjusted gross income strata: under \$30,000; \$30,000–\$49,999; \$50,000–\$99,999; \$100,000–\$499,999; above \$500,000 (where the percent who itemize are: 7, 21, 44, 80, and 93, respectively).

28. We use data from the Federal Reserve's Survey of Consumer Finance (SCF), which contains information on the average mean value of mortgages and home equity/home value from 2002 to 2015 for each Census region.

29. 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 percent inflation over the long run (e.g., see their policy statements on their website regarding 2 percent: 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 of inflation may not be unreasonable. For robustness, we consider alternative expectations of price later in the paper.

to the user cost literature (e.g., Himmelberg, Mayer, and Sinai 2005; Poterba and Sinai 2008), we vary this 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 in our calculation that follow the literature, we return to a discussion of these simplifications and ways to possibly create a more precise estimate in section 12.6.

12.4.2 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 the Zillow microdata. Because we have fine, transaction-level price data, we are able to first use actual market prices for P (when available and when it does not fail the outlier criteria discussed above). 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).[30] Turnover varies considerably by state and locality; approximately one third of properties in our dataset sold at least once within the window we study (2002–2015). For the value of the home in the following quarter, we posit that the price is simply the transacted price adjusted by the predicted price's appreciation/depreciation (discussed below). We use the same logic for the quarters following that sale until there is a new sale of that property.[31] Broadly, using more direct price data conforms most closely to the principles of valuation laid out by the SNA, where market prices are "the basic reference for valuation in the SNA" (United Nations et al. 2010, 22), and thus much of our aggregate calculation flows directly from millions of observed market prices underlying the housing stock.

As a more general principle of valuation, 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

30. The ACS has home price data with reasonably good coverage; however, these data come 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. Ideally, with linked microdata, eventually we would like to explore the differences between these datasets for use in the national accounts.

31. This method would likely be altered if it were implemented in national accounts over a longer time series because 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 this simplified approach.

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)" (United Nations et al. 2010, 22). Hence, for homes that did not sell during our sample period, we predict their prices based on transactions of similar homes that sold in each quarter using a hedonic model.[32] Conceptually, most of a home's value can be explained by its physical characteristics, location, and time (Rosen 1974); hence, our hedonic model uses sale prices of similar homes along these dimensions to estimate an imputed market valuation for each home in our dataset.[33] 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 a predicted sale price, $\hat{P}$, based on a hedonic model for each state by quarter separately for home i in quarter t in location j:

$$\begin{aligned} Sale\ Price_{ijt} = \alpha + \sum \beta X_{it} + \gamma LOCATION_{jt} + \sum \delta sqft_{it} \\ * LOCATION_{jt} + \sum \varphi acreage_{it} * LOCATION_{jt} + \varepsilon_{it}, \end{aligned}$$

where X is a set of physical characteristics (bedrooms, bathrooms, age of the structure, living area measured by square feet, lot size measured by the natural log of 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 the natural log of acreage, respectively.[34] For practicality in estimation, we initially use Census tract fixed effects, although we obtain similar estimates using finer-level geo-

32. Within-quarter hedonic regressions allow for all coefficients in the model to change across quarters, accounting for changing tastes and preferences for location or for each housing attribute in the model.

33. Aside from the voluminous literature in real estate, hedonic valuation is not uncommon in the national accounts and price index literatures. For example, see Pakes (2003) or Benkard and Bajari (2005) for applications with personal computers.

34. While the Zillow ZTRAX data contain a lot more information about individual properties that would help with valuation, we chose the variables with extensive coverage across all states in the dataset. 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, bathrooms) was missing, we imputed the number based on the size of the home, based on the rest of our sample. For SFRs with missing bedrooms, we replaced 1, 2, 3, 4, and 5 bedrooms for the following square footage buckets: < 500, 500–999, 1000–1999, 2000–3000, and 3000+. For non-SFRs and urban properties with missing bedrooms, we replaced 1, 2, and 3 bedrooms for the following square footage buckets: < 600, 600–999, 1000+. For all units, we replaced missing bathrooms with a full bathroom per each 1,000 square feet up to 3 bathrooms. Overall, the results are not sensitive to dropping these observations with missing characteristics entirely, but our coverage in some states/counties where this is more systematic would raise issues of representativeness if we drop them.

graphic fixed effects like Census block groups or blocks.[35] To avoid making predictions with thin cells, we specify that a given tract have at least 10 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) 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.[36] For non-SFRs, which we estimate separately from detached SFRs, we omit acreage and other SFR-specific characteristics from the hedonic model.[37] In addition, we estimate price predictions for urban single-family homes with very small lots (less than one tenth of an acre) with non-SFRs; and, we separately estimate rural homes, which we define as having between 1 and 100 acres. In both cases, we do this only to generate better price predictions for these properties, as we eventually aggregate all SFRs together by number of bedrooms by Census Division, which we discuss more below in section 12.4.4.

12.4.3 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 percent, whereas Hawaii and Alabama have average rates of 0.32 percent and 0.48 percent, respectively.[38] Even within states there is con-

35. 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.

36. This approach is used commonly in the hedonic valuation literature for housing and land. See, for example, Kuminoff and Pope (2013). For some of the larger states like California, this approach yields too many interaction terms that bump up against the limit for number of variables that can be used in a single regression for many statistical software packages, which required us to run substate samples (Northern CA versus Southern CA, for example). This allows noninteracted coefficients to vary within states.

37. 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 for our models using census tract fixed effects, 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:

$$\textit{Average Percent Error } (APE) = [(\textit{Predicted Price} - \textit{Actual Price})/\textit{Actual Price}] * 100$$

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. Overall, APE fell with ± 5 percent for the vast majority of quarters, with only a handful of quarters in the ± 5–10 range.

38. 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 US: https://www.usatoday.com/story/money/personalfinance/2017/04/16/comparing-average-property-taxes-all-50-states-and-dc/100314754/.

siderable 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 are collected primarily from local tax assessor office databases, the coverage of property taxes is quite good. We use individual tax data to 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.[39]

We made this choice for a couple reasons. First, regarding the denominator, the assessment value is often much lower than the market value, so applying the rate based on the assessed value to the market value of P in the user cost calculation 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, as we describe in section 12.4.1 above, we allow an estimated average benefit varying by home type, region, and home size, as household income (and therefore marginal tax rate and likelihood of itemization) varies tremendously across the US, which we capture to some extent with this approach.[40]

12.4.4 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 corresponding quantities of the housing stock by location/region, type of home (SFR versus non-SFR), and number of bedrooms. We use the weighted unit counts of the housing stock from the ACS for each year of our sample, which provides a yearly count

39. We currently have one year of tax amount data from Zillow but updating these 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), we winsorized 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.

40. Our ambition is to eventually use linked administrative data to back out a more precise, idiosyncratic estimate of the tax benefits to owning a home. In addition, 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 versus renters. We return to this point in section 12.6.

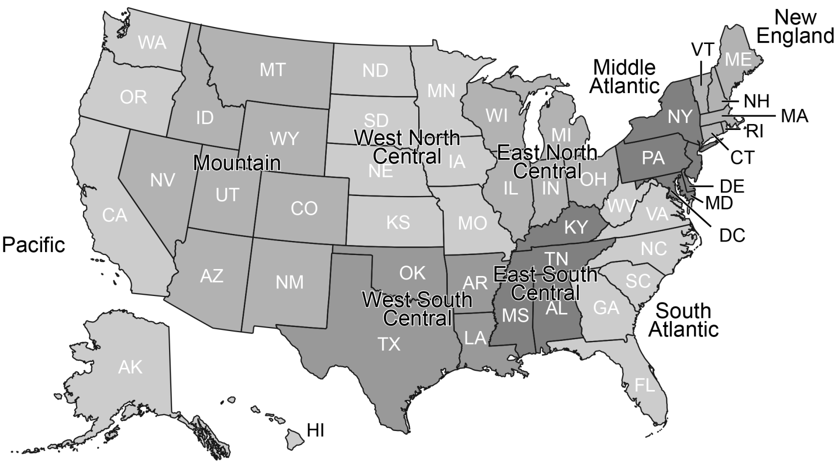

Fig. 12.3 Census divisions
Source: https://www.census.gov/geo/reference/webatlas/divisions.html.

of the aggregate number of residential housing units by Census Division, depicted in figure 12.3. Because the BEA's current method treats vacant homes differently than tenant- or owner-occupied homes, we omit these from our aggregation, reconstructing estimates according to the BEA's current method and using the same quantity of homes from the ACS such that the difference between the two series is independent of quantities used (labeling this "Quantity Adjusted PCE Housing" to reflect this difference from the official series). For illustrative purposes, refer to table 12.1 below, where we show the calculation of our national estimate for Q4 of 2015. For each Census Division or region of the US, we multiply the average user cost for each type of home (SFR versus non-SFR) for each bedroom category.[41]

This method of aggregation assumes that the nonmissing data are reasonably representative of the missing data. For example, Indiana's sale prices are missing from the ZTRAX dataset, as it is among the nondisclosure states that do 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

41. 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 12.1. 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 than averages by physical characteristics.

Table 12.1 **User cost aggregation—Example quarter**

Total User Cost Calculation (Default Specification) for 2015 Quarter 4							
		SFR			Non-SFR		
Division	Bedrooms	Avg. user cost	Q	P ∗ Q (billions)	Avg. user cost	Q	P ∗ Q (billions)
1	0 or 1	14,565	79,713	1	30,133	761,608	23
	2	16,669	491,998	8	32,612	1,006,532	33
	3	20,603	1,603,041	33	23,622	533,706	13
	4	29,814	838,816	25			
	5+	41,131	204,366	8			
2	0 or 1	11,749	142,736	2	17,386	2,599,754	45
	2	10,635	1,027,587	11	17,580	2,624,879	46
	3	15,848	3,614,253	57	28,243	2,174,197	61
	4	24,420	2,234,490	55			
	5+	38,896	579,746	23			
3	0 or 1	7,239	220,172	2	7,245	1,751,404	13
	2	6,887	1,946,805	13	10,839	2,480,621	27
	3	10,251	6,553,425	67	9,393	937,491	9
	4	16,547	2,979,940	49			
	5+	24,727	668,551	17			
4	0 or 1	9,682	143,659	1	10,554	769,223	8
	2	9,749	1,051,504	10	12,062	952,057	11
	3	12,754	2,678,916	34	14,576	351,747	5
	4	16,979	1,522,571	26			
	5+	20,061	470,828	9			
5	0 or 1	9,631	197,364	2	7,303	2,037,536	15
	2	8,813	1,922,406	17	9,670	3,258,601	31
	3	11,897	7,526,960	90	15,778	1,869,658	29
	4	20,120	3,739,500	75			
	5+	29,923	1,091,405	33			
6	0 or 1	7,300	94,430	1	6,881	443,190	3
	2	6,123	739,063	5	7,384	691,375	5
	3	7,685	2,895,377	22	10,281	246,935	3
	4	12,386	1,059,573	13			
	5+	18,240	243,589	4			
7	0 or 1	11,302	212,743	2	4,329	1,461,312	6
	2	5,616	1,315,520	7	7,323	1,449,698	11
	3	8,589	5,129,666	44	8,339	475,987	4
	4	13,350	2,283,730	30			
	5+	18,331	435,305	8			
8	0 or 1	15,553	127,213	2	10,601	779,253	8
	2	14,278	759,204	11	10,698	1,068,443	11
	3	14,736	2,597,256	38	13,958	428,687	6
	4	21,199	1,580,893	34			
	5+	28,338	623,233	18			

Table 12.1 (cont.)

Total User Cost Calculation (Default Specification) for 2015 Quarter 4

		SFR			Non-SFR		
Division	Bedrooms	Avg. user cost	Q	P * Q (billions)	Avg. user cost	Q	P * Q (billions)
9	0 or 1	17,924	314,491	6	23,344	2,515,810	59
	2	23,840	1,575,736	38	31,575	2,884,457	91
	3	25,817	5,077,243	131	36,109	1,132,319	41
	4	34,382	2,928,474	101			
	5+	43,812	755,755	33			
			Subtotal (SFR)	1,216		*Subtotal (non-SFR)*	618
							Total user cost: 1,216 + 618 = 1,834

Wisconsin) reflect the Indiana market.[42] Missing data itself is not a prohibitive limitation for constructing national accounts (statistical agencies always have limited data); the issue is rather the representativeness of the data we do have. While many of these states are reasonably represented by their neighboring states' housing markets (e.g., Indiana), one exception might be Texas (the largest state for which we have missing price data).[43]

12.4.5 Varying Ex Ante Expected Price Appreciation/Depreciation

Finally, for robustness, we vary the E[π] term of ex ante expected price appreciation. Our default specification assumes a very long-run view of home price inflation of a constant 2 percent 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. Rather than assuming that homeowners take a *constant long-run, national* view of price expectations, we can instead consider that that they take a *variable short-run*, *local* view of price expectations. Thus, our alternative specification supposes that homeowners expect ex ante price appreciation to be their local (county-level) average yearly price inflation from the prior two years (quarter $t - 8$ to $t - 5$ and $t - 4$ to $t - 1$). This is

42. Recall that one of the limitations of this dataset is that there are 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.

43. If this method (or similar) were to be adopted by the BEA or others, supplemental data would be required to verify these assumptions or to reweight 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 developing alternative housing estimates. 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 and is a significantly smaller sample.

calculated by taking the average percent change of the median predicted price by county over the previous eight quarters from our hedonic model estimates discussed above.[44]

In this alternative specification, we also limit appreciation (depreciation) expectations to 5 percent (−5 percent) to avoid substantially negative user costs and excessive volatility based on expectations. One can think of this specification as price appreciation being expected to cover or offset (approximately) the maintenance, physical deterioration of the property, and owner risk premium (which itself may fluctuate in proportion to price expectations). 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.[45]

12.5 Results

Our full set of results for all years and quarters in our sample appears in table 12.2, which shows both the total and average user cost estimates of housing services as well as the corresponding estimates by housing type (SFR versus non-SFR) by quarter. A visual of these data is shown in figures 12.4 and 12.5. Specifically, figure 12.4 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 total user costs of detached SFRs are consistently higher than non-SFRs, as one would expect.

The key figure of the paper is figure 12.5, where we compare our average yearly user cost measure of housing services with the BEA's yearly estimate of housing services from PCE, using the ACS to adjust the quantity of the stock of housing in each year to be equal across both series. Note that we compare the full estimates of aggregate housing services because we are

44. 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 nominal results.

45. 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 into 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 and is not uncommon in the academic literature.

Table 12.2 Housing user costs by quarter from 2002 through 2015

	Full Sample		SFR		Non-SFR	
	Total user cost ($B)	Avg. user cost	Total user cost ($B)	Avg. user cost	Total user cost ($B)	Avg. user cost
2002q1	1,489	14,876	1,051	15,773	438	13,088
2002q2	1,577	15,711	1,114	16,653	463	13,829
2002q3	1,498	14,884	1,051	15,665	447	13,322
2002q4	1,461	14,476	1,022	15,172	439	13,081
2003q1	1,481	14,638	1,032	15,261	450	13,383
2003q2	1,505	14,813	1,052	15,486	453	13,455
2003q3	1,677	16,445	1,169	17,119	508	15,080
2003q4	1,712	16,727	1,183	17,247	528	15,669
2004q1	1,711	16,657	1,184	17,177	526	15,595
2004q2	1,957	19,001	1,354	19,587	603	17,804
2004q3	1,947	18,848	1,340	19,345	606	17,833
2004q4	1,916	18,501	1,305	18,787	611	17,919
2005q1	1,961	18,885	1,322	18,980	640	18,692
2005q2	2,048	19,698	1,382	19,799	666	19,492
2005q3	2,139	20,545	1,446	20,655	693	20,319
2005q4	2,217	21,272	1,492	21,271	725	21,273
2006q1	2,280	21,848	1,533	21,800	747	21,948
2006q2	2,489	23,799	1,683	23,851	807	23,692
2006q3	2,458	23,440	1,659	23,451	798	23,417
2006q4	2,381	22,654	1,596	22,498	784	22,979
2007q1	2,415	22,922	1,624	22,821	791	23,131
2007q2	2,513	23,816	1,693	23,755	820	23,942
2007q3	2,460	23,263	1,662	23,277	797	23,234
2007q4	2,256	21,294	1,517	21,210	738	21,469
2008q1	2,051	19,326	1,378	19,235	673	19,517
2008q2	2,083	19,606	1,409	19,652	674	19,511
2008q3	2,027	19,051	1,374	19,153	653	18,841
2008q4	1,779	16,697	1,202	16,743	577	16,602
2009q1	1,591	14,912	1,078	15,006	513	14,719
2009q2	1,742	16,296	1,189	16,526	553	15,823
2009q3	1,786	16,666	1,221	16,942	565	16,100
2009q4	1,741	16,216	1,189	16,472	552	15,690
2010q1	1,771	16,461	1,209	16,720	562	15,931
2010q2	1,747	16,221	1,202	16,608	546	15,429
2010q3	1,572	14,578	1,083	14,956	490	13,808
2010q4	1,566	14,501	1,073	14,817	493	13,857
2011q1	1,651	15,278	1,132	15,631	519	14,558
2011q2	1,633	15,076	1,125	15,503	508	14,209
2011q3	1,461	13,455	1,006	13,832	455	12,691
2011q4	1,355	12,451	928	12,735	427	11,875
2012q1	1,344	12,318	922	12,627	422	11,693
2012q2	1,353	12,395	929	12,726	424	11,726
2012q3	1,332	12,189	916	12,542	416	11,477
2012q4	1,353	12,376	921	12,614	432	11,896
2013q1	1,407	12,853	956	13,086	451	12,385

(*continued*)

Table 12.2 (cont.)

	Full Sample		SFR		Non-SFR	
	Total user cost ($B)	Avg. user cost	Total user cost ($B)	Avg. user cost	Total user cost ($B)	Avg. user cost
2013q2	1,488	13,561	1,015	13,875	473	12,934
2013q3	1,687	15,339	1,146	15,632	541	14,755
2013q4	1,697	15,394	1,150	15,662	547	14,858
2014q1	1,729	15,656	1,162	15,813	567	15,343
2014q2	1,776	16,041	1,196	16,247	580	15,633
2014q3	1,773	15,979	1,190	16,149	583	15,643
2014q4	1,723	15,493	1,150	15,575	573	15,331
2015q1	1,675	15,030	1,111	15,027	564	15,036
2015q2	1,793	16,066	1,195	16,143	598	15,914
2015q3	1,832	16,390	1,220	16,456	612	16,259
2015q4	1,834	16,380	1,216	16,369	618	16,401

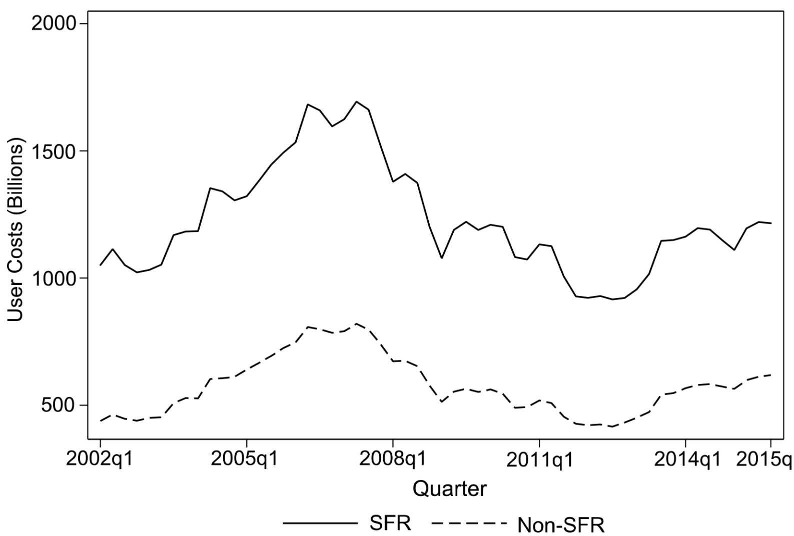

Fig. 12.4 Total quarterly user costs by SFR/non-SFR (default specification)

estimating the user cost for all residential homes in our sample, applying the same method to all homes whether they are owner-occupied or not.[46]

46. 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 include vacant 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 12.5. Finally, some states and municipalities had limited data in the early few years of this sample, which may not

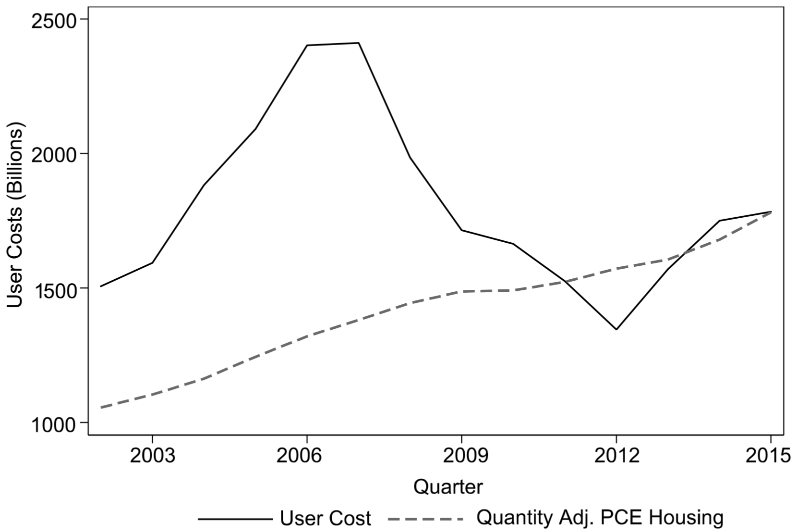

Fig. 12.5 Total yearly user cost (default) compared to PCE housing estimates

Our default aggregate measure of housing was initially much higher than the BEA's estimate in 2002, but this gap widened precisely when home prices throughout much of the US appreciated considerably during the run up to the financial crisis and Great Recession.

The more pronounced fluctuations in the path of the user cost-based estimate from 2002 through 2010, during the infamous bubble-bust years, bear a striking resemblance to national house price indices like Case-Shiller's, rising about $1 trillion from 2002 to the peak in 2006, 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 CPI measures using a user cost versus rental-equivalence approach. Indeed, this result is consistent with Ambrose et al. (2015) in that a notable drop occurs in the latter part of the decade. 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, consistent with the time series dynamics of the price indices in the figure we discussed in the introduction (figure 12.2).

One alternative specification of the user cost method, factoring in recent (eight quarters) and 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 12.6 shows a user cost closer to the rental-

have been random, as richer counties may have digitized these records earlier and more consistently, possibly explaining some of this difference in the first couple of years.

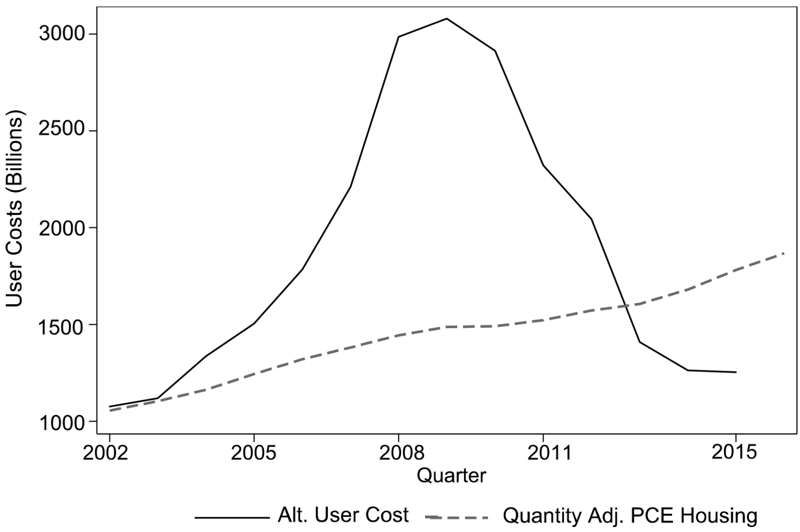

Fig. 12.6 Total alternative user cost compared to PCE housing

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. However, given that this specification is more aggressive in its price expectations assumption, this result should be interpreted with care, as it incorporates greater volatility into the series based on a very simple model of price expectations. 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).

For robustness, we vary some of the assumptions underlying the user cost formula, which we show in figure 12.7. First, rather than incorporate a fixed homeownership risk premium of 2 percent, 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 percent 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. Third, if we omit

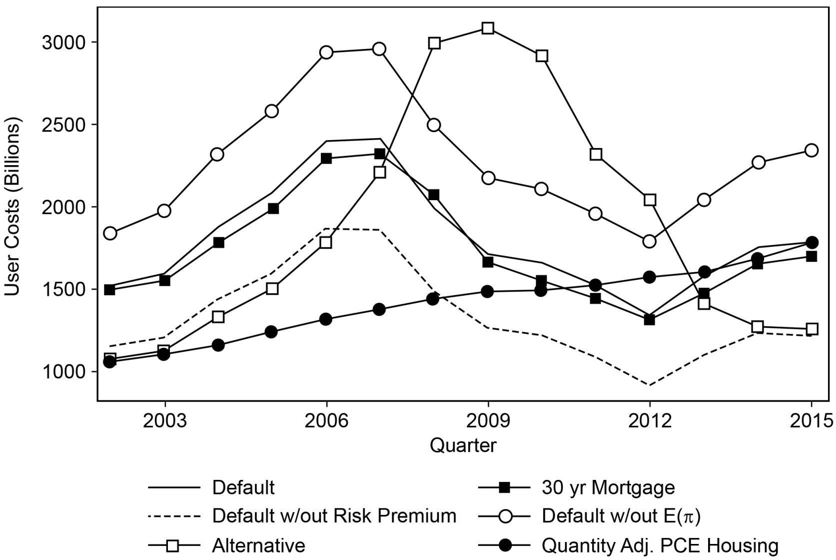

Fig. 12.7 Comparing different user cost methods with PCE

the risk premium entirely, we see an analogous downward shift in the 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 (where, even if one disagrees with the precise constant we use, if asset risk changes directly with price expectations, the choice of the constant becomes less relevant if they offset).

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 suffer to a greater extent from a thin cell problem. As an example, figures 12.8 and 12.9 show average user cost by region (Census Division) for SFRs and non-SFRs respectively, although the data easily allow us to provide measures at the county or tract level (except, of course, for states with missing price data). The estimates produce the expected result—that the Pacific region and New England have the highest average user costs of housing, 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 12.10 depicts average user cost per residential unit for three different specifications and the cor-

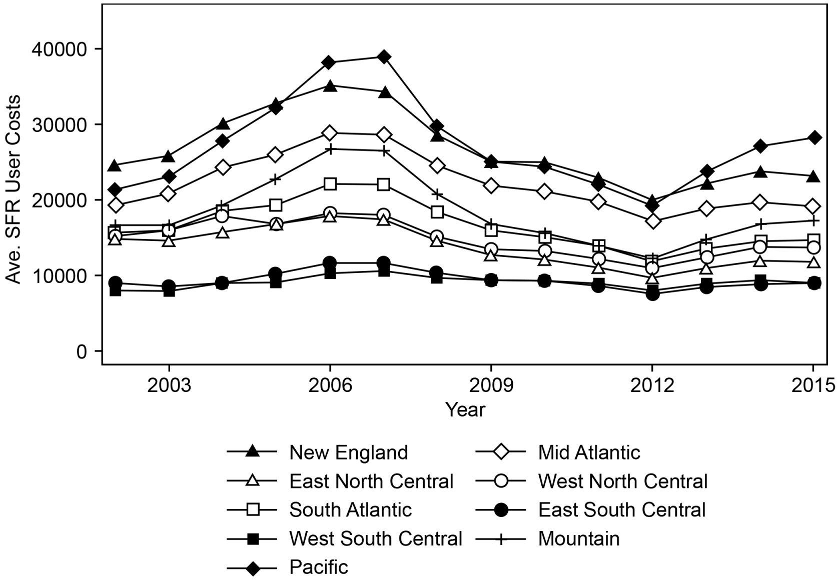

Fig. 12.8 Average yearly user costs for SFR by Census division

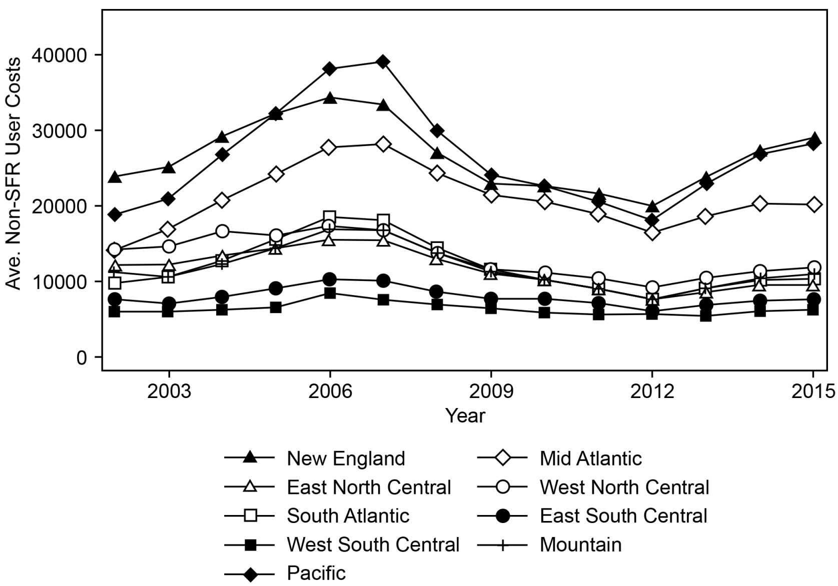

Fig. 12.9 Average yearly user costs for non-SFR by Census division

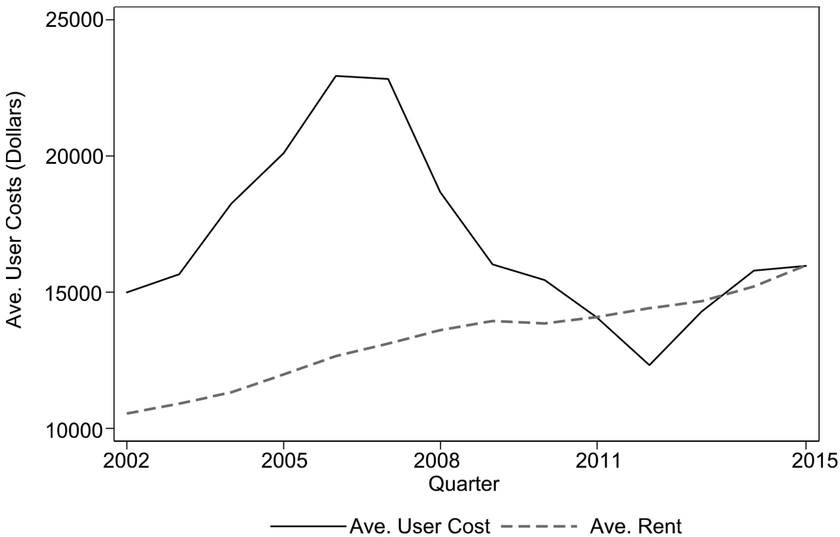

Fig. 12.10 Average user costs and PCE average rent

responding per unit space rent estimate (BEA). While the shape is identical to figure 12.7, the magnitudes may be helpful for assessing reasonability of the estimates, with the nominal average user cost and space rent both near $15,000 per year in the final couple of years in our sample period.

12.6 Discussion

Though for reasons discussed below the BEA is not adopting the user cost method, it is worth discussing a few caveats when comparing it to the current method and potential avenues for future research. 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 only for the most recent years we estimate, but the series behaves very differently over the bubble-bust period of the 2000s. Indeed, 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 the theoretically predicted 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 to policy makers?[47] These are foundational conceptual questions in the economic measurement literature (e.g., Alchian and Klein 1973; Gilchrist and Leahy 2002; and Goodhart 2001), which this paper does not attempt to settle.

We made a number of methodological simplifications and assumptions which, if adopted by a national statistical office at some point, would need to be explored further because some (likely small portion) of the differences may be attributable to these choices. Additional precision gained from refining these estimates may, at least in part, help bridge the aforementioned gap between user cost and rental-equivalence estimates (particularly in the post-bubble/bust years when the gap was not as large). For example, the mortgage interest and property tax deductions are highly idiosyncratic depending on a number of factors such as income, where the probability of itemization and marginal tax rates could be 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. Finally, linked administrative data may also help bridge 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 rental data and user cost estimates based on Zillow data, as this would show the most direct apples-to-apples comparison of the two methodological approaches. These linked data could also help us test or even construct better sample weights to ensure the composition of the sample accurately represents the characteristics of the entire stock of housing in the United States.

After considering a number of options, the BEA does not plan to adopt the user cost approach because it plans to modify its rent-based approach by incorporating new source data (Census ACS data) and updating its method to include a new owner-premium adjustment (see Aten 2018). The proposed modified rental-equivalence approach would be less volatile and more incremental compared to user cost-based estimates. Nevertheless, this research demonstrates the potential upside to incorporating new data and exploring new methods in the national accounts more generally, and housing in particular. Statistical agencies are continuously seeking ways to lower response

47. 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 policy makers? We leave this, however, to future research.

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. For example, survey respondents are asked to place a value on their own home. The kind of microdata used in this study could be used to update or even replace statistics that use these data (e.g., rent-to-value ratios or the housing stock quality measure used to adjust the BEA's current rental-equivalence method for owner-occupied housing). Linked Zillow-ACS data could provide an estimate for calculating an owner premium for owner-occupied housing, supplementing the (adapted) rental-equivalence method proposed by Aten (2018) by using market transaction values as opposed to survey-based values, which is currently being explored by BEA researchers. As another example, Big Data sources could also substantially improve precision for regional and type stratification, as linked data could provide additional details about individual homes (e.g., number of bathrooms, size of the home in square feet) that are not reported in a survey like the ACS, providing further potential for improving the economic measurement of housing services.

References

Albouy, D., and A. Hanson. 2014. "Are Houses Too Big or in the Wrong Place? Tax Benefits to Housing and Inefficiencies in Location and Consumption." *Tax Policy and the Economy* 28:63–96.

Alchian, A. A., and B. Klein. 1973. "On a Correct Measure of Inflation." *Journal of Money, Credit and Banking* 5 (1): 173–91.

Ambrose, B. W., N. E. Coulson, and J. Yoshida. 2015. "The Repeat Rent Index." *Review of Economics and Statistics* 97 (5): 939–50.

Anenberg, E., and E. Kung. 2014. "Estimates of the Size and Source of Price Declines due to Nearby Foreclosures." *American Economic Review* 104 (8): 2527–51.

Aten, Bettina. 2018. "Valuing Owner-Occupied Housing: An Empirical Exercise Using the American Community Survey (ACS) Housing Files." BEA Working Paper, Washington, DC. https://www.bea.gov/research/papers/2018/valuing-owner-occupied-housing-empirical-exercise-using-american-community.

Baldwin, A., A. O. Nakamura, and M. Prud'homme. 2009. "Different Concepts for Measuring Owner Occupied Housing Costs in a CPI: Statistics Canada's Analytical Series." In *Price and Productivity Measurement*, vol. 1, *Housing*, edited by W. E. Diewert, B. M. Balk, D. Fixler, K. J. Fox, and A. O. Nakamura, 151–60. Manchester, UK: Trafford Press.

Benkard, C. L., and P. Bajari. 2005. "Hedonic Price Indexes with Unobserved Product Characteristics, and Application to Personal Computers." *Journal of Business and Economic Statistics* 23 (1): 61–75.

Bernstein, A., M. T. Gustafson, and R. Lewis. 2019. "Disaster on the Horizon: The Price Effect of Sea Level Rise." *Journal of Financial Economics* 134 (2): 253–72.

Bian, X., B. D. Waller, and S. A. Wentland. 2016. "The Role of Transaction Costs in Impeding Market Exchange in Real Estate." *Journal of Housing Research* 25 (2): 115–35.

Braga, B., and R. I. Lerman. 2019. "Accounting for Homeownership in Estimating Real Income Growth." *Economics Letters* 174:9–12.
Bui, L. T., and C. J. Mayer. 2003. "Regulation and Capitalization of Environmental Amenities: Evidence from the Toxic Release Inventory in Massachusetts." *Review of Economics and Statistics* 85 (3): 693–708.
Bureau of Economic Analysis. 2018. Table 2.5.5: Personal Consumption Expenditures by Function. Accessed December 17, 2018. https://apps.bea.gov/iTable/iTable.cfm?reqid=19&step=2#reqid=19&step=2&isuri=1&1921=survey.
Campbell, J. Y., and J. F. Cocco. 2007. "How Do House Prices Affect Consumption? Evidence from Micro Data." *Journal of Monetary Economics* 54 (3): 591–621.
Coulson, N. E., and H. Li. 2013. "Measuring the External Benefits of Homeownership." *Journal of Urban Economics* 77:57–67.
Crone, T. M., L. I. Nakamura, and R. Voith. 2000. "Measuring Housing Services Inflation." *Journal of Economic and Social Measurement* 26 (3, 4): 153–71.
Crone, T. M., L. I. Nakamura, and R. P. Voith. 2009. "Hedonic Estimates of the Cost of Housing Services: Rental and Owner Occupied Units." In *Price and Productivity Measurement*, vol. 1, *Housing*, edited by W. E. Diewert, B. M. Balk, D. Fixler, K. J. Fox, and A. O. Nakamura, 51. Manchester, UK: Trafford Press.
Davis, M. A., A. Lehnert, and R. F. Martin. 2008. "The Rent-Price Ratio for the Aggregate Stock of Owner-Occupied Housing." *Review of Income and Wealth* 54 (2): 279–84.
Diewert, W. Erwin. 2003. "The Treatment of Owner Occupied Housing and Other Durables in a Consumer Price Index." Discussion Paper 03-08, Department of Economics, University of British Columbia, Vancouver, Canada. http://www.econ.ubc.ca/discpapers/dp0308.pdf.
———. 2006. "Conclusions and Future Directions." Summary paper of the OECD-IMF Workshop on Real Estate Price Indices, Paris. http://www.oecd.org/dataoecd/32/21/37848333.pdf.
———. 2008. "Conclusions and Future Directions." Paris OECD-IMF Workshop on Real Estate Price Indexes, Paris, November 6–7, 2006.
———. 2009. "Durables and Owner-Occupied Housing in a Consumer Price Index." In *Price Index Concepts and Measurement*, Studies in Income and Wealth, vol. 70, edited by W. E. Diewert, J. Greenlees, and C. Hulten, 445–500. Chicago: University of Chicago Press.
Diewert, W. E., and A. O. Nakamura. 2009. "Accounting for Housing in a CPI." In *Price and Productivity Measurement*, vol. 1, *Housing*, edited by W. E. Diewert, B. M. Balk, D. Fixler, K. J. Fox, and A. O. Nakamura, 7–32. Manchester, UK: Trafford Press. www.vancouvervolumes.com and www.indexmeasures.com.
Diewert, W. E., A. O. Nakamura, and L. I. Nakamura. 2009. "The Housing Bubble and a New Approach to Accounting for Housing in a CPI." *Journal of Housing Economics* 18 (3): 156–71.
Flavin, Marjorie, and Takashi Yamashita. 2002. "Owner-Occupied Housing and the Composition of the Household Portfolio." *American Economic Review* 92 (1): 345–62.
Garner, T. I., and R. Verbrugge. 2009. "Reconciling User Costs and Rental Equivalence: Evidence from the US Consumer Expenditure Survey." *Journal of Housing Economics* 18 (3): 172–92.
Gilchrist, S., and J. V. Leahy. 2002. "Monetary Policy and Asset Prices." *Journal of Monetary Economics* 49 (1): 75–97.
Gill, H. L., and D. R. Haurin. 1991. "User Cost and the Demand for Housing Attributes." *Real Estate Economics* 19 (3): 383–96.

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 *Housing Markets and the Economy: Risk, Regulation, and Policy: Essays in Honor of Karl E. Case*, edited by E. L. Glaeser and John M. Quigley. Cambridge, MA: Lincoln Institute of Land Policy.

Goeyvaerts, G., and E. Buyst. 2019. "Do Market Rents Reflect User Costs?" *Journal of Housing Economics* 44 (June): 114–30.

Goodhart, C., 2001. "What Weight Should Be Given to Asset Prices in the Measurement of Inflation?" *Economic Journal* 111 (472): 335–56.

Gordon, R. J., and T. vanGoethem. 2007. "Downward Bias in the Most Important CPI Component: The Case of Rental Shelter, 1914–2003." In *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, Studies in Income and Wealth, vol. 67, edited by Ernst R. Berndt and Charles R. Hulten, 153–95. Chicago: University of Chicago Press.

Haffner, M., and K. Heylen. 2011. "User Costs and Housing Expenses: Towards a More Comprehensive Approach to Affordability." *Housing Studies* 26 (4): 593–614.

Hill, R. J., and I. A. Syed. 2016. "Hedonic Price–Rent Ratios, User Cost, and Departures from Equilibrium in the Housing Market." *Regional Science and Urban Economics* 56: 60–72.

Himmelberg, C., C. Mayer, and T. Sinai. 2005. "Assessing High House Prices: Bubbles, Fundamentals and Misperceptions." *Journal of Economic Perspectives* 19 (4): 67–92.

Jorgenson, D. 1963. "Capital Theory and Investment Behavior." *American Economic Review* 53 (2): 247–59.

———. 1967. "The Theory of Investment Behavior." In *Determinants of Investment Behavior*, edited by Robert Ferber, 129–75. New York: NBER.

Katz, Arnold 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 *Price and Productivity Measurement*, vol. 1, *Housing*, edited by W. E. Diewert, B. M. Balk, D. Fixler, K. J. Fox, and A. O. Nakamura, 33–50. Manchester, UK: Trafford Press. www.vancouvervolumes.com and www.indexmeasures.com.

———. 2017. "Imputing Rents to Owner-Occupied Housing by Directly Modelling Their Distribution." BEA Working Paper WP2017-7, Washington, DC.

Komolafe, M., 2018. "Dwelling Services, with an Emphasis on Imputed Rent in the European Union." *Regional Statistics* 8 (1): 168–86.

Kuminoff, N. V., and J. C. Pope. 2013. "The Value of Residential Land and Structures during the Great Housing Boom and Bust." *Land Economics* 89 (1): 1–29.

Lebow, D. E., and J. B. Rudd. 2003. "Measurement Error in the Consumer Price Index: Where Do We Stand?" *Journal of Economic Literature* 41 (1): 159–201.

Linden, L., and J. E. Rockoff. 2008. "Estimates of the Impact of Crime Risk on Property Values from Megan's Laws." *American Economic Review* 98 (3): 1103–27.

Lowenstein, Lara. 2018. "Consumption of Housing During the 2000s Boom: Evidence and Theory." Working Paper, Federal Reserve Bank of Cleveland, Cleveland, OH.

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, Chicago, December 2007.

McCarthy, J., and R. W. Peach. 2010. *The Measurement of Rent Inflation*. Staff Report No. 425, Federal Reserve Bank of New York, New York.

McFadyen, S., and R. Hobart. 1978. "An Alternative Measurement of Housing Costs and the Consumer Price Index." *Canadian Journal of Economics* 11 (1): 105–12.

Mian, A., K. Rao, and A. Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128 (4): 1687–1726.

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.

Moulton, J. G., B. D. Waller, and S. A. Wentland. 2018. "Who Benefits from Targeted Property Tax Relief? Evidence from Virginia Elections." *Journal of Policy Analysis and Management* 37 (2): 240–64.

Moulton, J. G., and S. A. Wentland. 2018. "Monetary Policy and the Housing Market." Working Paper, Washington, DC.

Organisation for Economic Co-operation and Development. 2009. *Measuring Capital: OECD Manual 2009*. 2nd ed. https://www.oecd.org/sdd/productivity-stats/43734711.pdf.

Ozimek, A. 2014. "Sticky Rents and the CPI for Owner-Occupied Housing." Working Paper. http://gradworks.umi.com/35/95/3595699.html.

Pakes, A. 2003. "A Reconsideration of Hedonic Price Indexes with an Application to PC's." *American Economic Review* 93 (5): 1578–96.

Poterba, J., and T. Sinai. 2008. "Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income." *American Economic Review* 98 (2): 84–89.

Poterba, J. M. 1992. "Taxation and Housing: Old Questions, New Answers." *American Economic Review* 82 (2): 237–42.

Poterba, J. M., D. N. Weil, and R. Shiller. 1991. "House Price Dynamics: The Role of Tax Policy and Demography." *Brookings Papers on Economic Activity* 1991 (2): 143–203.

Rosen, S. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy* 82 (1): 34–55.

Sinai, T., and N. S. Souleles. 2005. "Owner-Occupied Housing as a Hedge against Rent Risk." *Quarterly Journal of Economics* 120 (2): 763–89.

United Nations, Commission of the European Communities, International Monetary Fund, Organisation for Economic Co-operation and Development, and World Bank. 2010. System of National Accounts 2008. New York: United Nations. 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): 671–99.

Wentland, S., B. Waller, and R. Brastow. 2014. "Estimating the Effect of Crime Risk on Property Values and Time on Market: Evidence from Megan's Law in Virginia." *Real Estate Economics* 42 (1): 223–51.