

Do People Respond to the Mortgage Interest Deduction? Quasi-Experimental Evidence from Denmark*

Jonathan Gruber, MIT

Amalie Jensen, University of Copenhagen

Henrik Kleven, Princeton University

July 2017

Abstract

Using linked housing and tax records from Denmark combined with a major reform of the mortgage interest deduction in the late 1980s, we carry out the first comprehensive long-term study of how tax subsidies affect housing decisions. The reform introduced a large and sharp reduction in the mortgage deduction for top-rate taxpayers, while reducing it much less or not at all for lower-rate taxpayers. We present three main findings. First, the mortgage deduction has a precisely estimated zero effect on homeownership. This holds even in the very long run. Second, the mortgage deduction has a sizeable impact on housing demand at the intensive margin, inducing homeowners to buy larger and more expensive houses. Third, the largest effect of the mortgage deduction is on household financial decisions, inducing them to increase indebtedness. These findings suggest that the mortgage interest deduction distorts the behavior of homeowners at the intensive margin, but is ineffective at promoting homeownership at the extensive margin and any externalities that may be associated with it.

*Gruber: gruberj@mit.edu; Jensen: asj@econ.ku.dk; Kleven: kleven@princeton.edu. We thank Jim Poterba for helpful discussions.

1 Introduction

Governments around the world have a common and expensive feature of their tax codes: tax subsidies to homeownership. In many countries, the deductibility of mortgage interest provides a significant subsidy to ownership. In the United States, this tax subsidy amounted to \$101 billion in 2014, which is more than 20% as large as the entire federal deficit in that year ([Gruber 2015](#)).

These subsidies are generally motivated by the perceived externalities of homeownership. As [Glaeser & Shapiro \(2003\)](#) write: “To its supporters, the home mortgage interest deduction is the cornerstone of American society. Homeownership gives people a stake in society and induces them to care about their neighborhoods and towns. By subsidizing property ownership, the deduction induces people to invest and then to have a stake in our democracy. Ownership makes people vote for long-run investments instead of short-run transfers.” [Glaeser & Shapiro \(2003\)](#) review a large body of empirical evidence that supports the existence of these positive externalities: homeownership is positively correlated with political activism and social connection, and homeowners take better care of their properties, leading to higher values for surrounding houses. While such correlations are suggestive, there is little causal evidence for the existence of externalities due to owning rather than renting.¹

Any externality-related argument for the mortgage interest deduction requires that the policy is able to increase homeownership. Yet, we have relatively little evidence on how these tax subsidies impact housing decisions. There is a sizeable literature, reviewed below, on the effects of tax subsidies on financial decisions such as indebtedness. But there is little evidence on how these expensive subsidies impact the decision to purchase a house and the characteristics of that house.

The lack of evidence on real housing responses is likely due to the fact that a convincing empirical estimate must meet three requirements. The first is exogenous variation in the mortgage interest deduction that is sufficiently large to be able to detect any effects. The second is micro data that matches tax records to real information on housing decisions. And the third is a sufficiently long time period to capture the long-run effect on homeownership, which tends to be a slow-moving outcome.

¹The most convincing evidence is in [Engelhardt *et al.* \(2010\)](#), who use randomized access to savings accounts that can be used for home downpayments as an instrument for homeownership. They find no evidence on a wide variety of measures of social capital, although there was some effect on home maintenance expenditures.

In this paper, we focus on a setting and natural experiment that meet these three criteria. The experiment is a major tax reform in Denmark in the late 1980s. This reform significantly reduced the subsidy to negative capital income (from mortgages and other borrowing) for taxpayers in the top bracket, while reducing it much less for taxpayers in a middle bracket and raising it slightly for taxpayers in a bottom bracket. The scaling back of the mortgage deduction raised the net-of-tax interest rate by about 80% for the top group, by about 30% for the middle group, and left it roughly unchanged for the bottom group.² These large and discrete changes provide an ideal setting for a difference-in-differences approach. Furthermore, we have for Denmark a unique set of data that match income tax records to housing records for the entire population over more than three decades. This allows us to carry out the first comprehensive long-term study of how tax policy affects housing decisions.

We use the matched panel data on taxes and housing to document the effects of tax subsidies on the extensive margin of homeownership, the intensive margin of housing demand (home size and home value), and finally the intensive margin of homeowner indebtedness. In each case, we present clear graphical evidence of the effects of the tax reform, and then follow that with regression-based estimates of elasticities with respect to the net-of-tax interest rate.

We draw three main conclusions. First, there is a tightly estimated and robust zero effect of tax subsidies on homeownership. Over multiple time periods, and considering multiple empirical strategies, we find no effect of the tax policy change on whether households own or rent. Second, there is a clear effect of tax subsidies on the size and value of homes. The scaling back of the subsidy lead to sizeable reductions in both square footage and appraised values, driven by households who move to different homes after the reform. The elasticity of home size with respect to the net-of-tax interest rate is about 0.1 in absolute value, while the elasticity of home value is about 0.2. Third, the largest effect of tax subsidies is on household financial decisions. The reduction in the subsidy induced homeowners to reduce total interest expenses by up to 20%, with an implied elasticity with respect to the net-of-tax interest rate equal to 0.4.

These results have important implications for the debate about tax subsidies to owner-occupied housing. In particular, they strongly dispute the notion that the tax subsidy is justified by the promotion of externalities through ownership. There are significant effects on home size and home

²At the time of the reform, the top bracket included about 10% of the population, the middle bracket included about 35% of the population, the bottom bracket included about 45% of the population, with the remaining 10% being below the exemption threshold.

value, but it is unclear what externalities are delivered through larger and more expensive houses.³ Rather, the intensive margin effects are most naturally viewed as distortions. The largest effect appears to be on financial, not real, household decisions for which there is also no externality argument. Taken together, our findings suggest that mortgage deductions distort housing and debt demand at the intensive margin, while being ineffective in promoting homeownership at the extensive margin and any externalities associated with it.

The paper proceeds as follows. Part 2 reviews the previous literature on taxation and homeownership. Part 3 presents the institutional setting, and describes our unique data. In Part 4, we estimate the effects of the mortgage interest deduction on homeownership, house size and value, and indebtedness. Part 5 concludes.

2 Literature Review

There is a long standing literature that discusses the mortgage interest deduction and its effects on household decision making. Early work in this area focused on discussing conceptually the effects of tax policy on housing and the income distribution ([Aaron 1970](#); [Rosen 1979](#)). Perhaps the earliest empirical study to address the behavioral consequences of the tax subsidy is [Rosen & Rosen \(1980\)](#). They studied the time series correlation between the tax subsidy and homeownership, estimating that 25% of the growth in homeownership in the U.S. was due to the mortgage interest subsidy.

Subsequent studies, however, have pointed out that over the past several decades there is no evidence that tax policy impacts homeownership (e.g. [Glaeser & Shapiro 2003](#)). This literature is reviewed in [Gale *et al.* \(2007\)](#). They show that over the forty years from 1965 to 2005 there were very wide swings in the tax price of homeownership, but very little movement in the time series of homeownership. They also review the cross-sectional international evidence, which suggests little correlation.

Other than such correlational evidence, there is no work that estimates the impact of the mortgage interest deduction on homeownership. There has been some work suggesting that more targeted tax policies could raise homeownership. Most notable is [Engelhardt *et al.* \(2010\)](#) who examine a randomized trial of a subsidized savings account for home purchases offered to low-income families in Oklahoma. This study finds that those randomized into the subsidized savings account

³In principle, there could be spillovers associated with home size and value. If a homeowner invests in a larger and nicer house, this could positively affect the surrounding houses by improving the neighborhood or negatively through competition. We have no clear evidence on such effects.

were 7-11% more likely to purchase a home.

While there is little causal evidence on the impact of the mortgage deduction on homeownership, there is a richer literature studying the impacts on house prices and household financial decision making. The seminal study of the impact of taxes on house prices is [Poterba \(1984\)](#), who used an asset price framework to show that the interaction of inflation with the tax system could explain much of the rapid rise in home prices in the 1970s. [Poterba \(1991\)](#) and [Poterba \(1992a\)](#) extend this framework to consider the impact of tax reforms in the 1980s on house prices.

The literature on mortgage interest deductions and household financial decision making is reviewed in [Poterba \(2002\)](#). He reviews evidence showing that household portfolio allocation is in general sensitive to tax prices, and that in particular household indebtedness is sensitive to the mortgage subsidies embedded in the U.S. tax code.

Over the past decade, there have been a number of “natural experiment” studies that use reforms in the tax treatment of mortgage interest to estimate impacts on household indebtedness. [Ling & McGill \(1998\)](#) and [Dunsky & Follain \(2000\)](#) both study the impact of the sizeable reduction in the mortgage interest deduction implied by the tax rate cuts of the U.S. Tax Reform Act of 1986. They find that there was a large financial response, with mortgage debt falling in response to the increase in the tax price of such debt.

On the other hand, studies of the financial impact of the mortgage interest deduction in other nations have been more mixed. [Jappelli & Pistaferri \(2007\)](#) study an Italian reform that changed the deduction from one linked to marginal rates to a flat deduction across all brackets. They find no effect on mortgage debt on either the extensive or intensive margins. [Kleven & Schultz \(2014\)](#) and [Alan *et al.* \(2016\)](#) study the change in Danish tax policy that is the focus of the current paper. [Kleven & Schultz \(2014\)](#) estimate an elasticity of negative capital income that vary between -0.10 and -0.13 , while [Alan *et al.* \(2016\)](#) find an elasticity of interest expenses of -0.07 with respect to the tax subsidy.

In summary, there is little evidence to date that tax policy has an impact on real housing decisions, although existing studies are limited to time series and cross-country approaches. There is more evidence that the mortgage interest deduction impacts household financial decisions, but the evidence varies across nations.

3 Institutional Setting and Data

3.1 Tax Treatment of Owner-Occupied Housing in Denmark

The Danish income tax has three tax brackets — a bottom, middle, and top bracket — along with an exemption threshold. Up until the major 1987-reform, individuals were taxed based on their total taxable income defined as the sum of labor and capital income minus deductions. Taxable capital income was a net income concept, with deductions for interest expenses on mortgages and other forms of debt. The top bracket tax rate was very high, 73% in the average municipality, and around 10% of the population was liable to pay it.⁴ The middle bracket tax rate was 62% (paid by about 35% of the population) and the bottom bracket tax rate was 48% (paid by about 45% of the population).

The 1987-reform changed this system in a fundamental way. It introduced a form of “dual income taxation” combining a progressive tax on labor income with a (roughly) proportional tax on capital income. The tax system continued to be based on capital income net of interest expenses, but capital income became subject to a much lower tax rate and with an asymmetry in the tax treatment of negative and positive net capital income. Specifically, negative capital income was taxed according to the bottom tax rate, whereas positive capital income was taxed according to either the bottom or the middle tax rate depending on the labor income of the taxpayer. Because most homeowners with mortgages have negative net capital income, the value of the mortgage interest deduction was in general given by the bottom rate after the reform.⁵

The implications of the 1987-reform for marginal tax rates on capital income are illustrated in Figure 1. Panel A shows marginal tax rates on negative capital income in the three brackets, while Panel B shows marginal tax rates on positive capital income. Focusing on the main scenario in Panel A, we see that the impact of the reform was quite dramatic. It set the capital tax rate equal to 51% in all brackets, creating large variation across brackets due to their different pre-reform tax rates. The tax price relevant for mortgage and housing demand is the net-of-tax interest rate, $(1 - \tau) R$. The reform increased this rate by 81% for the top group, while increasing it by only 29% for the middle group and reducing it by 6% for the bottom group. This is the identifying variation on which our paper is based.

⁴This tax rate includes a flat local income tax that varies somewhat across municipalities and counties.

⁵This is due to the fact that the return on retirement savings is tax-deferred and that, following the 1987-reform, the return on equity is taxed according to a separate progressive schedule. Excluding retirement savings and equity, net capital income is negative for around 80% of all homeowners (and a larger fraction of homeowners with mortgages).

Unlike some countries with mortgage interest deductions (such as the U.S.), Denmark taxed imputed rental income from owner-occupied housing at the time of the 1987-reform (this tax has since been replaced). Because of this, a reduction in the tax rate on capital income has two offsetting effects on housing demand. On the one hand, it reduces the mortgage interest deduction, which gives an incentive to own less housing. On the other hand, it reduces the taxation of rental income, which gives an incentive to own more housing. An “ideal” tax system would treat homeownership as a business and tax its real economic profits, namely the difference between true rental income and the user cost of housing (see e.g., [Poterba 1992b](#); [Poterba & Sinai 2008, 2011](#)). In this case, we would expect little or no effect of capital taxes on the intensive margin of housing demand, although there could still be effects on the extensive margin.

Importantly, the Danish tax on imputed rental income was artificially low and therefore did not correspond to the ideal system just described. The calculation of taxable rental income was based on a rental rate of only 2.5% of an assessed home value. Moreover, the assessed home value was based on a tax appraisal set below the true market value (as an implicit tax subsidy). The tax was further reduced by the presence of a deduction equal to 1% of the home value up to a cap. To conclude, the tax on imputed rents was small and therefore unlikely to have much of an impact on incentives.

Finally, as can be seen from [Figure 1](#), tax reforms in the mid- and late-1990s further reduced tax rate on negative capital income. However, these reforms introduced little variation across tax brackets, apart from some modest variation between the bottom and middle brackets. As a result, these reforms do not allow for the same clean difference-in-differences design made possible by the 1987-reform. We therefore focus on the 1987-reform in this paper.⁶

3.2 Data and Descriptives

Our analysis is based on administrative data for the full Danish population from 1980-2011. The data comes from several administrative registers, linked at the individual level through personal identification numbers. Each individual has a home address code and a property identification number, which can be linked to a housing register with information on homeownership and property characteristics. Our linked dataset includes detailed information on incomes, taxes, homeownership, house characteristics, and a wide range of socio-economic variables. We are able to link

⁶See [Kleven & Schultz \(2014\)](#) for a detailed description of all the different tax reforms and evidence on their impact on reported taxable income.

individuals to their spouses/cohabitants and children, which is important for constructing precise measures of tax treatment and homeownership as we describe below.

Tax variables related to housing wealth and mortgage debt are not available before 1987, so we cannot measure homeownership using the individual tax files. However, the housing register has information on whether the house is occupied by the owner. When matched to individuals, this enables us to construct our homeowner measure. Because individuals and information about their homes are matched via the home address, homeownership is defined at the household level.⁷

There can be several people living at the same address of an owner-occupied home, not all of whom are owners. This creates a complication for our homeownership measure based on information from the housing register. We deal with this measurement problem as follows. We start by dropping all children living with their parents from the dataset. Everyone living in a home not occupied by the owner are defined as renters. For those living in homes occupied by the owner(s), we distinguish between three cases. The first case is where only one person is living at the address. He or she is then defined as the owner. The second case is where a married or cohabiting couple is living at the address. They are both defined as homeowners. These two (clean) cases account for almost 90% of individuals living in owner-occupied homes. The third and more complicated case is where the home is inhabited by two people who are not cohabiting, or by three or more people. To assign ownership in this case, we use the information about owned property available in the individual tax files from 1987 onwards. When assigning ownership in the years prior to 1987, we check whether a person at the address is still living at that address in 1987 and owns property according to the tax files. If so, we define the person as a homeowner in the earlier year. If the person has a spouse/cohabitant, then he or she is defined as a homeowner as well. Any remaining people at the address are defined as renters. Using this method, we are able to construct a homeowner measure for an additional 8-9% of the individuals living in owner-occupied homes. This leaves us with a small group of individuals for whom we cannot assign owner/renter status. We drop these individuals from our analysis sample.

Figure A.I in the appendix validates our homeownership measure by comparing it to the ideal tax-based measure available from 1987. The figure shows three homeowner rate series: our baseline measure using the housing records as described above, our measure using the housing records

⁷Even with access to individual information about homeownership, we would want to construct a household-based measure. The reason is that the norm for how to register a home among couples has changed over time. While a home used to be registered in the name of the husband, it became more common during the 1980s to register the home in the name of both spouses. For an individual-based measure, this increase in dual-registration would create an artificial upward trend in homeownership.

but focusing only on the “clean” cases (single person or a cohabitating/married couple at the address), and the ideal measure using the individual tax records on housing wealth. We see that the two measures using housing records are essentially identical, and that both of them are very close to the precise tax-based measure in terms of both the trend and the level. The level of homeownership is somewhat lower when using the housing register, because by construction this measure does not capture owners who are not residing in their owned property. The most common example in Denmark are people who rent in their main residence, but own a holiday home.⁸

When studying responses at the intensive margin, we consider home size, home value, and interest expenses as our outcomes. Home size is defined as the area of the house (in square meters) used for habitation. The home value is based on an appraisal by the Danish Tax Authorities done for the purposes of assessing imputed rental income and taxable wealth. Interest expenses include interests paid on mortgage debt, consumer debt and other types of loans, all of which are deductible. We are not able to observe mortgage interest payments separately over the entire analysis period, and so we focus on total deductible interest expenses (for homeowners) as the outcome.

Our main analyses are based on a semi-balanced panel of individuals who we observe in every year between 1980-1996, and who are between the ages of 20 and 80. The sample is only ‘semi-balanced’ due to the age restriction: individuals drop out of the sample after they turn 80, and they enter only after they turn 20. These restrictions give us a sample of about 40 million individual-year observations.

Figure 2 shows descriptive evidence on homeownership over time. Panel A plots the aggregate homeowner rate in the adult population between 1980-2011. Just under 60% of the population are homeowners, a number that has stayed remarkably constant over time. The stability of the homeowner rate over three decades with large variation in tax subsidies to housing (as shown in Figure 1) provides a first indication — but does not prove — that homeownership is unresponsive to tax incentives. Such time series evidence is consistent with findings for the U.S. and other countries (e.g. Glaeser & Shapiro 2003; Gale *et al.* 2007). While there is no aggregate time trend in homeowner rates, Panel B shows that there is a strong life-cycle pattern. Homeownership increases strongly in early adulthood, stays relatively flat during middle age, and decrease gradually in later life. This life-cycle pattern has shifted to the right over time, consistent with a general shift in the timing of education, marriage, parenthood, and retirement.

⁸Indeed, the increase in the gap between the tax-based measure and our measure of homeownership during the first part of the 2000s may be due to increased purchases of holiday homes (by renters) during the housing boom. The gap declines again after the Great Recession.

4 Effects of the Mortgage Interest Deduction

In this section we present evidence on the main margins of behavior that can respond to the mortgage interest deduction: homeownership, home size and home value among owners, and interest expenses among owners. While there is existing work on interest/borrowing responses to tax subsidies, the evidence on homeownership and home size is novel. We are not aware of any clean evidence on how real housing investments at the intensive and extensive margins respond to the mortgage interest deduction.

4.1 Empirical Specification

The analysis is based on a difference-in-differences (DD) design exploiting the 1987 tax reform. As described above, the reform sharply reduced the mortgage interest deduction for taxpayers in the top bracket, while reducing it much less for taxpayers in the middle bracket and leaving it almost unchanged for taxpayers in the bottom bracket. This type of variation introduces a trade-off when choosing treatment and control groups. A comparison between the top and middle groups has the advantage that treatments and controls are relatively similar (strengthening the parallel trends assumption), but the disadvantage that both groups are treated to a significant extent (requiring an assumption of homogeneous responsiveness). Conversely, a comparison between the top and bottom groups might be more tenuous in terms of the parallel trends assumption, but does not require the homogeneity assumption. Of course, the observed pre-reform trends will provide a direct indication of the validity of the parallel trends assumption for the different comparisons. We consider all types of specifications in the empirical analysis. We first show results based on comparing those in the top bracket to everyone below the top bracket (i.e., combining the two possible control groups), and then show results based on the two possible control groups separately.

Although the analysis will be implemented on a sample of individuals, it is in reality household-based. There are two reasons for this. First, our measure of homeownership is household-based. As described in section 3, an individual is defined as a homeowner if she and/or her spouse are registered as homeowners. While there were measurement reasons for doing it this way, it is also economically meaningful to view homeownership as a household decision. Second, our measure of tax treatment is also household-based. Even though the Danish tax system relies on individual filing, it contains certain elements of jointness and this is particularly relevant for the tax treatment of capital income. The highest-bracket individual in the family tend to determine the degree to

which the household benefits from the mortgage deduction. Hence, we define tax treatment as follows: a taxpayer is assigned to the top bracket if either she or her spouse are in the top bracket, a taxpayer is assigned to the middle bracket if either she or her spouse are in the middle bracket and none of them are in the top bracket, while a taxpayer is assigned to the bottom bracket if both she and her spouse are in this bracket.

The analysis is based on a balanced panel of individuals who we observe in every year over a specified time window. To begin with, we consider the time window 1980-96, giving us 7 pre-reform years and 10 post-reform years. We will consider a longer time window as well. The assignment of treatment status is based on pre-reform tax status. The treated are those located in the top bracket prior to the reform, while the controls are those located in the middle or bottom brackets prior to the reform. We assign treatment status using three pre-reform years, requiring treatments and controls to be in a given bracket in every year between 1984-86. Using three years to assign treatment status ensures that the groups are relatively stable in terms of their tax status.⁹

The balanced panel approach implies that people in our estimation sample become older over time. This is a potential issue due to the strong lifecycle effects on homeownership (see Figure 2 above). Homeownership increases strongly until the middle ages and then falls gradually over the rest of the lifecycle. This poses a challenge for a DD based on variation across tax brackets: lower tax brackets have more young people as well as more old people than higher tax brackets, and so the groups are differentially affected by age effects. This creates nonparallel trends in the raw data. To deal with this issue, we focus on DD specifications that control nonparametrically for age effects. We show graphically that this strategy does very well in terms of absorbing differential trends.

We present graphs based on the following regression specification, which we run separately for the treatment and control groups:

$$Y_{it}^g = \sum_y \alpha_y^g \cdot \mathbf{I}(y = t) + \sum_a \beta_a^g \cdot \mathbf{I}(a = \text{age}_{it}) + \nu_{it}^g, \quad (1)$$

where Y_{it}^g is the outcome for individual i in group g in year t , and where $\mathbf{I}[\cdot]$ is an indicator function. The first term on the right-hand side includes a full set of year dummies, while the second term on the right-hand side includes a full set of age dummies. Without the age dummies, the coefficient α_t^g would correspond simply to the average outcome for group g in year t . With the age dummies,

⁹Taxpayers who move across brackets during the three pre-reform years 1984-86 are therefore dropped from the estimation sample. This restriction makes us drop around 13% of the sample.

the time profile of α_t^g captures changes in the average outcome for each group, having controlled nonparametrically for the fact that individuals in the group become older over time. We will present graphs in which we plot mean predicted outcomes in each year based on 1986 (pre-reform) age effects, i.e.

$$\hat{Y}_{t,86\text{-age}}^g \equiv \hat{\alpha}_t^g + \text{E} \left[\sum_a \hat{\beta}_a^g \cdot \mathbf{I}(a = \text{age}_{i1986}) \right]. \quad (2)$$

This gives the mean outcome in each group over time assuming the age composition stays as it was in the pre-reform year.

When trends are parallel (which we verify based on pre-trends), the effect of the reform can be obtained based on a simple DD. The effect in post-reform year t_1 relative to pre-reform year t_0 is equal to $(\hat{\alpha}_{t_1}^T - \hat{\alpha}_{t_0}^T) - (\hat{\alpha}_{t_1}^C - \hat{\alpha}_{t_0}^C)$ where T, C denote treatment and control. When showing DD estimates, standard errors are clustered at the household level.

For some of the outcomes and control groups we consider, the pre-trends are not fully parallel after absorbing lifecycle trends through the age dummies. In those cases, we estimate a group-specific time trend using pre-reform data and residualize the outcome variable based on the estimated time trend. Specifically, we first regress the outcome Y_{it}^g on a linear time trend $\theta^g \cdot t$ and the full set of age dummies using data from 1980 to 1986. We then run the specification (1) on the full period using the residualized outcome $Y_{it}^g - \hat{\theta}^g \cdot t$. As we shall see, allowing for linear pre-trends (where necessary at all) is in general sufficient to eliminate differential trends.

4.2 Extensive Margin: Homeownership

We start by investigating the effect of the mortgage interest deduction on the decision to owning vs renting. The traditional argument for such tax subsidies is that, without them, homeownership would be inefficiently low due to the presence of externalities (see e.g., [Glaeser & Shapiro 2003](#)). This policy prescription presumes that subsidies are able to increase homeownership, thus yielding the supposed Pigouvian welfare benefits. While it is difficult to estimate the potential externalities from homeownership, the effects of tax subsidies on owning can be estimated with our quasi-experiment and data. If this effect is small or zero, the externality argument is irrelevant.

Our first set of results are presented in [Figure 3](#). The figure shows the homeowner rate over time for individuals in the top bracket ('treatments') and individuals below the top bracket ('controls'). Each series is based on the specification in equations (1)-(2) in which we absorb lifecycle trends through group-specific age dummies. The homeowner rate for the control group has been

normalized to the level of the treatment group in the pre-reform year 1986. The figure shows two different time horizons in the top and bottom panels.

Panel A shows effects on homeownership in the medium run using our baseline sample: a balanced panel of individuals observed in every year between 1980-1996. We see that the treatment and control series track each other almost perfectly during this period. There is no sign of any effect of the sharp reduction of the mortgage interest deduction in 1987: the pre-reform trends are perfectly parallel and they stay perfectly parallel in the years after the reform.

It could be argued that 10 post-reform years are not enough to detect any effect of tax subsidies on homeownership, because this is a very slow-moving outcome. The slow-moving nature of homeownership may be due to the fact that most households make this decision at relatively young ages when they marry and have children (consistent with the pattern in Figure 2B) and do not often re-optimize later in the lifecycle. If so, homeownership is largely pre-determined later in the lifecycle, thus attenuating the short-run effects of the reform in the full population. This motivates the analysis in Panel B in which we consider the longest-possible period (1980-2011) for a balanced panel of individuals who we observe in every year over this period. By focusing on a balanced panel over such a long period, we are considering a sample of individuals who were relatively young at the time of the reform. These households experienced the reform during the phase of their lifecycle in which homeownership decisions tend to be made.

The results in Panel B are striking: the homeowner rates of the two groups track each other almost perfectly over more than three decades. While pre-trends become less indicative of post-trends as we move farther out, the combination of perfectly parallel pre- and post-trends is more than strongly suggestive. For there to be an effect on homeownership in this figure, there would have to be confounders that precisely offset the effect of the reform in every year over such a long period. Such a knife-edge scenario seems unlikely. Rather, the natural conclusion from Figure 3 is that the reduction of the mortgage interest deduction had a zero impact on homeownership.

In the appendix, we investigate the robustness of our results to the choice of control group. Instead of using everyone below the top tax bracket as controls, Figure A.II and A.III use either those in the middle bracket or those in the bottom bracket as controls. The top panel in each of these figures is based on our default specification (1)-(2) and the 1980-96 balanced panel. We see that the treatment and control series are not completely parallel when considering the middle and bottom groups separately. The controls follow a slightly steeper trend in one case, and a slightly flatter trend in the other case. Still, the graphs do not indicate any effect of the reform: the trends

are stable throughout the period, with no sign of anything different happening after the reform. To make this clear, the bottom panel considers the specification in which we net out group-specific linear time trends estimated on pre-reform data, as described in section 4.1. When allowing for such linear pre-trends, the treatment and control series track each other very closely over time, confirming the previous zero-effect on homeownership.

Table 1 presents DD estimates of the homeownership effect across a range of specifications: The three control group definitions considered above are shown in different panels, while different sets of controls are shown in the columns. The results in column (3) in which we include age fixed effects that vary by group corresponds to the graphical specification, while in the other columns we consider more parsimonious specifications without any age controls or with common age controls. To be precise, the estimates in the richest specification in column (3) are based on the following DD regression:

$$Y_{it}^R = \alpha_0 + \alpha_1 \cdot post_t + \alpha_2 \cdot treat_i + \alpha_3 \cdot post_t \cdot treat_i + \sum_a \beta_a \cdot \mathbf{I}(a = age_{it}) + \sum_a \beta_a^T \cdot \mathbf{I}(a = age_{it}) \cdot treat_i + \nu_{it}. \quad (3)$$

Here the outcome Y_{it}^R is residualized using the group-specific linear pre-trend as described above. We control for linear pre-trends in every specification for consistency, although in some specifications there is essentially no pre-trend as we have seen in the graphs. The estimated coefficient $\hat{\alpha}_3$ is the average 10-year effect of the reform on the homeowner rate of the treatment group relative to the control group, which we label “effect of reform” in the table. We also show estimates of the semi-elasticity of the homeowner rate with respect to the net-of-tax rate on negative capital income, i.e.

$$\varepsilon = \frac{\hat{\alpha}_3}{\Delta \log(1 - \tau^T) - \Delta \log(1 - \tau^C)}, \quad (4)$$

where $\Delta \log(1 - \tau^g)$ is the reform-induced change in the net-of-tax rate for group g between 1986 and 1987 as illustrated in Figure 1.

The following results in Table 1 are worth highlighting. First, the effect of the reform on homeownership is extremely small across all specifications. In fact, the estimates are often slightly positive, the opposite of what we would expect given that the reform made homeownership more costly. The largest negative effect on the homeowner rate is found in Panel B and columns (2)-(3), and even there the effect is less than a percentage point. Second, because the underlying tax variation is very large (see Figure 1), the small reform effects translate into very small elasticities. The

estimated elasticities vary between -0.016 and 0.017 across all specifications. Third, standard errors are very small due to the large sample size, making the effects very tightly estimated. We can therefore rule out elasticities of homeownership with respect to $1 - \tau$ in excess of 0.02 in absolute value.

To conclude, the evidence presented in this section shows that the large reduction in the tax subsidy to mortgage interest had no economically relevant effect on homeownership either in the short or the long run. As far as we are aware, this provides the first conclusive evidence on the impact of the mortgage interest deduction on homeownership.

4.3 Intensive Margin: Home Size and Home Value

Having found no effect of tax subsidies on the extensive margin of owning vs renting, we now turn to the effects on the intensive margin. Because there are no clear externalities associated with intensive margin responses, they are most naturally viewed as distortions. To study these distortions, we focus on home size (square footage) as our main outcome variable. While this is just one dimension of the ‘amount’ of housing that people own, it has the key advantage of being precisely measured in the administrative records for every home since 1980. Studying a broader intensive margin outcome — such as home values — is associated with measurement problems, but we will consider such an outcome as well.

Our first set of results are presented in Figure 4. This figure compares log home size over time for those in the top bracket (‘treatments’) and those below the top bracket (‘controls’). As before, the time series of each group is based on the specification in equations (1)-(2). Panel A shows results for the full sample, while Panel B shows results for movers only. In this analysis we do not condition on being a homeowner, because such a sample restriction is associated with potential selection problems.¹⁰

In Panel A the two series track each other perfectly in the years leading up to the reform, but begin to diverge immediately after the reform as taxpayers in the top group reduce their home

¹⁰For example, we could restrict the sample to individuals who are homeowners in the pre-reform year(s), naturally using 1984-86 as we do to assign treatment status. This restriction implies that we start from a homeowner rate of 100% at the time of the reform, and over time the homeowner rate declines as some individuals go from owning to renting. These transitions are heavily selected on home size, as they mostly include older individuals and divorcees who naturally downsize as they move into a rental. Other ways of restricting the sample (such as conditioning on being a homeowner in every year, both pre- and post-reform) give rise to other selection problems. Furthermore, households who are currently renting are potentially treated as they may transition to owning in the future at which point they face an incentive to buy a larger house. Only the “always-renters” are untreated. When we drop those who rent in every year, we obtain very similar results to those we show here.

size relative to taxpayers in the below-top group. The effect on home size is growing over time and does not reach a steady state within the 10 years following the reform. This dynamic pattern seems very natural. Leaving aside major renovations (which are costly and take time), changing the size of the house requires moving house. Since only 5-10% of people move in any given year, most households are effectively 'untreated' in the short run. Therefore, while the pattern in Panel A is clearly consistent with the presence of intensive margin responses, this is unlikely to capture the full long-run effect of the reform.

This motivates the analysis in Panel B in which we consider only movers. In each year, we plot the average home size among households who move in that year. This is a repeated cross-section selected from our balanced panel. As expected, the effect of the reform on home size is considerably larger in Panel B. Immediately after the reform, movers in the top group substantially reduce their home size relative to movers in the lower group. The effect grows for about 6-7 years and then stabilizes. Ten years after the reform, the divergence between the treatment and control groups is about 0.1 log points, an effect of about 10% on home size.¹¹

While Panel B of Figure 4 captures longer-run effects by considering movers ('re-optimizers'), another possibility is obviously to extend the time window. We do this in appendix Figure A.IV, which is based on our long panel (1980-2011). This figure shows that the home size effect in the full population increases gradually for a very long time, whereas the home size effect among movers is stable around 0.1 log points for a long time. Even though the long-run graph looks compelling, it may be argued that DD strategies in general are tenuous so long after the reform experiment as confounding shocks become increasingly likely over time.¹² For this reason, we prefer to obtain long-run effects based on the mover sample within a shorter time horizon.

Figures A.V and A.VI investigate robustness to the choice of control group. When using either the middle bracket or the bottom bracket as a control group, we obtain the same qualitative results as for the combined control group. The magnitude of the effect is smaller when comparing to the middle group and larger when comparing to the bottom group. This is consistent with the incentives created by the reform, because the middle group also experienced a sizeable reduction in

¹¹The sizeable intensive margin response by movers naturally raises the question of whether the moving decision itself is responsive to the tax reform. That is, are homeowners in the treated group more likely to move (buying a smaller house when they do)? We have analyzed this margin and find no clear effect of the tax reform on the propensity to move. This suggests a model in which moving decisions are determined by major life events (such as child birth, marriage, or job changes), while the mortgage interest deduction matters for the type of house that is purchased, conditional on moving.

¹²Note that this critique did not apply to the same extent to the long-run analysis of homeownership, because there we found a precisely estimated zero effect. As argued there, it would require an unlikely knife-edge case to reconcile the observed zero (in every year) with confounding shocks.

their deduction (albeit smaller than for the top group) while the bottom group experienced a small increase in their deduction. Hence the underlying tax variation is much larger when comparing to the bottom group, explaining the larger effects in this specification.

Table 2 shows estimates of the home size effect for movers across a range of different specifications. The table is constructed in the same way as the previous table for homeownership, with different control groups in panels A-C and different age controls in columns (1)-(3). Our baseline estimates (with group-specific age controls) in column (3) are based on the DD specification in equation (3) and the elasticity definition in equation (4).¹³ Because the outcome variable is now in logs, these are 'real' elasticities rather than semi-elasticities. The table shows that the results are very robust across specifications. With group-specific age controls, the home size elasticity is always close to 0.1 (in absolute value). Across all specifications in the table, the elasticity varies between 0.09-0.17 and with very small standard errors.

Our home size findings represent the first causal evidence of the impact of the mortgage interest deduction on real housing demand. It is of obvious interest to go from home size to a more comprehensive measure of home values. Households may respond to the mortgage deduction by changing other dimensions than square footage. The administrative tax records include information on home appraisals used for the purposes of calculating imputed rental income as well as taxable wealth (as Denmark had a wealth tax up until 1996). However, these appraisals are known to provide noisy measures of true home values: They are 'desk appraisals' that do not account for idiosyncratic improvements and repairs, they are not updated every year for every house, and they tend to be set below the true market value (as an implicit tax subsidy). For these reasons, the time variation in the appraisal of a given house will not capture the evolution of its real value very precisely. But the cross-sectional variation in appraisals is nevertheless likely to be informative: if a person moves to a less valuable house, it will typically have a lower tax-related appraisal. Hence, the mover analysis is potentially feasible with the home appraisal variable.

In Figure 5, we carry out such an analysis. This graph plots log home value for those in the top bracket and those below the top bracket in each year, conditional on moving in the given year.¹⁴

The graph suggests a clear impact on home values: the treatment and control series are roughly

¹³This implies that we control for linear pre-trends as described above (in this table and in all other tables). For the home size effect, this tends to make little difference to the estimates as pre-trends are similar between treatments and controls (as shown by the graphs).

¹⁴Because renters have zero housing wealth, the analysis of home values based on tax appraisals effectively conditions on homeownership in each year. This is different from the analysis of home sizes in which we made no such sample restriction, as discussed above.

parallel up until the time of the reform and then diverge. After 10 years, home values in the top group have dropped by about 0.15 log points (or about 15%) compared to home values in the lower group. This is a larger effect than on home size, consistent with the idea that homeowners adjust other quality margins (by more than square footage) in response to taxes. We repeat the analysis when splitting out the control group in Figures A.VII and A.VIII. These graphs confirm the effect on home values, with the size of the effect being larger when comparing to the bottom group than when comparing to the middle group. Again, these differences in effects are consistent with the differences in tax variation across comparison groups.

Table 3 presents regression results across specifications in the same way as the previous tables. In our baseline specification in column (3), the elasticity of home value with respect to the net-of-tax rate lies between 0.18 and 0.25 across the different control group specifications. This is about twice as large as the elasticity of home size.

To conclude, we have presented compelling evidence of intensive margin responses to the mortgage interest deduction. The effects of the reform on both home size and home value are sizeable. Because the tax variation created by the reform is very large, the implied elasticities of housing demand with respect to the net-of-tax interest rate are still fairly modest, about 0.1 for home size and 0.2 for home value.

4.4 Intensive Margin: Interest Expenses

In this section we investigate if the reform of the mortgage interest deduction affected the interest expenses (i.e. borrowing) of homeowners. As reviewed in section 2, there is existing work on how interest expenses respond to tax subsidies, including work on the Danish 1987-reform (Kleven & Schultz 2014; Alan *et al.* 2016). While the findings in this section are therefore less novel, it is important to confirm if the reform had an impact on the borrowing margin using our estimation sample and empirical approach. In particular, if we were to find no effect on interest expenses, it would be harder to believe our results on the housing demand margin presented above. This is because the main mechanism by which homeowners would want to reduce home size and home value in response to a higher after-tax interest rate is the incentive to reduce borrowing and thus (before-tax) interest expenses.

The Danish tax system offers deductability of all interest expenses, not just mortgage interest. Our data do not contain separate information on mortgage interest expenses and other types of interest expenses back in the 1980s, only total deductible interest expenses are recorded. Hence,

our analysis is forced to consider total interest expenses as the outcome variable. This is arguably not a strong data limitation given that, for homeowners, mortgage interest generally represents the bulk of their interest expenses. We also note that, because interest expenses are third-party reported by lenders, there is essentially no scope for misreporting (see [Kleven *et al.* 2011](#)). Our results in this section should therefore be interpreted as reflecting real borrowing responses, not evasion.

Figure 6 shows our first set of results for total interest payments among homeowners. Panel A considers our baseline specification (1)-(2) without adjusting for pre-trends, while Panel B adjusts for linear pre-trends. We focus on the time period from 1983 onwards in these graphs, as opposed to 1980 onwards in all previous graphs. This is because nominal interest rates fell sharply between 1982-83 in Denmark, a phenomenon that is typically attributed to the transition from a regime with frequent exchange rate devaluations (under the Social Democratic government until 1982) to a regime with a fixed exchange rate commitment (under the Conservative government from 1982). The fall in interest rates in the early 1980s had a differential impact on the interest expenses of the treatment and control groups, which might confound our analysis when including the full time period.

The graph shows clear and sizeable effects on interest expenses. In Panel A the pre-trends are not entirely parallel, but there is a clear break in the relative trends immediately after the 1987-reform. Taxpayers in the top bracket reduce their reported interest payments relative to the control group from 1987 onwards. In Panel B the pattern is even more striking: Here the two groups evolve similarly until the reform and then diverge sharply. The effect on interest expenses grows for about 3-4 years after the reform and then stabilizes at about 0.2 log points until the mid-1990s.¹⁵

To check the robustness of our results, we consider alternative control group specifications in appendix. The evidence in Figure A.IX in which we compare the top and middle groups is particularly compelling, because the pre-trends in Panel B of that figure are essentially perfect. The overall conclusion from the different graphs is that there is a clear and substantial effect on interest payments, consistent with the mechanism we have in mind for the housing demand responses studied above.

Table 4 summarizes our DD estimates across specifications. The preceding discussion implies that, for this outcome, it is most informative to focus on the results in Panel B in which we use the

¹⁵In the mid-1990s, the divergence between the two groups begins to grow further, but this could be due to confounding shocks occurring later.

middle group as control. In this panel and for our baseline specification in column (3), we estimate an elasticity of interest expenses with respect to the net-of-tax rate equal to 0.4. This is considerably larger than the elasticities of real housing demand reported above.

To conclude, we find that the mortgage interest deduction has the largest effect on interest expenses (borrowing), a smaller but still sizeable effect on real housing demand, and a zero effect on homeownership.

5 Conclusion

Across the world, investments in owner-occupied housing are heavily subsidized through the tax system and various other policies. In countries such as the U.S., the largest subsidy comes from the deductability of mortgage interest payments combined with the absence of a tax on the imputed rents. While these subsidies have been motivated by the possibility of positive externalities from owning, there is no conclusive evidence on the existence of such externalities. As a result, many economists argue that the current policy regime leads to overinvestments in housing as well as excessive borrowing by homeowners. Furthermore, the subsidy is inequitable, because higher-income households tend to own more housing and therefore benefit more from the subsidy.¹⁶

Despite having debated the mortgage interest deduction for decades, we have relatively little evidence on the impact of this policy on real housing demand and none that live up to the current empirical standards in the empirical literature. Does the deduction increase homeownership in the long run? Does the deduction increase housing demand conditional on ownership, and if so, by how much? Without persuasive answers to these questions, we cannot fully assess the policy and the arguments for abolishing it. In this paper we provide such answers using Danish administrative data and a historical experiment that drastically reduced the mortgage interest deduction for top-rate taxpayers. The data and experiment allow us to look at long-run effects.

We find that the tax subsidy has a precisely estimated zero effect on homeownership, even in the very long run. This finding eliminates the externality argument for having the subsidy: although the optimal Pigouvian subsidy is in general equal to the marginal externality benefit (which we do not estimate), the welfare gain of introducing the subsidy depends on the elasticity of homeownership and goes to zero when the elasticity is zero. We also find that the tax subsidy increases

¹⁶To put this more precisely using the language of optimal taxation, housing does not justify subsidization on Atkinson-Stiglitz grounds, because housing demand is positively correlated with earnings ability, conditional on labor income (see [Gordon & Kopczuk 2014](#)).

housing demand, conditional on owning. In response to the scaling back for the subsidy, homeowners reduce the square footage and value of their houses. These effects require moving house and therefore take a long time to play out, as we show using our long panel and historical experiment. The long-run effects on housing demand are quite sizeable, corresponding to elasticities of about 0.1 for home size and 0.2 for home value. Finally, we find large effects on borrowing (interest expenses), consistent with existing evidence in the literature.

These findings suggest that the mortgage interest deduction, in the absence of full taxation of imputed rents, reduces both efficiency and equity. The policy does not increase housing demand at the extensive margin (and therefore yields no externality gain from homeownership), and it creates sizeable distortions on housing and debt demand at the intensive margin. These distortions are not offset by distributional gains as the deduction tends to favor higher-income households.


References

- AARON, HENRY. 1970. Income Taxes and Housing. *American Economic Review*, **60**, 789–806.
- ALAN, SULE, LETH-PETERSEN, SOEREN, & MUNK-NIELSEN, ANDERS. 2016. Tax Incentives and Borrowing. *Economics Letters*, **145**, 162–164.
- DUNSKY, ROBERT M., & FOLLAIN, JAMES R. 2000. Tax-Induced Portfolio Reshuffling: The Case of the Mortgage Interest Deduction. *Real Estate Economics*, **28**, 683–718.
- ENGELHARDT, GARY V., ERIKSEN, MICHAEL D., GALE, WILLIAM G., & MILLS, GREGORY B. 2010. What Are the Social Benefits of Homeownership? Experimental Evidence for Low-Income Households. *Journal of Urban Economics*, **67**, 249–258.
- GALE, WILLIAM G., GRUBER, JONATHAN, & STEPHENS-DAVIDOWITZ, SETH. 2007. Encouraging Homeownership Through the Tax Code. *Tax Notes*, **115**, 1171–1189.
- GLAESER, EDWARD L., & SHAPIRO, JESSE M. 2003. The Benefits of the Home Mortgage Interest Deduction. *Chap. 2, pages 37–82 of: POTERBA, JAMES (ed), Tax Policy and the Economy, vol. 17*. MIT Press.
- GORDON, ROGER H., & KOPCZUK, WOJCIECH. 2014. The Choice of the Personal Income Tax Base. *Journal of Public Economics*, **118**, 97–110.
- GRUBER, JONATHAN. 2015. *In: Public Finance and Public Policy, Fifth Edition*. Worth Publishers.
- JAPPELLI, TULLIO, & PISTAFERRI, LUIGI. 2007. Do People Respond to Tax Incentives? An Analysis of the Italian Reform to the Deductibility of Home Mortgage Interests. *European Economic Review*, **51**, 247–271.
- KLEVEN, HENRIK, KNUDSEN, MARTIN, KREINER, CLAUS, PEDERSEN, SØREN, & SAEZ, EMMA-NUEL. 2011. Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark. *Econometrica*, **79**, 651–692.
- KLEVEN, HENRIK JACOBSEN, & SCHULTZ, ESSEN ANTON. 2014. Estimating Taxable Income Responses Using Danish Tax Reforms. *American Economic Journal: Economic Policy*, **6**, 271–301.
- LING, DAVID C., & MCGILL, GARY A. 1998. Evidence on the Demand for Mortgage Debt by Owner-Occupants. *Journal of Urban Economics*, **44**, 391–414.


- POTERBA, JAMES M. 1984. Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach. *Quarterly Journal of Economics*, **99**, 729–752.
- POTERBA, JAMES M. 1991. House Price Dynamics: The Role of Tax Policy and Demography. *Brookings Papers on Economic Activity*, **22**, 143–204.
- POTERBA, JAMES M. 1992a. Tax Reform and the Housing Market in the Late 1980s: Who Knew What, and When Did They Know It? *Conference Series, the Federal Reserve Bank of Boston*, **36**, 230–261.
- POTERBA, JAMES M. 1992b. Taxation and Housing: Old Questions, New Answers. *American Economic Review, Papers and Proceedings*, **82**, 237–242.
- POTERBA, JAMES M. 2002. Taxation, Risk-Taking, and Household Portfolio Behavior. *Chap. 17, pages 1109–1171 of: AUERBACH, ALAN J., & FELDSTEIN, MARTIN (eds), Handbook of Public Economics, vol. 3*. North Holland: Elsevier Science Publishers.
- POTERBA, JAMES M., & SINAI, TODD. 2008. Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income. *American Economic Review, Papers and Proceedings*, **98**, 84–89.
- POTERBA, JAMES M., & SINAI, TODD. 2011. Revenue Costs and Incentive Effects of the Mortgage Interest Deduction for Owner-occupied Housing. *National Tax Journal*, **64**, 531–564.
- ROSEN, HARVEY S. 1979. Housing Decisions and the U.S. Income Tax: An Econometric Analysis. *Journal of Public Economics*, **11**, 1–23.
- ROSEN, HARVEY S., & ROSEN, KENNETH T. 1980. Federal Taxes and Homeownership: Evidence from Time Series. *Journal of Political Economy*, **88**, 59–75.

Figure 1: Marginal Tax Rates on Capital Income in Denmark, 1980-2017

A: Marginal Tax Rates on Negative Net Capital Income

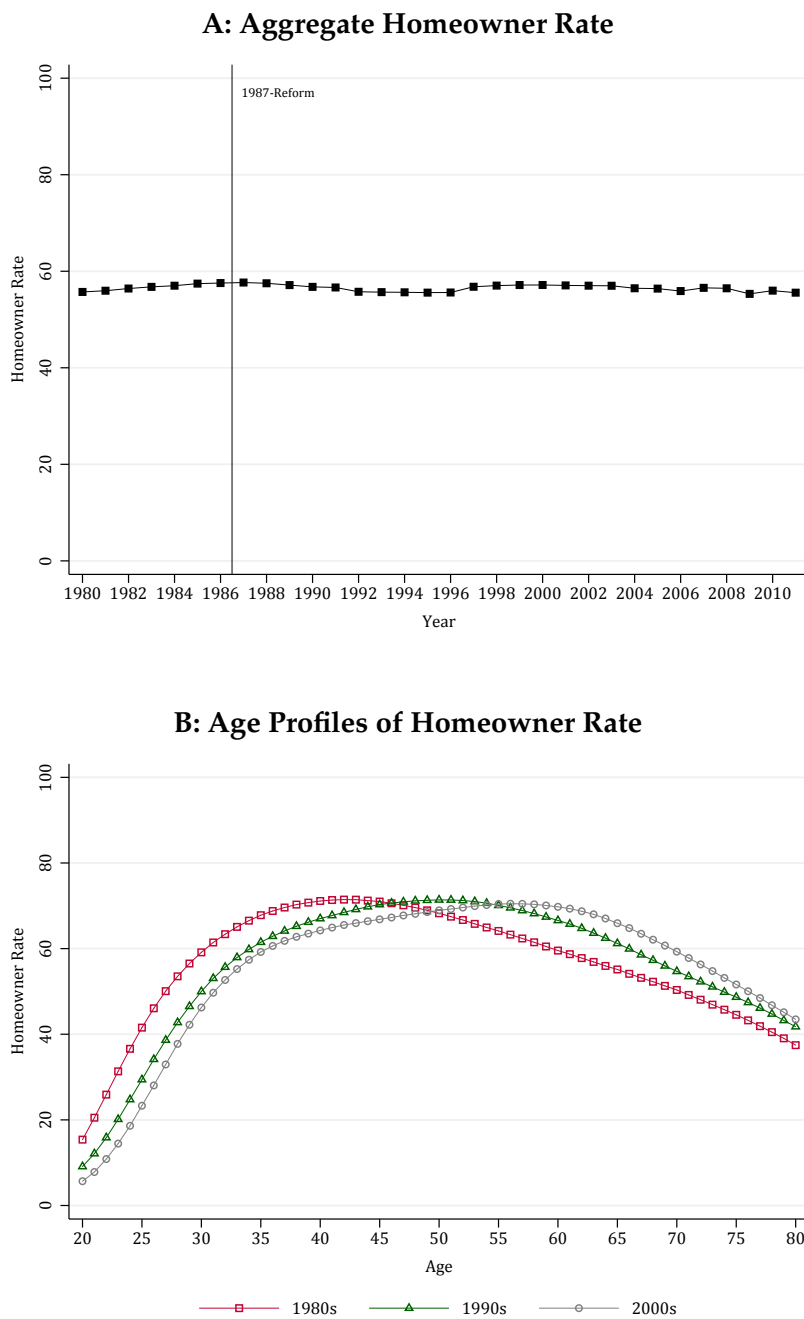


B: Marginal Tax Rates on Positive Net Capital Income



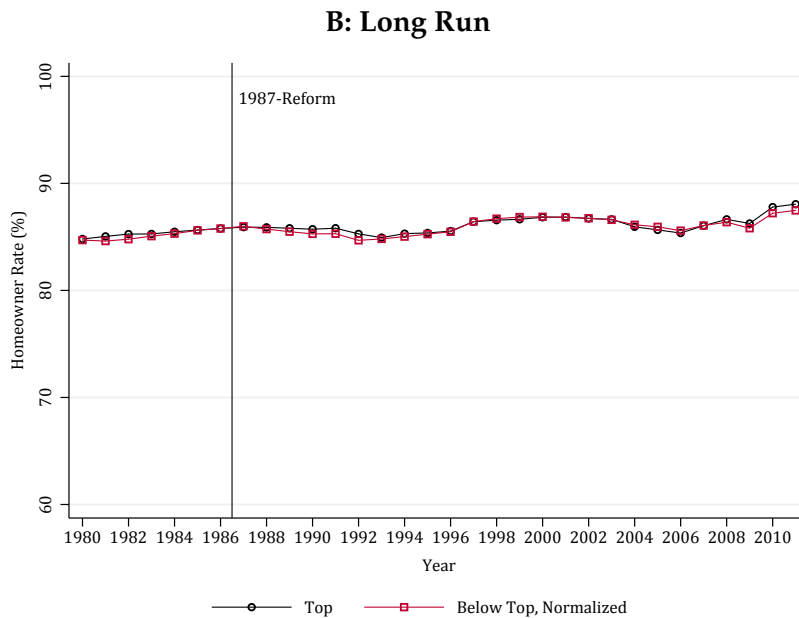
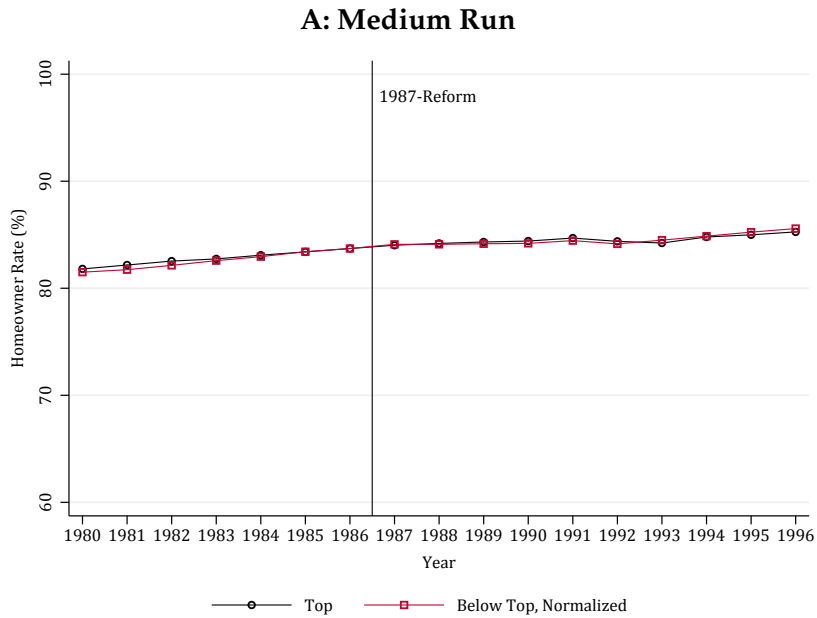
Notes: The figure is based on data from the Danish Tax Authorities (SKAT). The two panels show the marginal tax rates over time on negative and positive capital income, respectively, for different tax brackets.

Figure 2: Long-Run Evolution of Homeownership



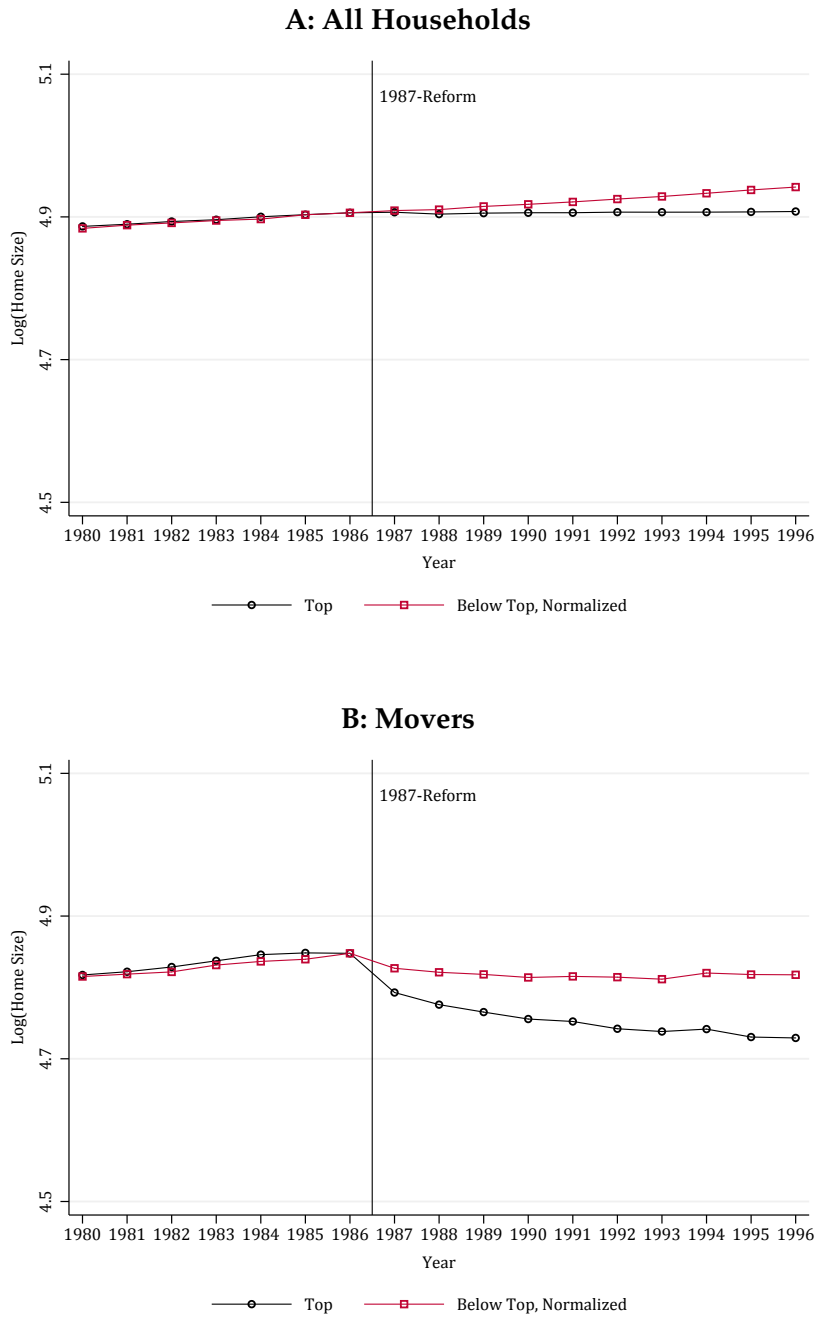
Notes: Panel A shows the homeowner rate over time for the Danish population above age 20. Panel B shows the age profile of the homeowner rate for three different decades. Homeownership is defined as described in Section 3.2.

Figure 3: Effect on Homeowner Rate




Notes: The graphs show mean predicted homeowner rates in each year for taxpayers in the top bracket (treatments) and taxpayers below the top bracket (controls) around the 1987-reform. Taxpayers are assigned to groups based on their status in 1984-1986, and we only include individuals who hold the relevant tax status in all three years. The mean predicted outcome is defined in equation (2) and is obtained from the regression in equation (1). This regression controls non-parametrically for group-specific age effects on homeownership. The results in Panel A are based on balanced panel of individuals observed all years between 1980-1996, while the results in Panel B are based on a balanced panel of individuals observed all years between 1980-2011.

Figure 4: Effect on Home Size (Square Feet)



Notes: The figure is constructed in the same way as Panel A of Figure 3, but with log(home size) as the outcome variable. Home size is defined as the area of the house used for habitation (in square feet). To avoid the results being affected by outliers, we have trimmed the home size variable at the 5th and 95th percentile. Panel A shows home size for all households (our full balanced panel), while Panel B shows home size in each year for those who move in that year (repeated cross-section picked from the balanced panel).


Figure 5: Effect on Home Value Among Movers




Notes: The figure is constructed in the same way as Panel A of Figure 3, but with log(home value) as the outcome variable. Home value is based on an appraisal made by the Danish Tax Authorities. The graph shows home value in each year for those who move in that year, similar to the previous graph on home size.

Figure 6: Effect on Interest Expenses

A: Baseline Specification



B: Controlling for Pre-Trends



Notes: Panel A is constructed in the same way as Panel A of Figure 3, but with $\log(\text{interest expenses})$ as the outcome variable. Interest expenses include interests paid on mortgage debt, consumer debt and other types of loans. Panel B is similar to Panel A, except that we control for a linear pre-reform time trend. To be specific, we estimate a group-specific linear time trend in $\log(\text{interest expenses})$ using only pre-reform years 1983-1986. When estimating the mean predicted outcome for the full period 1983-1996 (based on equations (1)-(2)), we have residualized the outcome using the estimated pre-trend from the first stage.

Table 1: Effect on Homeowner Rate

	(1)	(2)	(3)
A: Below-Top Group as Controls			
Effect of Reform	0.010 (0.001)	0.005 (0.001)	0.007 (0.001)
Elasticity wrt. $1 - \tau$	0.015 (0.001)	0.007 (0.001)	0.011 (0.002)
Observations	----- 34,440,869 -----		
B: Middle Group as Controls			
Effect of Reform	0.008 (0.001)	-0.008 (0.001)	-0.007 (0.001)
Elasticity wrt. $1 - \tau$	0.015 (0.002)	-0.016 (0.002)	-0.013 (0.002)
Observations	----- 15,980,104 -----		
C: Bottom Group as Controls			
Effect of Reform	-0.001 (0.001)	0.011 (0.001)	0.015 (0.001)
Elasticity wrt. $1 - \tau$	-0.001 (0.001)	0.012 (0.001)	0.017 (0.001)
Observations	----- 14,875,499 -----		
<i>Fixed Effects:</i>			
Age		×	×
Age x Treatment			×

Notes: The table shows estimates of the effect on the homeowner rate across a range of specifications. Panels A-C show three different control groups (below-top bracket, middle bracket, and bottom bracket), while columns (1)-(3) show different specifications of the age fixed effects. The baseline specification (3) with group-specific age dummies is presented in column (3). All estimates in the table are adjusted for group-specific linear pre-trends as described in the text. The elasticity with respect to the net-of-tax rate is defined in equation (4). The effects are estimated on a balanced panel of individuals observed in every year between 1980-1996, and they represent average effects over the 10 post-reform years.

Table 2: Effect on Home Size (Square Feet) Among Movers

	(1)	(2)	(3)
A: Below-Top Group as Controls			
Effect of Reform	-0.097 (0.002)	-0.095 (0.002)	-0.065 (0.003)
Elasticity wrt. $1 - \tau$	-0.142 (0.003)	-0.140 (0.003)	-0.096 (0.004)
Observations	----- 2,318,225 -----		
B: Middle Group as Controls			
Effect of Reform	-0.048 (0.002)	-0.057 (0.002)	-0.044 (0.003)
Elasticity wrt. $1 - \tau$	-0.094 (0.004)	-0.111 (0.004)	-0.085 (0.005)
Observations	----- 837,318 -----		
C: Bottom Group as Controls			
Effect of Reform	-0.146 (0.002)	-0.128 (0.002)	-0.090 (0.003)
Elasticity wrt. $1 - \tau$	-0.166 (0.003)	-0.146 (0.003)	-0.103 (0.003)
Observations	----- 899,309 -----		
<i>Fixed Effects:</i>			
Age		×	×
Age × Treatment			×

Notes: The table shows estimates of the effect on log(home size) among movers across a range of specifications. Home size is defined as the area of the house used for habitation (in square feet). Panels A-C show three different control groups (below-top bracket, middle bracket, and bottom bracket), while columns (1)-(3) show different specifications of the age fixed effects. The baseline specification (3) with group-specific age dummies is presented in column (3). All estimates in the table are adjusted for group-specific linear pre-trends as described in the text. The elasticity with respect to the net-of-tax rate is defined in equation (4). The effects are estimated on the sample of movers in each year (a repeated cross-section selected from the balanced panel between 1980-1996), and they represent average effects over the 10 post-reform years.

Table 3: Effect on Home Value Among Movers

	(1)	(2)	(3)
A: Below-Top Group as Controls			
Effect of Reform	-0.181 (0.007)	-0.158 (0.007)	-0.119 (0.008)
Elasticity wrt. $1 - \tau$	-0.266 (0.010)	-0.232 (0.010)	-0.175 (0.011)
Observations	----- 879,750 -----		
B: Middle Group as Controls			
Effect of Reform	-0.134 (0.008)	-0.106 (0.008)	-0.110 (0.008)
Elasticity wrt. $1 - \tau$	-0.260 (0.015)	-0.205 (0.015)	-0.214 (0.016)
Observations	----- 252,188 -----		
C: Bottom Group as Controls			
Effect of Reform	-0.318 (0.008)	-0.259 (0.008)	-0.215 (0.009)
Elasticity wrt. $1 - \tau$	-0.364 (0.009)	-0.297 (0.009)	-0.246 (0.010)
Observations	----- 195,449 -----		
<i>Fixed Effects:</i>			
Age		×	×
Age × Treatment			×

Notes: The table shows estimates of the effect on log(home value) among movers across a range of specifications. The home value is based on an appraisal by the Danish Tax Authorities. Panels A-C show three different control groups (below-top bracket, middle bracket, and bottom bracket), while columns (1)-(3) show different specifications of the age fixed effects. The baseline specification (3) with group-specific age dummies is presented in column (3). All estimates in the table are adjusted for group-specific linear pre-trends as described in the text. The elasticity with respect to the net-of-tax rate is defined in equation (4). The effects are estimated on the sample of movers in each year (a repeated cross-section selected from the balanced panel between 1980-1996), and they represent average effects over the 10 post-reform years.


Table 4: Effect on Interest Expenses

	(1)	(2)	(3)
A: Below-Top Group as Controls			
Effect of Reform	-0.196 (0.004)	-0.199 (0.004)	-0.172 (0.004)
Elasticity wrt. $1 - \tau$	-0.288 (0.006)	-0.293 (0.006)	-0.253 (0.006)
Observations	----- 13,615,363 -----		
B: Middle Group as Controls			
Effect of Reform	-0.192 (0.004)	-0.209 (0.004)	-0.204 (0.005)
Elasticity wrt. $1 - \tau$	-0.372 (0.008)	-0.405 (0.008)	-0.397 (0.009)
Observations	----- 7,248,682 -----		
C: Bottom Group as Controls			
Effect of Reform	-0.204 (0.005)	-0.132 (0.005)	-0.103 (0.005)
Elasticity wrt. $1 - \tau$	-0.233 (0.006)	-0.151 (0.006)	-0.118 (0.006)
Observations	----- 5,067,554 -----		
<i>Fixed Effects:</i>			
Age		×	×
Age × Treatment			×

Notes: The table shows estimates of the effect on log(interest expenses) across a range of specifications. Interest expenses include interests paid on mortgage debt, consumer debt and other loans. Panels A-C show three different control groups (below-top bracket, middle bracket, and bottom bracket), while columns (1)-(3) show different specifications of the age fixed effects. The baseline specification (3) with group-specific age dummies is presented in column (3). All estimates in the table are adjusted for group-specific linear pre-trends as described in the text. The elasticity with respect to the net-of-tax rate is defined in equation (4). The effects are estimated on a balanced panel of individuals between 1983-1996, and they represent average effects over the 10 post-reform years.


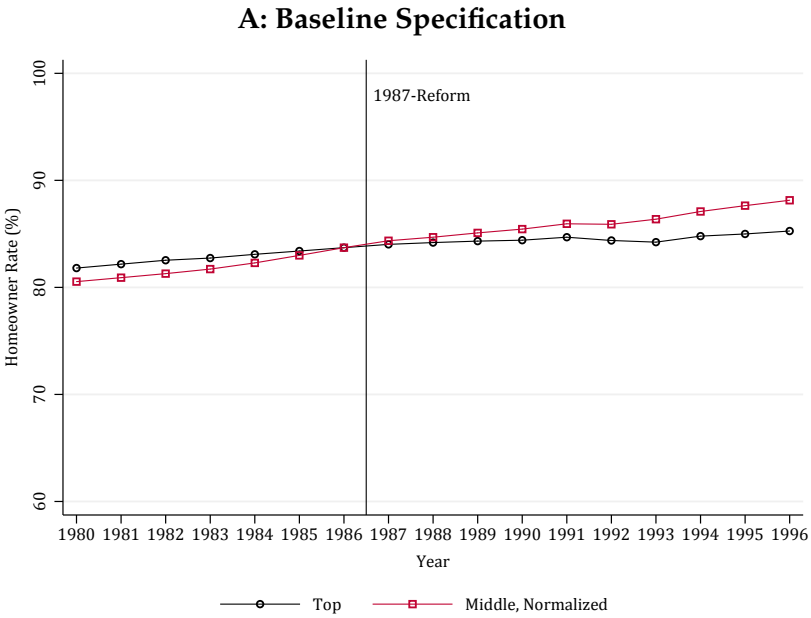
A Appendix

Figure A.I: Validation of Homeownership Measure



Notes: The figure compares our homeownership measure based on housing records to the ideal measure based on tax records available from 1987. The figure shows three homeowner rate series for the Danish population above age 20: our baseline measure using information from the housing records on living in an owner-occupied home, our measure using the housing records but including only "clean" cases (see text), and the ideal measure using information from the tax records on housing wealth. The two measures using housing records are essentially identical, and both of them are very close to the precise tax-based measure in terms of both the trend and the level. The level of homeownership is somewhat lower when using the housing register, because by construction this measure does not capture owners who are not residing in their owned property.


Figure A.II: Effect on Homeowner Rate Using Middle Group as Controls




Notes: Panel A is constructed in the same way as Panel A of figure 3, but using the middle tax group as controls. Panel B controls for pre-trends as follows: we estimate a group-specific linear time trend in the homeowner rate using only pre-reform data between 1980-1986. When estimating the mean predicted outcome for each group over the full period 1980-1996 (based on equations (1)-(2)), we residualize the outcome variable using the estimated pre-trend from the first stage.

Figure A.III: Effect on Homeowner Rate Using Bottom Group as Controls

A: Baseline Specification




B: Controlling for Pre-Trends



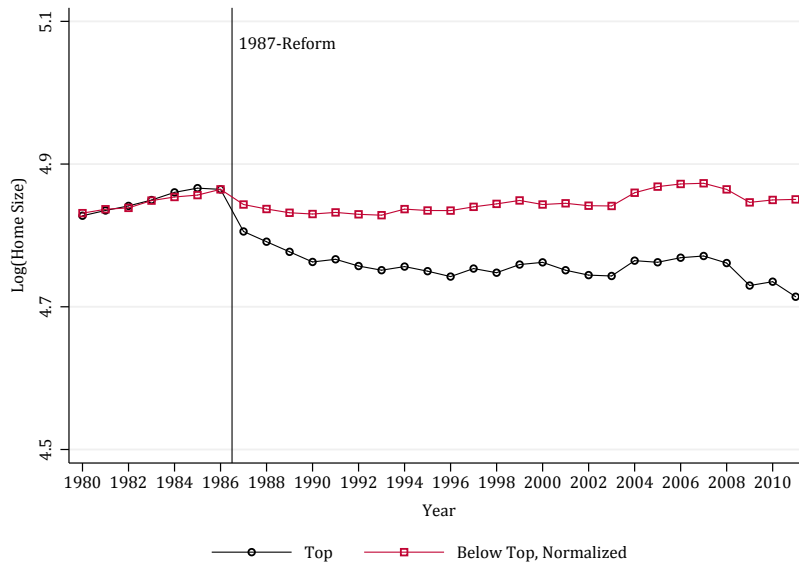
Notes: Panel A is constructed in the same way as Panel A of figure 3, but using the bottom tax group as controls. Panel B controls for pre-trends as follows: we estimate a group-specific linear time trend in the homeowner rate using only pre-reform data between 1980-1986. When estimating the mean predicted outcome for each group over the full period 1980-1996 (based on equations (1)-(2)), we residualize the outcome variable using the estimated pre-trend from the first stage.

Figure A.IV: Effect on Home Size in the Very Long Run

A: All Households


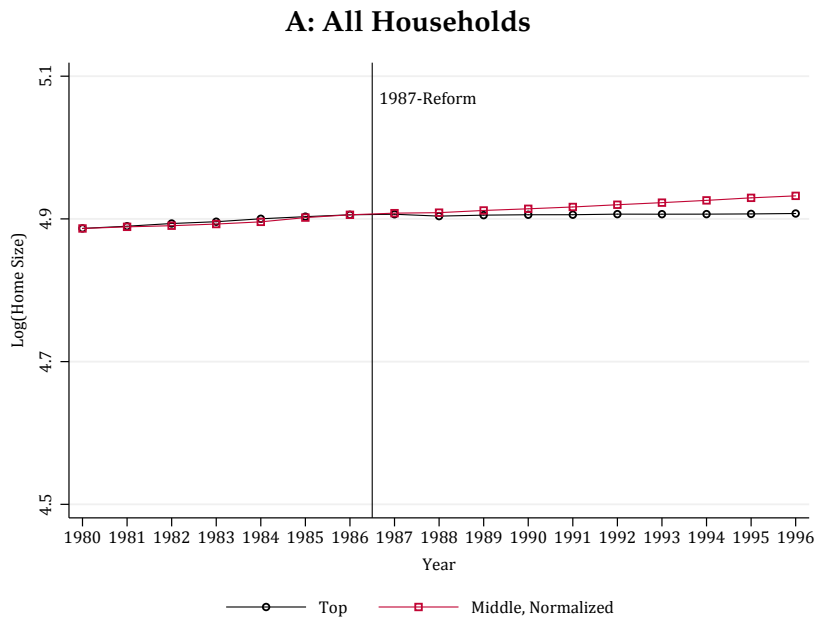


B: Movers





Notes: The figure is constructed in the same way as Figure 4, but using the longer balanced panel of individuals observed in all years between 1980-2011.

Figure A.V: Effect on Home Size Using Middle Group as Controls




Notes: This figure is constructed in the same way as Figure 4, except that here we use the middle tax group as controls.

Figure A.VI: Effect on Home Size Using Bottom Group as Controls




Notes: This figure is constructed in the same way as Figure 4, except that here we use the bottom tax group as controls.

Figure A.VII: Effect on Home Value Using Middle Group as Controls



Notes: This figure is constructed in the same way as Figure 5, except that here we use the middle tax group as controls.


Figure A.VIII: Effect on Home Value Using Bottom Group as Controls




Notes: This figure is constructed in the same way as Figure 5, except that here we use the bottom tax group as controls.

Figure A.IX: Effect on Interest Expenses Using Middle Group as Controls

A: Baseline Specification



B: Controlling for Pre-Trends



Notes: This figure is constructed in the same way as Figure 6, except that here we use the middle tax group as controls.

Figure A.X: Effect on Interest Expenses Using Bottom Group as Controls

A: Baseline Specification



B: Controlling for Pre-Trends



Notes: This figure is constructed in the same way as Figure 6, except that here we use the bottom tax group as controls.