

Tax Policy and the Heterogeneous Costs of Homeownership*

Kelly C. Bishop, Arizona State University

Jakob Dowling, Amazon.com

Nicolai V. Kuminoff, Arizona State University & NBER

Alvin D. Murphy, Arizona State University

October 24, 2023

Abstract

The real economic cost of homeownership depends on an intricate system of taxes and subsidies that vary over time and across the United States. We incorporate the key features of this system into a framework for measuring the annual user-cost of housing and we use it to document how housing costs and subsidies varied over time, across space, and with household demographics in 2016-2017. Then we examine how the Tax Cuts and Jobs Act of 2017 subsequently reduced subsidies and increased the relative cost of housing. We report how these changes varied by geography, homeownership, race, and voting behavior.

Key Words: user cost of housing, subsidies, homeownership, TCJA

JEL Classification Numbers: R28, R38, H22

*We especially thank Daniel Feenberg for developing a customized version of TAXSIM software to support this study. We are also grateful for insights and suggestions from Leah Brooks, Ed Coulson, Mike Eriksen, Fernando Ferreira, Eva de Francisco, Andra Ghent, Aaron Hedlund, Erik Hembre, Jaren Pope, Raven Molloy, Jessica Shui, and participants in seminars and conferences at the ASSA Annual Meetings, WEAI annual meetings, Arizona State University, Urban Economics Association Annual Meetings, and Innovations in Housing Affordability conference. Finally, we thank Raghav Warriar for excellent research assistance. Dowling's contributions to the paper were made prior to joining Amazon.

1 Introduction

The United States is one of several countries that subsidize homeownership through the tax code.¹ The US government transferred over \$95 billion to homeowners in 2017 alone by allowing them to deduct mortgage-interest and property-tax payments on their tax returns (Joint Committee on Taxation, 2017). Additionally, these transfers are unequally distributed. Renters are excluded, as are homeowners who do not itemize deductions on their tax returns.² Further, among itemizing homeowners, the subsidies are larger for households who face higher marginal income tax rates, own more expensive houses, and live in areas with higher property taxes. Thus, the US tax code ensures that different households would be charged different prices to live in the same house in the same year. These disparities are substantial: they can easily adjust the annual cost of home-ownership by 5% to 10%. Documenting these disparities is important for understanding the distributional consequences of tax policy and for understanding residential sorting and the demand for housing.³

The first contribution of this paper is to calculate the tax subsidy to homeownership and the real economic cost that would be paid by any household to own a house at any location in the US for one year. Calculating these location- and time-varying measures requires predicting whether a household will choose to itemize deductions when filing a federal tax return and accounting for other market forces that create variation in the real cost of ownership. These forces include the opportunity cost of capital invested in a property, the expected capital gains and risk premia from the investment, mortgage interest rates, and depreciation. We incorporate these features by merging American Community Survey public-use micro data describing 3.3 million households and the houses they occupied in 2016 through 2019 with several other public data files. We use NBER’s TAXSIM model for mapping household income,

¹Similar policies exist in Belgium, Ireland, the Netherlands, Switzerland, and Sweden and previously existed in Canada and the United Kingdom (Binner and Day 2015).

²Coulson and Li (2013) discusses the benefits of homeownership versus renting and the efficiency of subsidies.

³See, for example, Epple and Sieg (1999), de Bartolome and Rosenthal (1999), Sieg et al (2004), Bajari and Kahn (2005), Bayer, Ferreira, and McMillan (2007), Banzhaf and Walsh (2008), Wong (2013), Bayer, McMillan, Murphy, and Timmins (2016), Mangum (2017), Han, Han, and Zhu (2018), Bishop and Murphy (2019), Ouazad and Ranciere (2019), Bibler and Billings (2020), Caetano and Maheshri (2021), Ahlfeldt, Heblich, and Seidel (2022), Anenberg and Ringo (2022), and Han, Ngai, and Sheedy (2022).

expenditures, and tax filing strategies into tax burdens. We assume that households choose tax-filing strategies to minimize their tax burdens conditional on their income and expenditures. This process yields household-by-house-specific measures of the annual cost of housing as well as the tax subsidy to homeownership during 2016-2019. Interested readers can explore these data using interactive maps or download annual means for approximately 2,400 Public-Use Microdata Areas (PUMAs) defined by the U.S. Census Bureau at this paper’s website: www.housingusercost.org.

The high geographic and temporal resolution of our approach allows us to evaluate the distributive impacts of policies affecting the tax treatment of homeownership. This feature is useful for at least three reasons. First, it can support research on incorporating measures of equity into program evaluations (Banzhaf et al. 2019, Brouillette et al. 2022, Akbar et al. forthcoming). Second, it can help to disentangle the extent to which residential segregation is driven by demographic variation in the real cost of homeownership versus homophily (Aliprantis et al. 2022, Davis et al. 2023). Finally, it can enable analysts to meet new federal guidelines for extending regulatory analyses to include distributive impacts of policies by geography, wealth, race and other attributes (Biden 2023, US Office of Management and Budget 2023, Cronin et al. 2023).

The second contribution of this paper is to evaluate the distributive impacts of major changes to the tax treatment of homeownership made by the Tax Cuts and Jobs Act (TCJA) in 2017 and its scheduled expiration in 2025. We first analyze how tax subsidies and real housing costs changed in 2018 and 2019, the first two years after the TCJA took effect. Then we perform counterfactual simulations to calculate the subsidies and housing costs that would have prevailed in the absence of some (and all) of the TCJA’s changes to the tax treatment of homeownership. This exercise previews what may happen if Congress allows the TCJA’s individual tax provisions to expire in 2025. We use our framework to analyze how the effects of the TCJA’s tax treatment of housing are distributed by geography, tax filing behavior, homeowner status, race, and political party.

Our main empirical findings can be summarized by three broad conclusions. First, we provide new evidence on the extent to which the real cost of homeownership in the US varies across fine definitions of space and across households demographics

at a point in space. We estimate very large spatial variation in mean ownership costs across PUMAs, arising from local property taxes, expected capital gains, and residential sorting by income and other household attributes that affect tax liability. For example, consider a house worth \$300,000. The average annual cost of owning this house was \$18,300 at the 90th percentile of PUMAs, compared to only \$9,600 at the 10th percentile. Both the level and dispersion of these costs was significantly driven by tax policy. We find that the average homeowner received subsidies equal to 6.6% of their annual housing costs in 2016 and 2017.⁴ At the household level, the subsidy is increasing in income, property value, property taxes and, of course, homeownership, all of which correlate with other demographics such as race. As a result, we estimate that the mean subsidy among Black household-heads was about half the mean subsidy among White household-heads which, in turn, was about half the mean subsidy among Asian household-heads.

The second broad conclusion is that the TCJA's changes to tax policy made homeownership less affordable relative to other goods. Beginning in 2018, the TCJA drastically reduced the mortgage interest and property taxes recouped by tax subsidies. In 2018 and 2019, the average homeowner received subsidies equal to just 2.1% of ownership costs. We estimate that this figure would have been 6.9% had the pre-TCJA tax code stayed in place. Decomposing the overall subsidy into federal and state subsidies highlights how stark the changes were at the federal level – the TCJA had relatively small impacts on the state subsidy, but reduced the federal subsidy by over 80%.

Our third conclusion is that there were significant heterogeneity in the extent to which the TCJA reduced housing affordability. Homeowners in affluent PUMAs in coastal states that had the largest subsidies prior to the TCJA also saw the largest reductions. Since the TCJA was a partisan Republican bill, we examine how its effects correlate with political affiliation. We find that Democrat-voting PUMAs lost \$600 per household in annual subsidies, compared to \$340 for Republican-voting PUMAs.⁵ In terms of race, Asian household-heads lost \$923 on average, compared to

⁴As a comparison, the average homeowner with a mortgage received subsidies equal to 9.1% of their annual housing costs in 2016 and 2017.

⁵We designate areas as Democratic or Republican based on which candidate received the majority of votes in the 2016 presidential race between Hillary Clinton (Democrat) and Donald Trump (Republican).

\$528 for White household-heads, and \$220 for Black household-heads. We conclude that these effects will be reversed if the TCJA’s major tax provisions expire in 2025.

Our paper is closely related to a number of studies that develop methods for measuring implicit housing subsidies and the real cost of homeownership (e.g., Poterba 1984, 1992, Himmelberg, Mayer, and Sinai 2005, Harding, Rosenthal, and Sirmans 2007).⁶ We incorporate the insights of this literature and account for the rates at which households pay off their mortgages and account for non-linearity in the tax code. In particular, we predict how homeownership affects the likelihood that households minimize their tax burdens by choosing to itemize deductions. This step is quantitatively important for our policy implications.⁷ We validate this step by comparing our predictions for tax filing behavior to the most granular data on itemization rates reported in public IRS files. Our predictions closely match IRS data before and after the TCJA policy shift. For example, among households with adjusted gross incomes of \$1 million or less, the IRS itemization rates were 30.4% in 2017 and 11.2% in 2018; our model-based predictions are 30.4% and 10.9%. Our predictions also replicate how rates vary by income bin and PUMAs. For income bins the correlation between predicted and actual itemization rates is 0.99 and for PUMAs it is 0.95.

Our paper also relates to studies investigating the TCJA’s effects on housing markets. Ambrose et al. (2022) finds that the TCJA reduced tax subsidies to homeownership and Hembre and Dantas (2022) finds that the TCJA reduced ownership rates. Coen-Pirani and Sieg (2019) finds that the TCJA incentivized older, high-productivity households to relocate to lower-cost cities. Martin (2018), Rapoport (2019), and Sommer and Sullivan (2021) all conclude that the TCJA caused house prices to decline, and Li and Yu (2021) suggests heterogeneous effects on price

⁶Numerous papers have applied the user-cost concept to analyze housing policy and the costs of homeownership. Select examples include Glaeser and Shapiro (2003), Gyourko and Sinai (2003), Hilber and Turner (2004), Sinai and Gyourko (2004), Glaeser (2013), Albouy and Hanson (2014), Binner and Day (2015), Martin and Hanson (2016), DeFusco and Paciorek (2017), Knoll, Schularick, and Steger (2017), Sommer and Sullivan (2018), Davis (2019), Garriga, Manuelli, and Peralta-Alva (2019), Blouri, Buchler, and Schoni (2021), Fuster and Zafar (2021), and Gruber, Jensen, and Kleven (2021).

⁷Not owning a house may lead a household to minimize its tax burden by taking the standard deduction instead of reducing itemizations by the amount of the mortgage interest paid. Ignoring this non-linearity yields significantly higher estimates of subsidies because only 21% of currently itemizing homeowners would still itemize if they did not own a house.

growth. Unlike most of these studies, we choose to not explicitly model the TCJA’s effects on equilibrium outcomes like migration rates, homeownership rates, or housing prices. Instead, we add to the literature by examining how the TCJA’s effects on households vary by race, geography, tax filing behavior, mortgage tenure, and political party, and by designing counterfactual experiments to inform the distributional outcomes of allowing some or all of the TCJA’s tax provisions to expire in 2025.

Finally, the new database that we build to describe the real cost of housing for actual and counterfactual owners has potential to advance knowledge of markets for housing, labor, and public goods. In particular, the literature on residential and labor market sorting that builds on Tiebout (1956) and Roback (1982) considers how agents choose spatial locations based, in part, on annual housing costs (Kuminoff, Smith, and Timmins 2013). Our database improves the accuracy of existing measures for how these costs vary across metro areas (e.g., Albouy 2009, Bayer et al. 2009, Diamond 2016) and it provides new, quantitatively important evidence on how the cost of owning a particular house varies across households (e.g., Sieg et al. 2004, Bayer et al. 2016, Epple, Quintero, and Sieg 2020, Ma 2019). These data are also used in calculating expenditures on public goods and amenities (e.g., Albouy 2016, Bieri, Kuminoff, and Pope 2022) and they are crucial for understanding the decision to rent or own and its implications for wealth and welfare (e.g., Tracy, Schneider, and Chan 1999, Coulson and Li 2013, Binner and Day 2015, Gruber, Jensen, and Kleven 2021). The idea that the real cost of ownership can vary for the reasons we emphasize is well known, but its quantitative implications have been underexplored due to lack of data.⁸ We remove this limitation by developing a comprehensive US database using transparent methods that rely entirely on recurrent public data.

The next section presents our framework for calculating tax subsidies to homeowners and the real cost of ownership. Section 3 presents validation tests of our framework’s predictive accuracy. Section 4 summarizes our estimates for tax subsidies and ownership costs in 2016-2017. Section 5 summarizes how these subsidies and costs changed, on average, after the TCJA took effect in 2018-2019 and how the changes varied by demographics, and Section 6 concludes.

⁸Gindelsky, Moulton, and Wentland (2020) develops and analyzes national data on the user cost of housing using data and methods that could have been replicated by other researchers who had access to confidential micro data from Zillow before September 30, 2023 when Zillow terminated all data-sharing agreements with external researchers (Zillow 2022).

2 The Annual Economic Cost of Homeownership

The economic cost of owning a house depends on far more than the direct cost of a parcel of land and the structures that are built on it. The cost of homeownership also depends on the opportunity cost of capital invested in a property, the expected capital gains and risk premia from that investment, mortgage interest rates, annual depreciation of the property and, importantly, the tax code. On one hand, property taxes add to the cost of ownership. On the other hand, since 1913 the federal tax code has included two important subsidies that reduce the cost of owning a house relative to other consumption.⁹ First, homeowners can deduct certain state and local taxes, including property taxes, on their federal tax returns. Second, homeowners can deduct mortgage interest payments. The cumulative effect of this large set of taxes and capital costs on the real economic cost to a particular household of owning a particular house in a particular year can be measured by a single statistic – the user cost rate.

2.1 Defining the User Cost Rate

The user cost rate [henceforth UCR] is the key statistic for measuring spatial and temporal variation in the real costs of homeownership and for assessing how those costs are affected by policy. To fix ideas, let P_{ijt} denote the value of a house owned by household i in location j in year t . The annualized cost of homeownership for this house is specific to household i and is denoted \tilde{r}_{ijt} . This annualized cost can be expressed as a fraction of the house's value:

$$\tilde{r}_{ijt} = P_{ijt} \cdot UCR_{ijt}, \quad (1)$$

where UCR_{ijt} denotes the UCR that also varies by household, location, and year.

Following Poterba (1988, 1992) and Himmelberg, Mayer, and Sinai (2005), the

⁹These subsidies are increasing in a household's marginal tax rate, as they are driven by deductions, and may only be collected if homeowners itemized their expenses on their federal tax returns.

UCR can be expressed as:

$$UCR_{ijt} = (1 - ltv_{ijt}) \cdot rf_{ijt} + ltv_{ijt} \cdot rm_t + \omega_{jt} + \delta + \epsilon - \gamma_{jt} - s_{ijt} \quad (2)$$

In Equation (2), ltv is the loan to value ratio, rf is the risk-free after-tax rate of return on capital, rm is the mortgage interest rate, ω is the property tax rate, δ is the rate of depreciation, ϵ is the owner’s risk premium, and γ is the expected capital gains.¹⁰ The last term, s , is the tax refund of property taxes and mortgage interest obtained by a homeowner who itemizes their deductions, expressed as a fraction of the house value. We discuss definitions of the total subsidization of housing below.¹¹

In principle, every input to the UCR can vary by household, location, and time. The lack of i , j , and/or t subscripts on some of the inputs to the UCR formula in Equation (2) is designed to preview the fact that some inputs to the UCR formula do not vary at these levels or are measurable at coarser levels due to constraints on data granularity. We limit the potential for these constraints to influence our conclusions by aggregating our results by PUMA, year, and household type. Specifically, we leverage the linearity of Equation (2) by replacing (unobserved) household-specific measures for certain inputs with their corresponding (estimated) PUMA-by-year-by-type means. Thus, the subscripts denote the levels at which we calculate each component of the UCR formula.

2.2 Calculating the User Cost Rate

We start with data on all current homeowners in the American Community Survey (ACS) IPUMS 1% annual samples for 2016 through 2019 (Ruggles et al. 2022). We make two sample cuts. First, we drop approximately 6% of observations for which the home is a less-traditional dwelling such as a mobile home, trailer, boat, van, tent, or unspecified structure. Then we drop 1% of observations where the occupant self-reports a value for the home that is an extreme outlier (more than 6 standard

¹⁰We follow prior literature by specifying UCR to not be a function of price. In principle, some determinants of UCR could vary with price. For example, buying a more expensive house could affect the mortgage interest rate. During the time period of our study, the jumbo-conventional spread was small and negative (Fisher et al. 2021).

¹¹The formula for s is presented in Equation (4) and the TCJA’s impact on s is discussed in detail in Section 5.

deviations from the PUMA median and/or below \$10,000). These two sample cuts leave us with data describing 3,266,870 households that comprise 93% of the ACS homeowner sample. We summarize the data and procedures that we use to calculate each input to Equation (2) for the 93% ACS sample in the remainder of this section and provide additional details in Appendix A.

Loan-to-Value Ratio (ltv_{ijt})

We calculate each household’s current loan-to-value ratio, ltv_{ijt} , using the household’s responses to ACS questions regarding their mortgage financing. For the 36% of homeowners that report not having an active mortgage, we set ltv_{ijt} to zero. For the 64% of homeowners that report having an active mortgage, we derive an amortization schedule by combining their reported monthly mortgage payment with the reported year in which they purchased the home and an estimate for the mortgage interest rate (explained below). This amortization schedule allows us to impute ltv_{ijt} for each household with an active mortgage in each year.

The difficulty with deriving households’ amortization schedules is that we do not observe their individual mortgage terms and origination dates.¹² We address this information gap by calibrating the mortgage term so that our derived amortization schedule matches a closely-related data moment reported by Keys et al. (2016). That study reports that the average home loan had 23.4 years remaining in 2010. We reproduce this moment on our derived amortization schedule for 2016-2017 by assuming that homeowners select a 32-year term for a fixed-rate mortgage.¹³ While our calibration procedure recognizes that many households refinance their mortgages, our main results are robust to alternatively assuming a standard 30-year term for everyone.¹⁴ Finally, we use the derived amortization schedule and tenure in the home

¹²While 30-year fixed rate mortgages are very common, mortgage lengths can be shorter (e.g., via prepayments or shorter term lengths) or longer (e.g., via refinancing or home equity loans). Indeed, we observe that 30% of ACS households with 30+ years of tenure as homeowners still make mortgage payments.

¹³We use only the 2016 and 2017 ACS samples for this calibration to avoid any potential influence of the TCJA. The calibration procedure is described in more detail in Appendix A3.

¹⁴Keys et al. (2016) observe a random sample of outstanding mortgages and loan terms. Calibrating our assumed mortgage term to their sample can be viewed as a “reduced form” adjustment that integrates over heterogeneous loan types and refinancing behaviors. Any given household’s refinancing behavior (e.g., timing of refinancing or presence of cash-out refinancing) may, of course, differ from this mean. Appendix A3 explains why this adjustment has very little effect on our

to impute the current loan-to-value ratio.

After-tax Risk-Free Rate (rf_{ijt})

To calculate the household-specific after-tax, risk-free rate of the return to capital, rf_{ijt} , we begin with the (pre-tax) risk-free rate market, denoted by $ptrf_t$. To calculate $ptrf_t$, we use a rolling average over the prior ten years of the yield on U.S. Treasury securities at a 10-year constant maturity. We then use NBER’s TAXSIM software to obtain household-specific after-tax rates by calculating the retained fraction of investment returns after paying long-run capital gains tax. Using the after-tax rate follows the prior literature and allows the taxation of non-housing capital investment to reduce the UCR. We let σ_{ijt} denote the value of this reduction, where $\sigma_{ijt} = (1 - ltv_{ijt})(ptrf_t - rf_{ijt})$. We describe this procedure in more detail in Appendix A.3.1.

Mortgage Rate (rm_t)

We follow Himmelberg et al. (2005) and assume a common mortgage rate, rm_t , across households. We calculate rm_t by taking a rolling average over the prior ten years of interest rates on 30-year fixed rate mortgages across the US. Our choice for the term length is motivated by the fact that 30-year fixed rate mortgages account for over 90% of the mortgage market (Mortgage Bankers Association, 2022) and our decision to use a national average rate is motivated by evidence that geographic variation in rates is minimal (Bhutta, Fuster, and Hizmo, 2020). We choose a rolling average to address unobserved heterogeneity in household-specific origination dates. This has the desired effect of smoothing the UCR in response to fluctuations in year-specific interest rates.¹⁵

We cannot observe variation in mortgage rates across households in the ACS. Bhutta, Fuster, and Hizmo (2020) show that mortgage rates can vary across similar borrowers (by up to 54 basis points between the 10th and 90th percentiles). However, this cross-sectional variation is small relative to the temporal variation in rates that we capture with rm_t . For example, between 2009 and 2019, annual average interest

results relative to simply assuming a 30-year fixed rate mortgage for everyone.

¹⁵While mortgage rates were falling during the time period that we consider, the large adjustment costs associated with refinancing would lead many households not to refinance (Keys et al. 2016).

rates on 30-year fixed mortgages varied by 190 basis points (from 3.3% to 5.2%). The relevance of latent cross-sectional variation in mortgage rates for our UCR results is also dampened when we aggregate results to the PUMA-year level.

Property Tax Rate (ω_{jt})

We use ACS data to impute a property tax rate, ω_{jt} , for each PUMA-year which is the finest delineation allowed by ACS data. Specifically, we divide the sum of total reported property taxes paid by the sum of total reported property values in each PUMA and year.¹⁶ These are the most accurate property tax rate data that can be collected for the entirety of the US (Emrath, 2002) and this method has been used in several other contexts (Bieri et al., 2022; Cabral and Hoxby, 2012).

Depreciation (δ)

We set the annual depreciation rate of housing capital, δ , to be 2.49% per year. This is based on a repeat-sales model in Harding et al. (2007) that, to our knowledge, represents the most accurate national estimate of housing depreciation. As noted by Himmelberg et al. (2005), Halket, Nesheim, and Oswald (2020), and Head, Lloyd-Ellis, and Stacey (2023), the rate of depreciation may vary across house and geography due to heterogeneity in the ratio of land value to property value, a feature not captured here or in the prior literature that assumes a homogenous depreciation rate.

Owner's Risk Premium (ϵ)

We follow Flavin and Yamashita (2002) and Himmelberg et al. (2005) in using 2% as the owner's risk premium, ϵ . In principle, one could allow the risk premium to vary

¹⁶We restrict attention to households in owner-occupied houses and apartments. In the ACS, the variable for annual property taxes paid is reported in ranges. We calculate the midpoints of these ranges and assign the appropriate midpoint value to each household. The top range indicates that a household reports paying more than \$10,000 in annual property taxes. As we don't observe a midpoint for these households and we don't need all property-tax bins to estimate the single PUMA-specific property tax rate, these households are excluded from the calculation of property tax rates (but not from the subsequent analysis).

by PUMA.¹⁷ However, we believe that assuming a common housing risk premium provides a good approximation.

Expected Capital Gains (γ_{jt})

We calculate the location-specific nominal rate of expected future capital gains, γ_{jt} as the sum of two terms: expected real house price appreciation and expected inflation. We do not observe survey data on expected real house price appreciation, so follow Himmelberg et al. (2005) and estimate it under the assumption that households use historical real appreciation as a predictor of future real appreciation.¹⁸ We estimate time-invariant, historical real appreciation for 186 MSAs and the non-MSA markets of 49 states by using decennial Census and ACS data to calculate hedonic price indices for each of 235 distinct markets over 1990-2019.¹⁹ For the second term, expected inflation, we use the 10-year expected inflation rate from the Livingston Survey of professional forecasters, which varies by year. We do not include capital-gains taxes on expected capital gains as most households are exempt from housing-related capital-gains taxes.²⁰ We map the resulting market-specific measures back to PUMAs using a crosswalk provided by IPUMS. Appendix A.4 contains further details about this mapping and the construction of the price indices.

Tax Refund of Property Taxes and Mortgage Interest (s_{ijt})

Calculating s_{ijt} for household i in location j in year t presents two measurement challenges. First, s_{ijt} is only collected by households who choose to itemize deductions when filing their taxes, and this choice is not observable in the ACS or other public

¹⁷For example, one could use PUMA-specific variation over time to calculate PUMA-specific price variation and then use this to adjust the within-MSA variances estimated in Flavin and Yamashita (2002). Alternatively, one could potentially also use the small amount of geographic variation in offered mortgage rates to capture area-specific risk. Finally, we note that, in addition to affecting the UCR, volatility could also affect prices as shown in Amior and Halket (2014).

¹⁸Using data from the Federal Reserve Bank of New York’s Survey of Consumer Expectations, Armona, Fuster, and Zafar (2019) show that when offered information to help predict future house price growth, slightly more than half of participants chose historical growth over expert forecasts.

¹⁹All states except Rhode Island since every PUMA in that state is associated with an MSA.

²⁰The Taxpayer Relief Act of 1997 exempts from tax the first \$500,000 (\$250,000) of capital gains for married couples (individuals).

data.²¹ Second, s_{ijt} depends on several household characteristics including income, property tax payments, mortgage interest payments, other deductible expenses, and geographic location. We address both challenges by leveraging the richness of ACS data together with ancillary data on charitable giving from the Panel Study of Income Dynamics (PSID) and NBER’s TAXSIM 35 software to predict household-level itemization decisions and calculate their corresponding subsidy rates.²²

The calculation is performed using three simulated tax scenarios. In all three scenarios, one actual and two counterfactual, we assume that maximizing the household’s objective function corresponds to minimizing their tax burden. Households do this by choosing whether to itemize or take the standard deduction:²³

$$itemize^*(Z_{ijt}) = 1[\text{tax}(itemize|Z_{ijt}) < \text{tax}(std\text{deduction}|Z_{ijt})], \quad (3)$$

where Z_{ijt} captures all of the factors that determine a tax burden and $\text{tax}(\cdot)$ is the intricate, non-linear function that maps the household itemization decision and their Z_{ijt} into their tax burden. The tax determinants, Z_{ijt} , include income, age, number of dependents, marital status, state of residence, as well as deductible expenses such as local taxes, property taxes, mortgage interest, charitable giving, and medical expenses.²⁴

First, we use TAXSIM to estimate each household’s actual tax liability (state and federal). Importantly, we perform this step for all ACS households, not just homeowners, so that we can predict moments of the national distribution of filing behavior that are directly comparable with statistics published by the IRS. We present this comparison as a validation exercise in Section 3.2.

Second, for homeowners, we recalculate each household’s tax liability in a coun-

²¹The IRS does not provide data about itemization rates for homeowners separately from renters, which means additional data or assumptions are needed even to calculate *mean* subsidy rates.

²²The TAXSIM model can be accessed at taxsim.nber.org. Feenberg and Coutts (1993) provide an introduction.

²³See Poterba and Sinai (2008), Saez and Zucman (2016), Benzarti (2020), and Foote, Loewenstein and Willen (2021) for analyses of household decisions about whether to itemize or take the standard deduction.

²⁴Charitable donations and medical expenses are presumably jointly determined with other deductions and could respond to policy changes. While we do not explicitly model this co-dependency, we separately estimate charitable donations and medical expenses for pre and post policy, which allows the estimates to vary with policy. The validation exercise described in Section 3 suggests that this procedure works well.

terfactual scenario in which the return on housing equity is taxed. We do this by adding to the household’s income the hypothetical amount that their housing equity would return if invested in a non-housing asset. In this step, we assume they would obtain the risk-free rate of return on this investment, recognizing that this return would be taxed.²⁵

Third, we calculate what each homeownership household’s tax liability would be in another counterfactual scenario in which they no longer own a home. We do this by setting their deductible property tax and mortgage interest payments to zero before calculating their tax liability in TAXSIM. In this step, we also maintain the increase in taxable income from the return on housing equity as in the second step.

s_{ijt} then is calculated by the difference between the tax liabilities calculated in the second and third steps, divided by the household’s property value.

$$s_{ijt} = \frac{\text{ReducedTaxLiability}_{ijt}}{P_{ijt}} \quad (4)$$

where $\text{ReducedTaxLiability}_{ijt}$ is the reduced tax liability a homeowner faces due to deducting mortgage interest and property tax.

2.3 Defining the Housing Subsidy Rate

An important part of our analysis is describing heterogeneity in housing subsidization across time and geography. To do this, we must define what constitutes housing subsidization. By construction, a subsidy is defined relative to some counterfactual no-subsidy policy and we consider two different definitions of the housing subsidy rate.

Our preferred measure is to define the subsidy rate as s_{ijt} as shown in Equation (4) above. In this case, the assumed counterfactual world is the one in which mortgage interest and property taxes are not deductible in the US tax code. This definition has the benefit of considering a simple, commonly adopted counterfactual tax policy for the purpose of defining the subsidy.

The second measure is to define the subsidy rate relative to a counterfactual

²⁵This second step is what allows us to calculate household-specific after-tax opportunity cost of funds as described in Appendix A.3.

world in which imputed rental income (net of expenses) is taxed as Haig-Simons income.²⁶ Poterba and Sinai (2011) and Brueckner (2014) discuss the conditions under which the non-taxation of imputed rents can be mapped into the user cost function. In particular, Poterba and Sinai (2011) show that the non-taxation of imputed rents is equivalent in our framework to a subsidy of $s_{ijt} + \sigma_{ijt}$.²⁷

In our context, the distinction between these two subsidy definitions is not marked. First, the choice of subsidy definition does not impact our UCR calculations as defined in Equation (2). Second, we find that TCJA has little impact on σ_{ijt} , which is the only distinction between the definitions. Finally, Appendix A.3 presents results of our analyses using each definition of the housing subsidy rate, as well as explaining our exact procedures for calculating household-level tax liabilities.

3 Validation Test: Predicting Taxpayer Itemization Rates for 2016-2019

The validity of our measures for the real annual cost of homeownership depends on the accuracy of our predictions for the embedded tax subsidy, s_{ijt} . In principle, the ideal way to judge the accuracy of our predictions for s_{ijt} would be to compare them to the deductions taken on households' tax returns. Given the well-known barriers to obtaining administrative data on tax returns in the U.S., we perform a second-best validation test.

Our validation test exploits the way that the tax subsidy to homeownership is entangled with the decision to itemize. Tax filers must itemize to receive s_{ijt} and, all else constant, an increase in s_{ijt} increases the incentive to itemize. Thus, more accurate predictions for s_{ijt} should yield more accurate predictions for itemization behavior. With this in mind, we first use TAXSIM to predict whether each ACS household will minimize its tax burden by itemizing, given our predicted value for

²⁶According to Andrews, Caldera Sánchez, and Johansson (2011), imputed rent is taxed in Iceland, Luxembourg, the Netherlands, Slovenia, and Switzerland.

²⁷Using this definition, the subsidy rate is measured by the difference between the tax liabilities calculated in the first and third steps outlined in the description of s_{ijt} in Section 2.2. Poterba and Sinai (2011) also consider another effect, the option to prepay or default on a mortgage is implicitly subsidized, that we do not consider. We also abstract away from whether one should consider benefits received in return for local property taxes as discussed in Zodrow (2000).

s_{ijt} . Then we aggregate our predictions by income group and PUMA to compare them against annual itemization rates reported by the IRS from 2016 through 2019. We analyze predictive accuracy in the cross-sectional data for each year. Moreover, at the midpoint of the study period the federal tax code changed in ways that drastically reduced s_{ijt} and itemization rates. This quasi-experimental variation in tax policy allows us to judge the accuracy of our model-based predictions for how tax policy changes affect filing behavior and, thus, the annual cost of homeownership. The central benefit of having a validated model is that it allows us to predict counterfactual outcomes as well as actual outcomes at finer levels of aggregation than the publicly available IRS data allow.

Before presenting the results of a validation exercise that compares model-predicted and actual itemization rates before and after the implementation of the TCJA, we briefly outline the TCJA and its impact on incentives to itemize.

3.1 The TCJA Reduced the Incentive to Itemize

On December 22, 2017, President Trump signed the Tax Cuts and Jobs Act (TCJA) which went into effect for the 2018 tax year. The TCJA changed the tax code in several ways that reduced the incentive to itemize. Most importantly, the TCJA approximately doubled the standard deduction. For example, for a married couple filing jointly, the standard deduction increased from \$12,700 in 2017 to \$24,000 in 2018. This reduced the incentive for households to itemize and, thus, reduced their incentive to collect the tax subsidy to homeownership.

The TCJA further reduced the incentive to itemize by weakly reducing the tax subsidy to homeownership, s_{ijt} , that itemizers collect. First, the TCJA reduced the marginal income tax rates at which homeowners can deduct mortgage interest and property tax payments. Second, the TCJA reduced the maximum amount of indebtedness to which the mortgage interest deduction can be applied from \$1,000,000 to \$750,000. Third, the TCJA added a \$10,000 cap on the maximum amount of state and local taxes (SALT) that can be deducted. The fact that this SALT cap includes property taxes is particularly impactful because, prior to the TCJA, the SALT deduction could easily exceed the standard deduction in areas with high property values and/or high property taxes.

The net effect of these three reductions in the tax subsidy to homeownership varies across properties and, conditional on a property, it varies across owners. All else constant, the subsidy is reduced more for owners who have larger reductions in their marginal tax rates due to the TCJA, who have larger mortgage interest payments, and/or who have larger property tax payments. Consequently, we would expect the TCJA’s effect on itemization rates to vary by income group and geography. On aggregate, IRS data indicate that the number of itemizing households fell by over 60% between the 2017 and 2018 tax years.

3.2 Validation Test Results


In principle, our predicted itemization rates could diverge from the actual rates for three potential reasons. First, some households may file taxes in ways that fail to minimize their tax burdens. Second, some households may not accurately report all of their income-relevant information when responding to the ACS.²⁸ Third, the assumptions that we make in order to characterize a household’s filing options based on ACS data and TAXSIM software may introduce errors (e.g., our assumptions for the rates at which households pay off mortgages).

Despite these potential reasons for divergence, our predicted itemization rates are remarkably accurate. Figure 1 compares our predicted rates with actual rates in IRS data for 2016 through 2019. Before the TCJA, the IRS reports national itemization rates of 29.9% and 30.4% in 2016 and 2017, respectively. After the TCJA, the IRS reports rates of 11.2% and 10.8% in 2018 and 2019, respectively. Strikingly, our predictions never differ from the IRS data by more than one third of one percentage point.

The IRS also disaggregates itemization rates by income bin and geography. As a further validation check, we calculate the predicted and actual fractions of itemizers by income bin. We compare our estimates within 14 income bins ranging from adjusted gross income of $\leq \$0$ to \$500,000-\$1,000,000. The correlation between our predictions and IRS reported numbers is 0.99. The correlation only drops slightly, to

²⁸The first two failures have greater scope to affect high-income households who face more complicated tax situations due, for example, to owning a business, charitable giving, and decisions for how to amortize capital gains and losses.

Figure 1: Validation - Predicted versus Actual Itemization Rates



Note: The figure contrasts the fractions of all tax-filing households who choose to itemize according to IRS data with our predictions for the tax-minimizing filing strategy.

0.95, when we repeat the comparison using itemization rates for PUMAs.²⁹ Overall, these results increase our confidence that our framework produces accurate measures for the tax subsidy to homeownership, that it will produce accurate measures for the associated UCR, and that it will be capable of making accurate predictions for how counterfactual tax policies would modify the annual cost of homeownership through tax filing behaviors that can differ across demographic groups.³⁰

²⁹Figures B.1 and B.2 in the Appendix illustrate scatter plots and fitted regression lines between predicted and actual itemization rates by income bin and PUMA.

³⁰Our validation test effectively compares the model-predicted CDF of itemized deductions with the empirical CDF of itemized deductions at a single point, the standard deduction. In an ideal exercise, one would compare these CDFs at all possible points (or all points \geq the standard deduction). That being said, it is encouraging that the CDFs are virtually identical at the key point, the standard deductions, and that they continue to match well when using conditional CDFs that condition on income and geography.

4 User-Cost Rates and Subsidies


This section summarizes our estimates for the user-cost of housing and the tax subsidy to homeownership among heterogeneous household types in 2016 and 2017. We focus on these two years because they provide a baseline for evaluating the subsequent policy changes that we discuss in Section 5. Complete PUMA-by-year-specific measures of UCRs and tax subsidies for 2016-2019 can be downloaded or explored using an interactive tool on this paper’s companion website: www.housingusercost.org.

4.1 User-Cost Rates

User-Cost Rates Realized by Homeowners

Figure 2 shows our estimates of the average UCR among homeowners in each PUMA during 2016-2017. The average homeowner faced a UCR of 4.50%. The figure shows that UCRs vary greatly across the U.S. For example, at the 90th percentile of PUMAs, the mean UCR (6.1%) is nearly double the UCR at the 10th percentile (3.2%).

Figure 2: 2016-2017 Mean UCR by PUMA for Current Homeowners



Note: The figure shows the PUMA-specific mean UCR for current homeowners in 2016-2017. The figure is winsorized at the 1st and 99th percentiles.

The extent of spatial variation in UCRs in Figure 2 is striking. The large

variation across UCR septiles in Figure 2 implies that spatial variation in the UCR is likely to be a quantitatively important part of the implicit housing cost of consuming local public goods (Kuminoff, Smith, and Timmins, 2013). More broadly, by influencing the cost of residential sorting, the UCR influences the spatial allocation of labor, household wealth accumulation, and welfare (Albouy 2016, Diamond 2016). To decompose the geographic variation in UCR shown in Figure 2 into its determinants, Appendix Figures B.4 through B.8 show analogous maps to Figure 2 for each of the UCR determinants that vary across PUMAs. Finally, we note that our estimated UCR measures are positively correlated (correlation coefficient of 0.52) with rent-price ratios that are estimated using hedonic regressions that allow for spatially varying rent-price ratios at the PUMA level (Bayer et al., 2007).³¹

Implicit User-Cost Rates for Prospective Homebuyers

We also calculate implicit UCRs for hypothetical prospective homebuyers; i.e., the UCRs that households would face if they were to buy certain houses in certain areas. These implicit UCRs may differ from the UCRs realized by current owners for at least three reasons. First, for a loan of equal size, the share of monthly mortgage payments that goes towards interest is decreasing in time-since-origination. This means that, all else constant, a new buyer would get a larger tax subsidy, and thus face a lower UCR, than an owner who is further along in their mortgage repayment. Second, loan-to-value ratios are expected to be higher for new buyers. All else equal, this increases the UCR since it increases the loading in the user-cost formula on the mortgage rate, which is generally higher than the risk-free rate. Third, new buyers may differ from current owners in age, income, family structure, and other attributes that affect tax filing behavior, and thus the UCR.³² The net effect of these three differences on the UCR is ex-ante ambiguous and will depend, in part, on how prospective buyers sort themselves across housing markets.


³¹Figure B.9 in the Appendix illustrates a scatter plot and fitted regression line between rent-price ratios and UCR by PUMA.

³²In principle, new homebuyers and current homeowners could face different mortgage rates. For new buyers they could be lower than for current owners if rates are declining over time and there are fixed costs to refinancing. Analogously, if rates are increasing, we would expect current owners to have “locked-in” a lower mortgage rate than what new buyers face. In our estimates, we choose to use rm_t for both existing and prospective owners.

Our first approach to calculating the implicit UCR for prospective homebuyers assumes that buyers would replicate the residential sorting patterns that we observe for current owners. Specifically, we assume that the joint distribution of household and house characteristics for prospective buyers in each PUMA during 2016-2017 matches the distribution that we observe among owners who moved into that PUMAs during 2011-2016. We additionally assume that prospective buyers have initial loan-to-value ratios of 80%.³³

We find that the average prospective buyer faced an implicit UCR of 5.14% in 2016-2017. This statistic is 14% larger than the UCR realized by current owners in those years. Interestingly, the prospective buyers could have collected a 27% larger tax subsidy than current owners because a larger share of their monthly mortgage payments would have reflected deductible interest payments. However, the negative effect of this tax subsidy on the UCR is out-weighed by the positive effect of new buyers having higher loan-to-value ratios.

Figure 3: Distributions of User-Cost Rates



³³This assumption is conservative given that the average loan-to-value ratio for loans originated in 2017 was 85.6% (HMDA). For current homeowners we estimate an average current loan-to-value ratio of 57%.

Figure 3 shows the distribution of PUMA-specific mean UCRs realized by current owners in 2016-2017 as a solid line and the distribution of implicit UCRs for prospective buyers as a dotted line. The dotted line is constructed by focusing on households that recently purchased a house in the same PUMA (in the past five years) and calculating the counterfactual UCR that would apply to their house if they had originated a mortgage 2016-2017. Thus, this implicit UCR measure reflects the ways in which households sorted themselves across PUMAs by income and other demographics, as well as the ways in which they sorted over differently priced houses and neighborhoods within each PUMA.

As a second counterfactual exercise, we calculate the implicit UCR that would apply to an “average” American household in each PUMA, i.e., if there were no residential sorting. We do this by using the national distribution of households to repeatedly re-draw and re-assign households to houses at random, and thereafter recalculate their UCRs in their assigned houses.³⁴ The dashed line in Figure 3 shows the resulting UCR distribution for this “no-sorting” scenario. Without sorting, the distribution of UCRs narrows, yet the mean remains unchanged. We find that the variance shrinks by 45%, from 2.17% to 1.19% when we assign households to dwellings at random. This reveals that household characteristics such as income, as opposed to dwelling or local characteristics (e.g., property tax rates), account for a significant amount of the observed dispersion in the UCR distribution.

Finally, we illustrate the geographic variation in the determinants of the UCR by including separate maps for each of the spatially varying components ltv , rf , ω , γ , and s in Figures B4-B8 in Appendix B.4.

4.2 Housing Subsidies

Our estimates for the underlying components of the UCR imply that the average homeowner recouped 9.2% of their mortgage interest and property tax payments in 2016-2017 via reduced tax liability. This translates into an annual tax subsidy of \$1,032 for the average owner of a home worth \$300,000. To improve interpretability,

³⁴To implement this procedure, we draw 500 homeowning households and calculate what UCR they would face if they lived in each of the 3.2 million owner-occupied homes of our main sample. We then calculate the mean UCR for each home.

we calculate the percentage of annual homeownership costs that are subsidized: $(100 \cdot s_{ijt}) / (s_{ijt} + UCR_{ijt})$ where the denominator is the annual cost of owning the house in the absence of the subsidy, and the numerator is the subsidy. The advantage of using this statistic to describe subsidies is that it captures the ongoing annual subsidization as a percentage of the annual flow cost.³⁵ Depending on the context or economic question, the most relevant subsidy statistic may be $s_{ijt} \cdot P_{ijt}$ (the dollar value of a subsidy), $100 \cdot s_{ijt}$ (the percentage of the house’s value that is subsidized), or $(100 \cdot s_{ijt}) / (s_{ijt} + UCR_{ijt})$ (the percentage of annual homeownership costs that are subsidized).³⁶ Therefore, we make use of all three definitions below.

The average percentage of annual homeownership costs that are subsidized in 2016-2017 was 6.6%.³⁷ Thus, almost 7% of the average owner’s annual housing costs were subsidized by taxpayers. Decomposing this subsidy into its state and federal components reveals that 5.2% of housing costs were subsidized by the federal government, versus 1.4% by state.³⁸

Figure 4 shows the geographic heterogeneity in the percentage of annual homeownership costs that are subsidized during 2016-2017 by dividing PUMA-specific means of $(100 \cdot s_{ijt}) / (s_{ijt} + UCR_{ijt})$ into septiles. The spatial variation is striking. In the bottom septile, the average subsidy was less than 2.5% and in the top septile it exceeded 10%.

The percentage of annual costs that are subsidized is close to zero in areas where the majority of homeowners elected to take the standard deduction on their federal tax returns. This includes areas where incomes and house prices are relatively low and where homeowners have few deductible expenses. The subsidy is also lower in areas where loan-to-value ratios are lower because more households have paid off their mortgages. For example, areas with near-zero subsidy rates include low-income suburbs of Dallas, Miami, and Memphis that had mean adjusted gross

³⁵The fact that UCR_{ijt} is in the denominator highlights the importance of accurately estimating the UCR when analyzing subsidies. Without knowing the UCR, one cannot calculate the percentage of annual housing costs that is subsidized.


³⁶Given the presence of moving costs, one could also consider a subsidized definition that reflects exceptions over future subsidies, user costs, and duration: $E[\sum_{\tau=t}^{t+duration} (100 \cdot s_{ij\tau}) / (s_{ij\tau} + UCR_{ij\tau})]$.

³⁷In comparison, the subsidy rate was 9.1% for the subset of homeowners that had the option to take advantage of mortgage interest deduction because they had an active mortgage.

³⁸The corresponding figures for households with an active mortgage are 7.5% and 1.6%, respectively.

incomes under \$50,000 and mean home prices under \$150,000. In these areas, over 90% of homeowners took the standard deduction and many had paid off their mortgages as evidenced by mean LTVs around 30%. Subsidies also tend to be lower in states with no personal income tax. Indeed, 48 of the 50 PUMAs with the lowest mean subsidies were in three states with no personal income tax: Texas, Florida, and Tennessee.

Figure 4: Mean Housing Subsidies by PUMA in 2016-2017



Note: The figure shows the PUMA mean subsidy, defined as the percentage of the user cost that is subsidized by the federal and state governments in 2016-2017. The figure is winsorized at the 1st and 99th percentiles.

By contrast, the percentage of annual costs that are subsidized exceeds 10% in areas where incomes, taxes, and itemization rates are relatively high. In these PUMAs, mortgages are generally larger, with mean house prices often exceeding \$500,000. Many of the high-subsidy PUMAs are located in affluent cities in Western states such as California and Oregon, but also in rural parts of the Midwest (especially Minnesota and Wisconsin), and high-tax jurisdictions along the Eastern Seaboard.

Finally, we note that we choose to focus on estimating actual subsidies received. This measure combines the features of the tax code (that households take as given) with the endogenous decisions that households make that determine deduction amounts. Alternatively, one could calculate hypothetical subsidies for fixed

deduction amounts.³⁹

5 Policy Analysis: Distributional Impacts of the Tax Cuts and Jobs Act

As discussed in Section 3.1, the TCJA reduced the incentive for households to collect a tax subsidy for homeownership starting in 2018. It did this partly by reducing the size of the subsidy that could be collected by itemizers and partly by doubling the standard deduction collected by non-itemizers. Homeowners who switched from itemizing to taking the standard deduction lost their subsidies. In addition, those who continued to itemize saw the nominal value of their subsidies decline. Moreover, the real value of their subsidies further declined by the increase in the standard deduction, as increasing the standard deduction increased the portion of the household’s total deduction that was neutral to their homeowner status.⁴⁰

The TCJA tax provisions affecting households are set to expire in 2025. In the meantime, Congress must choose whether to let the tax provisions expire, to extend them, or to extend some provisions but not others. The most controversial provision is perhaps the SALT cap that limits the amount of state and local taxes that can be deducted. There have been several proposals to undo the SALT cap. For example, an early and unsuccessful version of the federal Build Back Better Act included a proposal to increase the cap from \$10,000 to \$80,000. Further, the states of New York, Connecticut, Maryland, and New Jersey unsuccessfully challenged the cap’s constitutionality in federal court.⁴¹

Motivated by this background, we begin by measuring how the TCJA changed subsidy rates and UCRs between 2016-2017 and 2018-2019. We then simulate the counterfactual effects of eliminating all of the TCJA’s individual tax provisions and of eliminating just the SALT cap. We focus on two outcomes: the subsidy rate and the UCR. Our analysis of the TCJA complements prior studies that estimated how

³⁹See Section 4.1 for an analysis of UCR under counterfactual populations.

⁴⁰We refer to these itemizable deductions below the standard deduction as “wasted deductions” following Follain and Ling (1991).

⁴¹In April 2022, the Supreme Court denied certiorari, letting stand the Second Circuit Court’s ruling that the SALT cap was constitutional.

the TCJA affected housing market outcomes without modeling the changes to tax subsidies and UCRs that may have driven those effects.

5.1 Estimation of the Impact of Policy Changes

Using our framework, one could estimate the impact of a policy by comparing the pre- and post-policy outcomes. However, similar to other policy analyses, it is often the goal to recover the causal impact of the policy by controlling for any confounding effects that would have occurred in the absence of the policy. Thus, a preferable estimate is to compare post-policy outcomes with predicted counterfactual outcomes in the post-period assuming that the policy had not taken place. Before presenting the results of this analysis in subsections 5.2-5.6, it is worth noting which behavioral and general equilibrium responses are incorporated.⁴²

As the estimated post-policy subsidy rate and post-policy UCR are calculated using actual post-policy data, any behavioral responses by households should be captured. For example, if the policy were to cause changes to itemization rates, mobility, loan-to-value ratios (ltv_{ijt}), purchase prices, or homeownership rates, this would be captured in the post-policy rates that we recover because we observe the post-policy distribution of household decisions/outcomes for these variables. Furthermore, if the policy were to affect the determined-in-equilibrium, time-varying inputs in the UCR formula (i.e., rf_{ijt} , rm_t , ω_{jt} , γ_{jt}), these equilibrium effects would be captured to the extent that we capture any time-varying changes in these factors.⁴³ Finally, some attributes are assumed to be time-invariant (i.e., δ , ϵ) and do not vary in our analysis.

In our primary counterfactual measures, we implement the pre-policy tax code (while adjusting for inflation and bracket creep) and we have a model of itemization to counterfactually predict what would have happened to itemization rates. Although we are addressing these features (which are arguably the most important features), we do not use a counterfactual model to predict what would have happened to ltv_{ijt} ,

⁴²Gervais (2002) estimates the impact of the preferential tax treatment of housing using a dynamic general equilibrium model that features heterogeneous agents but without geographically varying determinants of subsidies.

⁴³See Section 2.2 and Appendix A for detailed descriptions of how we estimate ltv_{ijt} , rf_{ijt} , rm_t , ω_{jt} , and γ_{jt} , including how the estimates of rf_{ijt} and rm_t incorporate temporal smoothing.

rf_{ijt} , rm_t , ω_{jt} , and γ_{jt} , and simply use their post-policy actual distributions. As a robustness exercise, we construct alternative counterfactual values of subsidies and UCR, which address these behavioral/equilibrium changes, but which require an additional assumption. Under this additional assumption, that the pre-policy trends in these rates would have continued absent a policy, one can use the pre-policy trends in outcomes to construct post-policy counterfactual outcomes.

5.2 The TCJA’s Effects on Housing Subsidies and UCRs

Repeating the analysis from Section 4 for the years after the TCJA went into effect reveals that, in 2018 and 2019, the average homeowner recouped 3.1% of their mortgage interest and property-tax payments via reduced tax liability, compared to 9.2% during 2016-2017. As a result, the average percentage of annual homeownership costs that are subsidized dropped from 6.6% to 2.1%. Importantly, the TCJA’s passage coincided with macroeconomic trends that modified other inputs to the UCR formula apart from the tax subsidy. We find that the national average UCR decreased by approximately 5 basis points (or 1.1%) because of changes to UCR components other than s_{ijt} , especially falling interest rates.⁴⁴

To provide a better measure of the causal impact of the reform, we construct counterfactual measures for the percentage of annual costs that are subsidized and the UCR that would have been realized in 2018-2019 had the 2017 tax code remained in place during those years amid realized income growth.⁴⁵ Figures 5 and 6 compare these “No-TCJA” counterfactual measures (dashed lines) with the actual measures (solid lines). We estimate the causal impact of the TCJA by the differences between the actual and counterfactual measures.

Figure 5 shows that had the 2017 tax code remained in place in 2018-2019,

⁴⁴The result that both mean subsidies and UCRs were lower in 2018-19 than in 2016-17 is due to the rapidly falling interest rates over the period.

⁴⁵To account for the fact that tax brackets and standard-deduction amounts are annually adjusted for inflation by the IRS, we use constant 2017 dollars in all years. However, even with this inflation adjustment, and in the absence of changes to the tax code, real growth in incomes increases subsidies through higher marginal tax rates due to the progressivity of the tax code, as well as through increased deductible expenses which tend to grow with income. Just as subsidies (and real GDP) grew by approximately 2% from 2016 to 2017, our counterfactual analysis indicates that this growth in subsidies would have continued if the 2017 tax code had remained in place and the TCJA not been implemented.

subsidies would have risen to offset 6.9% of annual homeownership cost in comparison to the realized rate of 2.1%. This shows that the causal effect is slightly larger than the raw time-difference would indicate. Thus, the TCJA caused the tax subsidy to homeownership to fall by 70% during the first two years after the TCJA’s implementation.⁴⁶ This large reduction in subsidies is particularly notable given the very limited historical changes in subsidies documented in Sinai and Gyourko (2014). Figure 6 shows that the TCJA-induced reduction in subsidies increased the average UCR by 0.25 percentage points (a 6% increase).⁴⁷ Using our alternative counterfactual measures of subsidies and UCR would lead to virtually identical estimates of the impact of the TCJA on subsidies (trivially smaller) and bigger estimates of the impact of the policy on UCR (a 9% increase versus 6% using our primary approach).

The dotted lines in Figures 5 and 6 show results from a second counterfactual simulation that eliminates the SALT cap but leaves all other TCJA provisions in place.⁴⁸ Comparing the dotted and solid lines shows that eliminating just the SALT cap would have relatively modest effects on mean subsidies and UCRs. For example, we estimate that eliminating the SALT cap would increase the average post-TCJA subsidy from 2.1% to 3.1%. This coincides with a reduction in the UCR of only 0.04 percentage points. However, we show below that these small changes to mean subsidies and UCRs reflect large transfers to a small share of homeowners.

⁴⁶While the TCJA was a federal policy, integrating over filing decisions in TAXSIM implies that the TCJA also caused the state-level portion of the tax-subsidy to fall from 1.4% to 1.1% by reducing the incentive to itemize. If we focus solely on the federal part of the subsidy, the reduction was even larger. We estimate that the TCJA reduced federal subsidies by 82%, from 5.5% to 1.0% of the annual cost of homeownership.

⁴⁷The relatively substantial reduction in UCR that is predicted in the absence of the TCJA is driven by falling mortgage and risk-free rates during this period. Section 2.2 documents how we calculate these measures, which include temporal smoothing to address issues like refinancing, and Table A.3 shows their values by year. In the absence of this smoothing, the decline would be larger.

⁴⁸We also used a novel version of NBER’s TAXSIM software that allows for hypothetical tax environments to compare a total elimination of the cap with an increase in the cap to \$80,000 based on an early version of the Build Back Better Act, H.R.5376, that passed the House of Representatives but was not voted on in the senate. Our results show that for over 99.9% of homeowners, these two policies would have identical effects on their homeownership subsidy and only around 60,000 homeowning households per year would additionally benefit from raising the cap above \$80,000.

Figure 5: Average Subsidy from 2016 to 2019






Figure 6: Average User Cost Rate from 2016 to 2019



5.3 Heterogeneity by Geography

Figure 7 summarizes geographic heterogeneity in the TCJA’s effect on housing subsidies in 2018-2019 by mapping the percentage-point reduction in the subsidy rate across PUMAs (i.e., the difference between the solid and dashed lines in Figure 5).⁴⁹ PUMAs that received larger subsidies prior to the TCJA (shown in Figure 4) generally saw larger reductions. Further, many areas lost almost the entirety of their pre-TCJA subsidies. In 10% of PUMAs, mean subsidies fell by over 90%.

Figure 7: Geographic Distribution of the Impact of TCJA on Subsidies



Note: The figure shows the percentage point reduction in the subsidy rate between 2016-2017 and 2018-2019 that we attribute to the Tax Cuts and Jobs Act. The figure is winsorized at the 1st and 99th percentiles.


The geographic heterogeneity in Figure 7 arises from interactions between several underlying factors. These include spatial variation in inputs to the UCR formula, spatial variation in household income and other taxpayer characteristics, interactions between federal and state tax codes, and how all of these factors interact to determine households’ tax-minimizing filing strategies. Appendix B.1 explains how federal and state tax codes interact.

Next, we turn to the heterogeneous effects of eliminating the SALT cap while leaving all other TCJA provisions in place. We estimate that only 15% of homeowners would benefit from this counterfactual policy. However, these beneficiaries would

⁴⁹Figure B.3 shows the disparate geographic impact of the TCJA on UCRs.

have large gains, with the average beneficiary experiencing a 6.6 percentage point increase in their homeownership subsidy. This represents a 24% increase in the subsidy level among those who receive subsidies.

Figure 8: Percent of Lost Subsidies that Would Be Regained if the SALT Cap Were Eliminated



Note: The figure shows PUMA-specific mean shares of the lost tax subsidy to homeownership that would have been returned to homeowners had there been no cap on the amount of deductible state and local taxes in 2018 and 2019. The figure is winsorized at the 1st and 99th percentiles.

Another way to measure the heterogeneous effects of eliminating the SALT cap is to calculate how much of the subsidy that was removed by the TCJA (Figure 7) would be returned if the SALT cap were eliminated. Figure 8 shows the geographic distribution of this measure. Three features stand out. First, the variation across PUMAs is substantial. PUMAs in the bottom three septiles regain less than 10% of their lost subsidies, whereas the top septile regains more than 30% of their lost subsidies. Second, PUMAs that see no effect are predominantly located in states that do not have income taxes (i.e., Florida, Nevada, South Dakota, Tennessee, Texas, Washington, and Wyoming). More broadly, the effects are smaller in lower-income and lower-tax areas. Finally, most states include some PUMAs that experience relatively large impacts.

5.4 Heterogeneity by Tax-Filing Behavior

To interpret the heterogeneity in Figure 7, and to explore the underlying mechanisms, we calculate the TCJA’s effects for three groups of taxpayer homeowners. First, we consider “always-itemizers” who would minimize their taxes by itemizing under both the pre-TCJA and post-TCJA tax codes. These are generally higher-income households and comprise 17% of homeowners. These households lose a portion of their housing subsidy, largely due to the increase in “wasted deductions.”

Second, we consider “switchers” who would minimize their taxes by itemizing under the pre-TCJA taxcode and by taking the standard deduction under the post-TCJA taxcode. These households comprise 31% of homeowners. While these households lose the entirety of their homeowner tax subsidy, the vast majority (91%) pay equal or lower taxes due to the increased standard deduction.

Finally, we consider “never itemizers” who would minimize their tax burden by taking the standard deduction under both the pre-TCJA and post-TCJA tax codes. These households comprise 51% of homeowners. These households see no subsidies under either tax code, but pay strictly lower taxes under the post-TCJA tax code due to the increased standard deduction. Although, we focus on homeowners, the vast majority of renters are effectively similar to never-itemizing homeowners – they received no housing subsidy before or after the TCJA, but pay strictly lower taxes under the post-TCJA tax code due to the increased standard deduction.

Table 1 summarizes the TCJA’s impact on housing subsidies for each group.⁵⁰ The always-itemizers see the biggest absolute reduction in their subsidy rate. However, this group also receives the largest subsidies under the pre-TCJA tax code. As a percentage change, however, the switchers lose more as they lose 100% of their federal subsidy.

⁵⁰There is additionally a fourth group who would minimize their tax burden by taking the standard deduction under the pre-TCJA tax code and by itemizing under the post-TCJA tax code. These households comprise only a very small group of homeowners (0.04%) who were affected by the TCJA’s changes to the cutoffs between certain tax brackets. Our analysis includes this group as well, although they are omitted from this discussion here for expositional purposes.

Table 1: 2018-2019 Mean Housing Subsidies by Itemization Status

Homeowner Type	% of Homeowners	Housing Subsidy (%)					
		Without TCJA			With TCJA		
		Federal	State	Total	Federal	State	Total
Never itemizers	51	0.0	0.9	1.0	0.0	1.0	1.0
Switchers	31	7.8	1.4	9.2	0.0	0.7	0.7
Always itemizers	17	17.8	3.0	20.8	5.9	2.2	8.1
All homeowners	100	5.5	1.4	6.9	1.0	1.1	2.1

Note: This table shows the 2018 to 2019 mean share of the annual cost of homeownership that is subsidized by federal and state governments for three types of homeowners distinguished by the impact that the implementation of the TCJA had on their itemization status. “Always itemizers” itemize under both the pre-TCJA and post-TCJA tax codes, “Switchers” are induced by the implementation of the TCJA to switch from itemizing their deductions to taking the standard deduction, and “Never itemizers” take the standard deduction under both tax codes. Subsidies are shown for these types under two alternative tax codes: first, a “Without TCJA” counterfactual that maintains the 2017 pre-TCJA tax code, and thereafter actual subsidies under the TCJA, as it was implemented.

5.5 Heterogeneity by Voting Behavior

Previous studies found that the TCJA benefited households in states decisively won by Donald Trump in the 2016 presidential election more than it benefited households in states won by Hillary Clinton. For example, Altig et al. (2020) finds that the TCJA increased the remaining lifetime spending of households in “red” states by 1.6% as opposed to only 1.3% in “blue” states. In a similar vein, we explore how the TCJA’s impact on tax subsidies to homeowners differed between red and blue counties, and how these areas would be differentially affected by the expiration of the TCJA’s individual tax provisions.

We merge our PUMA-level average UCR and subsidy measures with county-level data from the MIT Election Data and Science Lab describing the results of the 2016 presidential election. As PUMAs are designed to have roughly equal populations of around 100,000, their geographic size varies inversely with density. This means that in urban areas, counties often contain multiple PUMAs, whereas in rural areas, a single PUMA can span several counties. We therefore use the crosswalk from Bieri et al. (2022) to merge the datasets at the finest possible spatial resolution. This results in the aggregation of 2,351 PUMAs and 3,143 counties into 982 locations. Of these,

430 are metropolitan counties (aggregations of PUMAs) for which election results are directly available. In rural areas, we have 459 locations where a single PUMA contains multiple counties. There we calculate the election result by aggregating the votes cast in each constituent county. Finally, a relatively small number of PUMAs encompass parts of multiple counties. We merge all adjacent counties in such cases to create larger PUMA-county unions, for which we can calculate vote shares. There are 93 such unions.

Once households are linked to the election result in their area, we can calculate how various tax regimes would affect Republican- vs. Democratic-voting areas. We define an area as “Republican” if more votes were cast for Donald Trump than for Hillary Clinton in the 2016 election, and “Democratic” otherwise. Table 2 reports the subsidy obtained by the average household in 2018 and 2019 in each type of area, both as a dollar amount and as the percentage of the annual cost of homeownership that is subsidized. Since a constant change to subsidies will have larger impacts in areas with more homeowners, we include all households, both owners and renters. Renters do not receive any subsidy to homeownership, so they are given a value of zero in these tabulations.

Mean subsidies are calculated under three tax regimes. First, as described in Section 5.2, we calculate what subsidies would have been in 2018 and 2019 had the TCJA not been implemented. Without the TCJA, 4.7% of the annual cost of homeownership would have been subsidized for the average household in Democratic areas, compared with 3.4% for the average household in Republican areas. In dollar terms, this translates to annual subsidies of \$980 and \$480 respectively.

Second, we calculate the actual subsidies in 2018 and 2019 under the TCJA as implemented. The TCJA caused the subsidy gap between Republican and Democratic areas to shrink significantly; the average household in Democratic areas lost more subsidy dollars. In these areas, the average household lost \$626 (a 64% reduction) of subsidies per year due to the TCJA, while the average household in Republican areas lost \$344. We note, however, that, given their lower pre-TCJA subsidy amounts, households in Republican areas experienced a larger percentage loss compared with their 2016 and 2017 baseline subsidies (72%).

Finally, we consider the counterfactual tax regime where the SALT cap were

Table 2: Mean Subsidies by Election Result under Alternative Tax Regimes

	Tax Regime		
	No TCJA	TCJA	SALT Cap Elimination
<u>Dollar Amount of Subsidy</u>			
Republican	\$480	\$136	\$207
Democrat	\$980	\$354	\$609
All	\$748	\$253	\$423
<u>% of Annual Costs Subsidized</u>			
Republican	3.45%	1.01%	1.32%
Democrat	4.72%	1.51%	2.32%
All	4.13%	1.27%	1.86%

Note: This table shows mean subsidies to homeownership from federal and state governments in dollar amounts and as the share of the annual cost of homeownership that is subsidized under alternative tax regimes. Means are calculated over all households for 2018-2019 and include non-homeowners who do not receive any subsidy. The first column (No TCJA) represents a counterfactual tax regime where the 2017 pre-TCJA tax code stays in place. The second column (TCJA) shows subsidies under the actual tax code in place during 2018-2019 and the third column (SALT Cap Elimination) shows another counterfactual tax regime with subsidies that would have been received in 2018-2019 if all provisions of the TCJA except the SALT Cap were in place. Republican and Democratic areas are defined based on results from the 2016 presidential election at the finest possible spatial resolution (either PUMAs, counties, or their union; see main text for details).

eliminated, but all TCJA provisions are retained. This regime would almost double subsidies in Democratic areas while generating a more modest increase in Republican ones. Elimination of the SALT cap would thus represent a significant increase in subsidies to homeowners in Democrat-voting areas.

5.6 Heterogeneity by Race of the Household Head

TCJA provisions do not vary directly with race or ethnicity, but their correlation with income, geography, and homeownership may cause the TCJA provisions to have impacts that differ systematically across racial or ethnic groups. Understanding these distributive impacts may help to advance research on racial segregation in housing markets (Aliprantis et al. 2022, Davis et al. 2023) and on racial and ethnic gaps in economic outcomes (Banzhaf et al. 2019, Brouillette et al. 2022, Akbar et al. forthcoming). It may also be a requirement for future regulatory analyses (Biden

2023, US Office of Management and Budget 2023, Cronin et al. 2023).

With this in mind, Table 3 summarizes how the effects of each tax regime vary by the self-reported race of each ACS household head. We aggregate the ACS data on race into four categories. The Asian category combines responses that indicate Chinese, Japanese, or “Other Asian” ancestry. The Other category combines responses that indicate “American Indian or Alaska Native,” “Other race,” or multiple races. All four categories contain households who indicate Hispanic, Spanish, or Latino origin.

Panel A of Table 3 shows that the average Asian household receives a homeownership subsidy that is more than twice as large as the average White household under all three tax regimes. The subsidy to the average White household is again twice as large as the subsidy to the average Black household.

Table 3: Mean Subsidies by Race under Alternative Tax Regimes

	Tax Regime		
	No TCJA	TCJA	SALT Cap Elimination
<u>A. Dollar Amount of Subsidy</u>			
White	\$786	\$258	\$439
Black	\$346	\$117	\$162
Asian	\$1,545	\$622	\$423
Other	\$470	\$174	\$250
<u>B. % of Annual Costs Subsidized</u>			
White	4.44%	1.34%	1.97%
Black	2.32%	0.77%	0.95%
Asian	6.03%	2.11%	3.46%
Other	2.54%	0.87%	1.15%

Note: This table shows mean subsidies to homeownership from federal and state governments in both dollar amounts and as the share of the annual cost of homeownership that is subsidized under alternative tax regimes. Means are calculated over all households for the period 2018-2019 and include non-homeowners who get no subsidy. The first column (No TCJA) represents a counterfactual tax regime where the 2017 pre-TCJA tax code stays in place. The second column (TCJA) shows subsidies under the actual tax code in place during 2018-2019 and the third column (SALT Cap Elimination) shows another counterfactual tax regime with subsidies that would have been received in 2018-2019, if all provisions of the TCJA except the SALT Cap were in place. Race is self-reported for the head of each household.

Panel B shows the subsidy as the share of the annual cost of homeownership.

If the differences in Panel A were driven entirely by differences in house prices, then the shares in Panel B would be equal. The fact that they are not reveals that the racial disparities in housing subsidies are also driven by factors such as geographic location, income, and homeownership rates. Geographic location is the main driver of the large subsidies to Asian households. Homeownership rates are higher among White households than Asian ones (64% versus 58%), but Asian homeowners tend to have higher incomes and are more likely to live in large urban areas of high-tax, and thus high-subsidy, states such as California, New York, and New Jersey. Consequently, under the TCJA, Asian homeowners receive 12% of all subsidies to homeownership, while comprising only 5% of homeowners.

By contrast, Black households' relatively low subsidies are driven by their lower rate of homeownership. Though Black households are more likely to live in Southern states that have lower tax rates and lower subsidies, the main reason why the average White household receives a subsidy that is two times larger is because only 39% of Black households own their homes compared with 65% of White households. Focusing only on homeowners, the White-Black gap is much smaller with the average White homeowner receiving \$1,215 in subsidies compared with \$891 for the average Black homeowner under the "No TCJA" tax regime. Under the TCJA regime, this gap shrinks even further with subsidies of \$398 and \$302 for the average White and Black homeowners, respectively.

Comparing results across the three tax regimes reveals interesting distributional effects. The average household of each group lost similar fractions (60-70%) of their subsidies due to the TCJA. However, eliminating the SALT cap while maintaining the other TCJA provisions would have a more disparate racial impact. White and Asian households would see a 70% increase in their subsidies, whereas Black households would see only a 38% increase.⁵¹

⁵¹Since non-homeowners are unaffected and receive zero subsidies under all three tax regimes, these percentage changes, and thus the conclusions about disparate racial impacts, remain the same if we focus only on homeowners.

6 Conclusion

The real economic cost of homeownership is hard to measure due to non-linearity in the US tax code. We incorporate this non-linearity into a model of optimal tax-filing behavior and use rich micro-data to build and validate a novel database of user-cost rates and tax subsidies to homeowners across the US from 2016 to 2019. PUMA-by-year means of our estimates can be explored using interactive maps or downloaded at www.housingusercost.org. In addition, our use of recurrent publicly-available data makes it straightforward to update and extend our results.

It is important to develop accurate measures for user-cost rates because these rates are a key input to estimating the demand for housing, as well as the demand for any local public good or amenity that is capitalized into housing prices. Accurate measures for user-cost rates are also needed to evaluate the distributional effects of policies that affect housing markets. We demonstrate this by using our estimates to show how federal tax subsidies to homeownership disproportionately benefit certain demographic groups. On the extensive margin, renters are excluded. On the intensive margin, the subsidies are larger for households that face higher marginal income tax rates, own more expensive houses, and live in higher property-tax areas. These distortions are a significant source of variation in the real economic cost of housing and correlate with voting behavior, income, and race.

We also use our estimates to show that the TCJA reduced the mean subsidy rate to homeownership by 70% starting in 2018. The largest reductions occurred in Democrat-voting, affluent areas of coastal states that received the largest subsidies before the TCJA. Asian and White households also saw larger reductions, on average, than Black households. Further, we show that the TCJA increased the user-cost of homeownership disproportionately for new homebuyers.

Many of the TCJA's provisions, including the controversial cap on SALT deductions, are set to expire in 2025. We show that eliminating the SALT cap would have minimal impacts on the average subsidy to homeownership, with strongly heterogeneous effects. The vast majority of the benefits would accrue to homeowners in Democrat-voting areas and to Asian and White households.

Since we maintain a sharp focus on the TCJA's controversial tax provisions for

homeownership, our findings do not characterize the distributive welfare implications of the entire TCJA, or federal housing policy in general. Rather, our analysis serves to demonstrate how our estimates of heterogeneous user-cost rates and subsidies can help to provide sharper answers to numerous economic questions. Future studies can employ our framework to analyze any tax policy that impacts housing costs or subsidies. It can also inform research on geographic inequality in housing costs and living standards. Finally, combining our user-cost estimates with housing-demand elasticities could help to identify how tax and housing policies affect home-purchase decisions and the long-run accumulation of wealth by income and race (Akbar et al. forthcoming).

References

- Ahlfeldt, G. M., Heblich, S., and Seidel, T. (2023). Micro-geographic property price and rent indices. *Regional Science and Urban Economics*, page 103836.
- Akbar, P. A., Hickly, S. L., Shertzer, A., and Walsh, R. P. (forthcoming). Racial segregation in housing markets and the erosion of black wealth. *The Review of Economics and Statistics*.
- Albouy, D. (2009). The unequal geographic burden of federal taxation. *Journal of Political Economy*, 117(4):635–667.
- Albouy, D. (2016). What are cities worth? Land rents, local productivity, and the total value of amenities. *Review of Economics and Statistics*, 98(3):477–487.
- Albouy, D. and Hanson, A. (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(1):63–96.
- Aliprantis, D., Carroll, D. R., and Young, E. R. (2022). What explains neighborhood sorting by income and race? *Journal of Urban Economics*, 103508.
- Altig, D., Auerbach, A., Higgins, P., Koehler, D., Kotlikoff, L., Terry, E., and Ye, V. (2020). Did the 2017 tax reform discriminate against blue-state voters? *National Tax Journal*, 73(4):1087–1108.
- Ambrose, B. W., Hendershott, P. H., Ling, D. C., and McGill, G. A. (2022). Homeownership and taxes: How the TCJA altered the tax code’s treatment of housing. *Real Estate Economics*, 50(5):1167–1200.

- Amior, M. and Halket, J. (2014). Do households use home-ownership to insure themselves? evidence across us cities. *Quantitative Economics*, 5(3):631–674.
- Andrews, D., Caldera Sánchez, A., and Johansson (2011). Housing markets and structural policies in OECD countries.
- Anenberg, E. and Ringo, D. (2022). Volatility in home sales and prices: Supply or demand? *FEDS Working Paper*.
- Armona, L., Fuster, A., and Zafar, B. (2019). Home price expectations and behaviour: Evidence from a randomized information experiment. *The Review of Economic Studies*, 86(4):1371–1410.
- Bajari, P. and Kahn, M. E. (2005). Estimating housing demand with an application to explaining racial segregation in cities. *Journal of Business and Economic Statistics*, 23(1):20–33.
- Banzhaf, H. S. and Walsh, R. P. (2008). Do people vote with their feet? An empirical test of tiebout. *American Economic Review*, 98(3):843–63.
- Banzhaf, S., Ma, L., and Timmins, C. (2019). Environmental justice: The economics of race, place, and pollution. *Journal of Economic Perspectives*, 33(1):185–208.
- Bartolomé, C. A. d. and Rosenthal, S. S. (1999). Property tax capitalization in a model with tax-deferred assets, standard deductions, and the taxation of nominal interest. *Review of Economics and Statistics*, 81(1):85–95.
- Bayer, P., Ferreira, F., and McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4):588–638.
- Bayer, P., Keohane, N., and Timmins, C. (2009). Migration and hedonic valuation: The case of air quality. *Journal of Environmental Economics and Management*, 58(1):1–14.
- Bayer, P., McMillan, R., Murphy, A., and Timmins, C. (2016). A dynamic model of demand for houses and neighborhoods. *Econometrica*, 84(3):893–942.
- Benzarti, Y. (2020). How taxing is tax filing? using revealed preferences to estimate compliance costs. *American Economic Journal: Economic Policy*, 12(4):38–57.
- Bhutta, N., Fuster, A., and Hizmo, A. (2020). Paying too much? Price dispersion in the us mortgage market. *FEDS Working Paper*.
- Bibler, A. and Billings, S. B. (2020). Win or lose: Residential sorting after a school choice lottery. *Review of Economics and Statistics*, 102(3):457–472.

- Biden, J. R. (2023). Executive order 14094: Modernizing regulatory review. *Federal Registrar*, 88(69):21879–21881.
- Bieri, D. S., Kuminoff, N. V., and Pope, J. C. (2023). National expenditures on local amenities. *Journal of Environmental Economics and Management*, page 102717.
- Binner, A. and Day, B. (2015). Exploring mortgage interest deduction reforms: An equilibrium sorting model with endogenous tenure choice. *Journal of Public Economics*, 122:40–54.
- Bishop, K. C. and Murphy, A. D. (2019). Valuing time-varying attributes using the hedonic model: When is a dynamic approach necessary? *Review of Economics and Statistics*, 101(1):134–145.
- Blouri, Y., Büchler, S., and Schöni, O. (2021). The geography of housing subsidies. *MIT Center for Real Estate Research Paper*, (21/05).
- Brouillette, J.-F., Jones, C. I., and Klenow, P. J. (2022). Race and economic well-being in the united states. *Working Paper*.
- Brueckner, J. K. (2014). Eliminate the mortgage interest deduction or tax imputed rent? leveling the real-estate playing field. *Cityscape*, 16(1):215–218.
- Cabral, M. and Hoxby, C. (2012). The hated property tax: Salience, tax rates, and tax revolts. *NBER WP 18514*.
- Caetano, G. and Maheshri, V. (2019). A unified empirical framework to study segregation. *Working paper*.
- Campbell, S. D., Davis, M. A., Gallin, J., and Martin, R. F. (2009). What moves housing markets: A variance decomposition of the rent–price ratio. *Journal of Urban Economics*, 66(2):90–102.
- Coen-Pirani, D. and Sieg, H. (2019). The impact of the tax cut and jobs act on the spatial distribution of high productivity households and economic welfare. *Journal of Monetary Economics*, 105:44–71.
- Coulson, N. E. and Li, H. (2013). Measuring the external benefits of homeownership. *Journal of Urban Economics*, 77:57–67.
- Cronin, J.-A., DeFilippes, P., and Fisher, R. (2023). Tax expenditures by race and hispanic ethnicity: An application of the us treasury department’s race and hispanic ethnicity imputation. *Working Paper*.
- Davis, M. (2019). The distributional impact of mortgage interest subsidies: Evidence from variation in state tax policies. *working paper*.

- Davis, M. A., Gregory, J., and Hartley, D. A. (2023). Preferences over the racial composition of neighborhoods: Estimates and implications. *Working Paper*.
- DeFusco, A. A. and Paciorek, A. (2017). The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit. *American Economic Journal: Economic Policy*, 9(1):210–40.
- Diamond, R. (2016). The determinants and welfare implications of us workers’ diverging location choices by skill: 1980-2000. *American Economic Review*, 106(3):479–524.
- Emrath, P. (2002). Property taxes in the 2000 census. *Housing Economics*, 50(12):16–21.
- Epple, D., Quintero, L., and Sieg, H. (2020). A new approach to estimating equilibrium models for metropolitan housing markets. *Journal of Political Economy*, 128(3):948–983.
- Epple, D. and Sieg, H. (1999). Estimating equilibrium models of local jurisdictions. *Journal of political economy*, 107(4):645–681.
- Feenberg, D. and Coutts, E. (1993). An introduction to the TAXSIM model. *Journal of Policy Analysis and management*, 12(1):189–194.
- Fisher, L. M., Fratantoni, M., Oliner, S. D., and Peter, T. J. (2021). Jumbo rates below conforming rates: When did this happen and why? *Real Estate Economics*, 49(S2):461–489.
- Flavin, M. and Yamashita, T. (2002). Owner-occupied housing and the composition of the household portfolio. *American Economic Review*, 92(1):345–362.
- Follain, J. R. and Ling, D. C. (1991). The federal tax subsidy to housing and the reduced value of the mortgage interest deduction. *National Tax Journal*, 44(2):147–168.
- Foote, C. L., Loewenstein, L., and Willen, P. S. (2021). Cross-sectional patterns of mortgage debt during the housing boom: evidence and implications. *The Review of Economic Studies*, 88(1):229–259.
- Fuster, A. and Zafar, B. (2021). The sensitivity of housing demand to financing conditions: evidence from a survey. *American Economic Journal: Economic Policy*, 13(1):231–65.
- Garriga, C., Manuelli, R., and Peralta-Alva, A. (2019). A macroeconomic model of price swings in the housing market. *American Economic Review*, 109(6):2036–72.

- Gervais, M. (2002). Housing taxation and capital accumulation. *Journal of Monetary Economics*, 49(7):1461–1489.
- Gindelsky, M., Moulton, J. G., and Wentland, S. A. (2019). Valuing housing services in the era of big data: A user cost approach leveraging zillow microdata. *Big Data for 21st Century Economic Statistics*. Univeristy of Chicago Press.
- Glaeser, E. (2012). Triumph of the city: How our greatest invention makes us richer, smarter, greener, healthier, and happier. *Penguin*.
- Glaeser, E. L. (2013). A nation of gamblers: Real estate speculation and american history. *American Economic Review*, 103(3):1–42.
- Glaeser, E. L. and Shapiro, J. M. (2003). The benefits of the home mortgage interest deduction. *Tax policy and the economy*, 17:37–82.
- Gruber, J., Jensen, A., and Kleven, H. (2021). Do people respond to the mortgage interest deduction? Quasi-experimental evidence from denmark. *American Economic Journal: Economic Policy*, 13(2):273–303.
- Gyourko, J. and Sinai, T. (2003). The spatial distribution of housing-related ordinary income tax benefits. *Real Estate Economics*, 31(4):527–575.
- Halket, J., Nesheim, L., and Oswald, F. (2020). The housing stock, housing prices, and user costs: The roles of location, structure, and unobserved quality. *International Economic Review*, 61(4):1777–1814.
- Han, B., Han, L., and Zhu, G. (2018). Housing price and fundamentals in a transition economy: The case of the beijing market. *International Economic Review*, 59(3):1653–1677.
- Han, L., Ngai, L. R., and Sheedy, K. D. (2022). To own or to rent? the effects of transaction taxes on housing markets. *discussion paper 17520, CEPR*. 3.
- Harding, J. P., Rosenthal, S. S., and Sirmans, C. (2007). Depreciation of housing capital, maintenance, and house price inflation: Estimates from a repeat sales model. *Journal of Urban Economics*, 61(2):193–217.
- Head, A., Lloyd-Ellis, H., and Stacey, D. (2023). Heterogeneity, frictional assignment, and home-ownership. *International Economic Review*.
- Hembre, E. and Dantas, R. (2023). Tax incentives and housing decisions: Effects of the tax cut and jobs act. *Regional Science and Urban Economics*, 95:103800.
- Hilber, C. A. and Turner, T. M. (2014). The mortgage interest deduction and its impact on homeownership decisions. *Review of Economics and statistics*, 96(4):618–637.

- Himmelberg, C., Mayer, C., and Sinai, T. (2005). Assessing high house prices: Bubbles, fundamentals and misperceptions. *Journal of Economic Perspectives*, 19(4):67–92.
- Joint Committee on Taxation (2017). Estimates of federal tax expenditures for fiscal years 2016–2020. Congress of the United States Washington, DC.
- Keys, B. J., Pope, D. G., and Pope, J. C. (2016). Failure to refinance. *Journal of Financial Economics*, 122(3):482–499.
- Knoll, K., Schularick, M., and Steger, T. (2017). No price like home: Global house prices, 1870-2012. *American Economic Review*, 107(2):331–53.
- Kuminoff, N. V., Smith, V. K., and Timmins, C. (2013). The new economics of equilibrium sorting and policy evaluation using housing markets. *Journal of Economic Literature*, 51(4):1007–62.
- Li, W. and Edison, G. Y. (2022). Real estate taxes and home value: Evidence from TCJA. *Review of Economic Dynamics*, 43:125–151.
- Ma, L. (2019). Learning in a hedonic framework: Valuing brownfield remediation. *International Economic Review*, 60(3):1355–1387.
- Mangum, K. (2017). The role of housing in carbon emissions. *Andrew Young School of Policy Studies Research Paper Series*, (17-05).
- Martin, H. (2018). The impact of the tax cuts and jobs act on local home values. *FEB of Cleveland Working Paper*.
- Martin, H. and Hanson, A. (2016). Metropolitan area home prices and the mortgage interest deduction: Estimates and simulations from policy change. *Regional Science and Urban Economics*, 59:12–23.
- Ouazad, A. and Rancière, R. (2019). City equilibrium with borrowing constraints: Structural estimation and general equilibrium effects. *International Economic Review*, 60(2):721–749.
- Poterba, J. and Sinai, T. (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. (1984). Tax subsidies to owner-occupied housing: an asset-market approach. *The Quarterly Journal of Economics*, 99(4):729–752.
- Poterba, J. M. (1992). Taxation and housing: Old questions, new answers. *The American Economic Review*, 82(2):237–242.

- Poterba, J. M. and Sinai, T. (2011). Revenue costs and incentive effects of the mortgage interest deduction for owner-occupied housing. *National Tax Journal*, 64(2):531–564.
- Rappoport W, D. E. (2019). Tax reform, homeownership costs, and house prices. In *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, volume 112, pages 1–35.
- Roback, J. (1982). Wages, rents, and the quality of life. *Journal of Political Economy*, 90(6):1257–1278.
- Ruggles, S., Flood, S., Goeken, R., Grover, J., Meyer, E., Pacas, J., and Sobek, M. (2022). IPUMS USA: Version 12.0 [dataset]. Minneapolis, MN. <https://doi.org/10.18128/D010.V12.0>.
- Saez, E. and Zucman, G. (2016). Wealth inequality in the united states since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, 131(2):519–578.
- Sieg, H., Smith, V. K., Banzhaf, H. S., and Walsh, R. (2004). Estimating the general equilibrium benefits of large changes in spatially delineated public goods. *International Economic Review*, 45(4):1047–1077.
- Sinai, T. and Gyourko, J. (2004). The (un)changing geographical distribution of housing tax benefits: 1980-2000. *Tax Policy and the Economy*, 18:175–208.
- Sommer, K. and Sullivan, P. (2018). Implications of US tax policy for house prices, rents, and homeownership. *American Economic Review*, 108(2):241–74.
- Sommer, K. and Sullivan, P. (2019). The effect of the tax cuts and jobs act on the housing market. In *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, volume 112.
- Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy*, 64(5):416–424.
- Tracy, J. S., Schneider, H. S., and Chan, S. (1999). Are stocks overtaking real estate in household portfolios? *Current Issues in Economics and Finance*, 5(5).
- US Office of Management and Budget (2023). Circular a-4. *Draft for Public Review*.
- Wong, M. (2013). Estimating ethnic preferences using ethnic housing quotas in singapore. *Review of Economic Studies*, 80(3):1178–1214.
- Zodrow, G. R. (2001). Reflections on the new view and the benefit view of the property tax. *Property taxation and local government finance*, pages 79–111.

Supplemental Appendix for **Tax Policy and the Heterogeneous Costs of Homeownership**

Kelly C. Bishop, Jakob Dowling, Nicolai V. Kuminoff, and Alvin D. Murphy

October 24, 2023

Contents

A Building a Database of User Cost Rates	2
A.1 Overview	2
A.2 Sample Construction	3
A.3 Estimating Subsidies, After-Tax Rates of Return, and Loan to Value Ratios	5
A.3.1 Defining Subsidies and After-Tax Rates of Return	6
A.3.2 Taxing Implicit Rents and Considering Alternative Subsidy Definitions	8
A.3.3 Using TAXSIM Software to Estimate Tax Liabilities and Filing Behavior	9
A.4 Estimating Expected Capital Gains	17
B Additional Results	22
B.1 Validation	22
B.2 Geographic Variation in the Impact of TCJA on UCRs	25
B.3 Geographic Variation in Baseline Measures	25
B.4 Correlation with Estimated Rent-Price Ratios	28

A Building a Database of User Cost Rates

A.1 Overview

This appendix provides additional details on how we calculate the heterogeneous tax subsidies to homeownership and the user cost rate (UCR) of housing. Utilizing data from the American Community Survey (ACS), we construct annual measures for UCRs and subsidies for 3.3 million households from 2016 through 2019. The methodology described here can be applied to construct updated estimates as additional years of data from the annual ACS are released.¹

The sample for which we calculate subsidies and UCRs is described in Section A.2. We separate the UCR into components that we separately measure, estimate, or obtain from the literature, at differing levels of heterogeneity and spatial resolution. As shown in Equation (2), which we replicate here for convenience, the UCR can be expressed as:

$$UCR_{ijt} = (1 - ltv_{ijt})rf_{ijt} + ltv_{ijt} \cdot rm_t + \omega_{jt} + \delta + \epsilon - \gamma_{jt} - s_{ijt}$$

In the equation, ltv is the loan to value ratio, rf is the risk-free after-tax rate of return on capital; rm is the mortgage interest rate; ω is the property tax rate, δ is the rate of depreciation, ϵ is the owner's risk premium, γ is the expected capital gains. The last term, s , is the subsidy comprised of property taxes and mortgage interest paid that a homeowner who itemizes their deductions obtains, expressed as a fraction of their house value.

Since the subsidy, s_{ijt} , and the risk-free after-tax rate of return, rf_{ijt} , both

¹Likewise, it would be straightforward to produce estimates of historical user cost rates and subsidies as far back as the ACS releases allow.

depend on each household’s unique tax situation (income level and sources, family structure, geographic location, etc.) they are jointly estimated in the TAXSIM-based procedure described in Section A.3 below. Estimates of household-specific loan to value ratios, ltv_{ijt} , are obtained in an intermediary step in this procedure, but also enter directly into the user-cost formula.

The remaining inputs of the formula are estimated separately (rm_t , ω_{jt} , and γ_{jt}) or obtained from the literature (δ and ϵ) as described in Section 2 of the main text. Section A.4 provides additional details on the estimation of expected capital gains, γ_{jt} .

A.2 Sample Construction

The sample for which we calculate subsidies and UCRs is comprised of nearly all home-owning households in the 2016 through 2019 annual 1% ACS data. The construction of this sample is described here.

Though our focus is on homeowners, we calculate tax-filing behavior for renters as well which allows us to validate our model against data provided by the Internal Revenue Service (IRS). We therefore start with all 5.6 million households in the four years of ACS data. Our first sample cut is to exclude the 126 households that report having more than 10 children under the age of 19. This is done to avoid exceeding the maximum number of dependents that TAXSIM can handle when performing tax calculations.

We also drop 2,279 households (2,152 of which are homeowners) that have estimated adjusted gross incomes (AGIs) greater than \$1 million.² These households

²Since the IRS reports statistics by AGI but it is (unadjusted) gross income that is reported in the ACS, we make our sample cuts based on AGI by first using TAXSIM to estimate AGI for the whole sample, and then making cuts.

account for 0.32% of income tax returns in our sample period, and 0.04% of households in our sample. We choose to exclude this small group of very high-earning households from our analysis since the ACS does not capture information from this group very well as evidenced by their under-representation in the sample. Furthermore, both income variables and home values are likely to be top-coded for this group, and they are likely to face more complicated tax situations that are not captured in survey responses. For these reasons, we choose to omit this group from our analysis, even though our main results are essentially unchanged when they are included. With these two data cuts, we have constructed our validation sample that we use for comparisons with IRS data.

Next, we construct our estimation sample by restricting attention to the subset of 3.5 million households that report owning their residence. Of these, we exclude 220 thousand households that report living in either a mobile home or trailer, in a boat, tent, or van, or for which this information is missing. Additionally, we exclude 20 thousand home-owning households for which self-reported home value is less than \$10,000 and an additional 20 thousand for which self-reported home value is more than 6 standard deviations larger or smaller than the median self-reported home value in their specific PUMA. Finally, we drop 4 households with an active mortgage who have missing information about the length of their tenure in the home. Note, however, that all of the homeowners that are excluded from the estimation sample are still included in the validation sample. Table [A.1](#) summarizes the sample construction.

Table A.1: Validation and Estimation Sample Construction

	#HHs	Change
<u>Validation Sample</u>		
all HHs in the ACS, 2016-2019	5,613,645	–
exclude HHs with 11+ children under 19	5,613,519	-126
exclude HHs with adjusted gross income > \$1m	5,613,519	-2,279
<u>Estimation Sample</u>		
all owner-occupiers in Validation Sample	3,525,457	-2,085,783
exclude non-traditional properties	3,302,860	-222,597
exclude homes with value < \$10k	3,283,506	-19,354
exclude homes with value > 6 st.dev. of PUMA median	3,266,874	-16,632
exclude HHs with missing tenure for active mortgage	3,266,870	-4

Note: Non-traditional properties are mobile homes, trailers, boats, tents, vans, other, or missing.

A.3 Estimating Subsidies, After-Tax Rates of Return, and Loan to Value Ratios

To estimate the housing subsidy that a homeowner receives, we compare their actual total tax liability (federal and state) to their counterfactual total tax liability if the household did not own their home. Similarly, comparing the total tax liability of a household in different states of the world allows us to calculate after-tax rates of return. Finally, the process of estimating a household’s tax liability involves imputing the amount of mortgage interest they paid during the year. That imputation also yields an estimate of the household’s loan to value ratio, which is needed both to calculate subsidies and because it enters into the user cost rate formula directly.

This remainder of this section proceeds as follows. First, we define housing subsidies and after-tax rates of return in terms of differences in total tax liabilities. Thereafter, we describe how we use TAXSIM to estimate total tax liabilities.

A.3.1 Defining Subsidies and After-Tax Rates of Return

Although our subsidy measure is calculated as the straightforward difference in tax liabilities when the household owns their house and when they rent, calculating the counterfactual tax liabilities requires precise definitions and additional assumptions. For example, to calculate the counterfactual tax liability for a homeownership household, we need to make assumptions about what would happen to the equity they have invested in their home. Furthermore, we need to consider the rate of return on that investment as well whether income from that investment would be subject to taxation. We formalize these assumptions below.

Consider a homeownership household's total tax liability (ttl) in three different scenarios:³

ttl_0 Actual tax liability as a homeowner.

ttl_1 Counterfactual tax liability with *increased* income from return on invested housing equity.

ttl_2 Counterfactual tax liability with the same increased income as in ttl_1 but where the household no longer has any deductible housing expenses (i.e., mortgage interest payments and property taxes paid are set to zero).

Using these total tax liabilities, we then calculate:

$$s = (ttl_2 - ttl_1)/P,$$

$$\sigma = (ttl_1 - ttl_0)/P.$$

Note that because tax liability increases with income but ttl_0 is fixed, σ in-

³We define total tax liability as the sum of federal and state income tax liabilities.

creases with the rate of return on investment in a non-housing asset. However, since both tll_1 and tll_2 increase with this rate, s is less sensitive.⁴

The user cost rate of housing can then be written as:

$$ucr = ucr_0 - \sigma - s,$$

where $ucr_0 = (1 - ltv)ptrf + ltv \times rm + \omega + \delta + \epsilon - \gamma$. Following the main text, rm , ω , δ , γ , and $ptrf$ are the mortgage rate, property tax rate, depreciation rate, expected capital gains, and pre-tax risk-free rate, respectively, and where we have dropped subscripts for convenience. The after-tax rate is obtained by adjusting the pre-tax rate, $ptrf(1 - ltv)$, by σ , such that the after-tax rate is defined as:

$$rf = ptrf - \frac{\sigma}{(1-ltv)}$$

In our primary subsidy definition, s describes the subsidy rate that enters into the UCR formula, whereas σ is an opportunity cost adjustment term. In dollar terms, the amount of subsidy a household receives for being a homeowner is worth sP . However, we define the housing subsidy as the percentage of the annual cost of homeownership that is subsidized. This statistic is given by:

$$\text{Percentage of Costs Subsidized} = s/(ucr + s) = s/(ucr_0 - \sigma),$$

where the denominator describes the annual cost of owning the house in the absence of the subsidy and the numerator is the subsidy rate. The advantage of describing subsidies using this statistic is that it captures the ongoing subsidization as a percentage of the flow cost.

⁴Due to notches in the tax system, tll_1 and tll_2 might increase by different amounts for a constant increase in the rate of return on investment.

A.3.2 Taxing Implicit Rents and Considering Alternative Subsidy Definitions

Our primary subsidy definition above simply reflects the difference in tax liability for the same household when owning and renting under the current actual tax system. This definition does not compare the current tax system to a hypothetical tax system where imputed rent is taxed. As discussed in the text, an alternative subsidy definition would additionally count σ as a subsidy instead of as an opportunity-cost-adjustment term. This alternative definition would then be:

$$\text{Alternative Percentage of Costs Subsidized} = (s + \sigma)/(ucr + s + \sigma) = (s + \sigma)/ucr_0,$$

and, in dollar terms, the subsidy would be worth $(s + \sigma)P$.

This alternative definition will, by construction, always yield larger estimates of the tax subsidy to housing. Since the TCJA mainly affected s , the relative change in subsidies due to the TCJA will be smaller under this definition.⁵ Table A.2 quantifies the subsidies for the mean homeowner by year under the two definitions. Finally, we note that our estimates of user cost rates do not depend on the choice of subsidy definition. In either case, $ucr = ucr_0 - \sigma - s$, and the choice of subsidy definition is simply a labeling decision that determines whether σ is considered part of the subsidy or whether it is an adjustment-term used to obtain the after-tax opportunity cost of funds.

The rate of returns on capital is an input into the estimation of housing subsidies and user-cost rates. Since our UCR formula includes a risk-premium term, we believe the appropriate benchmark rate is the risk-free rate of return, *ptrf*. Table A.2 explores the sensitivity of our estimates to this assumption. The first panel shows

⁵The σ term is only affected through changes in income tax rates and brackets.

Table A.2: Subsidy for the Mean Homeowner under Alternative Specifications

	sP	σP	$(s + \sigma)P$	$\frac{s}{ucr+s}$	$\frac{s+\sigma}{ucr+s+\sigma}$
<i>ptrf</i>					
2016	\$1,217	\$790	\$2,008	6.6%	10.5%
2017	\$1,227	\$770	\$1,996	6.6%	10.6%
2018	\$428	\$782	\$1,210	2.2%	6.3%
2019	\$422	\$776	\$1,198	2.1%	6.3%
<i>rm</i>					
2016	\$1,231	\$1,334	\$2,565	6.9%	13.5%
2017	\$1,241	\$1,342	\$2,582	7.0%	13.8%
2018	\$429	\$1,353	\$1,782	2.3%	9.4%
2019	\$423	\$1,368	\$1,791	2.2%	9.6%

Table A.3: Pre-Tax Risk-Free Rates of Return and Mortgage Interest Rates

	<i>ptrf</i>	<i>rm</i>
2016	2.82%	4.59%
2017	2.59%	4.35%
2018	2.52%	4.20%
2019	2.41%	4.09%

our main estimates using the risk-free rate of return. The second panel repeats these estimates using the mortgage interest rate, rm . Table A.3 shows these rates by year.

A.3.3 Using TAXSIM Software to Estimate Tax Liabilities and Filing Behavior

Regardless of which definition we use, we need to be able to estimate total tax liabilities under any candidate tax scenario, for any household in the ACS. For validation purposes, and to explore mechanisms through which tax policies affect subsidies, we

are also interested in estimating whether or not a household itemizes their taxes or takes the standard deduction.

To do so, we use a novel version of NBER’s TAXSIM software. This version is identical to version 35 of TAXSIM which is publicly accessible online at <https://taxsim.nber.org/taxsim35/>, except for a modification that allows us to impose counterfactual tax policies such as a change in, or elimination of, the SALT cap.

To use TAXSIM to estimate the tax liability and filing behavior of a household we need information about the household’s income, age, number of dependents, marital status and deductible expenses. The ACS provides reasonably detailed information about all but the last component. Deductible expenses require some imputations and can, for our purposes broadly be divided into four categories: property taxes paid, mortgage interest paid, charitable giving and eligible medical expenses. Together with state and local income taxes (for which TAXSIM automatically calculates eligible deductions), these four categories account for most itemized deductions. Below we describe how we construct these inputs required to run TAXSIM.

Income, Age, Dependents and Marital Status

These inputs are directly observable in the ACS. However, the variables that TAXSIM uses and the variables available in the IPUMS ACS are often coded differently. We re-code the ACS variables to make them TAXSIM-compatible. For each person in every household, the ACS contains the following income variables:

- INCWAGE (Wage and salary income)
- INCBUS00 (Business and farm income)
- INCINVST (Interest, dividend, and rental income)

- INCSS (Social Security income)
- INCRETIR (Retirement income)
- INCWELFR (Welfare or public assistance income)
- INCSUPP (Supplementary Security Income)
- INCOTHER (Other income)
- INCEARN=INCWAGE+INCBUS00 (Total personal earned income)
- INCTOT (Total personal income), which is the sum of all income variables above.

We then allocate these income variables for each person in the household to the income categories in TAXSIM as follows:

- pwages (wage income of primary taxpayer): we allocate the sum of INCEARN for the household head and the sum of INCOTHER for the entire household.
- swages (wage income of spouse): will be zero if the household head is not married, otherwise the spouse's INCEARN is allocated.
- ltcg (long term capital gains): we allocate the sum of household INCINVST. Note that since ACS does not specify the source of investment income, it is not clear how to allocate this variable in TAXSIM as it could just as well be allocated to interest received or dividends for example. However, TAXSIM does not accept negative values for dividends or interest received, but does for capital gains. Since INCINVST is negative for some households, we allocate it as long-term capital gains. In practice this makes little difference and results are quantitatively similar under alternate allocations.

- pensions (taxable pensions and IRA distributions): we allocate the sum of household INCRETIR.
- gssi (gross social security benefits): we allocate the sum of household INCSS and INCSUPP.
- transfers (other non-taxable transfer income): we allocate the sum of household INCWELFR.

In addition to income, several features of the tax code depend on the taxpayer's age. We use the ACS head of household's age and, if applicable, their spouse's. The number of and age of dependents also determine things such as personal exemption calculations and eligibility for child tax credits. Since ACS does not include information about eligible childcare expenses, we do not include these when running TAXSIM. For all other categories of dependents, we simply count the number of people in each household who fall into each age group.

Finally, a household's filing status is determined by their marital status. This can be either single (or head of household) for unmarried taxpayers, joint (married), or separate (married). TAXSIM's user instructions note that filing as married-separate is not usually desirable under US tax law. We therefore assume that head-of-households who are separated (or coded as married with spouse absent) in the ACS are filing jointly. This assumption affects only a small number of households and does not substantially change our estimates of mean subsidies or user cost rates. We thus assume all households in the ACS that are married (with spouse present or absent) or separated but not divorced, choose the tax filing status "Married, filing jointly." All other households are coded as filing as "Single or head of household."

Property Taxes Paid

Property taxes paid are reported in ranges in the ACS. We calculate the midpoints of these ranges and assign the appropriate midpoint value to each household. The top range, however, is open-ended so we need to impute payments for households who report paying \$10,000 or more in property taxes.

For this group, we impute their property taxes paid as follows. First, we calculate the effective property tax rate for each PUMA as described in the main text while excluding any households in the top property tax range. We do this by dividing total taxes paid by total reported property values in the PUMA. Second, we apply this PUMA-specific imputed tax rate to the home value of each household in the top property tax range to obtain these households' annual tax payments. If this imputed payment is larger than \$10,000, we use the imputed payment, otherwise we assign the household as paying \$10,000 in property taxes annually.

Under this method, 50% of households in this top group get assigned \$10,000 in property tax payments. However, since this group is a small share of homeowners, using imputed payments directly (instead of $\max\{\$10,000, \text{imputed payment}\}$) does not substantially change our estimates of mean subsidies or UCRs.⁶

Mortgage Interest Paid (and Loan to Value Ratios)

Mortgage payments are reported in the ACS as monthly dollar amounts for first and (potential) second mortgages. These are top-coded at the 99.5th percentile in the state where the household resides. Higher amounts are expressed as the state means of values above the listed top-code value for that specific year.

Respondents are also asked whether these monthly payments include property taxes and insurance, both of which are also reported in the ACS. This enables us

⁶In principle, one could apply a similar imputation to all households, which would effectively be a shrinkage estimator, where all property payments are shrunk towards the PUMA-specific mean rate.

to calculate net annual mortgage payments for each household by first subtracting property taxes paid (if property taxes are included in payments) from the gross annual mortgage payments. Then, if insurance is included, we subtract the insurance payments (which are reported in dollar amounts).

However, if a household says property taxes are included in their mortgage payments but claim property taxes paid exceed 100% of the reported annual mortgage payments, we assume they have misreported and that in fact property tax payments were not included in the mortgage payments. We also do the same thing for the insurance payments. Only a small fraction of households are considered misreporting under this procedure.

This results in net annual mortgage payments for all homeownership households in the sample. However, it is only the share of annual mortgage payments that are interest payments that is deductible. This share will depend on the interest rate of the mortgage and where in the repayment schedule the household is, i.e., the term of the loan and time elapsed since origination. For mortgage rates, we assume a homogenous mortgage rate across households set equal to the 10-year average of 30-year fixed rate mortgages in the US. Results are not sensitive to the choice of mortgage rate used when calculating mortgage-interest paid and loan to value ratios.⁷

With interest rates and each household's annual (net of taxes and insurance) payments in hand, we can calculate an amortization schedule as long as we know either the initial loan amount, or the total number of monthly payments required (i.e., the mortgage term). Since knowing the initial loan amount requires making

⁷For example, replacing the mortgage rate used when calculating mortgage-interest paid and loan to value ratios with the 2009-2019 range of 30-year mortgage rates referenced in Section 2.2 yields UCR that are 98%-104% of the baseline estimated UCR. In contrast, the UCR is naturally sensitive to, rm , the choice of mortgage rate used in the direct calculation of UCR, as can be seen in Equation (2).

assumptions about house-specific appreciation rates and households' initial down payments, we will instead make an assumption about the total number of payments. Once we have this amortization schedule and information about when a household moved into their home, we can calculate the fraction of mortgage payments that go toward interest.

For each household, the ACS reports when they moved into the housing unit they are interviewed in. Like property taxes, this variable is binned into ranges and we assign each household the midpoint of the range. Households in the top bin – those that report moving in 30 or more years ago – are assumed to have moved in 35 years ago.

While this choice is arbitrary, the resulting imputations are not particularly sensitive to this choice. First, only 9.4% of households with a mortgage are in this group. Second, and more importantly, these households are likely to be near the end of their mortgage terms, which means the fraction of payments that go toward interest will be small. In the extreme case that everyone got a 30-year mortgage and made on-time payments with no refinancing, this entire group should have already paid off their mortgages. This case, however, is rejected in the data given that 30% of homeowners who moved in more than 30 years ago report making mortgage payments.

For this reason, and to account for households' option to refinance, we do not assume that all households get 30-year mortgages. Instead, we choose the length of the mortgage so that the average number of years remaining for repayments in our 2016-2017 sample matches the comparable number reported in Keys et al. (2016).⁸ In their CoreLogic sample that covers 85% of mortgages active in 2010, the average

⁸We use only the 2016 and 2017 ACS samples for this matching to avoid any potential influence of the TCJA.

loan has 23.4 years remaining. In our ACS sample, under the assumption that each household with a mortgage got their loan when they moved in, the average number of years remaining depends on both the mortgage term, as well as how long we assume households in the “30+ years”-group on average have been in their home. Under our assumption that this group of households has stayed in their home for 35 years, we need to assume that all households get a mortgage that is 32.5 years long in order to match the average years remaining. We note that because the group of households with a tenure longer than 30 years is a small fraction of our total sample of households with a mortgage, and because a small fraction of this group’s mortgage payments go toward interest, the necessary mortgage length to match the average years remaining is not sensitive to how long we assume this group has been in their homes. For example, assuming they all had been in their house for 30, 40 or even 50 years instead of 35, at most changes the implied mortgage length by a couple of months.

With interest rates, net monthly payments, and the total number of payments required, we can then calculate an amortization schedule for each household. With information about their tenure, we can then see where in this schedule they are, and thus what share of their payments go toward interest and principal, respectively. We then use these shares to calculate their annual (deductible) interest paid. This amortization schedule also predicts the current loan value, which together with self-reported property values are used to calculate each household’s loan to value ratio which then enters the user-cost formula directly.

Charitable Giving and Medical Expenses

Since the ACS does not contain information about deductible charitable giving or medical expenses, we use information from the 2017 and 2019 waves of the PSID

which ask households about their tax-deductible giving to charities and about their medical expenses.

We use quintile limits from the Census Bureau (based on CPS ASEC data) to group the PSID households into quintiles of the national income distribution. For each household, we then sum their total (eligible) giving to charitable organizations. For medical expenditures, households can deduct only those medical expenses that exceed 7.5% of their adjusted gross income. We therefore compare each household in the PSID’s reported medical expenses to their household income. If expenses exceed 7.5%, we include them in that household’s deductible expenses. We then calculate mean deductible expenses for each income quintile.

Finally, we group the households in our main ACS sample into quintiles of their year-specific income distributions and assign each household the quintile-specific mean deductible expenses calculated from the PSID.⁹

A.4 Estimating Expected Capital Gains

To obtain expected capital gains for each PUMA, we begin by separately estimating hedonic price indices for 235 distinct markets (186 MSAs and the non-MSA areas of all states except Rhode Island) over the 1990-2019 period. To do this we use Census and ACS data. For 1990 and 2000, we use the decennial Census 5% state samples. After 2000, the annual ACS 1% samples are available from 2005 onward. For each market, j , we estimate the following model:

$$\log(y_{it}) = \beta_t + X_{it}\alpha + u_{it}$$

where $t \in \{1990, 2000, 2005, 2006, \dots, 2019\}$ and y_{it} is the self-reported home price

⁹Unfortunately, it is not possible to use the PSID to get the geography-specific mean deductible expenses.

of household i in year t and X_{it} is a vector of dwelling characteristics. The X_{it} vector includes indicators for access to kitchen and plumbing facilities, indicators for the number of rooms and number of bedrooms, indicators for the age of the structure and whether the house is a single-family home or contains multiple units, and if so, how many.

By exponentiating the estimates of β_t , we obtain a nominal hedonic-price index. For the years in which we lack annual data (1991-1999 and 2001-2004), we assume a constant growth rate, e.g., the price growth in 2003 is given by $\left(\frac{P_{2005}}{P_{2000}}\right)^{\frac{1}{5}}$. The nominal indices are converted to real indices by deflating with the consumer price index (CPI), excluding shelter.

Finally, we calculate the nominal capital gains parameter, γ_{jt} , as the sum of the average growth rate of this real price index between 1990 and 2019 and the 10-year inflation expectations obtained from the Livingston Survey of professional forecasters. This leaves us with 235 market-specific measures which we map back to PUMAs using a crosswalk provided by IPUMS and the procedure outlined below.

Metropolitan Area Definitions and PUMAs

All PUMAs that are fully located either within or outside an MSA simply get assigned their market-specific expected capital gains parameter. For PUMAs where some, but less than 100%, of the PUMA-population lives within an MSA, we calculate the PUMA-specific term as a population-share weighted average of the MSA and the state non-MSA measures.

Metropolitan areas are identified using the geographically constant MET2013 variable in IPUMS from 2000 onward. From 1990 until 2011, metro areas can also be identified by the METAREA variable. It should be noted that METAREA is

not a geographically constant variable as it is contingent on varying delineations of metro areas across time and on variations in available geographic information and in confidentiality restrictions among samples. These varying definitions of metro areas notwithstanding, we match each METAREA from the 1990 census sample with a MET2013 variable from the later sample. This matching procedure is described below.

Matching METAREA with MET2013

There are 295 distinct MET2013 codes in the data, and 334 METAREA codes. We first join the METAREA codes with the MET2013 codes based on an exact match of their labels. For example, METAREA code “8” has the label “Akron, OH,” as does the MET2013 code “10420.” Since the label “Akron, OH” is exactly the same for both the METAREA and MET2013, these codes are matched in the first step. 127 of the metro areas are matched in this manner.

This leaves 375 unmatched codes: 207 in METAREA and 168 in MET2013. We manually match these remaining variables by inspecting their labels. For example, the MET2013 code “12260” has the label “Augusta-Richmond County, GA-SC.” We match this with the METAREA code “60” which corresponds to “Augusta-Aiken, GA/SC.”

Some METAREA codes are matched with multiple MET2013 areas. This occurs when more than one MET2013 area corresponds to a METAREA, e.g. the “Salt Lake City-Ogden, UT” METAREA gets matched with both the “Ogden-Clearfield, UT” and the “Salt Lake City, UT” MET2013 areas. This also occurs for the “San Francisco-Oakland-Vallejo, CA” METAREA which gets mapped to both the “Vallejo-Fairfield, CA” and the “San Francisco-Oakland-Hayward, CA” MET2013 codes. Since the analysis will be based on MET2013-areas (because these

easily map into PUMAs), this means that for 1990, the houses in “San Francisco-Oakland-Vallejo” will be used for both the “San Francisco-Oakland-Hayward” and the “Vallejo-Fairfield” regressions. After 2000, however, these metro areas will be allowed to diverge.

After this hand-matching procedure, 32 MET2013 and 78 METAREA codes remain unmatched. Out of the matched areas, we then discard areas which are not observed in all 16 samples (1990, 2000, 2005-2018). This occurs mostly because a MET2013 area has been created or discontinued, e.g., in 2011, both “Flint, MI,” and “Hammond, LA” lost their MET2013-status while “Ithaca, NY” and “Florence, SC” were created.

State Non-Metro Areas

For each state, we also consider all the areas that are not within an MSA or where the MSA is not identified as a single market. We therefore estimate capital gains for each state's non-MSA market as well. These parameters are then mapped to the user cost of households in all PUMAs within each state that are not within one of the matched MET2013-areas. Reasons a house will be included in the state non-metro sample are:

- The MET2013 or METAREA codes are 0, indicating “Not identifiable or not in an MSA”
- The MET2013 and METAREA variables were not matched
- The MET2013 and METAREA variables were matched, but this area is not identified in all 16 samples


In total, we end up with 186 consistent metro areas and 49 state non-metro areas (all states except Rhode Island). All homes in the data are thus located in one of these 235 markets.

B Additional Results

B.1 Validation


Figures [B.1](#) and [B.2](#) contrast the fraction of all tax-filing households that choose to itemize according to data from the Internal Revenue Service (IRS) with predictions from our empirical design for the tax-minimizing filing strategies. Specifically, in [Figure B.1](#), we compare across households across fourteen, equally-spaced, income bins ranging from adjusted gross income of less than \$0 to \$500,000-\$1,000,000. In [Figure B.2](#), we compare across households across approximately 2,400 PUMAs. Both Figures indicate a very strong correlation between predicted and actual itemization rates. Additionally, intercepts and slopes of the fitted regression lines are close to zero and one, respectively.

Figure B.1: Comparing Predicted and Actual Itemization Rates by Income



Note: Each observation represents a particular income-year bin. The regression line is fitted with equal weights on each bin. Results are very similar when weighting by income-year bin size.


Figure B.2: Comparing Predicted and Actual Itemization Rates by PUMA



Note: Each observation represents a particular PUMA-year. The regression line is fitted with equal weights on each PUMA-year. Results are very similar when weighting by PUMA-year population size.

B.2 Geographic Variation in the Impact of TCJA on UCRs


Figure B.3: Change in User Cost Rates by PUMA



Note: The percentage-point increases are winsorized at the 1st and 99th percentiles.


B.3 Geographic Variation in Baseline Measures

Figure B.4: Loan-to-Value Ratios for Homeowners by PUMA, 2016-2017




Note: The loan-to-value ratios are winsorized at the 1st and 99th percentiles.

Figure B.5: After-Tax, Risk-Free Rates for Homeowners by PUMA, 2016-2017




Note: The rates are winsorized at the 1st and 99th percentiles.

Figure B.6: Property-Tax Rates for Homeowners by PUMA, 2016-2017




Note: The rates are winsorized at the 1st and 99th percentiles. The variation is driven by variation in households' marginal tax rates.

Figure B.7: Expected Capital Gains for Homeowners by PUMA, 2016-2017



Note: The expected capital gains are winsorized at the 1st and 99th percentiles. Expected capital gains are defined by the expected growth rate less expected inflation, as explained in Section A.4.


Figure B.8: Subsidy Rate for Homeowners by PUMA, 2016-2017



Note: The subsidy rates are winsorized at the 1st and 99th percentiles.

B.4 Correlation with Estimated Rent-Price Ratios

Figure B.9: Comparing Rent-Price Ratios and User-Cost Rates by PUMA



Note: Each observation represents a particular PUMA-year. The regression line is fitted with equal weights on each PUMA-year. Results are very similar when weighting by PUMA-year population size.