The Distributional Impact of Mortgage Interest Subsidies: Evidence from Variation in State Tax Policies

Matthew Davis*
The Wharton School, University of Pennsylvania

JOB MARKET PAPER
January 15, 2019
Click to access most recent version

Abstract

Mortgage interest tax deductions are a widespread, expensive, and regressive tax expenditure, so understanding the distribution of the policy's costs and benefits is a question of first-order economic importance. This paper combines a sufficient-statistics approach with direct estimates of the induced effect on house prices to measure the policy's economic incidence, as distinct from its statutory incidence. I start with a flexible economic framework that expresses the policy's distributional impact in terms of a key parameter: the capitalization effect, or the extent to which the deduction increases house prices. I then directly estimate this parameter, drawing on a national database of housing transactions and exploiting sharp variation in tax rates and itemization rules at state borders. Comparing the sale prices of observationally identical homes purchased on either side of a border, I find that a one percentage point increase in the tax rate applied to mortgage interest increases house prices by 0.8%, which is sufficient to erase the tax savings for a first-time borrower when their loan-to-value ratio is under 60%. Finally, I combine the empirical result and the derived incidence expressions to show the distribution of the policy's impacts among new home-buyers. Accounting for non-itemization rates indicates that average buyers at most incomes do not benefit from the MID, though there is some heterogeneity across income levels and housing markets.

*I am grateful to Fernando Ferreira, Joe Gyourko, Ben Keys, and Todd Sinai for their guidance during this project. Amanda Chuan, Gilles Duranton, Ryan Fackler, Joe Harrington, Ben Hyman, Hal Martin, Madison Priest, Marco Tabellini, Maisy Wong, and participants in the Wharton BEPP student seminar, the Wharton Urban/Real Estate Lunch, the Ohio State Real Estate and Housing Conference, the Wharton Applied Economics Workshop, and the Urban Economics Association Meetings also provided valuable feedback. Financial support from a National Science Foundation Graduate Research Fellowship and the Lincoln Land Institute's C. Lowell Harris Dissertation Fellowship is gratefully acknowledged. Correspondence can be addressed to the author at mattda@wharton.upenn.edu. All errors are my own.

1. Introduction

The 2008 financial crisis alerted many economists and policy-makers to the costs faced by households who have heavily invested in the housing market, often taking on significant debt in the process (Mian and Sufi 2011; Mian, Rao, and Sufi 2013; Ganong and Noel 2018). While subsequent years have seen a robust discussion of possible remedies (Piskorski and Seru 2018), an underemphasized policy factor is the fact that governments often subsidize home purchases through the tax code. Indeed, mortgage interest deductions (MIDs) are found throughout developed economies despite their considerable fiscal burdens and oft-maligned distributional properties (Bourassa et al., 2013).[1] While often understood as a tax benefit for middle-income homeowners, economists and policy analysts have long argued that the MID is massively regressive; after all, tax deductions are most valuable to those in high tax brackets, and no tax relief is provided to renters or non-itemizing homeowners, who are less wealthy on average. Hence, understanding the distribution of the policy's impacts – not just the total economic distortion – is a question of first-order economic importance, yet one for which we have very little direct evidence.

This paper measures the economic incidence of the mortgage interest deduction by combining a flexible theoretical framework with direct empirical estimates of the induced change in house prices. I start by deriving incidence expressions from a model that places very few restrictions on the structure of household decision-making while allowing for a rich characterization of the household's cost of capital. These formulas show that the distribution of policy impacts depends critically on a key sufficient statistic: the capitalization effect, or the extent to which the policy increases house prices. Next, I directly estimate capitalization, deploying a border design that relies on variation in state tax rates and itemization policies. Finally, I bring the theoretical and empirical portions together to map out the distribution of costs and benefits across the population, drawing on a large and detailed housing transactions dataset.

The theory serves as a guide to tracking the policy's costs and benefits among various agents in the economy. Importantly in this setting, the impacts are likely to be unequally distributed across households, so the set-up allows for different tenure choices, credit constraints, and debt levels. The results show that, as with any tax or subsidy, the deduction's economic incidence is the sum of two effects: a direct monetary transfer and the ensuing change in equilibrium prices. The relative magnitudes of these two terms determines the ultimate distribution of benefits. In other

[1] In the United States, the mortgage interest deduction (MID) provided homeowners with $77 billion of tax relief in 2015, more than four times the amount allocated to housing vouchers for low-income households.

words, the capitalization response is a sufficient statistic for the policy's effect on the distribution of resources (Chetty and Finkelstein 2009).

Measuring capitalization has proved challenging for empiricists, owing to challenges in both research designs and data availability. Prior work has either relied on structural general equilibrium models (Sommer and Sullivan 2018) or expressed incidence in terms of existing estimates of the elasticities of supply and demand for homes (Rappoport 2016, Hanson and Martin 2016). The use of these indirect approaches largely reflects a dearth of credible research designs for directly estimating the level of capitalization. Appropriate policy experiments are rare; in the United States' case, the federal MID policy has not substantially changed since its inception.[2] At the most basic level, analysis requires plausibly exogenous, market-level variation in the level of the subsidy, combined with data on market prices under these differing policies.[3]

I overcome the lack of federal policy experiments by turning to a new source of identifying variation. State tax rules vary substantially both in the cross-section and over time, and my identification strategy exploits the fact that different tax treatments induce sharp variation in borrowing costs at state borders, due to either differences in personal income tax rates or state-specific itemization policies.[4] Specifically, I construct a sample of census tract-pairs that share a common border but lie in different states. Within these pairs, I compare the sale prices of observationally identical houses that transact in the same time period on either side of the state border. Relating the difference in prices to the difference in effective tax rates yields an estimate of the capitalization of the deduction.

The border design is made possible by a national database of housing transactions, notable for both its breadth and its detail. The broad data coverage ensures that I have sufficient sample near

[2]While time series variation in federal tax rates creates variation in the value of the deduction, such policy changes are confounded by a host of macroeconomic and time-series factors that prevent credible identification of the capitalization effect. Glaeser and Shapiro (2003) also show that state-level homeownership rates are not correlated with tax rates.

[3]Some papers use household-level variation in tax incentives to estimate causal effects on outcomes other than from market-level prices, the effect of interest here. Gruber et al. (2017) consider a policy change that had heterogeneous affects across the income distribution in Denmark, and Hilber and Turner (2014) show that households are not more likely to purchase a home after moving to a state with a larger tax subsidy. Both find a null effect on overall homeownership. Finally, Berger et al. (2000) estimate the price impacts of an interest subsidy tied to specific buildings in Sweden, finding results consistent with full capitalization.

[4]The level of state income taxes is highly variable, ranging from a high of 12.3% for high-income filers in California to zero in 17 states. (Because state tax payments may be deducted from one's federal tax return, the true rate is less than the statutory rate during my sample period. I use effective tax rates that account for federal deductibility throughout my analysis. See Section 4.2 for details.) Differences in rates generate variation in the after-tax cost of mortgage interest, since reducing one's taxable income is more valuable in a high-tax state. Furthermore, differences in state itemization rules induce variation in the specific tax rate applied to mortgage interest. In 2015, ten states collected income taxes but did not allow residents to deduct mortgage interest on their state return. Three other states placed a cap on the value of total deductions, effectively eliminating the subsidy at the margin for taxpayers claiming significant deductions.

state borders to execute the border design; by 2015, the data includes deeds records from over 90% of U.S. counties and a total of 4.2 million residential transactions. In addition, the detailed house-level characteristics allow me to construct house-type-by-geography fixed effects that ensure the estimates are driven by like-for-like comparisons.

Finally, I deploy a simulated instruments strategy to ensure that my estimates are driven by policy variation rather than to differences in the local tax base (Currie and Gruber 1996, Gruber and Saez 2002). The key "treatment" variable in my regressions is the average marginal tax rate, which is a function of both tax rules and the composition of the tax-paying population. The latter may respond directly to tax policy, due to sorting or other economic effects of tax variables. Solving this issue requires instruments that predict local rates but are not related to demographics. I therefore use the National Bureau of Economic Research (NBER) TAXSIM programs to compute statutory tax rates that are not influenced by population sorting. By using these simulated rates to instrument for the local average rate, I eliminate the endogeneity associated with changes in the tax affecting the non-statutory determinants of local effective rates.

The results show that the mortgage interest deduction induces large changes in house prices. Specifically, a one percentage-point increase in the tax rate applied to mortgage interest results in a 0.8 percent increase in house values. These estimates are consistent across specifications using a variety of strategies to control for local price trends, regional variations in the valuation of different housing amenities, and the interaction between the two.

This level of capitalization significantly reduces the tax benefits for itemizing homeowners. For a typical itemizer who finances her first home purchase at 80% LTV and faces a 25% marginal tax rate, the price increase wipes out 84% of the subsidy value. Because roughly half of homeowners do not itemize on their tax returns – and therefore cannot claim the deduction – the policy does not help new homebuyers, on net. Existing homeowners are largely insulated from the price effects, as any capital gains or losses are offset by an increase in future housing costs.

Other parties that do not claim deduction are nevertheless affected by the price increase. I do not directly estimate impacts on rents, but under plausible assumptions about housing market efficiency renters suffer a proportional increase in housing costs. Housing developers see an increase in profits; rising prices imply that they capture a portion of the subsidy value.

The capitalization estimates are robust to several potential threats to the border strategy. First, one naturally worries about unobserved factors that might affect house prices. Focusing on variation within small geographic areas allays many of these concerns. The border design effectively partials out local economic conditions, provided that they do not also vary sharply at the state border. Furthermore, directly controlling for measures of local economic performance does not substantively change the point estimates. Exposure to certain amenities and fiscal policies is more likely to coincide with state borders, however. I therefore estimate models that control for education spending, state fiscal conditions, and inter-governmental transfers. These approaches show similar levels of capitalization, suggesting that the original specification successfully isolates variation in deduction value rather than these potential confounders.

Using income tax rates for identification raises two other concerns. First, it is possible that households may sort in response to tax rates, and that the resulting redistribution of households could change house prices. Second, higher tax rates lead to lower after-tax income, which could directly shift housing demand.[5] To address these concerns, I estimate models that rely solely on variation induced by state itemization policies, rather than variation in income tax rates. The resulting estimates are consistent with earlier results. Unless households are sorting between states to take advantage of itemization rules, the two concerns discussed here are alleviated.

My most robust models use state fixed effects to identify capitalization through a limited number of changes in state tax policies. These difference-in-difference style models place significant strain on the housing transactions dataset, which contains a sufficiently broad sample near state borders for a relatively short duration. Nevertheless, the estimates are similar to those in the cross-sectional border design (though they are less precise). Furthermore, they survive state-level economic and policy controls – suggesting that endogenous rate changes are not driving the estimates – and are similar when using time-series variation in itemization policies instead of tax rates.

The final step in the analysis uses the incidence expressions to take the capitalization estimates to the data. I add demographic information by combining the housing transaction dataset with data collected under the Home Mortgage Disclosure Act (HMDA). Crucially, HMDA filings report household income, which I use to impute household-level marginal tax rates and the likelihood that households itemize. Accounting for observed itemization rates and larger estimated capitalization effects for more expensive homes shows that incidence is negative for homebuyers across the income distribution, though the estimates are for the most part not statistically differentiable from zero.

[5]Note that this second effect is distinct from the "income effect" commonly associated with analyses of tax changes. At issue is a change in income independent of changes in prices and/or tax rates, rather than a change in effective spending power due to a change in prices.

These findings contribute direct empirical evidence of capitalization effects to the extensive theoretical literature on the effects of the MID, which has grown significantly since Poterba's (1984) user-cost model. Many papers have extended this framework to analyze the impacts of various MID policy reforms,[6] though the degree of price sensitivity varies widely with assumptions about the supply response. Hanson and Martin (2016) show that adding a supply side reduces the price sensitivity by 64% from the full-capitalization case. Rappoport (2016) demonstrates theoretically that the MID might hurt borrowers in closed, inelastically-supplied housing markets (though his calibration suggests that it is rare in practice). Sommer and Sullivan's (2018) structural estimates predict that eliminating the MID would *increase* homeownership by reducing prices to the point where more potential buyers could afford down payments. The empirical results in this paper do not require assumptions about the supply side or the elasticity of demand; they estimate price sensitivity directly using natural policy variation.

I also contribute to a literature that emphasizes the important role that mortgage payments and housing debt play in household financial well-being. Mian and Sufi (2011) and Mian, Rao, and Sufi (2013) show that households with higher mortgage debt suffer greater consumption declines and increased default rates after a housing downturn. Ganong and Noel (2018) identify cash flow constraints as a key contributor to mortgage default for underwater borrowers. MID capitalization likely contributes to this effect, as higher prices result in higher monthly payments, but the tax benefit is not realized until the following year.

Finally, this paper's findings are relevant to the broader literature on the effect of capital market conditions on asset prices. Garrett et al. (2017) find municipal bond prices over-capitalize differences in state tax rates, finding evidence of more-than-complete capitalization, and Argyle et al. (2018) show that the benefits of favorable car loan terms are often subsumed by induced price increases.[7] My results show that a similar form of capitalization effects is vital to understanding a $29 trillion asset class owned by more than 60 percent of households in the United States.

The remainder of this paper is structured as follows. Section 2 outlines notable features of the MID in the United States and walks through a theoretical framework that traces how price changes affect various parties participating in the housing market. Section 3 explains my approach

[6]Poterba and Sinai (2008) find that eliminating the MID would reduce user costs of capital by 40 basis points on average, with larger effects for younger and wealthier households. Other papers analyze the spatial distribution of the deduction's benefits (Gyourko and Sinai 2003), the deduction's effect on household portfolio decisions (Poterba and Sinai 2011), and the deduction's role in determining household location decisions, (Albouy and Hanson 2014).

[7]See Glaeser, Gottlieb, and Gyourko (2013); Adelino, Schoar, and Severino (2012); and Bhutta and Ringo (2017) for discussions of the relationship between house prices and mortgage interest rates.

to estimating the price response, and Section 4 describes the data I use. I present my results and robustness checks in Section 5. Section 6 combines these estimates with national data on characteristics of borrowers and owners to analyze the distribution of the policy's impacts, and Section 7 concludes.

2. Background and a Framework for Distributional Analysis

2.1 MID Basics

When the income tax was first instituted in the United States, individuals could deduct all interest on personal debts from their taxable income. This status quo persisted until the Tax Reform Act of 1986, which eliminated the preferential tax treatment for all forms of personal interest aside from mortgage interest. The MID also survived 2017's Tax Cuts and Jobs Act (TCJA), though the law reduced the maximum deductible balance from $1 million to $750,000. More importantly, the bill doubled the standard deduction, making the MID less attractive for many homeowners.[8]

Today, the MID enjoys a privileged position in American politics. The New York Times labeled the deduction "sacrosanct" after 93 percent of respondents to a 2011 poll described it as either "very important" or "somewhat important." It is widely praised by industry representatives and leaders of both political parties, who frequently emphasize the importance of middle class tax relief.[9]

It should be noted that the economics literature takes a somewhat different view of the true source of the subsidy to owner-occupied housing (perhaps unsurprisingly). As I discuss at greater length in the following subsection, housing economists have long emphasized the non-taxation of imputed rent, i.e. the "dividend" produced by housing assets (Hendershott and Slemrod, 1982). My focus is somewhat different than these papers, however, and more in line with current policy

[8]CBO (2018) and JCT (2018) estimate that the number of itemizers will fall by roughly 60% in response to the reform, and JCT (2018) projects an equal decline in the total MID tax expenditure. The decline is primarily driven by the increase in the standard deduction, effectively zeroing out the benefits for many middle-income claimers. Some high-income claimers see somewhat smaller tax reductions due to the reduced cap.

[9]When asked asked about tax benefits for homeowners at an event for constituents, Democratic Senate Minority Leader Chuck Schumer recently promised "no reductions in your deductions." Both the National Association of Realtors (NAR) and the National Association of Homebuilders (NAHB) highlight their lobbying on the issue on their websites, with the NAR claiming that "to even mention policy changes that would reduce the tax benefits of homeownership could endanger property values." The NAHB site also claims "the mortgage interest deduction helps make the tax code more progressive and primarily benefits middle class taxpayers ... 82 percent of households who benefit from the mortgage interest deduction have incomes of less than $200,000." The NAR's site notes that "almost two-thirds (64 percent) of the families who claim the mortgage interest deduction have household incomes between $50,000 and $200,000, and 42 percent have incomes of less than $100,000."

proposals. Specifically, this paper analyzes the impact of reducing the deductibility of mortgage interest in isolation, rather than reforming the tax code to treat owner-occupiers like landlords. Hence, I often use the word "subsidy" to refer to the tax expenditure created by the MID, rather than the non-taxation of the housing dividend.

Despite the political rhetoric surrounding the policy, several features of the MID's design combine to concentrate the policy's benefits among the wealthy. First, at a very basic level, the deduction only provides value to homeowners, who are wealthier than renters on average. Second, among homewners, the deduction's value increases with the size of one's mortgage. Then, because the subsidy is a tax deduction – not a credit – its value increases with the filer's marginal tax rate, as reductions in taxable income generate larger decreases in taxes in higher tax brackets. Finally, the option to claim the standard deduction prevents many lower-income homeowners from benefiting. In 2015, the last year of my sample, tax filers only benefited from itemizing if they could claim at least $6,300 (single filers) or $12,600 (joint filers) in total deductions. The TCJA's near-doubling of the standard deduction will only further stratify the benefits from 2018 forward.

In other words, the MID is not designed to target the middle class. To this point, Figure 1 compares the distribution of taxpayer income to the distribution of federal MID value (i.e. the reduction in tax obligations generated by the deduction) using IRS Statistics of Income data from 2015. Filers with Adjusted Gross Incomes (AGI) over $200,000 collect roughly one third of the benefits, despite the fact that they file only 4.5 percent of tax returns. Filers with AGIs under $100,000 also collect one third of the total benefits, despite comprising over 85% of tax units.

Figure 2 depicts how the distribution of deduction value has changed in recent years. Since the trough of the housing cycle in 2008, the benefits have increasingly concentrated among high earners, as house prices in high-income areas have increased and new mortgage issuance to middle and lower-middle income households has stalled.

Static calculations like those in Figures 1 and 2 tell only part of the distributional story; after all, statutory incidence is not economic incidence. Like any tax or subsidy, the primary distributional impact is the sum of two effects: the direct transfer of funds (in this case in the form of a reduced tax bill) and the ensuing change in market prices (Kotlikoff and Summers 1987). The next subsection provides a theoretical framework for tracing the effects of price changes through the population in order to analyze the distributional impacts of the subsidy.

2.2 Incidence

The primary goal of the empirical part of this paper is to measure the sensitivity of house prices to the tax rate applied to mortgage interest. While this parameter is interesting and important in its own right, it is also a necessary input for analyzing how the deduction's impacts are divided among various participants in the housing markets.

This section presents a straightforward, public finance-style model that makes this connection explicit. Making few structural assumptions, this sort of theory allows me to translate these reduced-form estimates into quantities that are informative about the policy's welfare impacts. These "sufficient statistics"-style models are common in the empirical public finance literature (see Chetty and Finkelstein (2009)), though this setting differs in two important respects. First, only certain consumers are subsidized, as renters and non-itemizers receive no tax benefits. Second, the tax benefit is applied in the capital markets, rather directly to final goods. Thus, even though much of the intuition follows from workhorse treatments of economic incidence (e.g., Kotlikoff and Summers (1987)), it is necessary to formally model the relevant decisions and institutional factors.[10]

2.2.1 The User Cost of Owner-Occcupied Housing

The decision to purchase a home depends on a host of prices, tax rates, and features of the mortgage market. Accordingly, calculating the true economic cost of homeownership is a non-trivial exercise. The dominant method among housing and public finance economists is to group the various prices, institutional factors, and non-monetary costs associated with a home purchase into a single flow cost known as the user cost of housing. As shown by Poterba (1984), the user cost formulation facilitates analysis of various aspects of the tax regime on the true economic cost of homeownership. In the ensuing years, empirical analyses of the cost of homeownership have largely converged on an agreed specification of the user cost, discussed in detail in Himmelberg, Mayer, and Sinai (2005), Poterba and Sinai (2011) and Albouy and Hanson (2014).

Following those authors, the economic cost of purchasing an additional dollar's worth of hous-

[10] Rappoport (2016) uses a similar set-up to derive incidence expressions for itemizing home-owners. The approach here differs in several respects. First, it incorporates a more fully developed treatment of the user cost of capital, as described in the following subsection. Second, it explicitly accounts for impacts on renters and non-itemizers. Finally, as emphasized earlier, capitalization effects are estimated directly, not calibrated.

ing under can be written as follows:

$$c^{itemize} = [1 - \tau_y \phi_m \lambda - \tau_y (1 - \lambda)] r_f - \phi_m \tau_y \lambda (r_m - r_f) + (1 - \phi_p \tau_y) \tau_p + m + \sigma - \pi_e \qquad (1)$$

where r_f and r_m are the interest rates on a risk-free asset and mortgages, respectively; λ denotes the loan-to-value (LTV) ratio; τ_y and τ_p are income and property tax rates; ϕ_m and ϕ_p indicate the proportion of mortgage interest and property taxes that are deductible; m stands for combined maintenance and depreciation, assumed to be a fixed proportion of house value; σ is a risk premium; and π_e denotes the expected rate of capital gains on housing.

This formulation allows for separate marginal tax rates on income and the deductions. The marginal tax rate on interest income τ_y is assumed to be equal to the rate on labor income. To allow for restrictions on deductions, I introduce a parameter $\phi_m \in [0, 1]$, which governs the portion of mortgage interest that may be written off. Hence, $(1 - \phi_m \tau_y)$ denotes the effective after-tax price of mortgage interest, which some authors refer to as the mortgage subsidy rate (Hilber and Turner 2014). The joint cost of debt and equity is therefore discounted by an LTV-weighted average of the two tax rates. This gives rise to the first term in (1).

Mortgages include the option to default or pre-pay if the economic environment changes, and as a result borrowers are charged a premium above the risk-free rate. This option is a benefit to buyers, and its value should be accounted for in the user cost. I follow Poterba and Sinai (2011) and assume that those options are priced fairly, implying that interest expenses above the risk-free rate are exactly offset by the value of the option. Hence, the second term reflects only the the tax savings generated by interest deductibility: $\phi_m \tau_y \lambda (r_m - r_f)$.

Households pay property taxes equal to a fraction τ_p of house value, and some portion of those taxes may be deducted from income tax returns. I allow property tax deduction policy to differ from mortgage interest deduction policy (i.e. ϕ_p may not equal ϕ_m), though in practice they are usually the same. Maintenance m includes both physical depreciation and the costs of upkeep, and the risk premium σ reflects the costs of additional risk created by a larger housing position in the portfolio. Expected housing capital gains π_e are presumed to be untaxed, as is the case for the vast majority of American homeowners.[11]

Note that when mortgage interest is fully deductible ($\phi_m = 1$), Equation (1) implies that there

[11]Capital gains on owner-occupied homes are not taxed until they exceed $250,000 (for single filers) or $500,000 (for married filers). Accordingly, most papers in this literature make the simplifying assumption that housing capital gains are untaxed (Poterba and Sinai 2011, Albouy and Hanson 2014).

is no tax-advantage to debt financing relative to equity. Because the cost of equity is taxed at the same rate as debt, changes to marginal tax rates affect each side of the capital structure identically. This has led housing economists to conclude that the true subsidy to owner-occupied housing is the non-taxation of imputed rent – i.e. the value of of housing services provided by the home. As this value is the dividend produced by the housing asset, it would be taxed under a system that treated housing like other income-producing assets (or, equivalently, a system that treated owner-occupiers like landlords). Under this view, the value of the total tax subsidy to owner-occupied housing is computed by comparing the expression in Equation (1) to the user cost under such an alternative system (see Hendershott and Slemrod (1983), Gyourko and Sinai (2003), and Himmelberg, Mayer, and Sinai (2005) for examples and further discussion.)

This paper focuses on a different margin of policy adjustment, and I therefore use the term subsidy to refer to a different quantity than this strand of the literature. Specifically, I am ultimately interested in the tax subsidy to mortgage interest – or, more precisely, the reduction in user costs attributable to variation in ϕ_m holding other tax quantities constant. I focus on interest deductibility primarily because it is the more policy-relevant margin of adjustment. Calls for reducing or eliminating the MID have become more frequent in recent years, and indeed the recent TCJA moved policy along this dimension. Conversely, few if any mainstream political actors lobby for the taxation of homeowners' imputed rental income.

Thus far, the user cost formulation has focused on homeowners who itemize. When filing taxes, however, roughly half of households claim the standard deduction instead of itemizing. These filers do not benefit from the tax subsidy to mortgage interest or property taxes, though they may feel their impacts via capitalization effects. The user cost for these households can be expressed as follows:

$$c^{non-itemize} = [1 - \tau_y(1 - \lambda)]r_f + \tau_p + m + \sigma - \pi_e \tag{2}$$

It should be noted that Equations (1) and (2) reflect the marginal user cost – i.e. the price of the last unit of housing consumption – rather than the average user cost. This is a deliberate choice; the ultimate goal of introducing the user cost framework is to allow me to compute how small changes in various tax parameters affect borrowers, so the relevant variation is at the margin. In general, though, the average user cost may be different from the marginal user cost. Computing average user costs would require netting out the portion of mortgage interest below the standard deduction, accounting for the progressivity of the income tax schedule, and recognizing that households may choose a different debt-equity mix at different levels of housing expenditure. These considerations become relevant when modeling the impacts of large overhauls of the tax system; see Martin (2018) and Follain and Ling (1991) for further discussion.

Finally, when analyzing the effects of changing tax parameters in the following sections, I assume that pre-tax interest rates are not affected. This amounts to assuming that the impacts of small changes in the tax treatment of United States mortgages are negligible in the face of global capital markets, as is commonly asserted in empirical work on the MID.[12] Furthermore, I do not explicitly solve for induced changes in households' preferred capital structure. The literature has not reached a consensus on the magnitude of these responses (see Poterba and Sinai (2011) for an extended discussion). In my setting, the impacts of changes to mortgage balances are negligible due to envelope conditions that I discuss in Section 2.2.3. One might also worry that other features of the mortgage market, such as loan-to-value restrictions or the availability of pre-payment options, might be respond to changes in tax policy. This consideration is mitigated by restricting the analysis to small perturbations of current policy parameters.

2.2.2 The Household Problem

With this formulation of the cost of purchasing housing in hand, I can now turn to the individuals' decision. In addition to housing, households consume a numeraire consumption good c to maximize an increasing, concave, and twice-differentiable flow utility function $u(C_t, H_t)$ with a positive cross-partial derivative. Housing services can be either rented or owned. Renters pay a per-unit price of ρ_t for H_t^R units of rental housing, while owners procure H_t^O of housing by paying $c^i(M_t)p_t$. c^i, with $i \in \{\text{itemize, non-itemize}\}$ denotes the user cost of housing as defined in the last subsection, which may vary with the level of mortgage debt M_t. Owner-occupied and rental housing are assumed to be perfect substitutes in the utility function.

Household balance sheets are comprised of mortgage debt, owned housing, and liquid savings denoted by D_t. Financial allocations are subject to two constraints. First, mortgage debt cannot exceed a fixed fraction of housing wealth, denoted by κ. Second, deposits must be positive, preventing households from borrowing at the lower risk-free interest rate. All borrowing occurs through the mortgage market.

[12] Poterba and Sinai (2011), Albouy and Hanson (2014), Rappoport 2016, Martin (2018), and and Sommer and Sullivan (2018) all assume pre-tax interest rates are not affected by taxes or deductions.

Formally, then, the household solves the following problem:

$$\max_{C_t, H_t^O, H_t^R, M_t, D_t} u(C_t, H_t^O + H_t^R) + \beta \mathbb{E}_t \left[V_{t+1}(D_t, M_t, H_t^O) \right]$$

subject to:

$$C_t + c^i(M_t) p_t H_t^O + \rho_t H_t^R + D_t \leq y_t(1-\tau_y) + (1 + r_f(1-\tau_y)) D_{t-1}$$

$$M_t \leq \kappa p_t H_t^O$$

$$D_t \geq 0$$

In this setting, a household may choose to rent a home instead of owning for two reasons. First, renting may be a cheaper option for certain individuals. User costs c^i are heterogeneous across individuals; they decrease with marginal tax rates and access to credit (via lower offered mortgage rates).[13] Hence, lower-income households face a higher cost of financing a home purchase. Second, some households might prefer to buy a home at an LTV greater than κ, but are instead bound by the credit constraint. These households may build up savings and purchase a home in the future.

The continuation value $V_{t+1}(D_t, M_t, H_t)$ is left unspecified for tractability. For the key results in this section to hold, I do not need to impose additional structure on expectations, the evolution of income, or terminal budget constraints. One cost of this simplification is that these results only apply to the first year after a policy change.

2.2.3 The Incidence of Mortgage Interest Deductibility for Households

In this framework, the first-order welfare impact of a small change in the mortgage subsidy rate on household welfare is as follows. Precisely, I consider the impact of a increasing the mortgage deductibility parameter ϕ_m by an amount $(1-\tau_y)d\phi_m$, which produces a change in the total mortgage subsidy rate of $d(1-\tau^{MID})$. The following Proposition details the resulting changes in buyers' value functions:

[13]For parsimony, I have written the budget constraint using the user cost notation developed in the previous subsection. The full budget constraint for itemizers is: $C_t + p_t H_t^O + \rho_t H_t^R - M_{t-1}(1 + r_m \phi_m (1-\tau_y)) + D_t + \tau_p \phi_p (1-\tau_y) \leq y_t(1-\tau_y) + (1 + r_f(1-\tau_y)) D_{t-1} + M_t(1 + r_m \phi_m (1-\tau_y)) + p_{t-1} H_{t-1}^0 (1 - \delta - m - \pi_e)$ For non-itemizers, ϕ_m and ϕ_p are set to zero. Note that the risk aversion σ and the value of the default option are not explicitly included in the decision problem, but I include them in all empirical user cost calculations.

Proposition 1 *In the consumer maximization problem described in the previous subsection, a change to the mortgage subsidy rate $(1-\tau^{MID})$ induced by a small change in deductibility ϕ_m has the following effect on the welfare of itemizing buyers when the capital constraint does not bind:*

$$\frac{dV_t}{d(1-\tau^{MID})} = \ell_t \left[M_t r_m + c^{itemize} H_t^O \frac{dp}{d(1-\tau^{MID})} \right]$$

where ℓ_t denotes the Lagrange multiplier on the budget constraint and $c^{itemize}$ is defined in Equation 1. For non-itemizing buyers, the first term in brackets is dropped and $c^{itemize}$ is replaced with $c^{non-itemize}$, defined in Equation 2.

When the borrowing constraint is binding, incidence is increased by the amount $\ell_t(r_m - r_f)H_t^O \frac{dp_t}{d(1-\tau^{MID})}$

The two terms in brackets convey the subsidy's key tradeoff. $M_t r_m$ is the total mortgage interest, i.e. the target of the subsidy. This term captures the direct effect of the change in the subsidy rate. The second term is the user cost of housing multiplied by the capitalization effect. Because the strength of the subsidy increases as $(1-\tau^{MID})$ decreases, this term is negative. Therefore, analyzing the relative magnitudes of these two terms is key to understanding how the policy affects buyers.

When the LTV limit binds, capitalization imposes an additional cost on borrowers who are not able to select their optimal debt level. These buyers are forced to take on even more debt to cover the purchase when prices rise. For unconstrained buyers, envelope conditions ensure that the cost of this additional debt is second-order. However, these conditions do not hold for constrained buyers; they would increase their borrowing if the constraint were loosened. They therefore suffer a first-order cost associated with the need to acquire additional financing at a sub-optimal LTV. This cost is proportional to the spread between costs of debt and equity, as shown in the term in the sentence of the proposition.

Finally, the budget constraint multiplier ℓ_t is the marginal value of an additional dollar of after-tax income at the optimum. It converts the term inside the brackets – the total change in income generated by the policy change – into utility terms.

Notably, the incidence expression does not show the impacts of many possible dimensions of adjustment, such as changes in housing consumption, debt, or tenure. This is a direct consequence of the envelope theorem. Because they are chosen optimally, small changes in quantities have no first-order effect on welfare; their marginal benefit is exactly offset by their marginal cost. Prices, however, are not chosen optimally, and thus changes in p have first-order bite. Put differently, this result shows that the effect of subsidies on prices is a sufficient statistic for the first-order welfare impact, and we can make well-grounded welfare statements without specifying the full structure of the model.

2.2.4 Other Parties

To connect rents and prices, I assume a competitive rental market in which absentee landlords can borrow at rate r_l, pay property taxes at rate τ_p, and incur efficiency costs equal to γ percent of house value. A zero-profit condition for landlords therefore implies that $\rho = p(r_l + \tau_p + \gamma)$. Hence, the effect on rents can be inferred from the change in prices. Because renters are not subsidized, their incidence only reflects the increase in rental costs $H_t^R \frac{d\rho}{d(1-\tau^{MID})}$.

Finally, I assume that new houses are supplied by a competitive home construction sector with increasing and convex costs of construction. Incidence for suppliers is simply the change in housing prices multiplied by the total magnitude of new housing services produced that period.

2.2.5 Summary

To summarize the key results of this section, Table 1 reports the first-order effects of a small change in the magnitude of the tax rate, commonly referred to as the economic incidence, for various parties. To ease interpretation, it is useful to normalize each term by the value of local housing and to suppress the individual's marginal utility of wealth. The normalization serves two purposes. First, dividing out the value of housing converts the capitalization effects to semi-elasticities. This form facilitates comparisons to the rest of the literature and is has attractive empirical properties, as logging observed house prices reduces their considerable dispersion. It also eases interpretation of empirical estimates. Once the capitalization semi-elasticity has been estimated, the degree of incidence can be inferred from information on mortgage rates and user costs. In particular, the parameters in the main expressions are all on similar scales.

The first two rows show the normalized incidence effects for first-time homebuyers, presuming that they are not subject to the borrowing constraint. Row 1 begins with the incidence for itemizing first-time buyers. Home purchasers suffer from financing more expensive houses (the first term in their formula) but this cost is offset by the value of the subsidy (the second term). By contrast, homebuyers who do not itemize, however, are subjected to the price increase without a corresponding increase in the subsidy. To the extent that buyers are constrained by borrowing limits, they suffer an additional cost of being driven further away from their ideal debt-equity mix.

Table 1: Summary of Incidence Expressions

Itemizing First-Time Buyers	$c^{itemize}\frac{d\ln(p)}{d(1-\tau^{MID})} + r_m * LTV$
Non-Itemizing First-Time Buyers	$c^{non-itemize}\frac{d\ln(p)}{d(1-\tau^{MID})}$
Additional Cost for Constrained Buyers	$(r_m - r_f)\frac{d\ln(p)}{d(1-\tau^{MID})}$
Itemizing Existing Owners	$r_m * LTV$
Non-Itemizing Existing Owners	0
Renters	$\frac{d\ln(\rho)}{d(1-\tau^{MID})} = \frac{d\ln(p)}{d(1-\tau^{MID})}$
Suppliers	$\frac{d\ln(p)}{d(1-\tau^{MID})}$

Notes: The first five rows show the incidence of a small change to the mortgage subsidy rate $(1-\tau^{MID})$, normalized by house values. Renters' and suppliers' incidence is normalized by rents and the total value of housing supplied, respectively.

The magnitude of this quantity is difficult to assess. Explicit down-payment requirements likely vary across lenders and individuals, but I do not observe these terms in my data. Fortunately, this term is likely small relative to the primary capitalization effect, as it is proportional to both the interest rate spread and the probability that a borrower is constrained.

Existing homeowners, on the other hand, are insulated from the price increase. While they may benefit from an increase in the price of their asset, the change is offset by the increased cost of purchasing their next home. Hence, itemizers are exclusively affected by the change in subsidy value, while non-itemizers are completely unaffected. I assume existing homeowners are unlikely to be LTV-constrained. The vast majority of US mortgages amortize over time, which would lead borrowers away from the constraint.

Finally, renters and suppliers are exposed to the price effect but do not receive any direct subsidy payments. Developers and homebuilders capture some of the subsidy value when prices increase.[14] While I can not observe rents directly, I can infer renter incidence from the estimated price effect. Under the assumption that the landlord sector is competitive and the supply of capital is inelastic, the price-to-rent ratio is constant. Thus, the increase in log rents equals the increase in log prices.

From an empirical perspective, the most important feature of the Table 1 is the prominence of the capitalization effect $\frac{d\ln(P)}{d(1-\tau^{MID})}$. In particular, it is the only quantity that cannot be directly

[14] The specific implications of the MID for developers and homebuilders is an interesting topic that I defer to future research.

inferred from data. Credible estimates of this parameter are therefore essential for understanding distribution of the subsidy's impacts. While prior work has obtained values for the price effect either by imposing significantly more structure on the problem (Sommer and Sullivan 2018) or by drawing existing supply and demand elasticities from the literature (Rappoport 2016), this project generates the first direct evidence of its magnitude.

It should be note that the incidence results in the previous subsection apply to a closed housing market, in which individuals must purchase or rent a home without considering alternative communities. In practice, of course, residents can choose to move in response to policy or price changes. The degree to which residents can opt out of the local market has important implications for our interpretation of capitalization effects.

For intuition, consider the limiting case of perfect mobility. If we believe that individuals are indifferent between living in either of two housing markets, then the law of one price will hold across these areas. Houses in either environment must sell for the same after-tax price, or everyone would live in the cheaper market. Thus, perfect mobility would seem to guarantee full capitalization, regardless of the elasticities of supply and demand.

I formalize this intuition in Appendix A.2 via a local labor markets model in which homeowners can choose to live in various cities, which may have different tax policies that affect the user cost of financing home purchases. Idiosyncratic tastes for locations generate limited mobility, as homeowners with a preference for a certain city will not immediately move in response to a favorable tax policy in another location. The model recovers the intuitions outlined above. In the closed-economy (zero-mobility) case, we recover identical predictions of capitalization rates as in the closed-economy model. As location preferences become less important, however, capitalization effects increase. Specifically, capitalization decreases with the variance of the location taste shock, which governs the extent to which mobility frictions reduce individuals' ability to arbitrage tax differences across markets by moving.

3. Empirical Approach

The primary empirical goal of this paper is to estimate the degree to which more generous mortgage interest subsidies drive up house prices – often referred to as capitalization. A simple approach to estimating capitalization might start with the following regression, in which the price P_{it} of house i transacting at time t is regressed on the state's net-of-tax rate applied to mortgage interest $(1 - \tau_{st}^{MID})$ and a vector of house characteristics X_i:

$$\ln(P_{ist}) = \theta(1 - \tau_{st}^{MID}) + \beta X_i + \varepsilon_{ist}$$

Depending on the contents of the control vector X_i, this equation nests the hedonic and repeat-sales specifications common in empirical studies of house prices. These designs are most likely to be valid when they examine variation within a single small geographic area. In these cases, it is reasonable to assume that other local determinants of house prices, such as access local public goods or weather, are common to all observations.

My sample spans much of the contiguous United States, however. At this level, there is surely heterogeneity in the local attributes that affect housing prices, and one might reasonably expect that such attributes are correlated with state tax policies. I therefore require an empirical approach that holds local confounders constant. A natural approach to controlling for unobserved local confounders is to add fixed effects for small geographic units. Restricting identifying variation so that if falls within, say, census tracts ensures that differences in prices is not being driven by amenities or economic conditions that are approximately constant within tracts.

Of course, MID policy does not vary within census tracts, as these tax rules are set at the state level. Thus, capitalization effects are not identified in models with tract, city, or county fixed effects. Any control strategy that relies on geographical fixed effects to control for unobserved local confounders must examine housing markets that lie in multiple states. In other words, plausible identification requires looking closely at variation near state borders.

I therefore exploit cross-border variation to estimate the impact of tax-induced differences in borrowing costs. By focusing on properties close to borders, these specifications are robust to local economic and geographic shocks that might be correlated with state tax policies. To begin, I construct all pairs of census tracts that border each other but lie on either side of a state border. To ensure that estimates are not affected by differences in the housing stock, I also also identify all houses within each pair that transact in the same quarter and match on observable characteristics.[15] Indexing matched house types by m, bordering county pairs by b, and states by s, I estimate variants of the following regression model:

[15]Specifically, I sort houses into coarse age bins before matching exactly on the age bin, number of bedrooms, number of bathrooms, condominium status, and whether the house is in a subdivision. I control directly for a quadratic polynomial in square footage in the regression.

$$\ln(P_{ihsbt}) = \theta(1 - \tau_{st}^{MID}) + \gamma_{hb} + \delta_{bt} + \varepsilon_{ihsbt} \quad (3)$$

The log-linear specification allows us to interpret θ as a semi-elasticity, consistent with the empirical literature estimating price responses to changes in interest rates (Adelino, Schoar, and Severino 2012; DeFusco and Paciorek 2017). γ_{hb} is a house-type-by-market effect, allowing the implicit price of individual housing characteristics to vary non-parametrically across markets. δ_{bt} is a flexible, market-level time trend that absorbs any local shocks common to each side of the border. Standard errors allow for arbitrary correlation within state-year combinations (the level at which treatment varies).[16] See Kroft et al. (2017) and Ljungqvist and Smolyansky (2014) for similar approaches to estimating the effects of local taxes using a border design.

Because marginal tax rates vary widely across the population, the specific construction of the dependent variable of interest τ_{st}^{MID} deserves some discussion. First, the log-linear specification allows me to focus exclusively on state tax policies, as the time-varying fixed effects control for changes in federal tax policy. Hence, τ_{st}^{MID} is the state subsidy rate above and beyond the federal rate. Second, the marginal tax rate includes all relevant information for the states' tax treatment of mortgage interest. In addition to the state marginal tax rate, I incorporate information on state rules governing itemization, federal tax deductibility, and various caps and clawbacks in state tax codes. In the simplest case, this implies that states that do not allow the deduction of mortgage interest are assigned a value of zero.[17]

Finally, the ideal treatment variable would be the marginal effective tax rate for the marginal homebuyer in each market. The characteristics of the true marginal buyer are not known, however, so I approximate her rate by calculating the average marginal tax rate in each state, the smallest

[16]When estimating the panel models discussed in Section 5.3, I cluster standard errors at the state level, per the recommendation of Bertrand, Duflo, and Mullainathan (2004).

[17]Because the empirical variation in τ^{MID} combines both variation in tax rates and variation, the estimated parameter is not necessarily identical to the theoretical capitalization effect discussed in Section 2.2.3, which considered the impact of varying itemization exclusively. It is straightforward to show the effect of varying the marginal income tax rate τ_y equals the original incidence expression plus several additional terms. Precisely, $\frac{dV}{d(1-\tau_y)} = \frac{dV}{d(1-\tau^{MID})} + \lambda_t \left[-r_f(p_t H_t^O - M_t) + y_t + \frac{dg_t}{d(1-\tau_y)}\right]$. The first term in brackets reflects changes in the after-tax price of equity (the term is negative because increasing the tax rate reduces the opportunity cost of housing equity). The final two terms are the direct increase in income due to the tax change (y_t) and the change in lump-sum government transfers g_t, which will be negative as long as the government balances its budget.

If the sum of all three terms is non-zero, they represent a wedge between the estimated quantity and the parameter derived from theory. While the magnitude of this wedge is difficult to assess theoretically, its relevance can be tested empirically. I assess its importance by estimating models that rely solely on variation in itemization rules. Fortunately, these specifications deliver similar results, suggesting that the primary specifications effectively targets the theoretical parameter of interest.

geographic identifier available in the TAXSIM/SOI files. I discuss how I compute local average marginal rates in Section 4.2.

Figure 3 depicts the identifying variation I exploit to measure capitalization. The map plots the combined (state plus federal) top marginal tax rate applied to mortgage interest for all counties in my housing dataset that border a state with a different mortgage tax rate in 2015. The cross-border difference can result from either a difference in state tax rates or itemization rules. Redder (or darker) counties have higher effective tax rates; states with white counties either have no state income tax or do not allow their residents to claim the MID. As the map shows, there is considerable variation in tax policies at state borders.

Local effective tax rates are a function of both policy rules and the local population, since marginal rates depend on income and household size. If local populations sort to one side of the border, either in response to tax rules or other state-specific factors, then the local average tax rates will reflect both policy rules and endogenous population responses. Accordingly, I estimate instrumented specifications that rely solely upon policy-induced variation.

My primary specification uses simulated state tax rates computed for fixed populations to instrument for the true average marginal rate of the states' residents. The characteristics of these populations do not vary across states. As a result, any variation in these instruments is driven purely by differences in policy. I choose three populations to target different moments of the income tax schedule: a taxpayer in the highest tax bracket, a stylized first-time homebuyer, and the entire national population. I discuss the the construction of these variables in more detail in Section 4.2

With these simulated instruments in hand, the first-stage equation is as follows:

$$\begin{aligned}(1-\tau_{st}^{MID}) =& \beta_1(1-\tau_{st}^{HighIncome}) + \beta_2(1-\tau_{st}^{Homebuyer}) \\ &+\beta_3(1-\tau_{st}^{NationalAvg.}) + \eta_{bt} + \alpha_{hb} + \nu_{ihsbt}\end{aligned} \tag{4}$$

This equation is over-identified, though I will also show estimates that use each instrument individually. All over-identified models are estimated via efficient two-stage generalized method of moments (GMM), though the two-stage least squares estimator produces almost identical results.

A separate concern is that household sorting might directly impact housing prices. There is considerable evidence that the characteristics of one's neighbors affects house values; see Wong (2013) and Bayer et al. (2007) for examples and discussion. In my setting, one might expect richer households to sort to the low-tax side of the border, and these changes might well be capitalized into house prices. I address this threat to identification by using a state's itemization policies – which include no direct information about the state tax rates – as an instrument for the mortgage subsidy rate in a two-state least squares specification. This approach purges estimates of equation (1) of any potential endogeneity associated with the sorting based on tax rates. These specifications will deliver consistent results as long as households do not choose their home state to take advantage of itemization rules.

Even in the instrumented model, the necessary identifying assumption is that MID deductibility does not correlate with any local-area unobservables. To address this concern, I take advantage of the fact that, in addition to varying spatially, state tax policies also change during my analysis period. Panel A of Figure 4 shows a different source of variation: changes in policy over time. The six red triangles indicate changes in the mortgage subsidy resulting from phasing out the state deduction. Changes to state income tax rates are more frequent, but often small. The seven labeled dots show changes of at least 2.5 percentage points to the top state income tax rate. Panel B highlights all of the counties exposed to change in tax rules, due to a policy change in their own state or a border state.

Adding a state fixed effect to Equation 3 soaks up any unobserved local effects and uses these policy changes to identify θ.[18] Because I continue to include border-pair fixed effects, this procedure estimates the change in the cross-border price differential after either state changes effective tax rates. The cost of this approach is statistical power, since much of the potentially useful cross-sectional variation is absorbed. I therefore expand my sample to include all border counties (rather than tracts) when estimating specifications with state fixed effects. While this weakens the case that across-border houses are true substitutes, the difference-in-differences estimator is consistent under the weaker assumption that identical houses in adjacent counties exhibit parallel trends in prices.

4. Data and Descriptive Statistics

[18] I obtain similar results if I use county or tract fixed effects instead of state effects, but the computations are considerably more burdensome and the estimates are slightly less precise.

4.1 Housing Transaction Data

The primary data source for this project is a detailed dataset of residential property transactions assembled from public deeds records by CoreLogic. While historical coverage varies geographically, by 2015 the dataset records transaction details for properties in 2,853 out of the 3,141 counties in the United States. The dataset provides basic descriptive information about the property, including square footage, number of bedrooms, number of bathrooms, etc.; detailed geographic information; and details of any transactions occurring during the coverage period, including sale price, date, and the details of any liens placed on the property.

To add demographic information about borrowers, I merge the transaction records with data reported under the Home Mortgage Disclosure Act (HMDA). The loan-level HMDA data are anonymized, but, following methods in Ferreira and Gyourko (2011) and DeFusco (forthcoming), I am able to determine most homeowners' race and household income as reported on the mortgage application by matching on the identify of the lender, loan amount, origination year, and the property's census tract. Most importantly, the income variable allows me to impute marginal tax rates and the likelihood that a household itemizes.

The dataset offers an unusual amount of breadth and detail, both of which are necessary for my research design. The full dataset contains over 4.2 million residential transactions in 2015 alone, dwarfing smaller survey datasets used in previous analyses (Hilber and Turner (2014), Hanson (2012)). Furthermore, the detailed transaction information – in particular the recording of exact sale prices, loan balances, and transaction dates – is not available in large public datasets like the American Community Survey, and its transaction-level granularity allows much more detailed study than the county-level averages published in the IRS Statistics of Income database.

The top panel of Table 2 reports summary statistics for the housing transactions sample. To facilitate comparisons between the border sample and states' interiors, Column (1) shows averages for non-border tracts, and (2) reports values for tracts on state borders. As reflected in their housing stock, border tracts are somewhat less well off than non-border tracts. The average house in a border region sells for $21,800 less than the average interior home, and it is 135 square feet smaller. Border homes are also slightly older and less likely to have three or more bathrooms. These differences are statistically significant, but their magnitudes do not raise grave concerns about representativeness.

My empirical approach requires a strict filter on data quality, particularly with respect to characteristics that affect the value of housing. I therefore do not include housing transactions for which I do not observe the number of bedrooms, the number of bathrooms, the house's age, or its square footage. 50.4% of transactions in border tracts are affected by this screen, reducing my final analysis sample to 425,017 transactions. Column (3) shows descriptive data for this sample. Houses in the final sample are noticeably larger than those in the full border sample (1,920 square feet on average, relative to 1,562), but other characteristics are well balanced.

4.2 Tax Data

I collect the details of state tax policies from TAXSIM, a software package that allows users to calculate federal and state income tax burdens for a wide range of taxpayer characteristics.[19] Importantly, the calculations incorporate much of the complexity of state tax codes that is not captured by state income tax rates alone. This includes information about mortgage interest deductibility as well as adjustments to effective tax rates attributable to the deductibility of state taxes on federal tax returns. Hence, for a given set of taxpayer characteristics, I can compute a reliable measure of the true after-tax price of mortgage interest.

In addition to the program that calculates tax obligations, I make use of the underlying SOI data files that detail the tax-filing characteristics of a representative sample of taxpayers. The finest level of geographic information is state of residence, so I can only compute averages within states, rather than the tract or county.[20] By varying the amount of mortgage interest reported by a small amount, I can infer the effective subsidy rate for all taxpayers in the sample.

Average marginal tax rates reflect characteristics of local populations that are potentially endogenous. It is therefore important for me to separate variation stemming from differences in tax bases from differences in tax policies, which are plausible sources of exogenous variation. I therefore use the TAXSIM programs to calculate simulated tax rates for three fixed populations: homeowners paying the top marginal tax rate, a stylized first-time homebuyer, and the national average. They then serve as useful instruments for the true average local tax rate.

For clarity, I can walk through the process of calculating the simulated tax rate for the national population. I start by collecting the full sample of taxpayer characteristics for a given year. Then, in the first step, I compute state and federal tax rates as if the full population lived in Alabama in

[19]Originally developed by Feenberg and Coutts (1993), the model and associated files are available at www.nber.org/~taxsim.

[20]State of residence is only explicitly recorded through 2008. In subsequent years, home states are imputed with the goal of matching the state-level distribution of key tax characteristics. I use these imputations in my main results, though my estimates are quite similar when I use 2008 population characteristics instead.

2008. After computing and storing the average rate for this population, I proceed to loop over the other 49 states and seven years in my sample, calculating average tax rates for the same population as if they lived in every state tax regime. I then repeat this entire process for a high-income taxpayer and a stylized first-time homebuyer, providing me with three instruments that provide variation along different sections of the income tax schedule.

Crucially, all taxpayer characteristics are held constant when I calculate each simulated rate. Any variation in the instruments is therefore driven by differences in state policies, not differences in state tax bases.

Summary statistics for tax variables can be found in Panel B of Table 2. Because the finest geographic information in the SOI microdata sample is the state of residence, I use publicly available aggregates to compute these averages. Specifically, tract averages are computed from the 2015 zip-code-level averages, using 2010 population weights to construct a crosswalk between zip codes and census tracts. The results reinforce the conclusion from Panel A that border tracts are somewhat poorer than average. Average Adjusted Gross Income was $65,288 in the border tracts (relative to $69,517 in the interior), and filers in border areas are roughly three percentage points less likely to itemize and claim the MID. Because I am measuring average mortgage subsidy rates at the state level, I am likely overestimating the tax rate at the border. This will only bias my estimates if the extent to which borders are not representative is correlated with tax policies. Classical measurement error in the treatment variable will be removed by the instruments.

4.3 Other Data Sources

While my main estimates rely only on the data sources described in the previous subsections, I draw on several commonly used data sources to construct control variables for robustness tests. Annual State GDP and county income data come from the Bureau of Economic Analysis' (BEA) Regional Economic Accounts, and county unemployment data is from the Bureau of Labor Statistics' (BLS) Local Area Unemployment database. The National Center for Education Statistics' F-33 survey provides spending and enrollment data for school districts, which I assign to census tracts using a block-weighted crosswalk. Finally, I obtain state government spending and inter-governmental spending within states from the Census Bureau's annual State Government Finances survey, which are then converted to per-capita measures using the Census' state population estimates.

5. Results

5.1 Border Design Estimates of Capitalization

This section presents my main estimates of the capitalization effect of the MID, specified as the relationship between log prices and the average marginal net-of-tax rate in the local area (denoted by $(1 - \tau_{MID})$). As a first step, therefore, I first regress log house prices on the local average marginal income tax rate without any instrumentation. Panel A of Table 3 shows the results of estimating Equation 3 directly via OLS. The resulting estimate is small, imprecise, and not statistically differentiable from zero: 0.299 (0.321).

As previously discussed, regressing prices on average marginal tax rates invites a host of endogeneity concerns. Average tax rates in a state depend on both policy rules and characteristics of the local population, including income, household type, etc. These population characteristics may respond directly to policy, either through the direct effects of local taxes and spending, through sorting, or through the simultaneous determination of demographics and policy via the local amenity and productivity factors.

Purging these estimates of endogeneity requires measures of the tax burden that do not depend on the local population. I therefore construct three statistics that that rely exclusively on variation induced by state policies. This approach, often referred to as "simulated instruments," is common in empirical public finance (Currie and Gruber 1996). I use state tax rules to calculate average marginal tax rates for one of three fixed populations: the national average, an individual paying the the top marginal income tax rate with a large mortgage, and a representative first-time homebuyer.

Panel B shows the results of estimating Equation 3 using these measures of the local tax rate. The cost of using these simulated rates is that they may accurately measure the local tax rate, since they remove the influence of local demographics by design. Thus, these regressions are best thought of as reduced-form approximations to the true relationship. Consistent with the hypothesis that a more generous tax deduction increases house prices, the coefficients are all negative and statistically significant.

The magnitudes in Panel B are difficult to interpret, however. The ideal specification would combine the local measure used in the naive OLS specification in Panel A with the stronger exogeneity claim of the simulated instruments.

Therefore, Panel C shows specifications that use the policy variables to instrument for the true average marginal tax rate. Columns (5)-(7) show results using each of the three policy measures individually. The estimates range from -0.948 (0.243) to -1.315 (0.309), implying that a one percentage point increase in the tax rate applied to mortgage interest raises house prices by 0.9%-1.3%. To exploit the full variation from all three instruments, my preferred specification pools all three instruments together into a single first-stage. Estimating this over-identified model via two-step GMM produces a semi-elasticity of -0.830 (0.229), as seen in Column (7).

To put this magnitude in perspective, recall from Table 1 that the incidence for itemizing buyers is proportional $c^{itemize}\frac{d\ln(P)}{d(1-\tau^{MID})}+r_mLTV$. Under commonly accepted parameter values,[21] the user cost for an itemizing homebuyer facing a 25% marginal tax rate and making a 20% down payment is 4.78%. For this purchaser, then, the price response implied by the overidentified model negates 84% of the tax benefit provided by the MID. If this household's LTV falls below 60%, the incidence turns negative. Indeed, as I discuss further in Section 6, adding non-itemizers to this calculation implies that the average new homebuyer would prefer a reduction in deductibility.

There are several reasons that we might expect such strong capitalization of the MID. The first stems from the non-linearity of debt subsidies. As originally noted by Rappoport (2016) and echoed by Sommer and Sullivan (2018), the combination of inelastic housing supply and less-than-complete financing of home purchases can result in homebuyers being harmed by borrowing subsidies.[22] Second, the marginal buyer in certain areas may have different characteristics than the average buyer. To the extent that wealthier itemizers "set the market" for land and houses, local prices may be determined by purchasers with costs of capital than the average resident's. Finally, as previously discussed, capitalization effects may be stronger in the empirical setting analyzed here, where moving across state lines is a relatively cheap option. Accordingly, we should not be overly surprised to see large price differentials across state borders.

5.2 Robustness of Border Design

The estimates discussed thus far use a flexible parameterization to control for house quality and local shocks. Specifically, border-by-quarter fixed effects soak up any shocks to local housing markets, while house-type-by-border effects account for local valuations of different house characteristics. Identification stems from comparisons between nearby homes that transact in the same

[21] See Section 6.

[22] See Appendix A.1 for further discussion

period, with additive quality adjustments specific to each location.

Table 4 shows that these results are robust to a wide variety of other strategies to control for house quality and a local shocks. Column (1) replicates the GMM results from Table 3 for ease of comparison, and the next five columns report similar specifications using various combinations of fixed effects for house types, border pairs, and time periods. All estimates are within (or nearly within) a standard error of the main specification. In a fully-saturated model with border pair by house type by quarter fixed effects, the estimated coefficient is -0.734 (0.262).

The rightmost three columns report results from hedonic regression specifications. To ensure covariate overlap, columns (8) and (9) restrict the sample to transactions for which an observationally identical house sells across the state border in the same quarter or year, respectively. Rather than using fully-interacted house type fixed effects, these specifications control for separate effects of bedrooms, bathrooms, house age, condominium status, subdivision status, and a quadratic in square footage. The resulting point estimates are slightly smaller, ranging from -0.572 (0.231) to -0.626 (0.243), but still quite large and statistically significant.

Having shown that the capitalization estimates in Table 3 do not depend on the specific functional form used to control for local shocks, I turn to the possibility that other characteristics might also change sharply at state borders and confound my estimates. Recall that my key identifying assumption is that observationally identical houses on either side of a border are close substitutes. If local economic outcomes, amenities, or policies co-vary in the same way, my results could be biased.

I can test for such bias by adding controls for local characteristics and observing how the coefficients change. This intuition – formalized by Altonji, Elder, and Taber (2005) and Oster (forthcoming) – follows from the observation that, even if the perfect set of control variables is not available to the researcher, controlling for imperfect proxies nevertheless conveys information about the possible bias. For instance, if a coefficient suffers from omitted variable bias, then adding controls that are correlated with the confounders will cause the targeted coefficient to shift towards the true, unbiased value. Conversely, if the estimates remain stable as new controls are added, the identifying variation is not threatened by the hypothesized confounder.

For this exercise to be informative, it is important to select a broad set of market-level control variables that are capitalized into house prices. The canonical models of spatial equilibrium and local public finance offer some guidance here. Rosen (1979) and Roback (1982) suggest focusing on measures of economic productivity and local amenities. I use local unemployment rates and average incomes to measure productivity. While amenities are difficult to measure – and most are unlikely to change sharply at state borders – public school spending is both measurable and likely to change at state borders. Hamilton (1975, 1976) also highlights the importance of local taxes and public goods. School quality is a particularly salient and measurable local public good. Local tax rates are much more difficult to measure, but I can use state government revenues as a measure of state-provided government services and inter-governmental transfers to proxy for the extent to which state governments provide resources to local entities. Furthermore, because schools receive a large share of their funding through local property taxes, education spending will be closely related to property tax burdens.

Table 5 shows specifications that control for a variety of potential confounders, along with uncontrolled estimate in the leftmost column (1) for ease of comparison. Reassuringly, the capitalization estimates are strikingly stable. Panel A uses the locally measured control variables – unemployment, mean income, and school expenditures. The largest deviation comes from the specification controlling only for unemployment.[23] The point estimate increases in magnitude by 0.140 log points to -0.970 (0.255). Controlling for all three local measures produces an estimate of -0.904 (0.263).

Results in the bottom panel control for state-level economic and government expenditure variables. Adding all three controls results in an estimate nearly identical to the uncontrolled model: -0.832 (0.298). All of the estimates in Table 5 are less than one standard error from the original specification, leaving one to conclude that unobserved confounders are unlikely to drive the estimates.

Finally, it is possible that household sorting could bias my estimates. Variation in the MID subsidy is driven by both itemization rules and local tax rates, and evidence suggests that households – particularly high earning ones – react to state tax rates when choosing residences (Moretti and Wilson 2017). If the characteristics of these households affect local housing prices, then my estimates could be biased.

To assess this concern, I estimate models that do not rely on variation in tax rates. Instead, I construct a dummy variable reflecting whether each state allows its residents to claim the MID,[24]

[23]The negative and statistically significant coefficient on school expenditure deserves a brief comment. Because I do not observe property tax rates, the "effect" of school spending likely bundles the effects of educational investment and increased local taxes. As emphasized by Hilber (2017), regressions of house prices on either tax rates or public goods can be positive, negative, or zero under various plausible models of local public finance.

[24]Note that this requires me to drop states with no state income tax from my sample, as they do not have itemization policies to begin with.

and I use this variable to instrument for the local average marginal rate. These estimates are shown in Table 6. While the standard errors are considerably higher (reflecting the fact that I have drastically restricted the available identifying variation), none of the estimates are statistically distinguishable from my preferred specification. Models that control for all three local characteristics used in the previous table – useful for increasing statistical power – result in an estimated coefficient of -1.258 (0.544).

5.3 Panel Estimates

Finally, I estimate difference-in-difference style regressions that identify capitalization via changes in state tax policies. Because these models include both state fixed effects and border-pair fixed effects, they estimate the extent to which cross-border price differentials change after one of the states changes its tax policy. They are therefore robust to concerns that the cross-sectional variation in state tax rates is endogenous to economic shocks at the border. However, they will still be biased if changes in state tax rates are endogenous to border conditions, a possibility I address by including state-level economic controls.

Because these specifications are substantially more demanding of the data, I expand the sample to include all counties on state borders, rather than just census tracts. This increases my regression sample from 632,238 observations to over 21 million, which is necessary to provide sufficient statistical power for meaningful statistical inference in these models. While houses in neighboring counties are less comparable than houses in neighboring tracts, the addition of state fixed effects adds a layer of robustness to the estimates that likely makes this tradeoff worthwhile.

As expected, the estimates in Table 7 are less precise than the cross-sectional estimates. Panel A uses the three tax rate instruments used in Tables 3-5. Without controls, the estimated capitalization semi-elasticity is -1.416 (0.554). Adding controls results in similar results, suggesting that endogenous tax rate changes are not a major threat.

The bottom panel uses panel variation in itemization policies exclusively. Panel B of Figure 4 shows that four states have either eliminated or capped deductions since 2008, providing sudden variation in mortgage deduction value that is not related to local tax rates. These estimates are comfortably in the range of the cross-sectional results, ranging from -1.098 (0.447) in the uncontrolled specification to -1.001 (0.457) in the model with all three state controls.

The specifications in Panel B are, in a certain sense, the most robust estimates shown in this paper. Most (if not all) of the threats to robustness described in the previous section are assuaged by the combination of state fixed effects and the suppression of variation in tax rates. The cost, in this case, is the loss of power and the need to rely on only three policy changes in a short panel. Nevertheless, it is reassuring that these specifications are statistically quite close to the main results.

5.4 Heterogeneity

It is reasonable to suspect that the average capitalization effects estimated thus far mask heterogeneity across geographic areas or sectors of the housing market. One might expect capitalization effects to vary within different segments of the housing market, perhaps because higher-income households are more attuned to the tax implications of their mortgage and housing choices.

To estimate submarket-specific capitalization rates, I separate my analysis sample based on a measure of house quality. I first estimate a hedonic regression of log prices on observable house characteristics. The fitted values from this regression form a measure of quality that depends on fixed housing characteristics, not observed prices. I then sort houses based on these predictions and group them into four evenly-sized quartiles. By sorting on the fitted values, I use the hedonic price function to infer house quality without explicitly conditioning on the dependent variable. Splitting the sample based on observed prices could bias the process, since the capitalization effect would influence both the observed dependent variable as well as the assignment to subsamples.

Table 8 shows estimates of the main specification for different segments of the house quality spectrum. The effect is monotonically increasing in house quality. Among the lowest quality houses (those in quartile one), a one percentage point rise in the after-tax rate $(1 - \tau^{MID})$ decreasing prices by 0.699 (0.365) percent. The price response increases in magnitude to -0.714 (0.278), -1.184 (0.326), and -1.306 (0.416) in quartiles two through four, respectively.

The quartiles in Table 8 are calculated within bordering tract-pairs. As a result, all of the variation is within markets, rather than across markets with different housing valuations. The same pattern holds if I calculate quartiles using the full-sample quality distribution, however. Those results can be found in Table A2; estimates range from -0.552 in the first quartile to -1.306 (0.416) in the fourth.

Taken together, these results imply that houses at the highest market segment are more responsive to credit subsidies than those at the bottom of the market. In the following section, I incorporate the heterogeneity of the capitalization rate across housing sub-markets when calculating the distribution of MID incidence for new home-buyers.

6. Interpretation and Distributional Impact

The final step of the analysis moves beyond average effects to consider the distribution of MID incidence. Here, I make use of the merged CoreLogic-HMDA sample described in Section 4.1. HMDA surveys include information on household income, which allows me to impute marginal tax rates and the likelihood of itemizing for individual households.

I focus on the distribution of incidence for new homebuyers. These effects reflect the most interesting heterogeneity, as incidence depends on the joint distribution of tax rates, itemization status, and mortgage balances. Renters receive no benefits from the MID, and I do not observe them in the data. Existing homeowners are effectively hedged against capitalization effects, because any capital loss they experience is offset by savings on their next purchase.[25] While I cannot explicitly separate first-time buyers from repeat purchasers in the data, I impose sample restrictions meant to approximate the income and LTV distributions of first-time buyers.[26]

The parameters I use to compute user costs are detailed in Table 9. My housing transactions data and the available summary data extracts from the IRS both terminate in 2015, so I use the average 2015 interest rates in this analysis. Some variables, such as house prices and initial mortgage balances, are observed in the housing data. Elsewhere, I adopt the conventions of the literature when assigning values to other parameters of the user cost equation (Poterba and Sinai 2011, Martin 2018). Marginal tax rates and itemization rates are imputed based on the average values for households of the same income and county of residence using public IRS SOI tables.

I use the framework developed in Section 2.2 to compute the impacts of a one percentage point reduction in the deductible portion of mortgage interest. There are several reasons to consider a reduction in deductibility as opposed to an increase. The first is policy relevance. As previously discussed, few if any mainstream political actors advocate for additional tax preferences for mortgages. The second reason relates to the generalizability of my empirical findings. To the extent that residents are mobile across state lines, the capitalization effects that I measure may overstate the price response in other parts of the country. Supply responses dampen capitalization, but a market with readily available (unsubsidized) substitutes behaves much like an inelastically supplied market. This motivates the focus on a reduction in the subsidy, as the durability of the housing stock

[25] Capital gains/losses are more directly relevant for households who are close to retirement or who are considering substantial changes to their housing consumption. My dataset does not allow me to follow households over time, however.

[26] Because first-time buyers earn less and borrow more than the average buyer, I drop purchasers with incomes in the top 25% of their MSA's distribution or LTVs in the bottom 25%.

implies housing markets respond to negative demand shocks as if supply were inelastic (Glaeser and Gyourko, 2005).

Figure 5 plots the average change in deduction value, the average change in housing costs (i.e. price times user cost), and average incidence within twenty evenly sized income categories. The sample includes all residential housing transactions in 2015, subject to the income and LTV restrictions discussed above. Price effects are computed using the quartile-specific parameters shown in Table 8.

Among the lowest income households, the subsidy provides very little benefit, largely because few of these households itemize. Accordingly, the policy change benefits these households by lowering house prices. Moving up the distribution, the deduction value increases as itemization rates and tax rates increase. This value is roughly offset by the price effect, however, which is larger for more expensive homes (as discussed in Section 5.4). While the point estimates indicate that even wealthier households benefit from a weakened deduction, the total incidence is not statistically different from zero.

To supplement the national averages, Figure 6 shows similar results for four metro areas: New York, San Francisco, Cleveland, and Detroit. In addition to showcasing a diverse set of housing markets, analyzing this set of cities benefits from the additional robustness of effectively inelastic housing supply – New York and San Francisco because of strict housing regulations, Cleveland and Detroit due to long and on-going population declines. Certain households in the middle of the income distribution in New York and San Francisco, as well as the wealthiest San Francisco home purchasers, would be harmed by a reduced deduction (though the magnitudes are well within statistical range of zero). Otherwise, the patterns are similar to the national picture, subject to compositional differences due to the local income distribution. House price impacts are bigger in the more expensive metros, but these effects are offset by higher itemization rates.

Finally, Appendix Figures A1 and A2 show similar results over the LTV distribution, as opposed to income. Incidence is, surprisingly, relatively constant at different debt levels, even though the theory indicates that high-LTV buyers should benefit more than low-LTV buyers. That result applies when holding other factors constant. In this case, the greater tax benefit afforded to larger borrowers is offset by the facts that these borrowers are less likely to itemize, are in lower tax brackets (and hence have higher user costs), and live in less valuable homes.

7. Conclusion

This paper takes a modern approach to a central question in housing and public finance. I measure the economic incidence of the mortgage interest tax deduction using plausibly exogenous policy variation and a flexible, sufficient statistics economic framework, rather than deriving incidence effects from a calibrated model. The key empirical target – the capitalization of mortgage subsidies into house prices – is estimated by using a national housing transactions dataset and a research design that exploits sharp policy variation at state borders. Combining the estimates with the derived incidence formulas allows me to to map out the distribution of policy effects.

My preferred specifications indicate that a one percentage point increase in the tax rate applied to mortgage interest increases house prices by 0.8%. These estimates are robust to a variety of specification checks; controls for local economic and political factors; specifications that rely on itemization rules instead of tax rates; and difference-in-difference specifications that estimate capitalization exclusively via changes to state tax policies.

These magnitudes imply that all buyers purchasing with loan-to-value ratios below 60% would prefer a reduction in the subsidy rate. At 80% LTV, an itemizing buyer sees 84% of the tax benefit cannibalized by capitalization. These effects help explain why, per the most credible available evidence, mortgage interest deductions do not meaningfully increase homeownership (Gruber et al. 2017, Hilber and Turner 2014). To the extent that the subsidies are captured by homebuilders or existing homeowners, they provide little incentive to own instead of rent.

This paper focuses closely on the MID's incidence on homebuyers. Extending the analysis to other parts of the market using new data sources and research designs is a promising avenue for future research. In particular, understanding the policy impacts on market interest rates and incentives for supplying different types of homes would paint a fuller picture of the deduction's costs and benefits. Furthermore, as discussed earlier, elevated house prices and mortgage debt levels increase risk for both households and the macroeconomy. Quantifying the extent to which housing subsidies magnify such risks could improve the design of macroprudential policy.

While many policy discussions focus on the statutory regressivity depicted in Figure 1, accounting for house price effects points to a different driver of economic incidence. The estimated capitalization implies that homebuyers' economic benefits are at best small, even if they receive a tax reduction. The subsidy is largely captured by the supply side, ultimately accruing to either shareholders in the homebuilding industry or divesting homeowners. Thus, the benefits are indeed concentrated among a particularly well-off population, but one determined by asset holdings rather than income.

References

[1] Adelino, Manuel, Antoinette Schoar, and Felipe Severino (2012). "Credit Supply and House Prices: Evidence from Mortgage Market Segmentation." NBER Working Paper No. 17832.

[2] Albouy, David, and Andrew Hanson. 2014. "Are Houses Too Big or In The Wrong Place? Tax Benefits to Housing and Inefficiencies in Location and Consumption." In *Tax Policy and the Economy* Volume 28, Jeffrey Brown ed.

[3] Altonji, Joseph, Todd E. Elder, and Christopher R. Taber. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of Political Economy*, 113(1): 151–184.

[4] Argyle, Bronson, Taylor Nadauld, Christopher Palmer, and Ryan Pratt. 2018. "The Capitalization of Consumer Financing into Durable Goods Prices." NBER Working Paper No. 24699.

[5] Bayer, Patrick, Fernando Ferreira, and Robert McMillan. 2007. "A Unified Framework for Measuring Preferences for Schools and Neighborhoods." *Journal of Political Economy*, 115 (4): 588-638.

[6] Berger, Timothy, Peter Englund, Patric H. Hendershott, and Bengt Turner. "The Capitalization of Interest Subsidies: Evidence from Sweden." *Journal of Money, Credit, and Banking*, 32 (2): 199-217.

[7] Bertrand, Marianne, Esther Duflo, Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-In-Differences Estimates?" *Quarterly Journal of Economics*, 119(1): 249-275.

[8] Bhutta, Neil, and Daniel Ringo (2017). "The Effect of Interest Rates on Home Buying: Evidence from a Discontinuity in Mortgage Insurance Premiums." Finance and Economics Discussion Series 2017-086. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2017.086.

[9] Bourassa, Steven C., Donald R. Haurin, Patric H. Hendershott, and Martin Hoesli. 2013. "Mortgage Interest Deductions and Homeownership: An International Survey." *Journal of Real Estate Literature*, 21 (2): 181-203.

[10] Chetty, Raj, and Amy Finkelstein. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics*, 1(1): 451-488.

[11] Congressional Budget Office (CBO). 2018. "The Budget and Economic Outlook: 2018-2028." Available at https://www.cbo.gov/publication/53651

[12] Currie, Janet, and Jonathan Gruber. 1996. "Health Insurance Eligibility, Utilization of Medical Care, and Child Health." *Quarterly Journal of Economics*, 111(2): 431-466.

[13] DeFusco, Anthony. Forthcoming. "Homeowner Borrowing and Housing Collateral: New Evidence from Expiring Price Controls." *Journal of Finance*, forthcoming.

[14] DeFusco, Anthony, and Andrew Paciorek. 2017. "The Interest Rate Elasticity of Mortgage Demand: Evidence from Bunching at the Conforming Loan Limit." *American Economic Journal: Economic Policy*, 9(1): 210-240.

[15] Diamond, Rebecca. 2016. "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000." *American Economic Review*, 106(3): 479-524.

[16] Feenberg, Daniel Richard, and Elizabeth Coutts. 1993. "An Introduction to the TAXSIM Model." *Journal of Policy Analysis and Management*, 12(1): 189-194.

[17] Ferreira, Fernando, and Joseph Gyourko. 2011. "Anatomy of the Beginning of the Housing Boom: U.S. Neighborhoods and Metropolitan Areas, 1993-2009." NBER Working Paper No. 17374.

[18] Follain, James R., and David C. Ling. 1991. "The Federal Tax Subsidy to Housing and the Reduced Value of the Mortgage Interest Deduction." *National Tax Journal* 44 (2): 147–68.

[19] Ganong, Peter, and Pascal Noel. 2018. "Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." NBER Working Paper No. 24964.

[20] Garrett, Daniel, Andrey Ordin, James W. Roberts, and Juan Carlos Suarez Serrato. 2017. "Tax Advantages and Imperfect Competition in Auctions for Municipal Bonds." NBER Working Paper No. 23473.

[21] Glaeser, Edward, and Jesse Shapiro. 2003. "The Benefits of the Home Mortgage Interest Deduction." NBER Tax Policy and the Economy, 17:37-82.

[22] Glaeser, Edward, and Joseph Gyourko. 2005. "Urban Decline and Durable Housing." *Journal of Political Economy*, 113(2): 345-375.

[23] Glaeser, Edward, Joshua Gottlieb, and Joseph Gyourko. 2013. "Can Cheap Credit Explain the Housing Boom?" In *Housing and the Financial Crisis*, Edward Glaeser and Todd Sinai (eds.), University of Chicago Press.

[24] Gruber, Jonathan, Amalie Jensen, and Henrik Kleven. 2017. "Do People Respond to the Mortgage Interest Deduction? Quasi-Experimental Evidence from Denmark." NBER Working Paper No. 23600.

[25] Gruber, Jonathan, and Emmanuel Saez. 2002. "The Elasticity of Taxable Income: Evidence and Implications." *Journal of Public Economics*, 84 (1): 1-32.

[26] Gyourko, Joseph, and Todd Sinai. 2003. "The Spatial Distribution of Housing-Related Tax Benefits in the United States." *Real Estate Economics*, 31(4): 527-576

[27] Hamilton, Bruce W. 1975. "Zoning and Property Taxation in a System of Local Governments." *Urban Studies*, 12: 205–11.

[28] Hamilton, Bruce W. 1976. "Capitalization of Intrajurisdictional Differences in Local Tax Prices." *American Economic Review*, 66: 743–53.

[29] Hanson, Andrew. 2012. "Size of Home, Home Ownership, and the Mortgage Interest Deduction." *Journal of Housing Economics* 21(3): 195–210

[30] Hanson, Andrew, and Hal Martin. 2016. "Metropolitan Area Home Prices and the Mortgage Interest Deduction: Estimates and Simulations from Policy Change." *Regional Science and Urban Economics*, 59: 12–23.

[31] Hendershott, Patric, and Joel Slemrod. 1982. "Taxes and the User Cost of Capital for Owner-Occupied Housing." *Real Estate Economics*, 10 (4): 375-393.

[32] Hilber, Christian. 2017. "The Economic Implications of House Price Capitalization: A Synthesis." *Real Estate Economics*, 45 (2): 301-3

[33] Hilber, Christian, and Tracy Turner. 2014. "The Mortgage Interest Deduction and its Impact on Homeownership Decisions." *Review of Economics and Statistics*, 96(4): 618-637.

[34] Himmelberg, Charles, Christopher Mayer, and Todd Sinai. 2005. "Assessing High House Prices: Bubbles, Fundamentals and Misperceptions." *Journal of Economic Perspectives*, 19(4): 67-92.

[35] Joint Committee on Taxation. 2018. "Estimates of Federal Tax Expenditures for Fiscal Years 2017-2012. Report JCX-34-18, available at https://www.jct.gov/publications.html?func=startdown&id=5095

[36] Kline, Patrick and Enrico Moretti. 2014. "People, Places and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs." *Annual Review of Economics*, 6.

[37] Kotlikoff, Laurence J. and Lawrence H. Summers. 1987. "Tax Incidence," *The Handbook of Public Economics*, edited by Alan J. Auerbach and Martin Feldstein, pp. 1043-1092.

[38] Kroft, Kory, Jean-William Laliberte, Rene Leal-Vizcaino and Matthew J. Notowidigdo. 2017. "Quantifying the Welfare Gains of Variety: A Sufficient Statistics Approach." Mimeo, available at http://korykroft.com/wordpress/KKLN_variety_july2017.pdf

[39] Ljungqvist, Alexander and Michael Smolyansky. 2014. "To Cut or Not to Cut? On the Impact of Corporate Taxes on Employment and Income." NBER Working Paper No. 20753.

[40] Martin, Hal. 2018. "The Impact of the Tax Cuts and Jobs Act on Local Home Values." Federal Reserve Bank of Cleveland Working Paper No. 18-06.

[41] Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics*, 128 (4): 1687-1726.

[42] Mian, Atif and Amir Sufi. 2011. "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis." *American Economic Review*, 101 (5): 2132-2156.

[43] Mian, Atif, and Amir Sufi. "Who Bears the Cost of Recessions? The Role of House Prices and Household Debt. *Handbook of Macroeconomics*, Volume 2, 255-296.

[44] Moretti, Enrico and Daniel J. Wilson. "The Effect of State Taxes on the Geographical Location of Top Earners: Evidence from Star Scientists." *American Economic Review*, 107(7): 1858-1903.

[45] Oster, Emily. Forthcoming. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business and Economic Statistics*

[46] Piskorski, Tomasz, and Amit Seru. 2018. "Mortgage Market Design: Lessons from the Great Recession." BPEA Conference Draft, Spring, available at https://www.brookings.edu/bpea-articles/mortgage-market-design-lessons-from-the-great-recession/

[47] Poterba, James. 1984. "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach." *Quarterly Journal of Economics*, 4(1): 729-752.

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

[49] Poterba, James, and Todd Sinai. 2011. "Revenue Costs and Incentive Effects of the Mortgage Interest Deduction for Owner-Occupied Housing." *National Tax Journal*, 64(2): 531-64.

[50] Rappoport, David. 2016. "Do Mortgage Subsidies Help or Hurt Borrowers?" Finance and Economics Discussion Series 2016-081. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2016.081.

[51] Roback, Jennifer. 1982. "Wages, Rents, and the Quality of Life." *Journal of Political Economy*, 90: 1257-1278

[52] Rosen, Sherwin. 1979. "Wages-based Indexes of Urban Quality of Life." In *Current Issues in Urban Economics*, edited by P. Mieszkowski and M. Straszheim. Baltimore: John Hopkins Univ. Press.

[53] Sommer, Kamila and Paul Sullivan. 2018. "Implications of the U.S. Tax Policy for House Prices, Rents, and Homeownership." *American Economic Review*, 108(2): 241-74.

[54] Suarez Serrato, Juan Carlos and Owen Zidar 2016. "Who Benefits from State Corporate Tax Cuts? A Local Labor Market Approach with Heterogeneous Firms." 2016. *American Economic Review*, 106(9): 2582-2624.

[55] Wong, Maisy. "Estimating Ethnic Preferences Using Ethnic Housing Quotes in Singapore." *Review of Economic Studies*, 80(3): 1178-1214.

Figure 1: Distribution of Income and Federal Mortgage Interest Deduction Value (2015)

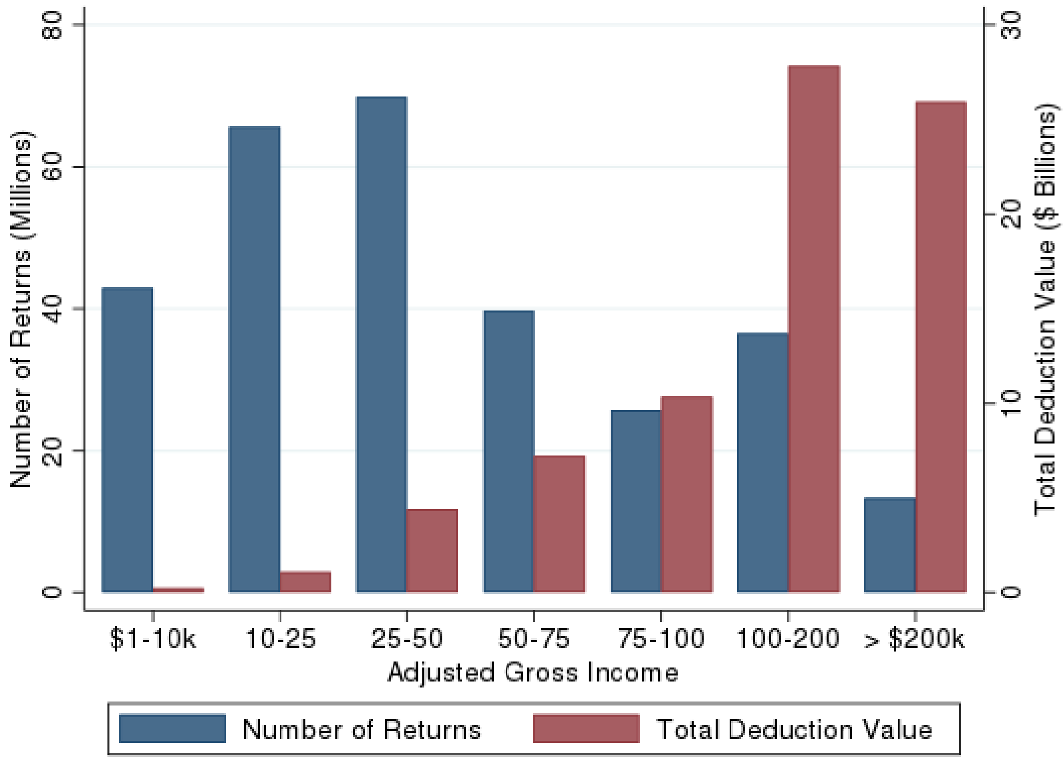

Note: Author's calculation using IRS Statistics of Income Data. Deduction value is calculated as the total amount deducted times the group's marginal tax rate, accounting for itemization rates and filing status. See Appendix for details.

Return to text.

Figure 2: Distribution of Federal Deduction Value Over Time

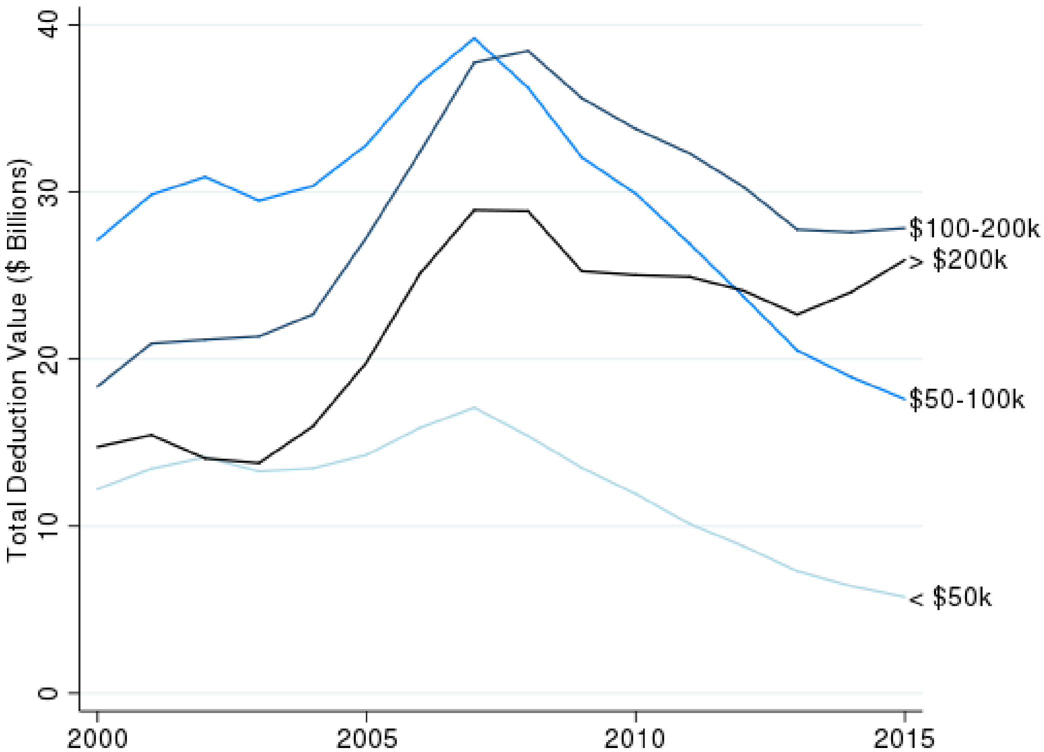

Note: Author's calculation using IRS Statistics of Income Data. Deduction value is calculated as the total amount deducted times the group's marginal tax rate, accounting for itemization rates and filing status. See Appendix for details.

Return to text.

Figure 3: Border Variation in Mortgage Tax Rates for Housing Sample

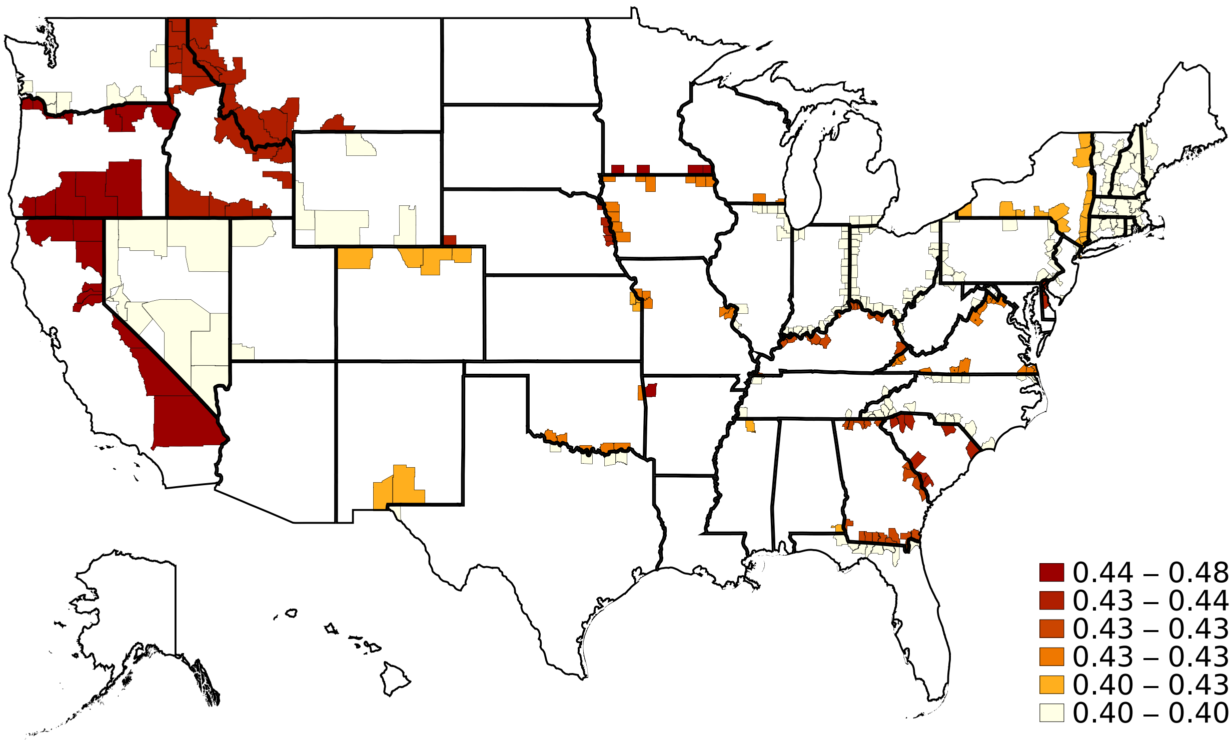

Note: The map plots all counties in the housing dataset that border a state with a different MID tax rate in 2015. The color denotes the top total (federal and state) marginal tax rate applied to mortgage income in 2015, accounting for all relevant clawbacks and deductions.
Return to text

Figure 4: Time-Varying Sources of Mortgage Tax Rate Variation

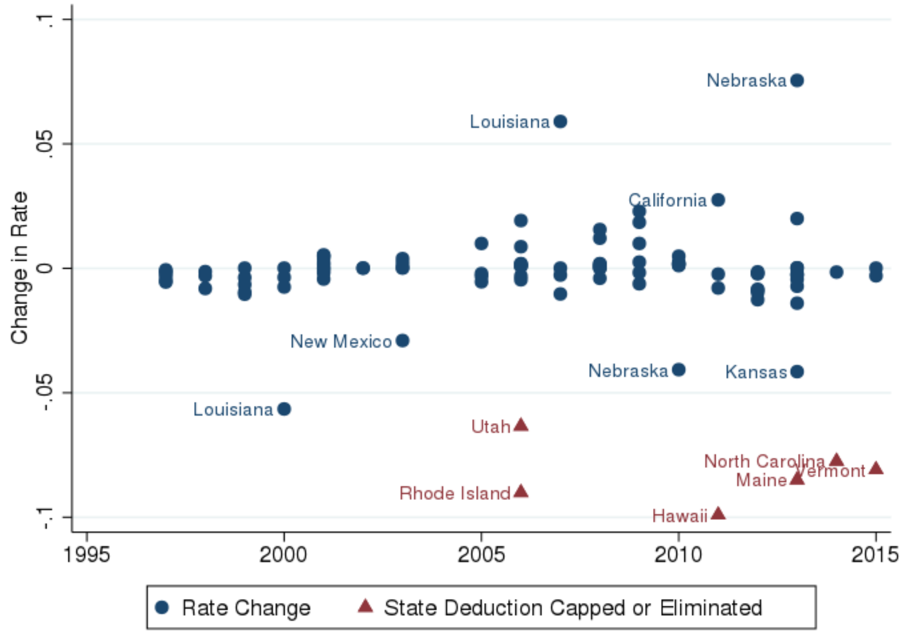

(a) Changes to Top State Mortgage Interest Tax Rates

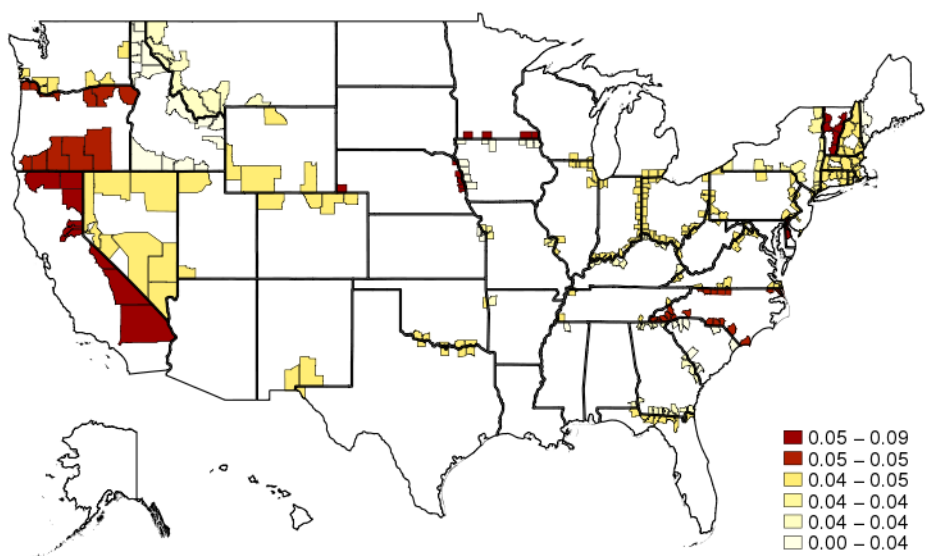

(b) Sample Counties Exposed to Rate Change, 2008-2015

Notes: Panel A plots all changes to the top state MID tax rate since 1996. Changes of at least 2.5 percentage points are labeled. Panel B plots all counties in the housing sample in states that change rates between 2008 and 2015, as well as counties neighboring those states. The color denotes the magnitude of the change in the rate.
Return to text.

Figure 5: Income-Specific Incidence of a Small Reduction in Mortgage Deductibility

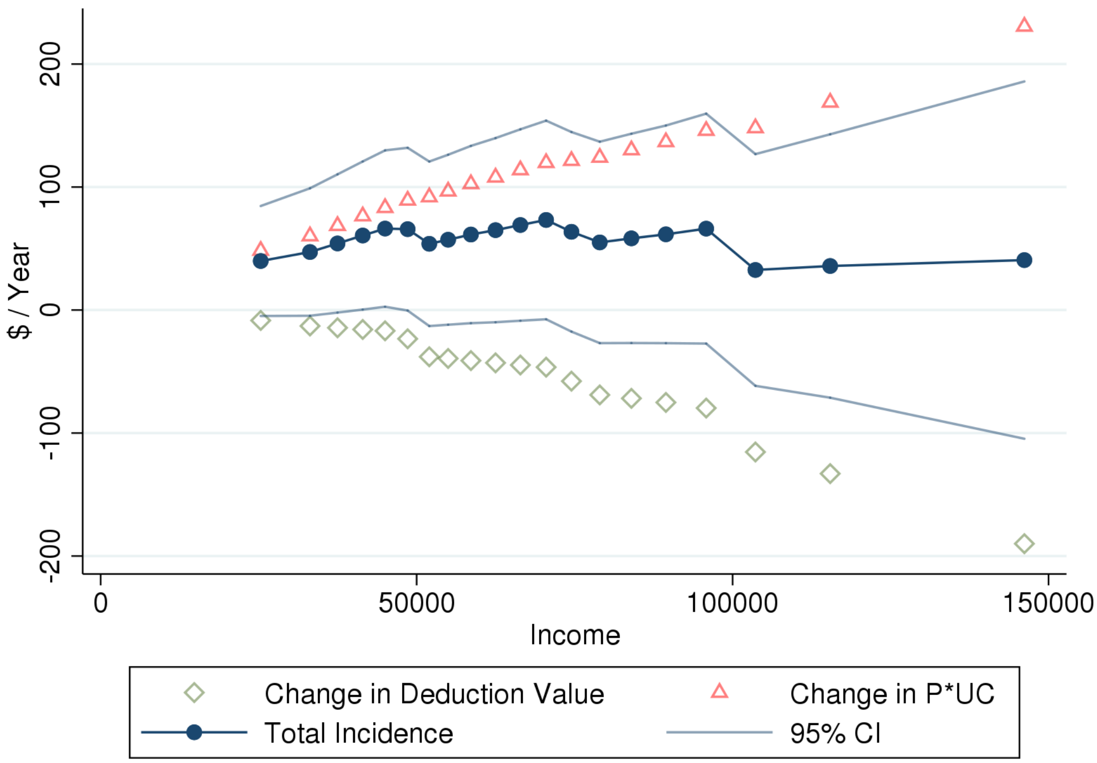

Notes: This figure reports the effect of a one-percentage point decrease in the tax rate applied to mortgage interest for new home-buyers. The sample includes all 2015 transactions in the merged HMDA-CoreLogic sample, dropping sales in the top 25% of their MSA's income distribution or the bottom 25% of their MSA' loan-to-value distribution. The plotted points show means computed within 20 evenly-sized income bins. Total incidence is the sum of the change in deduction value and the change in the flow cost of housing (Price * User Cost), and the associated confidence interval reflects statistical uncertainty in the value of the capitalization parameter.

Return to text.

Figure 6: Income-Specific Incidence of a Small Reduction in Mortgage Deductibility in Four Metro Areas

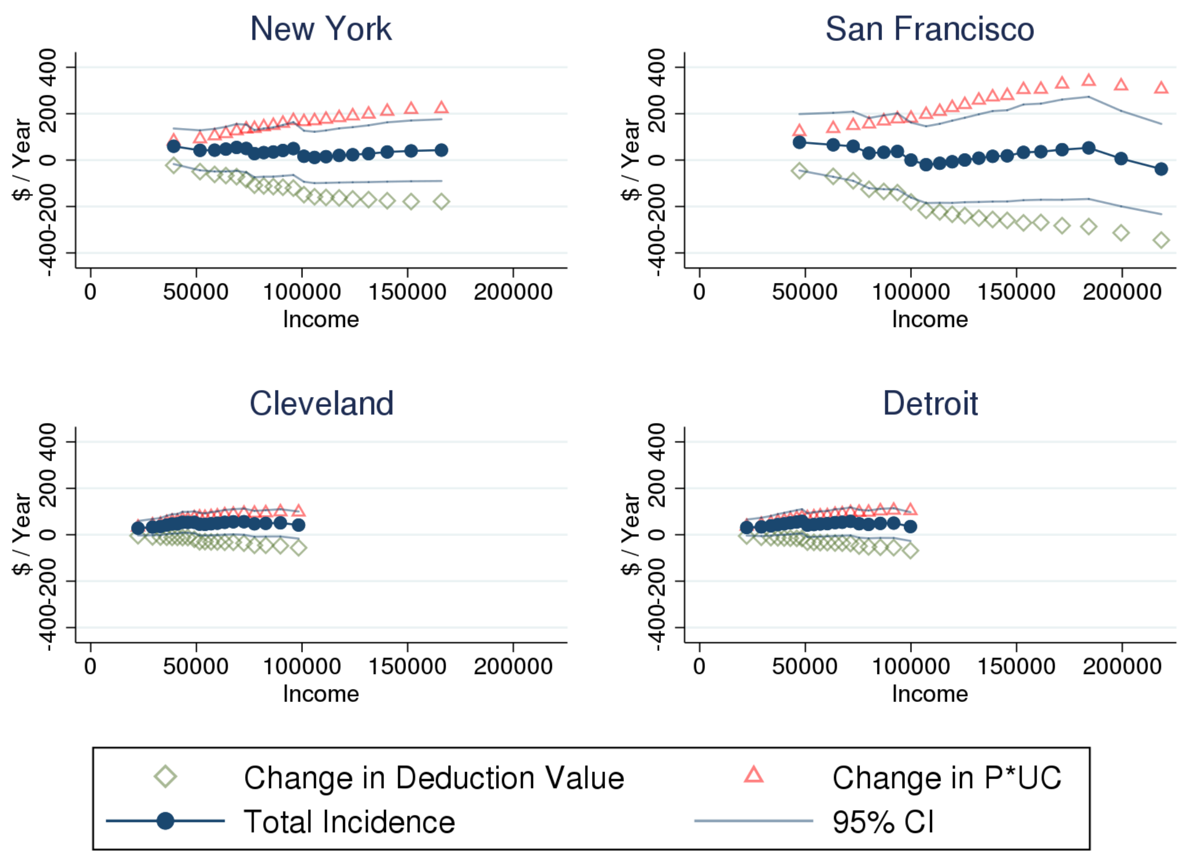

Notes: The four panels reproduce Figure 5 for specific metro areas.
Return to text.

Table 2: Summary Statistics

	Non-Border Tracts	Border Tracts	Analysis Sample
Panel A. Housing Transactions	(1)	(2)	(3)
Price (Thousands)	246.358	224.599	227.667
LTV at Origination	0.628	0.576	0.615
Square Footage	1,697	1,562	1,920
Three or More Bedrooms	0.782	0.789	0.791
Three or More Bathrooms	0.388	0.344	0.348
New Construction	0.049	0.038	0.044
Built in Last Decade	0.214	0.197	0.229
Condo	0.129	0.082	0.069
Subdivision	0.098	0.072	0.068
Missing Housing Characteristics	0.381	0.504	0.000
Number of Transactions	26,635,605	856,434	425,017
Panel B. Tax Filer Characteristics			
Adjusted Gross Income	69,517	65,288	-
Percent Itemizing	0.306	0.270	-
Percent Claiming MID	0.221	0.193	-
Avg. MID Per Claimer	8,624	7,724	-
Number of Tracts	69,705	3,744	-

Notes: Panel A shows average characteristics of housing transactions for various subsets of the CoreLogic sample. Column (1) contains all single-family residential transactions between 2008 and 2015. Column (2) restricts this sample to transactions in tracts that fall on state borders, and Column (3) includes only observations with non-missing housing characteristics (bathrooms, bedrooms, square footage, and age) and dropping extreme outliers (square footage over 10,000 or a price outside of [10,000 , 2,500,000]. Panel B shows average characteristics of tax units in these districts, drawn from the IRS Statistics of Income tables. These variables are measured at the zip code level and mapped to tracts via population-weighted averages.

Table 3: MID Capitalization Across State Borders

Panel A. Naive OLS	(1)			
$(1-\tau_{MID})$	0.299			
	(0.321)			
Panel B. Reduced Form	(2)	(3)	(4)	
$(1-\tau_{MID}^{NationalAvg.})$	-0.689***			
	(0.180)			
$(1-\tau_{MID}^{Homebuyer})$		-0.820***		
		(0.186)		
$(1-\tau_{MID}^{HighIncome})$			-1.037***	
			(0.281)	
Panel C. Instrumented Models	(4)	(5)	(6)	(7)
$(1-\tau_{MID})$	-0.948***	-1.315***	-1.165***	-0.830***
	(0.243)	(0.309)	(0.320)	(0.229)
IVs	$\tau_{MID}^{NationalAvg.}$	$\tau_{MID}^{Homebuyer}$	$\tau_{MID}^{HighIncome}$	All
Observations	632,238	632,238	632,238	632,238
Brdr-By-Type FE	X	X	X	X
Brdr-By-Qtr FE	X	X	X	X

Notes: This table reports the results of estimating Equation 3 using several independent variables and simulated instruments. The dependent variable in each column is log house price. The independent variable in Panels A and C is the average marginal net-of-tax rate applied to mortgage interest in each state and year $(1-\tau_{MID})$. The independent variables in Panel B are constructed by applying each state's tax rules in each year to one of three fixed populations: a nationally representative sample of taxpayers ($(1-\tau_{MID}^{NationalAvg.})$), a stylized first-time homebuyer ($(1-\tau_{MID}^{Homebuyer})$), and a high-income household paying the top marginal tax rate ($(1-\tau_{MID}^{HighIncome})$). Panels A and B are estimated by OLS; Panel C uses the independent regressors from Panel B as instruments for the true average marginal tax rate. The over-identified model in Column (7) is estimated via two-stop efficient GMM. The sample includes all housing transactions between 2008 and 2015 that occur in a census tract on a state border. Standard errors allow for clustering at the state-year level. Return to text.

Table 4: Robustness of Border Design Specification

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Fixed Effect Specifications						*Matched Hedonic Specifications*		
$(1-\tau^{MID})$	-0.830***	-0.632***	-0.590**	-0.646***	-0.734***	-0.880***	-0.626**	-0.572**	-0.581**
	(0.229)	(0.241)	(0.230)	(0.222)	(0.262)	(0.249)	(0.243)	(0.231)	(0.240)
Observations	632,238	650,819	101,046	657,289	385,552	550,462	650,819	101,046	181,133
R-squared	0.722	0.664	0.663	0.620	0.801	0.743	0.663	0.663	0.643
Match Sample	All	All	Qtr	All	All	All	All	Qtr	Yr
House Type FE		X		X					
Border Pair FE				X					
Brdr-By-Type FE	X								
Brdr-By-Qtr FE	X	X	X				X	X	X
Brdr-By-Year FE									
Brdr-By-Type-by-Qtr FE					X				
Brdr-By-Type-by-Year FE						X			
StateBorder-By-Qtr FE				X					
Hedonic Controls							X	X	X

Notes: This table reports the results of estimating Equation 3 using the over-identified model described the Table 3 with various subsets of controls and fixed effects. Sample and variable definitions follow the notes of Table 3 except where noted. Columns (3) and (8) restricts the sample to houses that transact in the same quarter as an identical house across a state border. Column (9) restricts the sample to matched transactions occurring in the same year. Hedonic controls include the number of bedrooms, number of bathrooms, condominium status, presence in a subdivision, house age, and a quadratic in square footage. Standard errors allow for clustering at the state-year level.

Return to text.

Table 5: Robustness of Border Design to Local Economic and Policy Controls

Panel A: Local Controls	(1)	(2)	(3)	(4)	(5)
$(1-\tau_{MID})$	-0.830***	-0.970***	-0.924***	-0.829***	-0.904***
	(0.229)	(0.255)	(0.224)	(0.226)	(0.263)
County Unemployment		-0.0377***			-0.0402***
		(0.00443)			(0.00601)
County Mean Income			0.000830		-0.00117
			(0.00140)		(0.00150)
Per-Pupil School Expenditure				-0.0142***	-0.0151***
				(0.00229)	(0.00219)
Observations	632,238	632,238	621,656	608,589	598,007
Panel B: State Controls	(6)	(7)	(8)	(9)	(10)
$(1-\tau_{MID})$	-0.830***	-0.833***	-0.846***	-1.023***	-0.832***
	(0.229)	(0.233)	(0.291)	(0.235)	(0.298)
State GDP		-0.000298			0.000323
		(0.00113)			(0.00131)
State Gov. Expenditure			-0.000727		0.0141
			(0.00753)		(0.0111)
Inter-Gov. Expenditure				-0.0504**	-0.0756**
				(0.0205)	(0.0328)
Observations	632,238	632,238	632,238	632,238	632,238

Notes: This table reports the results of estimating Equation 3 using the over-identified model described the Table 3 with various control variables. All regressions include border-x-quarter and border-by-house-type fixed effects. Sample and variable definitions follow the notes of Table 3. The units for all monetary variables are thousands of real 2016 dollars per person. Unemployment and mean income are measured at the county level. School spending is measured at the Census' Local Educational Area (school district) level; when tracts lie in multiple districts, they are assigned a weighted average of the overlapping districts' values. Standard errors allow for clustering at the state-year level.
Return to text.

Table 6: Border Design Using Only State Deductibility Instrument

	(1)	(2)	(3)	(4)	(5)
$(1 - \tau_{MID})$	-1.614**	-1.403**	-1.515**	-2.201***	-1.258**
	(0.651)	(0.648)	(0.620)	(0.622)	(0.544)
County Unemployment		-0.0473***			-0.0574***
		(0.00516)			(0.00757)
County Mean Income			-0.000727		-0.00414**
			(0.00167)		(0.00181)
Per-Pupil School Expenditure				-0.0186***	-0.0160***
				(0.00339)	(0.00280)
Observations	540,359	540,359	529,777	527,699	517,117

Notes: This table reports the results of estimating Equation 3 using a binary indicator for whether a state restricts the deductibility of mortgage interest as the only instrument. See Table 5 notes for variable definitions. All regressions include border-x-quarter and border-by-house-type fixed effects. The sample includes all housing transactions between 2008 and 2015 that occur in a census tract bordering a state with a different tax policy, after dropping states with no state income tax. Standard errors allow for clustering at the state-year level.
Return to text.

Table 7: Border Difference-in-Difference Estimates of MID Capitalization

	(1)	(2)	(3)	(4)	(5)
Panel A: Tax Rate Instruments					
$(1-\tau_{MID})$	-1.416**	-1.569***	-1.432**	-1.451**	-1.675***
	(0.554)	(0.400)	(0.564)	(0.557)	(0.432)
State GDP		0.00652***			0.00676***
		(0.00213)			(0.00223)
State Gov. Exp.			-0.00312		-0.0165*
			(0.00945)		(0.00929)
Inter-Gov. Exp.				0.0229	0.0237
				(0.0329)	(0.0324)
Observations	21,059,108	21,059,108	21,059,108	21,059,108	21,059,108
	(6)	(7)	(8)	(9)	(10)
Panel B: State Deductibility Instrument					
$(1-\tau_{MID})$	-1.098**	-0.851*	-1.088**	-1.108**	-1.001**
	(0.447)	(0.453)	(0.464)	(0.455)	(0.457)
State GDP		0.00651**			0.00690***
		(0.00247)			(0.00242)
State Gov. Exp.			0.00102		-0.0156
			(0.00978)		(0.00962)
Inter-Gov. Exp.				0.0171	0.0271
				(0.0338)	(0.0358)
Observations	21,059,108	21,059,108	21,059,108	21,059,108	21,059,108

Notes: This table presents instrumental variable esimates of equation 3, adding a state fixed effect. All regressions also include border-x-quarter and border-by-house-type fixed effects. Panel A uses the three instruments described in Table 3. Panel B instruments for $(1-\tau_{MID})$ with an indicator variable for states that remove or significantly restrict the deductibility of mortgage interest. The sample contains all counties that border a state with a different tax policy. Standard errors allow for clustering at the state level

Return to text.

Table 8: Capitalization Effects Across the Local House Quality Distribution

	Quartile 1	*Quartile 2*	*Quartile 3*	*Quartile 4*
	(1)	(2)	(3)	(4)
$(1-\tau_{MID})$	-0.699*	-0.714**	-1.184***	-1.306***
	(0.365)	(0.278)	(0.326)	(0.416)
Observations	139,263	134,133	132,838	127,780

Notes: This table shows estimates of the over-identified model discussed in Table 3 for subsamples of the data based on predicted house values. House quality is the predicted value from a hedonic regression. Specifically, it is the fitted value from a regression of log prices on the number of bedrooms, number of bathrooms, condominium status, presence in a subdivision, house age, and a quadratic in square footage. Quartiles are defined within border-pairs; therefore they reflect the house's rank in the local quality distribution. See Table A2 for estimates that define quartiles over the national quality distribution. See Table 3 notes for other details of the sample and specification.
Return to text.

Table 9: Parameter Values for User Cost Calculations

λ	Loan-to-Value Ratio	Observed in data
r_m	Mortgage interest rate	3.9%
r_f	Risk-free return (10-yr Treasury)	2.1%
$m + d$	Maintenance + depreciation	2.5%
σ	Risk premium for housing	2%
π	Expected capital gains	2%
τ_p	Property tax rate	1.5%
τ^{MID}	Tax rate applied to deductions	Inferred from income and county
τ_y	Tax rate on income	Inferred from income and county
ρ	Probability of itemizing	Inferred from income and county

Notes: Parameter values that are not observed or imputed are drawn from Poterba and Sinai (2011) and Martin (2018). r_m is the average mortgage rate for 30-year, fixed rate agency mortgages originated in 2015, the last year for which IRS SOI data is available. r_f is the average 10-year Treasury rate from 2015. τ^{MID}, τ_y, and ρ are imputed using average values for the appropriate county-income subgroup in IRS SOI county-level data files.
Return to text.

A. Appendices

A.1 Complete or Over-capitalization

The incidence results derived in Section 2.2.3 allow for a surprising possibility: new home purchasers can be made *worse* off by the interest deduction. This point deserves some discussion. Over-capitalization of the subsidy – when the price increase is larger than the subsidy payment – is possible even in a closed-economy model like the one above, and it is exacerbated when people can choose to enter or leave the market in response to policy changes.

First, consider the model outlined in Section 2.2.3, in which policy changes apply to a closed housing market. With competitive supply, we can differentiate the market-clearing condition $D(P, r) = S(P)$, apply the chain rule, and rearrange terms allows us to solve for the theoretical level of capitalization, $\frac{d\ln(P)}{\partial(1-\tau^{MID})} = \frac{r_m}{c}\frac{\varepsilon_D}{\varepsilon_S+\varepsilon_D}$. This takes a somewhat familiar form, with relative elasticities determining the division of the incidence between demand and supply, plus an additional term to account for the fact that the deduction applies only to mortgage interest rather than the entire cost of capital.[27]

Plugging the predicted capitalization into the buyers' incidence formula yields a result that is, at first, surprising: increasing the subsidy makes new home buyers better off if and only if $LTV < \frac{\varepsilon_D}{\varepsilon_S+\varepsilon_D}$. In other words, buyers with low levels of debt or those in inelastic markets would prefer to eliminate the MID.

This possibility, derived and emphasized by Rappoport (2016), is linked to the non-linearity of the subsidy. The market price is still set by the *marginal* cost of capital. However, the debt-equity mix of the marginal unit is not the same as earlier units. In particular, optimizing homebuyers finance later units of housing at higher loan-to-value ratios than initial units. Thus, the user cost of the marginal unit of housing is higher than the average unit, which determines the overall subsidy. Rappoport's empirical work suggests that the number of households harmed by the federal MID is small in practice, however. In Sommer and Sullivan's (2018) model, however, cutting the MID would benefit new home purchasers due to the strength of the price response.

A.2 A Model With Limited Mobility

Section 2.2.5 discusses how capitalization effects play out in models zero mobility or full mobility. Here, I present a model for the in-between case of limited mobility.

[27] Note that I have made the simplifying assumption that all residents finance home purchases at the same user cost c.

The set-up follows the applied literature on local labor markets, particularly Kline and Moretti (2014), Suarez-Serrato and Zidar (2016), and Diamond (2016). A unit mass of individuals choose to live in one of J cities, indexed by j. Utility is Cobb-Douglas in housing h and numeraire consumption c, with city-specific amenity values A_j entering multiplicatively.

Individuals also have idiosyncratic tastes for city j, denoted by ε_{ij}. The presence of location-specific shocks generates imperfect mobility. If a worker has a strong preference for a given city, she will not necessarily move even if the bundle of wages, house prices, and amenities in a different city would be preferred by the average individual.

$$U_{ij} = A_j h^{\alpha} c^{(1-\alpha)} \exp[\varepsilon_{ij}]$$

For simplicity, assume further that production is Cobb-Douglas; capital is perfectly mobile and inelastically supplied; and workers inelastically supply one unit of labor. Per Kline and Moretti (2014), this implies that the labor demand curve in each market is flat. That is, local wages w_j do not depend on the local population, and are instead pinned down by TFP and labor's share of production costs. This feature is not necessary for any of the qualitative results emphasized here, but it simplifies the exposition.

Consider the case in which all individuals purchase their home, and the cost of financing varies exogenously across cities, possibly because of local tax policies. We denote the user cost of housing capital in city j by u_j, and write:

$$\ln(V_{ij}) = \underbrace{\ln(A_j) + w_j}_{v_j} - \alpha \ln(u_j) - \alpha \ln(P_j) + \varepsilon_{ij}$$

Then, if idiosyncratic tastes are distributed Type I Extreme Value with variance parameter σ, we can solve for the population of city j:

$$N_j = \frac{\exp\left[\frac{v_j - \alpha \ln(u_j P_j)}{\sigma}\right]}{\sum_k \exp\left[\frac{v_k - \alpha \ln(u_k P_k)}{\sigma}\right]}$$

Here, we can see that the variance of the taste shocks σ governs the degree of mobility. If σ is small, local fundamentals v_j dominate location decisions and idiosyncratic tastes are inconsequential (i.e. perfect mobility). If σ is large, tastes dominate the location decision (little-to-no mobility)

To close the model, allow the housing supply elasticity ψ_j to vary locally, according to $H_j^S = B_j P_j^{\psi_j}$. Local housing demand is the local population times the individuals' housing consumption. In logs:

$$\ln(H_j^D) = C_j - \frac{\alpha+\sigma}{\sigma}\ln(u_j) - \frac{\alpha+\sigma}{\sigma}\ln(P_j) - \ln\left(\sum_k \exp\left[\frac{v_k - \alpha\ln(u_k P_k)}{\sigma}\right]\right)$$

Setting supply equal to demand implicitly defines prices, and we can take comparative statics. If J is large (and hence individual cities are small), we can ignore the effect of changes in u_j to the log-sum-exponential term in the demand equation. Put differently, we ignore the effects of changes to u_j on prices in other markets. Implicitly differentiating with respect to u_j and re-arranging yields the following:

$$\frac{\partial\ln(P_j)}{\partial\ln(u_j)} = -\frac{1+\frac{\alpha}{\sigma}}{1+\frac{\alpha}{\sigma}+\psi_j}$$

Interestingly, in the limiting case of perfect mobility, the degree of capitalization no longer depends on the elasticity of supply. Formally, we have the following result:

Proposition 2

$$\lim_{\sigma\to 0}\frac{\partial\ln(P_j)}{\partial\ln(u_j)} = 1$$

In other words, when households are perfectly mobile, changes in borrowing costs are perfectly capitalized into prices. regardless of the local supply elasticity.

The intuition follows from the logic outlined in Section 2.2. If households are indifferent between houses in two locations, then those houses must sell for the same price in equilibrium.

Note also that, with zero mobility, ($\sigma = \infty$), we obtain the same expression as we did in the closed economy model. This follows directly from the fact that, with Cobb-Douglass preferences, the elasticity of demand is one when $\sigma \to \infty$.

Figure A1: LTV-Specific Incidence of a Small Reduction in Mortgage Deductibility

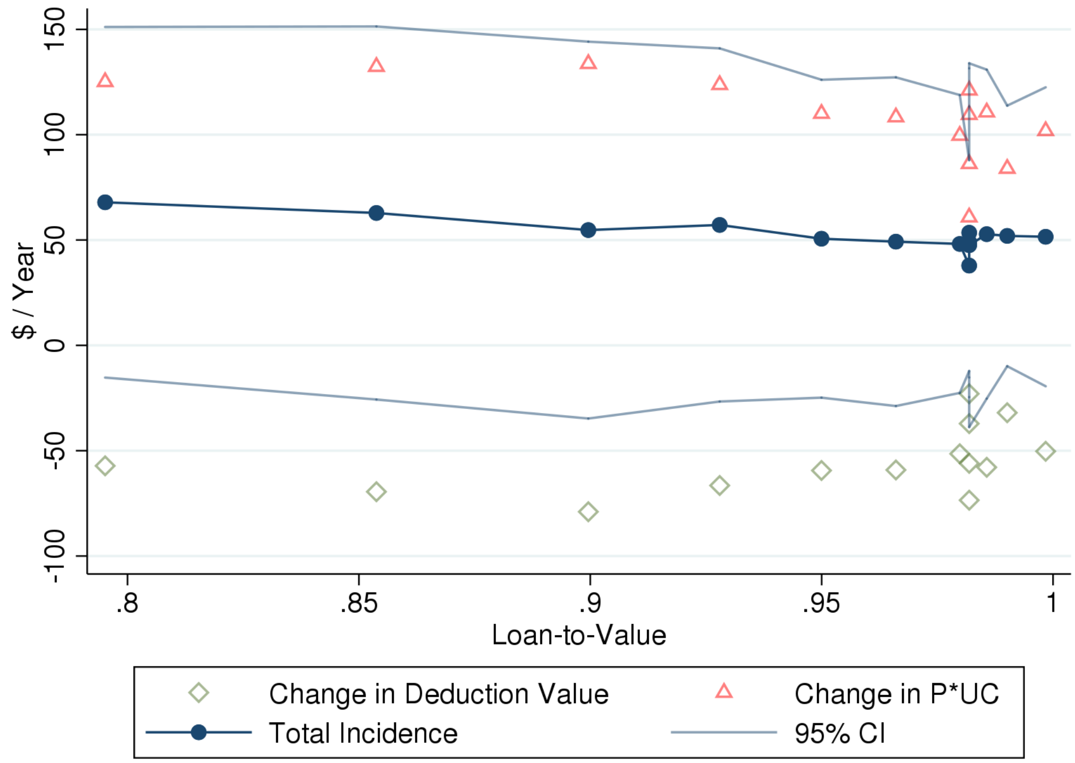

Notes:This figure reports the effect of a one-percentage point decrease in the tax rate applied to mortgage interest for new home-buyers. The sample includes all 2015 transactions in the merged HMDA-CoreLogic sample, dropping sales in the top 25% of their MSA's income distribution or the bottom 25% of their MSA' loan-to-value distribution. The plotted points show means computed within 20 evenly-sized loan-to-value bins. Total incidence is the sum of the change in deduction value and the change in the flow cost of housing (Price * User Cost), and the associated confidence interval reflects statistical uncertainty in the value of the capitalization parameter.
Return to text.

Figure A2: LTV-Specific Incidence of a Small Reduction in Mortgage Deductibility in Four Metro Areas

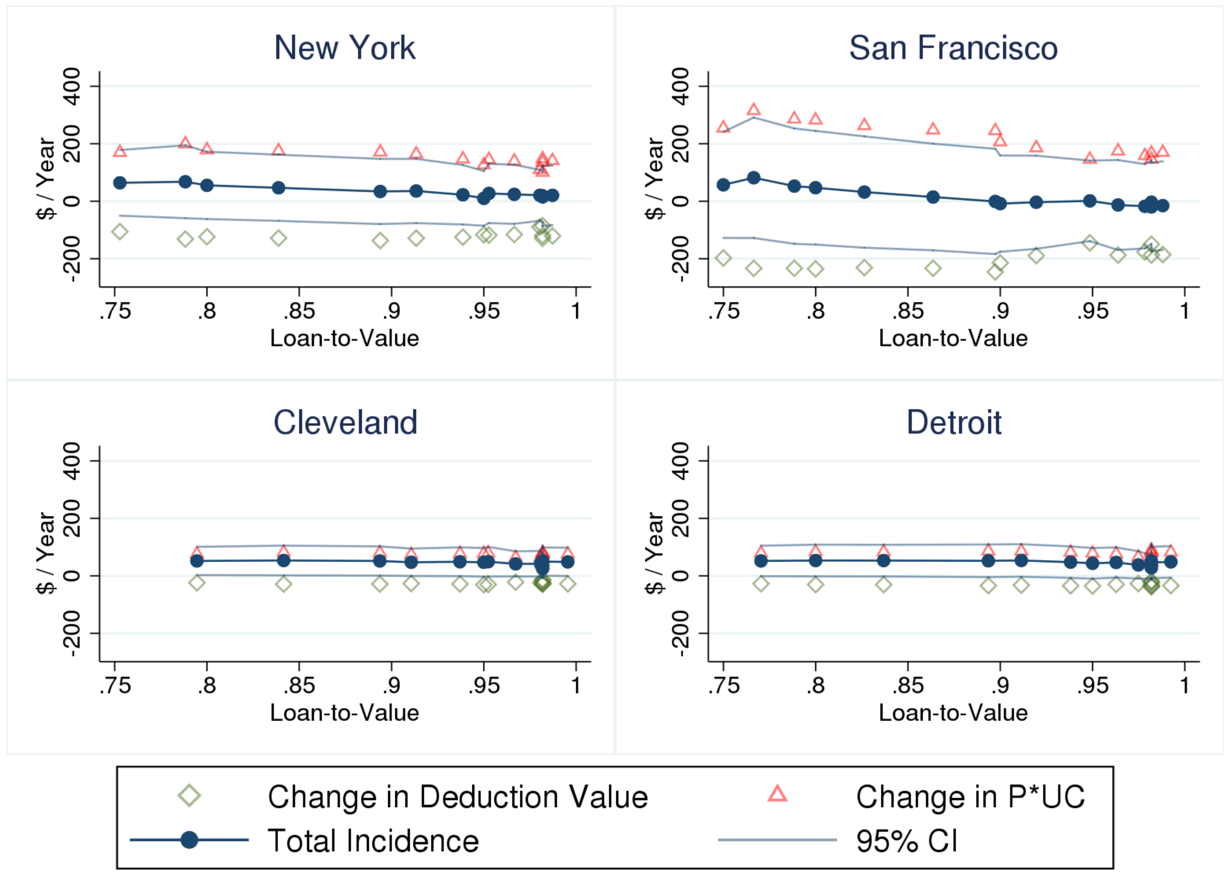

Notes: This figure reproduces Figure A1 for specific metro areas.

Return to text.

Table A1: First-Stage Results

	(1)	(2)	(3)
$(1-\tau_{MID}^{NationalAvg.})$	0.727***		
	(0.0314)		
$(1-\tau_{MID}^{HighIncome})$		0.897***	
		(0.0426)	
$(1-\tau_{MID}^{Homebuyer})$			0.623***
			(0.0289)
Observations	626,241	626,241	626,241
R-squared	0.975	0.966	0.967

Notes: This table reports the first-stage results underlying the estimates in Tables 3. See earlier table notes for details of the sample and specification.
Return to text.

Table A2: Capitalization Effects Across the National House Quality Distribution

	Quartile 1	*Quartile 2*	*Quartile 3*	*Quartile 4*
	(1)	(2)	(3)	(4)
$(1-\tau_{MID})$	-0.552	-0.715**	-0.878**	-1.409***
	(0.510)	(0.336)	(0.373)	(0.416)
Observations	139,935	134,188	130,925	135,656

Notes: This table shows estimates of the over-identified model discussed in Table 3 for subsamples of the data based on predicted house values. House quality is the predicted value from a hedonic regression. Specifically, it is the fitted value from a regression of log prices on the number of bedrooms, number of bathrooms, condominium status, presence in a subdivision, house age, and a quadratic in square footage. Unlike Table 8, quartiles are defined based on the full-sample distribution. See Table 3 notes for other details of the sample and specification.

Return to text.