American Economic Journal: Economic Policy 2021, 13(2): 273–303
https://doi.org/10.1257/pol.20170366

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

By Jonathan Gruber, Amalie Jensen, and Henrik Kleven*

Using a major reform that scaled back the mortgage interest deduction for middle- and high-income households in Denmark, we study how tax subsidies affect housing decisions. We present four main findings. First, the mortgage deduction has a precisely estimated zero effect on homeownership for high- and middle-income households. Second, the mortgage deduction has a clear effect on housing demand at the intensive margin, inducing homeowners to buy larger and more expensive houses. Third, the deduction has sizeable effects on household financial decisions, inducing them to increase indebtedness. Finally, the reduction of the tax subsidy lowered equilibrium house prices. (*JEL* G21, G51, H24, K34, R21, R31)

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 particularly large subsidy to ownership. In the United States, this subsidy alone amounted to $125 billion in 2020, which is more than 20 percent as large as the entire federal deficit in that year (Gruber 2021). Such subsidies have been motivated by the perceived externalities of homeownership (e.g., Glaeser and Shapiro 2003), an argument that relies on the ability of tax incentives to significantly increase homeownership.

Unfortunately, we have relatively little evidence on the effect of tax subsidies on real 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

* Gruber: MIT, Department of Economics, 50 Memorial Drive, E52-434 Cambridge, MA 02142, and NBER (email: gruberj@mit.edu); Jensen: University of Copenhagen, Department of Economics, Oester Farimagsgade 5 DK-1353, Copenhagen, Denmark (email: asj@econ.ku.dk); Kleven: Princeton University, Department of Economics, 20 Washington Road, Princeton, NJ 08540 (email: kleven@princeton.edu). Matthew Shapiro was coeditor for this article. We thank Jim Poterba and two anonymous referees for helpful comments.

† Go to https://doi.org/10.1257/pol.20170366 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

to capture the long-run effect on homeownership, which tends to be a slow-moving outcome.

In this paper, we focus on a setting and quasi-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 percent for the top group, by about 30 percent for the middle group, and left it roughly unchanged for the bottom group.[1] 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.

Based on this experiment and data, we estimate the effects of the mortgage interest deduction on the extensive margin of homeownership, the intensive margin of housing demand (home size and home value), the intensive margin of homeowner indebtedness, and house prices. 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 four main conclusions. First, there is a tightly estimated and robust zero effect of tax subsidies on homeownership for high- and middle-income households. Over multiple time periods, and considering multiple empirical strategies, we find no effect of the tax policy change on whether these 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 led to sizeable reductions in both square footage and appraised home values, driven by households who moved 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, there are sizeable effects of the mortgage interest deduction on household financial decisions. The reduction of the subsidy induced homeowners to reduce total interest expenses by almost 20 percent, with an implied elasticity with respect to the net-of-tax interest rate equal to 0.25. Finally, moving from our within-country setting to a cross-country setting, we provide suggestive evidence that the Danish tax reform had a negative impact on house prices. This conclusion is based on a difference-in-differences design that compares Denmark to a synthetic control country.

These results have implications for the debate about tax subsidies to owner-occupied housing. In particular, they suggest that such subsidies are ineffective at promoting externalities by increasing homeownership. There are significant effects on home size and home value, but it is unclear what externalities are

[1] At the time of the reform, the top bracket included about 10 percent of the population, the middle bracket included about 35 percent of the population, the bottom bracket included about 45 percent of the population, with the remaining 10 percent being below the exemption threshold.

delivered through larger and more expensive houses.[2] Likewise, the sizeable effects on financial decisions do not provide any positive externalities. While our findings speak against the desirability of tax subsidies to housing, it is important to bear in mind that only high- and middle-income households—roughly the upper half of the income distribution—are treated by the tax reform we study. It is possible that the effects are different at the bottom of the distribution.

The paper is organized as follows. Section I reviews the previous literature on taxation and homeownership. Section II describes the institutional setting and our administrative data. Section III develops a simple model of housing demand that facilitates interpretation of our empirical findings. Section IV presents evidence on the effects of the mortgage interest deduction on homeownership, home size and value, and indebtedness. Section V considers the potential general equilibrium house price effects of tax subsidies to homeownership. Section VI concludes.

I. Literature Review

A long-standing literature 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 and Rosen (1980). They studied the time series correlation between the tax subsidy and homeownership, estimating that 25 percent of the growth in homeownership in the United States 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 and Shapiro 2003). This literature is reviewed in Gale, Gruber, and Stephens-Davidowitz (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 were 7–11 percent more likely to purchase a home.

[2] In principle, there could be spillovers associated with home size and value. If a homeowner invests in a larger and nicer house, this could generate positive externalities on the neighborhood. However, our finding that the intensive-margin demand effects are driven entirely by movers—as opposed to stayers who improve their existing homes—goes against this interpretation. The best existing evidence on the possibility of such externalities is in Engelhardt et al. (2010), who used randomized access to savings accounts that can be used for home downpayments as an instrument for homeownership. They did find that the treatment group spent more on home maintenance—but only on the *inside* of the house, which generates private benefits, and not on the *outside* of the house, the part that generates social benefits.

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 US tax code.

There have been some quasi-experimental studies in the US that use reforms in the tax treatment of mortgage interest to estimate impacts on household indebtedness. Ling and McGill (1998) and Dunsky and Follain (2000) both study the impact of the sizeable reduction in the mortgage interest deduction implied by the tax rate cuts of the US 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 and 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 and Schultz (2014) and Alan, Leth-Petersen, and Munk-Nielsen (2016) study the change in Danish tax policy that is the focus of the current paper. Kleven and Schultz (2014) estimate an elasticity of negative capital income (such as from interest expenses on mortgages and other loans) that vary between -0.10 and -0.13, while Alan, Leth-Petersen, and Munk-Nielsen (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.

II. Institutional Setting and Data

A. *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 percent in the average municipality, and around 10 percent of the population was liable to pay it.[3] The middle bracket tax rate

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

was 62 percent (paid by about 35 percent of the population) and the bottom bracket tax rate was 48 percent (paid by about 45 percent 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.[4]

The implications of the 1987 reform for marginal tax rates on capital income are illustrated in Figure 1.[5] 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 percent 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. The reform increased this rate by 81 percent for the top group, while increasing it by only 29 percent for the middle group and reducing it by 6 percent for the bottom group. This is the identifying variation on which our paper is based.

Unlike some countries with mortgage interest deductions (such as the United States), 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 and Sinai 2008, 2011). However, the Danish tax system was not such an ideal system. The tax on imputed rental income was artificially low for several reasons: the calculation of taxable rental income was based on a rental rate of only 2.5 percent of an assessed home value, and the assessed home value was based on a tax appraisal set below the true market value. Moreover, the tax was reduced by the presence of a deduction equal to 1 percent 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-to-late 1990s further reduced the tax rate on negative capital income. However, these reforms introduced

[4] This 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 percent of all homeowners (and a larger fraction of homeowners with mortgages).

[5] Data on marginal tax rates over time are from the Danish Ministry of Taxation (1980–2017).

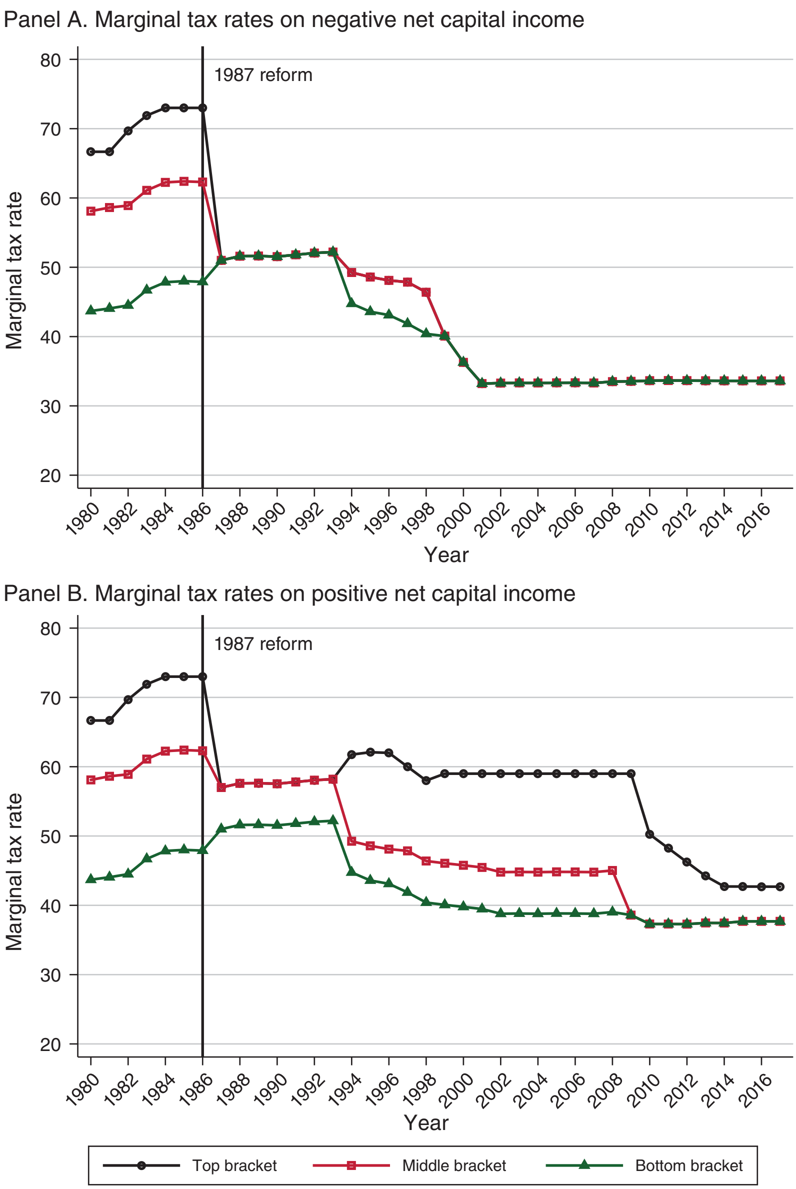

Figure 1. Marginal Tax Rates on Capital Income in Denmark, 1980–2017

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.

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]

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

B. *Data and Descriptives*

Our analysis is based on administrative data for the full Danish population from 1980 to 2011.[7] 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 socioeconomic variables. We are able to link 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.[8]

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 percent 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 onward. 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 percent 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.

[7] See Gruber, Jensen, and Kleven (2021) and Statistics Denmark (2020a, b, c, d, e, f).

[8] Even 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.

Figure A.I in the online 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 but focusing only on the "clean" cases (single person or a cohabiting/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.[9]

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 semibalanced panel of individuals who we observe in every year between 1980 and 1996, and who are between the ages of 20 and 80. The sample is only "semibalanced" 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 and 2011. Just under 60 percent of the population are homeowners, a number that has stayed remarkably constant over time. The stability of the homeowner rate over three decades that featured 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 United States and other countries (e.g., Glaeser and Shapiro 2003; Gale, Gruber, and Stephens-Davidowitz 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 decreases 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.

[9] Indeed, 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.

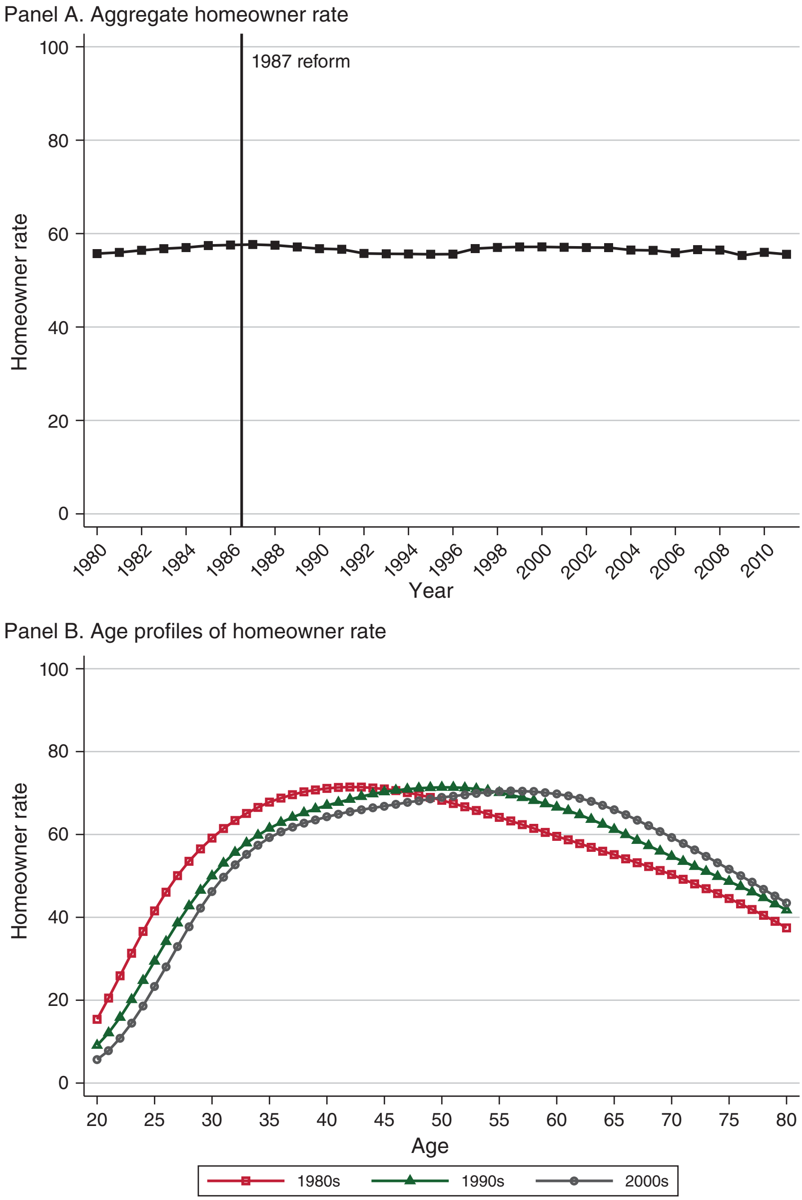

Figure 2. Long-Run Evolution of Homeownership

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 IIB.

III. Housing Market Model

To provide a lens for interpreting the empirical findings, this section develops a simple housing market model with intensive and extensive margin demand responses to taxes. Households derive utility from the consumption of a numeraire good c and housing h. They choose whether to become homeowners or renters (extensive margin) and how much housing to consume conditional on each (intensive margin). Each household lives in only one house, which they either own (in which case $h = h^o > 0$ and $h^r = 0$) or rent (in which case $h = h^r > 0$ and $h^o = 0$). To simplify the exposition, we consider the following parametrization of preferences:

$$u = c + \frac{A}{1 + 1/\varepsilon}\left(\frac{h}{A}\right)^{1+1/\varepsilon} + q \cdot I\{h^o > 0\}, \tag{1}$$

where A, ε, and q characterize housing preferences. As we shall see, the parameter ε determines the intensive-margin elasticity of housing demand. The parameter q captures a preference for owning over renting (which can be negative for those who prefer renting), which we allow to be heterogeneous in the population according to a smooth pdf $f(q)$ and cdf $F(q)$. The distribution of q is key for the extensive-margin elasticity of housing demand. We note that the quasilinear utility function eliminates income effects on housing demand as we will focus on the substitution effect of tax incentives.[10]

The budget constraint is given by

$$c + \left((1-\tau)m + \tau\tilde{r}\right) \cdot ph^o \cdot I\{h^o > 0\} + rh^r \cdot I\{h^r > 0\} = (1-\tau)y, \tag{2}$$

where y is exogenous labor income, τ is the income tax rate, m is the mortgage interest rate, r is the rental price, $\tilde{r}$ is imputed rent for homeowners, and p is the price per unit of housing.[11] Consistent with the Danish tax system, we allow for the taxation of imputed rental income, but note that the imputed rental price $\tilde{r}$ is set at a very low level as described above.

Let us first consider the intensive margin decision of how much housing to consume, conditional on either owning or renting. If owning, households maximize

$$u = (1-\tau)y - \left((1-\tau)m + \tau\tilde{r}\right) \cdot ph^o + \frac{A}{1+1/\varepsilon}\left(\frac{h^o}{A}\right)^{1+1/\varepsilon} + q, \tag{3}$$

which implies the following first-order condition:

$$h^o = A\left[\left((1-\tau)m + \tau\tilde{r}\right)p\right]^{\varepsilon}, \tag{4}$$

where $\varepsilon < 0$ is the elasticity of housing demand with respect to the user cost $((1-\tau)m + \tau\tilde{r})p$. Income taxation has an a priori ambiguous effect on demand conditional on owning, because of the offsetting effects from the mortgage interest deduction (positive effect) and the tax on imputed rental income (negative effect). However, when $\tilde{r}$ is sufficiently low, a higher income tax rate unambiguously increases housing demand among homeowners.

[10] As described in detail in Kleven and Schultz (2014), the Danish tax reform that we study combined reductions in marginal tax rates with base broadening, implying that changes in average tax rates were relatively small. As a result, income effects of the reform are likely to be small, and indeed Kleven and Schultz (2014) found little evidence of income effects in the context of taxable income responses.

[11] Note that while equation (2) includes taxes and interest in the user cost of housing, it ignores other costs (such as depreciation and maintenance) that are not pertinent to our analysis.

If renting, households maximize

$$u = (1-\tau)y - rh^r + \frac{A}{1+1/\varepsilon}\left(\frac{h^r}{A}\right)^{1+1/\varepsilon}, \tag{5}$$

which gives

$$h^r = Ar^{\varepsilon}. \tag{6}$$

We now turn to the extensive margin decision between owning and renting. A household decides to own if and only if the utility of owning is greater than or equal to the utility of renting, which puts a lower bound on the preference parameter q:

$$q \geq \left((1-\tau)m + \tau\tilde{r}\right)\cdot ph^o - rh^r + \frac{A}{1+1/\varepsilon}\left[\left(\frac{h^r}{A}\right)^{1+1/\varepsilon} - \left(\frac{h^o}{A}\right)^{1+1/\varepsilon}\right]. \tag{7}$$

Using the first-order conditions (4) and (6), this implies

$$q \geq \frac{A}{1+\varepsilon}\cdot\left(\left[\left((1-\tau)m + \tau\tilde{r}\right)p\right]^{1+\varepsilon} - r^{1+\varepsilon}\right) \equiv \bar{q}. \tag{8}$$

Households with $q \geq \bar{q}$ become homeowners such that the homeowner rate equals $1 - F(\bar{q})$. As with the intensive margin, the effect of income taxation on the extensive margin is in general ambiguous due to the potentially offsetting effects of the mortgage interest deduction and the tax on imputed rental income. When $\tilde{r}$ is sufficiently small, we have $d\bar{q}/d(1-\tau) > 0$, in which case a larger income tax rate unambiguously increases the homeowner rate. The extensive margin elasticity with respect to the net-of-tax rate $1-\tau$ can be defined as

$$\eta \equiv \frac{\partial\left(1-F(\bar{q})\right)}{\partial\bar{q}}\frac{d\bar{q}}{d(1-\tau)}\frac{1-\tau}{1-F(\bar{q})} = -f(\bar{q})\frac{d\bar{q}}{d(1-\tau)}\frac{1-\tau}{1-F(\bar{q})}. \tag{9}$$

It is useful to highlight two aspects of this elasticity. First, it is a partial equilibrium (micro) elasticity rather than a general equilibrium (macro) elasticity. The derivative $d\bar{q}/d(1-\tau)$ takes the rental rate r and the mortgage interest rate m as given, but these may respond to taxation in general equilibrium.[12] The focus on micro elasticities corresponds to what we identify empirically using our quasi-experimental approach. Second, the extensive margin elasticity depends on the density of homeowner preferences around the cutoff, $f(\bar{q})$. We have derived the elasticity for a particular income level y, and if there is heterogeneity in income, then the relevant density is a conditional one, $f_y(\bar{q})$. It is conceivable that, among high-income people, preferences for homeownership are sufficiently strong that the density is close to zero around the cutoff, $f_y(\bar{q}) \approx 0$, so that their extensive margin elasticity is zero. We will test this in the empirical analysis.

[12] We refer to Sommer and Sullivan (2018) for a recent calibration analysis of the general equilibrium effects of repealing the mortgage interest deduction.

The total demand for owner-occupied housing is given by

$$D = \int_y \int_{\bar{q}}^{\infty} h^o f_y(q) g(y)\, dq dy = A\Big[\big((1-\tau)m + \tau\tilde{r}\big)p\Big]^{\varepsilon} \cdot H, \tag{10}$$

where $g(y)$ is the density distribution of income and $H \equiv \int_y (1 - F_y(\bar{q}))g(y)\, dy$ is the aggregate homeowner rate in the population. This demand must be met by the available supply of housing, net of the housing stock that is rented out. For simplicity of exposition, we do not spell out a model of housing supply in the owner-occupied and rental markets. If housing demand responds to the mortgage interest deduction (at the intensive or extensive margins) and housing supply is less then perfectly elastic, then house prices will respond to taxes. In particular, in the short run where the housing stock is fixed, we would expect tax changes to capitalize into house prices (see, e.g., Poterba 1984). While our main objective is to provide evidence on demand responses at each margin (as captured by the elasticities ε and η), we will also provide suggestive evidence on general equilibrium price responses.

IV. Demand 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.

A. *Empirical Specification*

The analysis is based on a difference-in-differences (DiD) 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 start out by showing 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 comparing specifically to those in the bottom bracket who are (almost) entirely untreated. Both of these specifications capture effects on high-income households. For homeownership, we also show results from comparing middle- and bottom-bracket taxpayers, which capture effects on middle-income taxpayers.

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 II, 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; and 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 1980–1996, giving us seven pre-reform years and ten post-reform years. We will consider a longer time window as well. The assignment of treatment status is based on pre-reform tax status. Specifically, treatment status is determined using three pre-reform years, requiring treatments and controls to be in a given tax bracket in every year between 1984 and 1986.[13] Using three years ensures that the comparison groups are relatively stable in terms of their tax status. Even so, the bracket location of the households in our comparison groups is not perfectly persistent, implying that our estimates represent intention-to-treat (ITT) effects rather then treatment-on-the-treated (TOT) effects.

The balanced panel approach implies that people in our estimation sample become older over time. This is a potential issue due to the strong life-cycle effects on homeownership (see Figure 2 above). Homeownership increases strongly until the middle ages and then falls gradually over the rest of the life cycle. This poses a challenge for a DiD 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 DiD 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:

$$Y_{it}^{g} = \sum_{y} \alpha_{y}^{g} \cdot \mathbf{1}(y = t) + \sum_{a} \beta_{a}^{g} \cdot \mathbf{1}(a = age_{it}) + \nu_{it}^{g}, \tag{11}$$

[13] Taxpayers who move across brackets during the three pre-reform years 1984–1986 are therefore dropped from the estimation sample. This restriction makes us drop around 13 percent of the sample.

where Y_{it}^{g} is the outcome for individual i in group g in year t, and where $\mathbf{1}[\cdot]$ 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, 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 present graphs in which we plot the coefficient α_t^g for each year t plus a constant that ensures the plotted value in year 1986 (the last pre-reform year) matches the actual level in 1986.

When trends are parallel (which we verify based on pre-trends), the effect of the reform can be obtained based on a simple DiD. 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. For some of the outcomes and control groups, the pre-trends are not fully parallel after absorbing life-cycle trends through the age dummies. In those cases, we estimate a group-specific linear 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 1983 to 1986 (the last four pre-reform years). We then run the specification (11) 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.

B. *Extensive Margin: Homeownership*

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 equation (11) in which we absorb life-cycle 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 and 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 trends are perfectly parallel prior to the reform, and they stay perfectly parallel in the years after the reform.

It could be argued that ten 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 panel B of Figure 2) and do not often reoptimize later in the life cycle. If so, homeownership is largely predetermined later in life, 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)

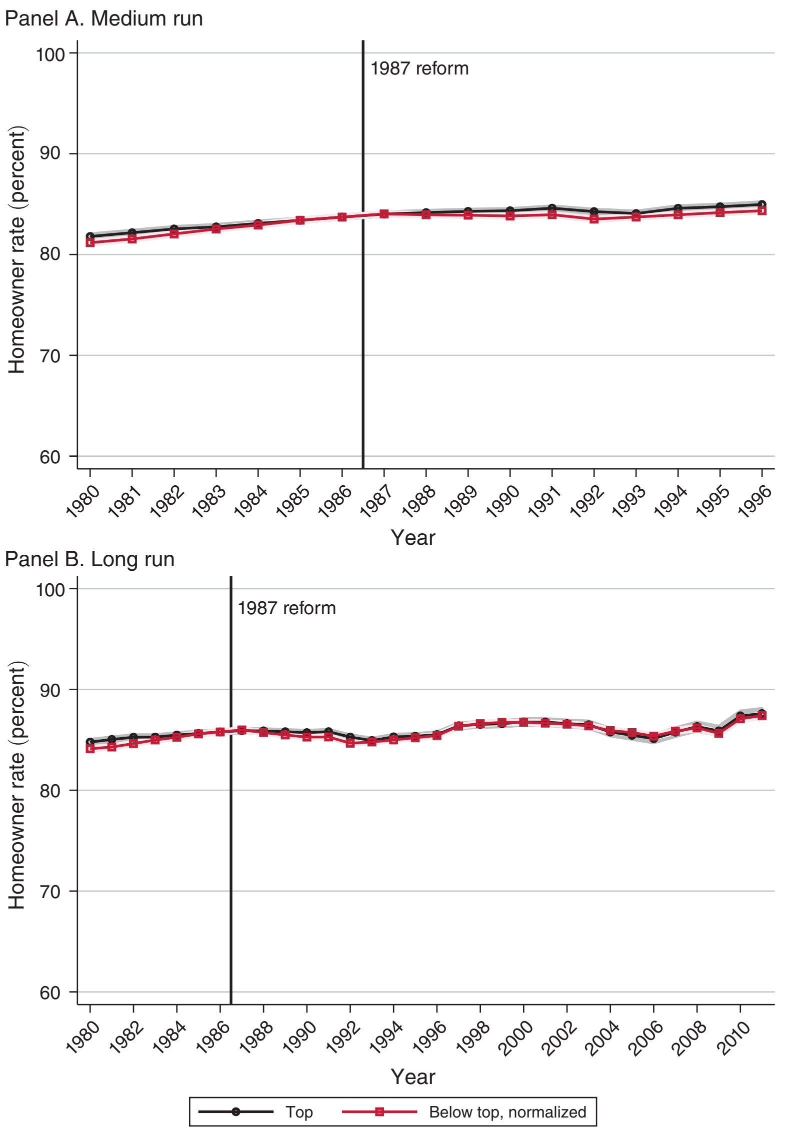

Figure 3. Effect on Homeowner Rate

Notes: The graph shows 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 tax status in 1984–1986. The plotted values are the year coefficients estimated from equation (11) plus a constant such that the 1986 values (in both groups) equal the actual homeowner rate in the top bracket in that year. The actual homeowner rates in 1986 were 83.7 percent for taxpayers in the top bracket and 62.0 percent for taxpayers below the top bracket. The results in panel A are based on balanced panel of individuals observed all years between 1980 and 1996, while the results in panel B are based on a balanced panel of individuals observed all years between 1980 and 2011. Confidence intervals are based on standard errors clustered at the household level.

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 life cycle 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 online Appendix, we investigate the robustness of our results to the choice of comparison groups. Online Appendix Figure A.II compares those in the top bracket to those in the bottom bracket. The top panel is based on our default specification (11) and the 1980–1996 balanced panel. We see that the treatment and control series are not completely parallel as the bottom group follows a slightly flatter trend than the top group. 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 online Appendix Section A. 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.

Online Appendix Figure A.III investigates if there are extensive margin responses further down the income distribution by comparing the middle and bottom groups. As described above, those in the middle bracket experienced a smaller, but still sizeable cut in their mortgage interest deduction. The figure is constructed in the same way as the previous one, with panel A showing our default specification and panel B adjusting for pre-trends. The analysis shows that, even in the middle part of the distribution, there are no effects of the reduced tax subsidy on homeownership. This is an important extension of our results: while it might be expected that high-income people are unresponsive along the extensive margin (due to their high baseline homeowner rates), this is certainly not true for people on moderate incomes. Taken together, our results imply that the entire upper-half of the income distribution—the top and middle brackets combined represented about 45 percent of the income distribution at the time of the reform—do not make homeownership decisions based on the housing subsidy. The tax reform we study does not allow us to identify responses (or their absense) in the lower half of the distribution.

Table 1 presents estimates of the homeownership effect and elasticities across specifications. Panels A and B show different control groups, while columns 1–3 show different sets of controls. The results in column 3 in which we include age fixed effects that vary by group corresponds to our preferred specification presented graphically, while the other columns consider more parsimonious specifications without any age controls or with common age controls. To be precise, the

Table 1—Effect on Homeowner Rate

	(1)	(2)	(3)
Panel A. Top versus below-top groups			
Effect of reform	0.011	0.004	0.006
	(0.001)	(0.001)	(0.001)
Elasticity with respect to $1-\tau$	0.016	0.005	0.008
	(0.001)	(0.001)	(0.002)
Observations		34,440,869	
Panel B. Top versus bottom groups			
Effect of reform	−0.009	−0.005	−0.002
	(0.001)	(0.001)	(0.001)
Elasticity with respect to $1-\tau$	−0.011	−0.006	−0.002
	(0.001)	(0.001)	(0.001)
Observations		14,875,499	
Fixed effects:			
Age		X	X
Age × treatment			X

Notes: The table shows estimates of the effect on the homeowner rate across a range of specifications. Panels A and B show two different control groups (below-top bracket and bottom bracket), while columns 1–3 show different specifications of the age fixed effects. The baseline specification (12) 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 (13). The effects are estimated on a balanced panel of individuals observed in every year between 1980 and 1996, and they represent average effects over the ten post-reform years. Standard errors are clustered at the household level.

estimates in the richest specification in column 3 are based on the following DiD regression:

$$Y_{it}^{R} = \alpha_0 + \alpha_1 \cdot post_t + \alpha_2 \cdot treat_i + \alpha_3 \cdot post_t \cdot treat_i + \sum_a \beta_a \cdot \mathbf{1}(a = age_{it}) + \sum_a \beta_a^T \cdot \mathbf{1}(a = age_{it}) \cdot treat_i + \nu_{it}. \tag{12}$$

Here the outcome Y_{it}^{R} is residualized using a 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 ten-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 convert this into a semielasticity of the homeowner rate with respect to the net-of-tax rate using the following formula:

$$\varepsilon = \frac{\hat{\alpha}_3}{\Delta\log\left(1-\tau^{T}\right) - \Delta\log\left(1-\tau^{C}\right)}, \tag{13}$$

where $\Delta\log(1-\tau^{g})$ is the reform-induced change in the net-of-tax rate on negative capital income 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 very small across all specifications. In fact, the estimates are sometimes 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—i.e., when comparing the top and bottom brackets. But even there the effect is less than one-quarter of a percentage point in our preferred specification that fully absorbs life-cycle trends. Second, because the underlying tax variation is very large, the small reform effects translate into very small elasticities. In our preferred specification, the estimated elasticities vary between -0.002 and 0.008. 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 $1 - \tau$ below -0.004.

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. These results apply to both top-rate and middle-rate taxpayers, representing roughly the upper-half of the income distribution in Denmark.

C. *Intensive Margin: Home Size and Home Value*

Having found no effect of tax subsidies on the extensive margin of owning versus renting, we now turn to the effects on the intensive margin. Because there is no evidence of any externalities associated with intensive margin responses, they are properly 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 equation (11). 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.[14]

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 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

[14] For example, we could restrict the sample to individuals who are homeowners in the pre-reform year(s), naturally using 1984–1986 as we do to assign treatment status. This restriction implies that we start from a homeowner rate of 100 percent 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.

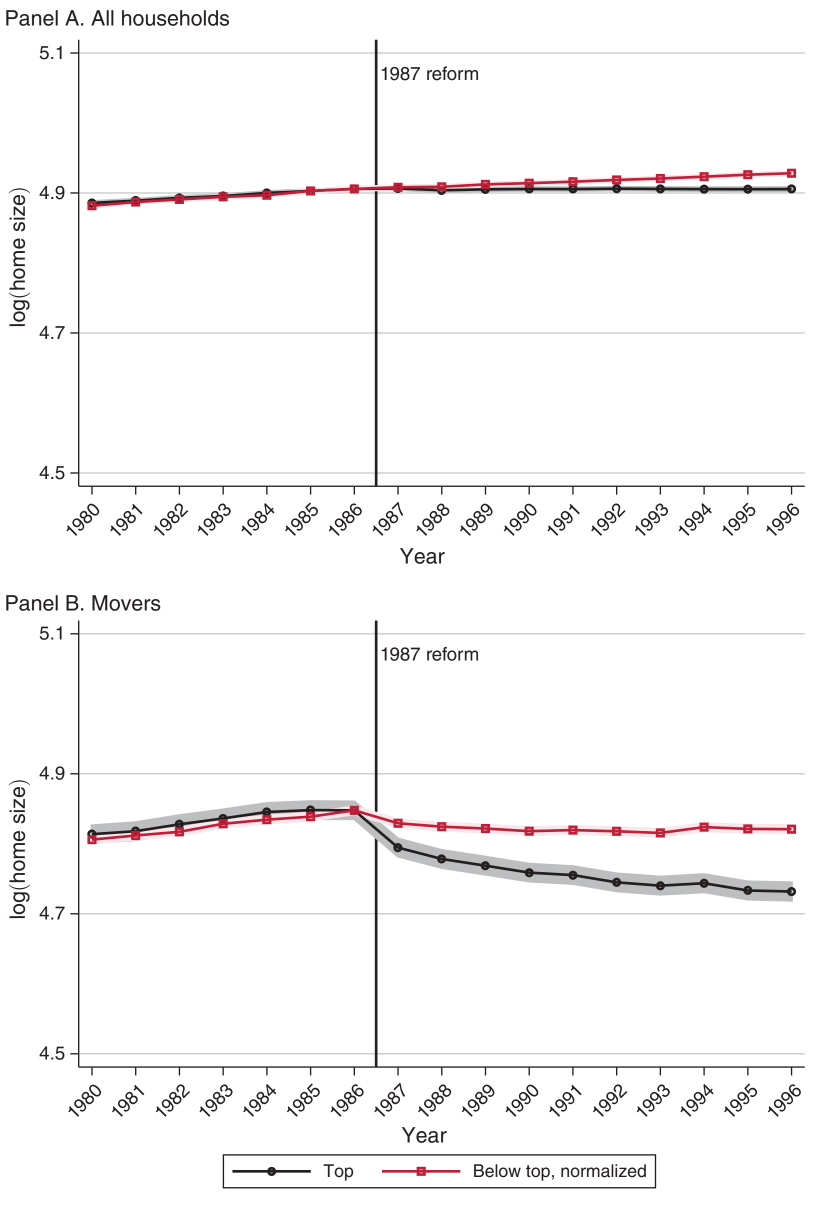

Figure 4. Effect on Home Size (Square Feet)

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 fifth and ninety-fifth 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).

ten 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 percent 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 below-top group. The effect grows for about six to seven 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 percent on home size.[15]

While panel B of Figure 4 captures longer-run effects by considering movers ("reoptimizers"), another possibility is obviously to extend the time window. We do this in online 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 DiD strategies in general are tenuous so long after the reform experiment as confounding shocks become increasingly likely over time.[16] For this reason, we prefer to obtain long-run effects based on the mover sample within a shorter time horizon.

Online Appendix Figure A.V investigates robustness to the choice of control group. When using taxpayers in the bottom bracket to form a control group, we obtain qualitatively similar patterns—but quantitatively larger effects—than when using all taxpayers below the top bracket. This is consistent with the incentives created by the reform, because the middle group also experienced a sizeable reduction in their deduction 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 specifications. The table is constructed in the same way as the previous table for homeownership, with different control groups in panels A and B and different age controls in columns 1–3. Our preferred estimates (with group-specific age controls) in column 3 are based on the DiD specification in equation (12) and the elasticity definition in equation (13).[17] Because the outcome variable is now in logs, these are "real" elasticities rather than semielasticities. The table shows that the negative effect on home size is robust across specifications, but that the magnitude is larger when

[15] The 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 childbirth, marriage, or job changes), while the mortgage interest deduction matters for the type of house that is purchased, conditional on moving.

[16] Note 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.

[17] This 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).

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

	(1)	(2)	(3)
Panel A. Top versus below-top groups			
Effect of reform	−0.066	−0.071	−0.063
	(0.002)	(0.002)	(0.003)
Elasticity with respect to $1 - \tau$	−0.096	−0.105	−0.092
	(0.003)	(0.003)	(0.004)
Observations		2,318,221	
Panel B. Top versus bottom groups			
Effect of reform	−0.127	−0.116	−0.105
	(0.002)	(0.002)	(0.003)
Elasticity with respect to $1 - \tau$	−0.145	−0.132	−0.120
	(0.003)	(0.003)	(0.003)
Observations		899,302	
Fixed effects:			
Age		X	X
Age × treatment			X

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 and B show two different control groups (below-top bracket and bottom bracket), while columns 1–3 show different specifications of the age fixed effects. The baseline specification (12) 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 (13). The effects are estimated on the sample of movers in each year (a repeated cross-section selected from the balanced panel between 1980 and 1996), and they represent average effects over the ten post-reform years. Standard errors are clustered at the household level.

comparing top-bracket taxpayers to bottom-bracket taxpayers than when comparing top-bracket taxpayers to everyone below the top bracket. This difference reflects partly that the tax variation is different in these comparisons, which can be seen from the fact that the elasticities vary much less. Focusing on our preferred specification in column 3, the elasticity of home size lies between 0.09 and 0.12 (in absolute value) depending on the comparison groups.

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.

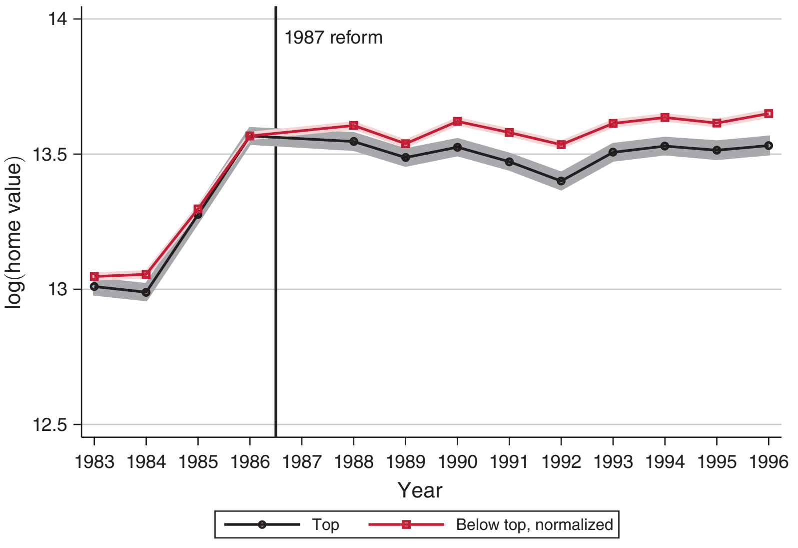

Figure 5. Effect on Home Value among Movers

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.

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.[18] The graph suggests a clear impact on home values: the treatment and control series are roughly parallel up until the time of the reform and then diverge. After ten years, home values in the top group have dropped by about 0.15 log points (or about 15 percent) 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. As for the other outcomes, the online Appendix repeats the graphical analysis when using only (untreated) taxpayers in the bottom bracket to form a control group. This graph confirms the effect on home values.

Table 3 presents regression results across specifications in the same way as the previous tables. In our preferred specification (column 3), the elasticity of home value with respect to the net-of-tax rate lies between 0.18 and 0.25 in absolute value across the different control groups. 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—around 0.1 for home size and 0.2 for home value.

[18] Because 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.

Table 3—Effect on Home Value among Movers

	(1)	(2)	(3)
Panel A. Top versus below-top groups			
Effect of reform	−0.181	−0.158	−0.119
	(0.007)	(0.007)	(0.008)
Elasticity with respect to $1 - \tau$	−0.266	−0.232	−0.175
	(0.01)	(0.01)	(0.011)
Observations		579,750	
Panel B. Top versus bottom groups			
Effect of reform	−0.318	−0.259	−0.215
	(0.008)	(0.008)	(0.009)
Elasticity with respect to $1 - \tau$	−0.364	−0.297	−0.246
	(0.009)	(0.009)	(0.01)
Observations		195,449	
Fixed effects:			
Age		X	X
Age × treatment			X

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 and B show two different control groups (below-top bracket and bottom bracket), while columns 1–3 show different specifications of the age fixed effects. The baseline specification (12) 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 (13). The effects are estimated on a balanced panel of individuals between 1983 and 1996, and they represent average effects over the ten post-reform years. Standard errors are clustered at the household level.

D. *Intensive Margin: Interest Expenses*

In this section, we investigate if the reform of the mortgage interest deduction affected interest expenses (i.e., borrowing) of homeowners. As reviewed in Section I, there is existing work on how interest expenses respond to tax subsidies, including work on the Danish 1987 reform (Kleven and Schultz 2014; Alan, Leth-Petersen, and Munk-Nielsen 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 deductibility 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 considers total interest expenses as the outcome variable. For homeowners, mortgage interest tends to represent 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

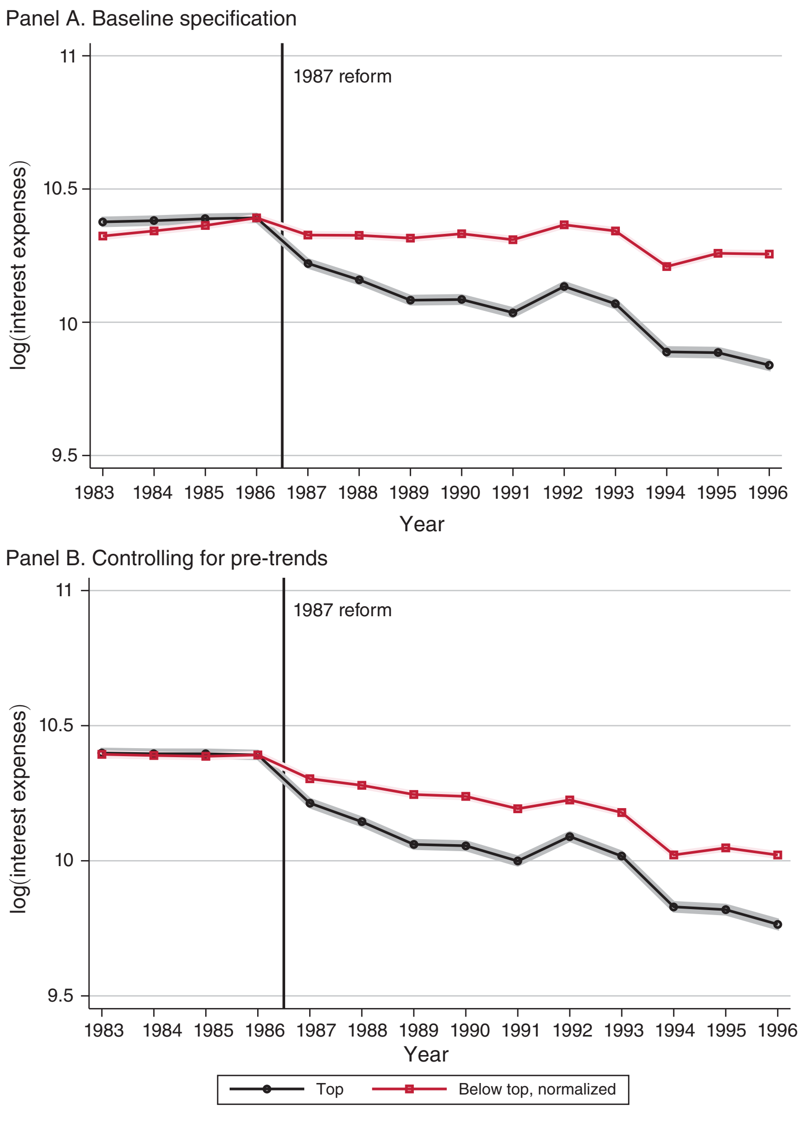

Figure 6. Effect on Interest Expenses

Notes: Panel A is constructed in the same way as panel A of Figure 3, but with log(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(interest expenses) using only pre-reform years 1983–1986. When estimating the mean predicted outcome for the full period 1983–1996 (based on equation (11)), we have residualized the outcome using the estimated pre-trend from the first stage.

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 (11) without adjusting for pre-trends, while panel B adjusts for linear pre-trends. We focus on the time period from 1983 onward in these graphs, as opposed to 1980 onward in all previous

Table 4—Effect on Interest Expenses

	(1)	(2)	(3)
Panel A. Top versus below-top groups			
Effect of reform	−0.196	−0.199	−0.172
	(0.004)	(0.004)	(0.004)
Elasticity with respect to $1 - \tau$	−0.288	−0.293	−0.253
	(0.006)	(0.006)	(0.006)
Observations		13,615,363	
Panel B. Top versus bottom groups			
Effect of reform	−0.204	−0.132	−0.103
	(0.005)	(0.005)	(0.005)
Elasticity with respect to $1 - \tau$	−0.233	−0.151	−0.118
	(0.006)	(0.006)	(0.006)
Observations		5,067,554	
Fixed effects:			
Age		X	X
Age × treatment			X

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 and B show two different control groups (below-top bracket and bottom bracket), while columns 1–3 show different specifications of the age fixed effects. The baseline specification (12) 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 (13). The effects are estimated on the sample of movers in each year (a repeated cross-section selected from the balanced panel between 1980 and 1996), and they represent average effects over the ten post-reform years. Standard errors are clustered at the household level.

graphs. This is because nominal interest rates fell sharply between 1982 and 1983 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 onward. 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 three to four years after the reform and then stabilizes at about 0.2 log points until the mid-1990s.[19] As before, we check the robustness of our results to the choice of control group in the online Appendix. When comparing the top and bottom groups in online Appendix Figure A.VII, the effect on interest expenses is somewhat smaller, but still clear and sizeable.

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

Table 4 summarizes our estimates across specifications. For our baseline specification in column 3, we estimate an elasticity of interest expenses with respect to the net-of-tax rate between 0.12 and 0.25 in absolute value across control groups. Overall, our findings show that the mortgage interest deduction has a clear effect on interest expenses (borrowing), consistent with the mechanism we have in mind for the housing demand responses studied above.

V. House Price Effects of the Mortgage Interest Deduction

Do house prices respond to the mortgage interest deduction? As discussed in Section III, we would expect the reduction of the tax subsidy to reduce equilibrium prices provided that housing demand responds and given housing supply is inelastic in the short run. We did not find any demand responses at the extensive margin, but the demand responses at the intensive margin may affect prices. This section provides suggestive evidence on house price effects using cross-country data and a synthetic control approach.[20]

A. *Data and Empirical Strategy*

To estimate general equilibrium house price effects, we compare the evolution of house prices in Denmark and other countries around the 1987 reform. We use data on quarterly residential property prices made available by the Bank for International Settlements (BIS).[21] Specifically, we focus on *real* property prices calculated by deflating nominal property prices with the Consumer Price Index.[22] Of the countries covered by the BIS dataset, we include Denmark along with 16 other developed countries that are relatively comparable to Denmark. These are the highest-income member countries of the OECD.[23]

Figure A.VIII in the online Appendix shows the raw house price series in Denmark and in the other countries. Panel A compares Denmark to the other Nordic countries, panels B–C compare Denmark to the rest of Europe, while panel D compares Denmark to English-speaking countries outside of Europe. The time series show that Denmark had a steeper house price trend than most other countries

[20] In our setting, there are two possible strategies for estimating house price effects of the mortgage interest deduction. These strategies are designed to uncover, respectively, local GE effects and national GE effects. For estimating local GE effects, we can divide Denmark into a set of regions that vary by their fraction of top-rate taxpayers and are therefore treated differentially by the reform. The housing market in regions with many top-rate taxpayers were hit hard by the reform, whereas regions with few top-rate taxpayers were hit less hard. We have explored such strategy, but find that it suffers from a fundamental problem: there is relatively little spatial heterogeneity in income levels in Denmark. Specifically, grouping municipalities into quintiles of the fraction of top-rate taxpayers, we find that, prior to the 1987 reform, the highest-income municipalities had about 20 percent top-rate taxpayers while the lowest-income municipalities had about 8 percent. Therefore, even if we compare only the top and bottom quintiles, the natural experiment would consist in a modest 12 pp difference in treatment intensity. This is a very small experiment to use for identification. As a result, we focus on estimating national GE effects using cross-country data and a synthetic control approach. While a cross-country strategy comes with its own set of challenges, it is based on a large experiment given that about half of the Danish population (those in the middle and top brackets) experienced substantial increases in the net-of-tax mortgage interest rate.

[21] See Bank for International Settlements (1983–1996).

[22] The results are very similar for nominal property prices.

[23] Specifically, the set of countries includes Australia, Belgium, Canada, Finland, France, Germany, Ireland, Italy, the Netherlands, New Zealand, Norway, Spain, Sweden, Switzerland, the United Kingdom, and the United States.

in the early 1980s. However, there is a marked slowdown in house price growth in Denmark (but not elsewhere) starting about one year before the 1987 reform. This slowdown occurred too early to be an effect of the reform, and, in fact, it is typically attributed to business cycle changes and macro stabilization policy in Denmark in the mid-1980s. After the reform, Denmark continues on a flatter trend than most other countries for a number of years. Notably, there is a small, but sharp downward shift in Danish house prices *immediately* after the reform, which suggests a negative house price effect of the reduced tax subsidy.

To estimate house price effects, we compare Denmark to a synthetic control country using the method by Abadie, Diamond, and Hainmueller (2010). The donor pool includes the 16 countries shown in online Appendix Figure A.VIII. To construct the weights that form the synthetic country, we match on house price levels in each of the pre-reform quarters. An important choice in this analysis is the length of the time horizon. While we consider different horizons below, we argue that a compelling analysis requires us to focus on a narrow time window. This is because, over longer periods, there are many confounding factors that impact house prices across countries and make it hard to interpret the patterns. Given the forward-looking nature of house prices, focusing on the short run should be informative for capturing the capitalization of taxes in prices. In our preferred specification, we consider house prices from four quarters before to four quarters after the Danish tax reform.

B. *Results*

We start by considering a long time window (1983–1996) in panel A of Figure 7. The graph shows broadly similar trends in the years leading up to the reform and a large gap opening up in the years following the reform. However, when inspecting the pre-trends more critically, we see that this specification does not match the trends very precisely. In particular, the synthetic control country does not feature the pre-reform slowdown in house price growth that we discussed above. In other words, the confounding house price changes in Denmark (due to business cycle changes and macro stabilization policy in the mid-1980s) are not captured by the synthetic control country. As a result, this analysis does not have a clear causal interpretation.

This motivates focusing on the narrower two-year window in panel B of Figure 7. This graph shows a more compelling picture: the house price trends are essentially identical in Denmark and the synthetic control country prior to the reform, and then they diverge immediately after the reform. The divergence is created by the sharp drop in Danish house prices just after the reform, as discussed above. The resulting gap implies that the 1987 reform reduced house prices in Denmark by about 12 percent. These patterns are suggestive and consistent with the intensive-margin demand effects estimated above.

VI. 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 United States, the largest subsidy comes from the deductability of mortgage interest

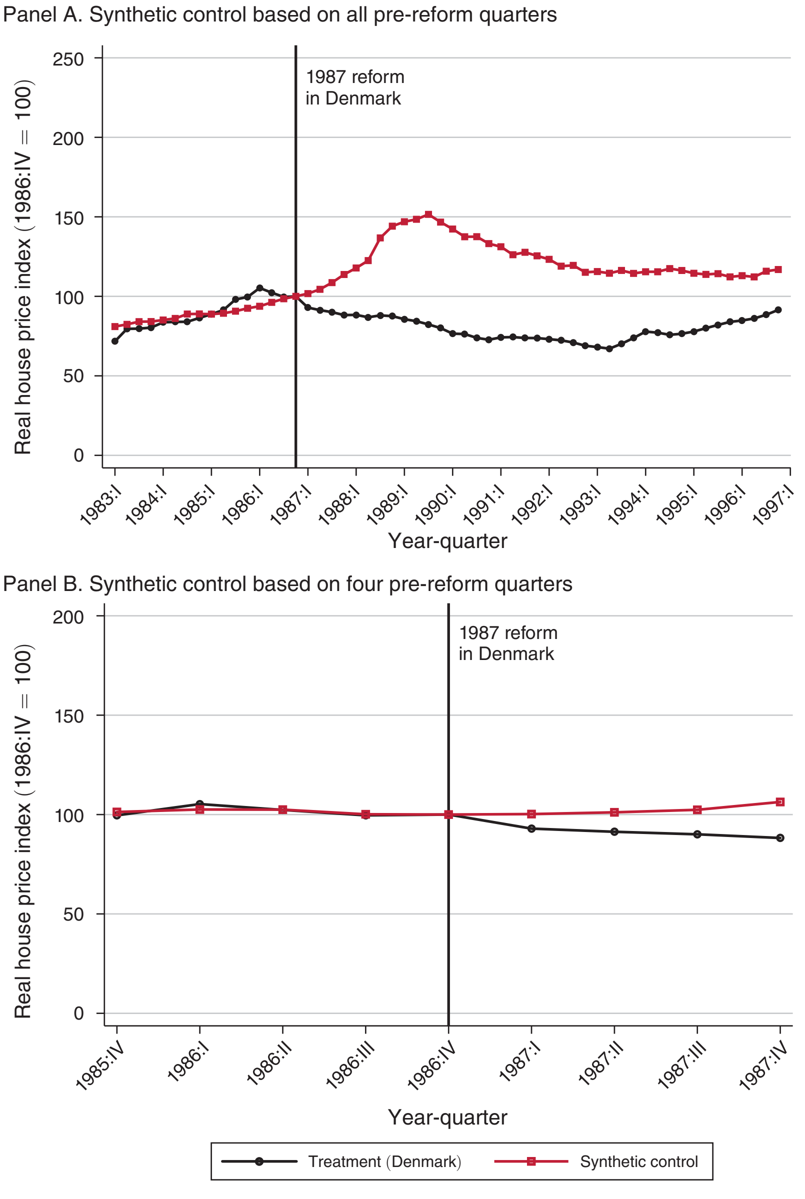

Figure 7. Effect on National House Prices

Notes: The graphs show real residential property prices in Denmark and a synthetic control country around the Danish 1987 reform, using data from the Bank for International Settlements' online database (http://www.bis.org/statistics/pp.htm). Each price series is normalized to 100 in 1986:IV, the last quarter before the reform. The price series for the synthetic control country represent a weighted average of other countries, where the weights are constructed by matching on house price levels in each pre-reform quarter using the Synth command in Stata. The donor pool for the synthetic country includes 16 countries: Australia, Belgium, Canada, Finland, France, Germany, Ireland, Italy, the Netherlands, New Zealand, Norway, Spain, Sweden, Switzerland, the United Kingdom, and the United States. Panel A considers a long time window (1983:I to 1996:IV) and here the synthetic control country is based on the following weights: Australia: 0.16 and United Kingdom: 0.84. Panel B considers a short time window (1985:IV to 1987:IV) and here the synthetic control country weights are: Australia: 0.5, Finland: 0.35, and Switzerland: 0.15.

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.[24]

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 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 strongly reduced the mortgage interest deduction for high- and middle-income 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. As a result, the policy does not generate any externalities associated with owning, at least not in the sample of households treated by the Danish reform. We do find that the tax subsidy affects 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. Consistent with the mechanism through which such demand responses should operate, we find sizeable effects of the mortgage interest deduction on borrowing (interest expenses). Finally, we provide suggestive evidence that the scaling back of the tax subsidy reduced house prices.[25]

These findings suggest that the mortgage interest deduction, in the absence of full taxation of imputed rents, may be harmful for 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.

[24] To 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 and Kopczuk 2014).

[25] In the public debate, the possibility of such house price effects is frequently used as an argument against removing the mortgage interest deduction. It is argued that, without this tax subsidy, house prices would fall and put homeowners under water. From an optimal tax perspective, this is not a compelling argument for keeping the mortgage interest deduction. First, the efficiency effects of taxation depend on real allocations rather than prices. The housing demand elasticities that we estimate represent the key sufficient statistics for measuring efficiency. Second, while any change in the taxation of housing capital may have implications for short-run economic fluctuations (due to the importance of the housing market for the macro economy) and the distribution of wealth (existing versus future homeowners, owners versus renters), this is hardly an argument for introducing or retaining a permanent subsidy to homeowners. Other policy instruments are available to deal with the business cycle and distributional consequences of removing the subsidy.

A limitation of our approach is that it only allows for estimating treatment effects in the upper part of the income distribution. The effects may be different among households further down, and this would affect the policy implications. We hope that future work will investigate if our findings are transportable to lower-income households and to other countries. This would put the policy implications on firmer ground.

REFERENCES

Aaron, Henry. 1970. "Income Taxes and Housing." *American Economic Review* 60 (5): 789–806.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association* 105 (490): 493–505.

Alan, Sule, Søren Leth-Petersen, and Anders Munk-Nielsen. 2016. "Tax Incentives and Borrowing." *Economics Letters* 145: 162–64.

Bank for International Settlements (BIS). 1983–1996. "National Sources, BIS Residential Property Price Database." BIS Statistics. http://www.bis.org/statistics/pp.htm (accessed August 16, 2019).

Danish Ministry of Taxation. 1980–2017. "Tax Rates over Time." Danish Ministry of Taxation Database. https://www.skm.dk/%20skattetal/satser/tidsserier/centrale-skattesatser-i-skattelovgivningen-2018-2020 (accessed April 6, 2020).

Dunsky, Robert M., and James R. Follain. 2000. "Tax-Induced Portfolio Reshuffling: The Case of the Mortgage Interest Deduction." *Real Estate Economics* 28 (4): 683–718.

Engelhardt, Gary V., Michael D. Eriksen, William G. Gale, and Gregory B. Mills. 2010. "What Are the Social Benefits of Homeownership? Experimental Evidence for Low-Income Households." *Journal of Urban Economics* 67 (3): 249–58.

Gale, William G., Jonathan Gruber, and Seth Stephens-Davidowitz. 2007. "Encouraging Homeownership Through the Tax Code." *Tax Notes Federal* 115: 1171–89.

Glaeser, Edward L., and Jesse M. Shapiro. 2003. "The Benefits of the Home Mortgage Interest Deduction." In *Tax Policy and the Economy,* Vol. 17, edited by James M. Poterba, 37–82. Cambridge, MA: MIT Press.

Gordon, Roger H., and Wojciech Kopczuk. 2014. "The Choice of the Personal Income Tax Base." *Journal of Public Economics* 118: 97–110.

Gruber, Jonathan. 2021. *Public Finance and Public Policy.* 7th ed. New York: Worth Publishers.

Gruber, Jonathan, Amalie Jensen, and Henrik Kleven. 2021. "Replication data for: Do People Respond to the Mortgage Interest Deduction? Quasi-Experimental Evidence from Denmark." American Economic Association [publisher], Inter-univeristy Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E119157V1.

Jappelli, Tullio, and Luigi Pistaferri. 2007. "Do People Respond to Tax Incentives? An Analysis of the Italian Reform of the Deductibility of Home Mortgage Interests." *European Economic Review* 51 (2): 247–71.

Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez. 2011. "Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark." *Econometrica* 79 (3): 651–92.

Kleven, Henrik Jacobsen, and Esben Anton Schultz. 2014. "Estimating Taxable Income Responses Using Danish Tax Reforms." *American Economic Journal: Economic Policy* 6 (4): 271–301.

Ling, David C., and Gary A. McGill. 1998. "Evidence on the Demand for Mortgage Debt by Owner-Occupants." *Journal of Urban Economics* 44 (3): 391–414.

Poterba, James M. 1984. "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach." *Quarterly Journal of Economics* 99 (4): 729–52.

Poterba, James M. 1991. "House Price Dynamics: The Role of Tax Policy and Demography." *Brookings Papers on Economic Activity* 22: 143–203.

Poterba, James M. 1992a. "Tax Reform and the Housing Market in the Late 1980s: Who Knew What, and When Did They Know It?" *Federal Reserve Bank of Boston Economic Research Conference Series* 36: 230–61.

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

Poterba, James M. 2002. "Chapter 17–Taxation, Risk-Taking, and Household Portfolio Behavior." In *Handbook of Public Economics*, Vol. 3, edited Alan J. Auerbach and Martin Feldstein, 1109–71. Amsterdam: Elsevier.

Poterba, James M., 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): 84–9.

Poterba, James M., 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.

Rosen, Harvey S. 1979. "Housing Decisions and the U.S. Income Tax: An Econometric Analysis." *Journal of Public Economics* 11 (1): 1–23.

Rosen, Harvey S., and Kenneth T. Rosen. 1980. "Federal Taxes and Homeownership: Evidence from Time Series." *Journal of Political Economy* 88 (1): 59–75.

Sommer, Kamila, and Paul Sullivan. 2018. "Implications of US Tax Policy for House Prices, Rents, and Homeownership." *American Economic Review* 108 (2): 241–74.

Statistics Denmark. 2020a. "Befolkningen (BEF, Population Demographics), 1980–2011." Danmarks Statistiks Forskningsservice (accessed November 2016).

Statistics Denmark. 2020b. "Boligtællingen (BOL, Homes), 1980–2011." Danmarks Statistiks Forskningsservice (accessed November 2016).

Statistics Denmark. 2020c. "Ejendomsvurderinger (EJVK, Property Assessments), 1980–2011." Danmarks Statistiks Forskningsservice (accessed November 2016).

Statistics Denmark. 2020d. "Familier (FAFA, Families), 1980–2011." Danmarks Statistiks Forskningsservice (accessed November 2016).

Statistics Denmark. 2020e. "Husstande og Familier (FAIN, Households), 1980–2011." Danmarks Statistiks Forskningsservice (accessed November 2016).

Statistics Denmark. 2020f. "Indkomst (IND, Income), 1980–2011." Danmarks Statistiks Forskningsservice (accessed November 2016).