Wealth Taxation and Wealth Inequality: Evidence from Denmark, 1980-2014*

Katrine Jakobsen (University of Copenhagen)
Kristian Jakobsen (Kraka)
Henrik Kleven (London School of Economics)
Gabriel Zucman (UC Berkeley and NBER)

April 23, 2017

Preliminary and Incomplete

Abstract

We use 35 years of administrative wealth data from Denmark combined with large changes in wealth taxation to estimate the effects of wealth taxes on wealth accumulation and wealth inequality. Denmark used to impose one of the world's highest marginal tax rates on wealth at the top of the distribution, but this tax was drastically reduced and ultimately abolished between 1989 and 1997. Because of specific features of how the wealth tax was designed, the tax cuts create large and compelling variation for understanding behavioral responses among the wealthiest segments of the population. Using a variety of empirical approaches, we show that the effects of wealth taxation on wealth accumulation and pre-tax wealth inequality are small, though not zero. When interpreting the wealth tax as a tax on the return to wealth, the implied savings elasticity with respect to the net-of-tax return is very small. Our results contribute to a nascent literature on wealth taxation, and they have implications for the optimal taxation of capital and labor income.

*Katrine Jakobsen: k.m.jakobsen@econ.ku.dk Kristian Jakobsen: krj@kraka.org; Henrik Kleven: h.j.kleven@lse.ac.uk; Gabriel Zucman: zucman@berkeley.edu. We thank Maxim Massenkoff for outstanding research assistance.

1 Introduction

What are the economic effects of taxing household wealth? While a large literature estimates the elasticity of labor supply and taxable income, less is known about how taxes affect the supply of capital. The lack of evidence makes it hard to assess the desirability of taxing household wealth, a proposal that has gained interest following Thomas Piketty's call for a global wealth tax (Piketty 2014) and new evidence of sharply rising wealth inequality in the United States (Saez & Zucman 2016). How would wealth taxes affect the saving and consumption decisions of the rich? Would they reduce wealth inequality, and by how much?

Answering these questions is difficult due to several major empirical challenges. First, while many countries collect data on taxable income — and have made such data available for research purposes — very few countries collect individual data on wealth. Second, because wealth accumulation is a dynamic and slow-moving process, estimating the effects of wealth taxes requires wealth data over a long time period. Third, because wealth is always very concentrated (much more than labor income, due to the cumulative processes that govern wealth accumulation), it is crucial to estimate behavioral responses for the wealthiest individuals. However, finding sources of variation allowing for such estimation has so far been elusive.

In this paper, we break new ground on these questions based on Danish wealth data and quasi-experimental variation in wealth taxation. Until 1997, Denmark taxed household wealth above an exemption threshold located around the 98th percentile of the household wealth distribution. Through to the 1990s, a dozen of OECD countries levied similar taxes (OECD 1988), but the Danish wealth tax was the largest of its kind. The marginal tax rate on taxable wealth was equal to 2.2% up until the late 1980s, which is a very high rate considering that it applies to the stock of wealth as opposed to only the return to wealth.[1] The Danish government implemented large changes to the wealth tax between 1989-1991 — cutting the marginal rate in less than half and doubling the exemption threshold for married couples — before eventually abolishing the tax in 1997. These policy changes probably represent the largest natural experiment with wealth taxation ever conducted. Besides the experiment, a key advantage of the Danish setting is that the authorities have been collecting high-quality, micro-level, population-wide data on wealth and income since 1980. Critically, wealth data are still collected after the abolition of the wealth tax, making it possible to track individual wealth over the long run. These data allow us to study wealth accumulation and wealth inequality during 8 years of high wealth taxation

[1] Assuming a return to wealth of e.g. 4.4% (which is in the ballpark of empirical estimates), a marginal wealth tax of 2.2% corresponds to a 50% tax on the flow of capital income.

(from 1980-1988), 7 years of more moderate taxation (from 1989-1996), and 15 years of zero wealth taxation (from 1997 onwards). Using these data, we make two contributions.

Our first contribution is to provide a comprehensive analysis of trends in wealth inequality in Denmark since 1980, contrasting the Danish experience with the U.S. experience. We combine the Danish registry data with national accounts and household balance sheet data in order to provide comprehensive, internationally-comparable estimates of Danish wealth inequality. The results show that wealth concentration in Denmark is lower than in the United States: the top-1% wealth share is currently around 20% in Denmark as opposed to 37% in the United States. What is more, the relative trends in wealth inequality between these two countries change strikingly around the late 1990s. Until that time, the two countries featured roughly similar trends of rising wealth inequality; if anything wealth inequality increased at a faster pace in Denmark that in the United States. However, wealth inequality has stabilized in Denmark since the late 1990s, while it has continued to grow sharply in the United States. Hence, the abolition of wealth taxes in Denmark was not followed by an upsurge in pre-tax wealth inequality: The quantiles of the wealth distribution that were affected by the wealth tax before its abolition grew at the same rate as the rest of the population after the abolition. This is the opposite pattern from what we would expect if wealth taxes were impeding wealth accumulation at the top, although these time series comparisons may of course be affected by a number of confounding factors.

This motivates our second contribution, which is to investigate the causal effect of wealth taxes on wealth accumulation at the micro level. To this end, we consider several difference-in-differences approaches using the large policy changes around 1990. We exploit three main sources of variation. First, we study the tax rate cuts in 1989-91 by comparing taxpayers above the exemption threshold (roughly the top 2% wealthiest) to taxpayers below the threshold. Second, among households located above the exemption threshold, some faced a zero marginal tax rate on wealth due to a ceiling on total tax liabilities (income plus wealth taxes) as a fraction of taxable income. Therefore, we also investigate the 1989-91 tax cuts by comparing taxpayers unaffected by the ceiling to taxpayers affected by it. This strategy allows us to identify behavioral responses among the very wealthiest. Third, we study the doubling of the exemption threshold for couples relative to singles implemented between 1989-1992. This reform eliminated wealth taxation for couples located roughly between the 98th and the 99.5th percentiles of the wealth distribution. We study responses to this experiment by comparing couples and singles in the affected range of the distribution.

Our findings suggest that the effect of wealth taxation on wealth accumulation is small, though not zero. We view our evidence as compelling in the sense that, in all three approaches just described, the pre-trends are completely parallel after making a non-parametric adjustment for pre-reform differences in portfolio composition.[2] We find the largest wealth effects when using the threshold change for couples, which is also the largest policy experiment. In this case, the wealth trajectories of treatments and controls clearly starts to diverge right after the reform. The DD estimate of the effect on annual wealth accumulation is about 3%. While this effect may seem large, the underlying tax variation driving it is also large. In fact, when interpreting our results in terms of the elasticity of savings with respect to the net-of-tax rate of return, the responses are very modest. This finding is relevant for calibrating models of optimal capital and labor taxation (see e.g., Saez & Stantcheva 2016).

Taken together, our results suggest that pre-tax wealth inequality was not strongly affected by the large cuts and eventual abolishment of wealth taxation in Denmark. The absence of strong effects on pre-tax wealth inequality implies that post-tax wealth inequality as well as consumption inequality increased as a result of the tax cuts. It is worth noting, however, that these quasi-experimental estimates capture only partial equilibrium effects of wealth taxation, not any general equilibrium effects.

Our paper contributes to a nascent literature on the behavioral responses to wealth taxes (Seim 2015; Brülhart *et al.* 2016; Zoutman 2015). Our main contributions compared to these studies is to document the evolution of wealth inequality based on comprehensive household-level wealth data, to consider a larger natural experiment with wealth taxes, to implement new estimation strategies that identifies behavioral responses at the very top of the distribution, and to follow taxpayers over many years to better capture the dynamic effects of wealth taxes.

The rest of the paper proceeds as follows. Section 3 describes the data, section 4 provides comprehensive evidence on the evolution of wealth concentration in Denmark since 1980 (and compares with recent estimates for the U.S.), section 5 presents our quasi-experimental study of household responses to wealth taxation, while section 6 concludes.

2 Data

We base our analysis on an administrative database maintained by the Danish Statistical Agency, DST. The database contains information on the entire Danish population between

[2]Without adjusting for pre-reform portfolio composition — for example differences in housing shares or equity shares across treatments and controls — there would be bias from non-tax related asset price movements.

1980-2014 in a range of registries. We are able to combine several different administrative registries using a unique personal identifier. These registries contain yearly information on wealth, earnings, occupation, family composition etc. The personal identifier allows us to link family members and track individuals and households over time. The data on income and wealth is not censored or top coded. Due to our focus on the top of the wealth distribution, this is a very important feature.

Wealth and income registries are based on tax return data from the Danish Tax Agency (SKAT). Wealth levels are stated by the end of the year. Most of the components measuring wealth are reported by third-parties such as banks, financial institutions, and government institutions, and is therefore considered to be very reliable. Until the abolishment of the wealth tax in 1997 end of year assets also include self-reported components. These are subject to auditing by the tax authorities. Accumulated pension funds are not part of the wealth data.

Assets includes third-party reported cash value of real estate, deposits, bonds, and shares. Value of stocks (listed and non-listed) are self-reported until 1997. Hereafter, the value of listed stocks is reported by banks, and other financial institutions. The asset value of real estate is based on valuations conducted by the tax authorities, which typically implies a valuation of real estate below the actual market value. Until the abolishment of the wealth tax in 1997 assets also include self-reported value of durables such as cars, boats and caravans and value of share certificates for housing cooperatives, premium bonds and cash-holdings. From 1997 the value of durables, cash holdings, non-listed stocks and shares outside deposits and private debt are not part of the wealth statement. From 1986-1996 taxable equity from self-employment is determined separately and added to the assets. Equity from self-employment is determined as taxable value (80 pct. in 1981, 70 pct. from 1983-1987, and 60 pct. in 1988-1996 of cash value or operational value) of buildings and operating equipment as well as inventory and financial assets net of debt. From 1997 equity from self-employment is no longer part of the wealth statement. However, the cash value value of buildings and financial assets from self-employment is included in the asset statement. From 1997, the asset statement is based solely on third-party reporting.

Liabilities include third-party reported values of debt in financial institutions, mortgage credit debt, credit and debit card debt, deposited mortgage debt, student debt and debt in the The Mortgage Bank (Public Institution), debt to financial corporations, debt to municipalities, and other liabilities such as unpaid taxes which are not deposited.

More details on Danish wealth data can be found in e.g. Boserup *et al.* (2014) Leth-Petersen (2010).

3 Wealth Inequality in Denmark, 1980-2014

3.1 Methods to Compute the Distribution of Household Wealth

We distribute 100% of the macroeconomic amount of household wealth at market value recorded in Denmark's household balance sheet, the stock part of Denmark's national accounts. Following internationally-agreed guidelines codified in the System of National Accounts, total household wealth includes all the non-financial and financial assets that belong to Danish resident households, minus debts. It includes in particular all funded pension wealth and excludes consumer durables and valuables. In 2012, Danish household wealth amounted to 379% of national income. Average household wealth per adult was US$ 242,000 (using market exchange rates to convert Danish kronor to dollars), a level comparable to that of the United States (US$ 234,000).

Our starting point to compute the wealth distribution is the detailed individual wealth records maintained by DST—the Danish statistical authority—since 1980. Although these data are generally of very high quality, they suffer from a number of limitations that we address as follows.

First, the DST data disregard funded pension assets, which are large in Denmark—43.5% of all Danish household wealth in 2012—and have grown a lot over time, see Online Appendix. We impute pension wealth as in Saez & Zucman (2016) by assuming that 40% belongs to wage-earners and 60% belongs to retirees. The pension wealth allocated to wage-earners is then allocated proportionally to wage income (winsorized at the 99th percentile), and the pension wealth allocated to retirees is allocated proportionally to the pension benefits paid out of pension funds. This imputation introduces a small noise at the micro-level which is irrelevant for our purposes of computing the distribution of wealth, as pensions are always a small fraction of wealth at the top where equities, bonds, and business assets dominate. Since 2012, estimates of individual pension wealth—provided by the main pension funds—are available in the DST data. As reported in the Online Appendix, the distribution of our imputed pension wealth is very close to that of (directly measured) pension assets in 2012.

Second, after the abolition of the wealth tax in 1997, the DST data do not capture forms of wealth that are not reported by third parties, most importantly unlisted corporate equities and non-corporate business assets. We impute unlisted equities by capitalizing dividends. The capitalization rate is equal to the market value of unlisted firms (as recorded in the national accounts) divided by the flow of dividends reported on tax returns. We similarly impute non-corporate

business assets by capitalizing business income in the same way. After these corrections we capture almost 100% of the total household wealth reported in the national accounts. We fill in the residual gap by multiplicatively scaling up registered assets, separately for each asset class—listed equities, fixed-income claims, housing assets, and debts. Importantly, since we include all forms of assets and liabilities at market value in our estimates, our distributional figures cover 100% of the aggregate wealth of Danish households. This makes our estimates comparable to the U.S. wealth inequality statistics of Saez & Zucman (2016) who similarly cover 100% of the aggregate wealth of U.S. households, including funded pension assets and unlisted businesses.[3]

For the purpose of computing the wealth distribution, our unit of analysis is the adult individual (aged 20 or above) with wealth equally split among married spouses. We disregard children below 20 who own negligible wealth. Dividing the wealth of married couples equally addresses the problem that jointly-owned assets are allocated to only one individual in the registry, even when they are owned in part by each spouse.[4]

3.2 The Evolution of Wealth Inequality Since 1980s: Denmark vs. the United States

Figure 1 shows the level of wealth inequality in Denmark and its evolution since 1980. There are three main findings. First, like in all other countries, Danish wealth is very concentrated, much more than labor income. The bottom 50% of the distribution owns a very small share of wealth, around 5% of the total. On aggregate, the assets of the bottom 50% are barely higher than its debts. The top 10% owns around half of total wealth. The middle 40%—the group of adults between the median and the 90th percentile that can be described as the middle class—owns about 45% of total wealth. Second, wealth inequality increased significantly in Denmark in the second half of the 1980s. The share of wealth owned by the top 10% increased from 42% in 1985 to 52% in 1990 (+ 10 percentage points), while the bottom 50% share fell by 7 percentage points and the middle 40% share by 3 percentage points. This evolution was driven by the dynamic of asset prices, in particular housing prices. Between 1985 and 1990, the nominal housing price index fell 13%, following an increase of 62% from 1980 to 1985. Because the share of housing in net wealth falls with wealth, housing slumps hurt the bottom 50% and the middle-class more

[3]One caveat is that we do not attempt to account for the unreported offshore holdings of Danish households, such as the bank deposits and portfolio of equities, bonds, and mutual fund shares held in Switzerland and similar tax havens. We refer the reader to Alstadsaeter *et al.* (2017) for an attempt at including offshore assets in the wealth distribution of Denmark and other Scandinavian countries.

[4]As reported in the Online Appendix, inequality among households is higher than inequality among equal-split adults. The difference, however, is negligible at the top of the distribution. When comparing our series to the United States, we use US series where wealth is also split equally between spouses.

than the top 10%, which leads to a rise in the top 10% share. Third, wealth inequality has remained roughly stable since the early 1990s. Although there have been cyclical short-run variation—again driven by asset prices—wealth inequality has not displayed any marked trend since 1990. If anything it has slightly decreased since then, despite the decline and eventual abolition of wealth taxation. The top 10% wealth share has been around 50% in recent years, a few percentage points below its level of the late 1980s.

Inequality is lower in Denmark than in the United States. As shown by Figure 2, the top 1% owns close to 40% of total wealth in the United States today, much more than in Denmark (about 20%). The gap between both countries has increased over time. Denmark was not very far from the United States in the 1980s and 1990s. In the early 1990s, the top 1% wealth share was close to 20% in Denmark (just like today), but the U.S. top 1% wealth share was significantly lower than it is today—around 27%. The gap between Denmark and the United States increased sharply in the 2000s and early 2010s: while the top 1% wealth share has remained stable in Denmark, it has increased 10 percentage points in the United States from 1995 to 2014.

The higher one moves up the wealth distribution, the more Denmark looked like the United States in the 1980s and 1990s. As shown by the top panel of Figure 3, the top 0.1% wealth share was only 2 to 3 points higher in the United States than in Denmark until the late 1990s. The top 0.01% wealth share were similar in both countries in the late 1980s and late 1990s (bottom panel of Figure 3). Wealth concentration followed the same rising trend in Denmark and the United States throughout the 1980s and 1990s. It rose in the second half of the 1980s, stabilized in the early 1990s, and rose again in the second half of the 1990s. The parallel evolution stopped in the late 1990s. Top shares fell a lot in Denmark during stock market crashes (in the early 2000s and then during the Great Recession), in contrast to the United States. As a result, the top 0.01% wealth share is twice higher in the United States than in Denmark today.

4 The Effects of Wealth Taxation

4.1 Tax Variation and Empirical Strategies

Denmark taxed wealth until 1996. Taxable wealth equalled the total net wealth of households, excluding pension wealth. Taxable wealth components thus included deposits, bonds, equities, housing, large durables (such as cars and boats) and business assets, net of any debts. A number of these components were third-party reported by financial institutions, leaving little scope for tax evasion. But some components were self-reported, namely unlisted equities, non-corporate

business assets, and durables.

The wealth tax had a (high) exemption threshold, making the tax progressive. During the time period we study the exemption threshold was located in between percentiles 96.5% and 99% of the household wealth distribution. Wealth falling below the threshold was not taxed, whereas wealth falling above the threshold was taxed at a flat rate. The tax rate was equal to 2.2% until 1988, but was substantially reduced and ultimately eliminated in two reforms. Between 1989-91 the tax rate was reduced from 2.2% to 1%, while in 1996-97 the wealth tax was abolished entirely. These tax rate changes are illustrated in Panel A of Figure 4.

In the analysis below we consider three different quasi-experimental strategies to estimate the effect of wealth taxation on wealth accumulation. The first strategy exploits the tax rate cuts in 1989-91, comparing taxpayers above the exemption threshold (treatments) to taxpayers below the threshold (controls). This difference-in-differences design relies on parallel trends in wealth accumulation between households above and below the threshold. We start by doing a repeated cross-section analysis comparing different percentiles of the wealth distribution over time, and then turn to a balanced panel analysis in which we follow the same individuals over time. In the panel analysis, households are assigned to the treatment or control group based on their pre-reform wealth levels. The reason we focus on the 1989-91 tax cuts (rather than the subsequent elimination of the tax) is a data limitation discussed earlier: Since abolishing the wealth tax in 1997, Statistics Denmark no longer records purely self-reported wealth. This data break makes the wealth tax abolishment difficult to study, and so we focus on the earlier tax cuts that do not have this limitation.[5]

The second strategy exploits the presence of a ceiling on total tax liability from all personal taxes (income taxes, social security taxes, and wealth taxes) as a fraction of taxable income. This ceiling – called *Det Vandrette Skatteloft* ("horizontal tax ceiling") – was in place to limit the total tax rate on households with large wealth relative to labor income ("rentiers"). According to the ceiling rule, if the sum of income taxes, social security taxes and wealth taxes exceed a certain fraction of taxable income, then tax liability is reduced by the excess amount. Wealth tax liability is reduced first (but by no more than 50%) and the income tax liability is reduced next (at the most to zero). The level of the tax ceiling changed over time and was equal to 78% just prior to the 1989 reform. Hence, the ceiling was binding for households whose total taxes were above 78% of income, but not by so much that meeting the ceiling would reduce wealth

[5]We will provide suggestive evidence on the wealth tax elimination using the market-value wealth series constructed in the first part of the paper. This series is consistently measured over time, but has the limitation that it does not exactly measure taxable wealth at the individual level.

tax liability by 50% and income tax liability by 100%. For households whose ceiling is binding, the marginal wealth tax rate is effectively zero. This motivates our second strategy in which we investigate responses to the 1989-91 cuts by comparing taxpayers unaffected by the ceiling (treatments) to taxpayers affected by it (controls). This difference-in-differences design relies on parallel trends between households with different wealth/income ratios among those who are above the exemption threshold.

Our third and final strategy exploits a large increase in the exemption threshold for couples relative to singles. Until 1988, singles and couples faced the same nominal exemption threshold for paying wealth taxes. This is difficult to justify on equity grounds, because a couple is less wealthy in per capita terms for a given level of household wealth. To rectify this issue, the exemption threshold for couples relative to singles was doubled between 1989-1992. These threshold changes are illustrated in Panel B of Figure 4. The implication of the reform is that couples in a certain range of the household wealth distribution became exempt from wealth taxation, whereas singles in the same range did not. Our differences-in-differences design here relies on parallel trends between couples and singles at the household wealth levels affected by the threshold move.

To summarize, we consider difference-in-differences strategies that rely on three different types of comparisons: the very wealthy vs the less wealthy (above vs below exemption threshold), the working rich vs rentiers (below vs above tax ceiling), wealthy couples vs wealth singles (affected vs unaffected by the threshold adjustment). As we shall see, the combination of all three provides a compelling picture of behavioral responses to wealth taxation.

Finally, it is worth noting that in principle there is a fourth quasi-experimental strategy: A bunching approach based on the budget set kink at the exemption threshold.[6] Bunching approaches in the wealth context are useful for uncovering evasion/avoidance responses, but not for uncovering real responses. While this is largely true even in the context of labor income responses (see Kleven 2016), it is particularly true for wealth responses. Taxable wealth depends not only on individual decisions, but also on asset price movements that occur continuously throughout the tax year. Given asset price uncertainty, it would be virtually impossible for a taxpayer to bunch at the exemption threshold through, say, real savings responses. While we find modest bunching at the kink (consistent with a small reporting response, probably coming from self-reported wealth), we focus on the difference-in-differences strategies that are better suited to uncover real responses.

[6] As mentioned earlier, a recent paper by Seim (2015) takes such an approach for Sweden.

4.2 Repeated Cross-Section Analysis

Before turning to our quasi-experimental analysis using panel data, this section presents simple and suggestive evidence based on repeated cross-sections. We compare the evolution of wealth in quantiles of the wealth distribution that are affected differently by wealth tax reforms. Specifically, we compare households in the top 0.5% of the wealth distribution (above the exemption threshold in all years) to households between percentiles 95 and 96.5 (below the exemption threshold in all years). As described above, the top 0.5% of the wealth distribution benefitted from large tax cuts through the 1989-91 reform and the 1996-97 wealth tax abolishment.

For this analysis we use the measure of market-value wealth constructed in the first part of the paper. Market-value wealth is fully comprehensive and includes assets (primarily pension assets) that are tax exempt. The key advantage of using this wealth concept is that it is consistently measured over time, thus allowing us to study the effects of both the 1989-91 and 1996-97 reforms. Taxable wealth, by contrast, is not precisely measured at the household level after the wealth tax abolishment in 1997, making it hard to study taxable wealth responses to this reform. In the panel analysis below, we use taxable wealth as our outcome variable and focus on the earlier reform where this outcome is measured both before and after.

The results are presented in Figure 5. It includes two panels each of which shows the evolution of market-value wealth in the treatment group (top 0.5%) and in the control group (P95-P96.5) between 1982-2012. The top panel focuses on total wealth, while the bottom panel excludes housing assets. The motivation for excluding housing assets is that housing is less important in the portfolios of the very wealthy than in the portfolios of the less wealthy. This implies that house price volatility (unrelated to wealth taxation) may act as a confounder when comparing the total wealth series of the two groups. We return to this issue in the panel analysis below.

The following insights are worth highlighting. First, when considering the 30-year period as a whole, there is essentially no difference in the growth of wealth between the two groups. Despite benefitting from large wealth tax cuts over the period, the top 0.5% do not experience significantly larger wealth growth. In fact, when leaving out housing, the two groups feature almost exactly the same wealth growth from 1982 to 2012. Second, the two series track each other very closely around the 1989 reform, especially when we exclude housing. This suggests little or no effect of the 1989 tax cuts. Third, wealth is strongly volatile in the top 0.5% from the late 1990s, making it harder to gauge the effects of the 1997 wealth tax abolishment. This volatility is driven by equity price swings around the dot-com bubble and the global financial

crisis. Given these events were clearly unrelated to the Danish wealth tax, the graph does not point to any clear wealth accumulation effects of eliminating the wealth tax.

4.3 Panel Analysis

4.3.1 Difference-in-Differences Using Rate Changes Above and Below Threshold

In this section we turn to panel analyses of wealth responses to wealth taxes. We first consider responses to the 1989-91 tax cuts by comparing taxpayers above and below the exemption threshold. We define treatments as those who are consistently above the threshold in the three pre-reform years 1986-1988, while controls are those who are consistently below the threshold in those three years. Assigning treatment status based on three years instead of just one year ensures that the groups are relatively stable and thus avoids bias from mean reversion. The analysis is based on a balanced panel of taxpayers who we observe with non-zero wealth in each year between 1982-96, and who are located in the top 5% of the wealth distribution in 1986-1988. Given the location of the exemption threshold during the pre-reform years (as shown in Figure 4, Panel B), we will be comparing households who are roughly above the 97.5th percentile to households who are roughly in the 95-97.5th percentiles.

The results are shown in Figure 6. Panel A shows three wealth series: log wealth in the treatment group (red series with dots), log wealth in the control group (grey series with squares), log wealth in the control group normalized to the level of the treatment group in 1988 (black series with squares). The following findings emerge from this panel. First, the treatment and control series evolve in a similar fashion overall: there is strong growth in wealth until 1986, then a sharp kink and relatively little growth thereafter. Second, pre-trends are not completely parallel between treatments and controls. The wealth of households above the exemption threshold grows faster than the wealth of households below the threshold. Third, this graph does not show any clear effects of the 1989-91 tax cuts. Households above the threshold experience faster wealth growth both before and after the reform, and the rate of divergence does not seem to increase after the reform. That is, the wealth growth advantage of the treated is not increasing following the tax cuts.

The raw difference-in-differences in Panel A is hampered by the absence of parallel trends. A potential reason for non-parallel trends when comparing households at different levels of wealth is heterogeneity in portfolio composition. With differences in asset portfolios, aggregate movements in asset prices (unrelated to taxation) will create differential trends. Specifically, a key difference between the very wealthy and the moderately wealthy lies in the relative importance

of housing assets. At the very top, housing is much less important than further down. This implies that the (large) house price cycles during the 1980s and 1990s may be confounding our comparison.

To deal with this issue, Panel B presents results from a specification in which we control non-parametrically for differential trends related to housing. For the treatment and control groups separately, we regress log wealth on a full set of year dummies along with a full set of interactions between year dummies and deciles of the 1988 share of housing in gross assets. The figure plots the estimated coefficients on the year dummies. Without the year-by-housing share dummies, the year coefficients would correspond to the raw averages depicted in Panel A. Including those interactions absorbs any differential wealth trend stemming from the fact that the two groups had different housing shares before the reform.

The results in Panel B are quite striking. Controlling for trends related to different pre-reform housing shares is sufficient to make pre-trends almost perfectly parallel between the two groups. The two groups track each other very closely during the pre-reform years, featuring the same growth leading up to 1986 and the same reversal of growth between 1986-1988. There is no short-run effect of the wealth tax cuts: The two series continue to track each other until 1991. After this time the two series diverge slightly – with faster growth among the treated – consistent with a medium-run response to the reform. However, this response is very small and only marginally significant.

4.3.2 Difference-in-Differences Using Tax Ceiling

We now investigate responses to the 1989-91 tax cuts by exploiting the tax ceiling rule described above. We will compare compare taxpayers for whom the tax ceiling is non-binding (treatments) to taxpayers for whom the tax ceiling is binding (controls). The treatments experience a reduction in their marginal wealth tax rate from 2.2% to 1%, while the controls experience either no change in their marginal tax rate or an increase in their marginal tax rate (from 0 to 1%). The latter scenario occurs for taxpayers whose ceiling is binding under the initial high wealth tax, but not under the reduced wealth tax. As before, we define treatment status using three pre-reform years 1986-1988, and consider a balanced panel of taxpayers who we observe with non-zero wealth in each year between 1982-96. We restrict the sample to the top 1% of the wealth distribution in 1986-1988.

Figure A.1 in appendix shows the fraction of households for whom the ceiling is binding (series in white dots) and the fraction for whom it is non-binding (series in black dots) by

quantile of the wealth distribution.[7] The graph implies that this strategy largely compares households at the bottom half of the top 1% to households at the top half of the top 1%. Hence, we are not just considering a different identification strategy to the one in the previous section, we are also estimating responses for different taxpayers: the very wealthiest in this section (top 1%) as opposed to more moderately wealthy in the previous section (top 2.5%).

The results are shown in Figure 7, which is constructed in the same way as Figure 6. Panel A shows that both the treatment and control series feature strong wealth growth until 1986 and much weaker growth thereafter (similar to the first analysis). The pre-trends are not parallel: the households with a non-binding ceiling (the treated) show weaker wealth growth in the period before the reform. The fact that post-trends do become parallel could be taken as evidence that the reform had an effect. That is, the wealth tax cut lead the treated taxpayers to increase wealth accumulation and thereby catch up with the wealth growth of untreated taxpayers. However, when inspecting the patterns closely, this convergence in wealth growth appears to start before the reform, which goes against such an interpretation.

To provide more conclusive evidence, we need to balance the groups in terms of their pre-trends. For the threshold DD, we did this by adjusting for differences in pre-reform housing shares. While heterogeneity in housing shares is important when comparing the very wealthy to the moderately wealthy as we did in the previous section, it is much less important for the comparison done here. This analysis is based on comparing two groups of very wealthy taxpayers who differ by their wealth/income ratios, which determine whether the tax ceiling is binding. A larger difference between these groups lies in the importance of equity in their portfolios, implying that fluctuations in stock prices are likely to confound our results.

Hence, in Panel B we control non-parametrically for time trends due to pre-reform differences in equity shares. For each group separately, we regress log wealth on year dummies and on interactions between year dummies and deciles of the 1988 equity share. The figure plots the estimated coefficients on the year dummies. The results are now very clear: the pre-trends are exactly the same in the two groups, and there is essentially no divergence happening after the reform. Hence, this strategy suggests a zero effect of wealth taxation on the very wealthiest taxpayers.

[7] The two fractions do not quite sum to 100%, because we are currently unable to precisely determine ceiling status for a small subset of households.

4.3.3 Difference-in-Differences Using Threshold Change for Couples

Our last strategy exploits the doubling of the exemption threshold for couples relative to singles between 1989-1992. Figure 8 investigates the impact of this reform by comparing couples located between the exemption threshold and 2 x the exemption threshold in 1988 (treatments) to singles located in that same range in 1988 (controls). The graph is otherwise constructed in the same way as the previous ones: Panel A shows the raw time series for treatments and controls, while Panel B shows non-parametrically adjusted time series that eliminate differential pre-trends. When comparing singles and couples, it becomes important once again to control for differences in the pre-reform share of housing. Therefore, as we did for the threshold DD above, Panel B absorbs the effect of differential trends across deciles of the 1988 share of housing in gross assets.

This approach yields our strongest evidence of wealth responses to wealth taxation. In Panel B, the two groups feature the exact same trends in the five years leading up to the reform, but begin to diverge immediately after the reform. The gap between the two groups builds up over time as annual wealth accumulation is consistently higher among couples who are no longer liable to pay wealth taxes. The DD estimate of the wealth effect is about 20% seven years after the reform, or just under 3% per year. While this effect may seem large, notice that the tax variation driving it is also very large. This is because the tax falls on the stock of wealth rather than only the return to wealth: Assuming a rate of return of e.g. 4.4%, eliminating a wealth tax of 2.2% corresponds to eliminating a 50% tax on the flow of income from wealth. Hence the elasticity of savings with respect to the net-of-tax return is actually quite modest.

To conclude, we have presented three difference-in-differences approaches to estimating behavioral responses to wealth taxes. These approaches rely on different sources of variation and capture responses for different groups of taxpayers. The evidence is compelling in the sense that, in all three approaches, pre-trends are completely parallel after making non-parametric adjustments for pre-reform differences in portfolio composition (that would otherwise create bias from non-tax related asset price movements). The overall finding is that the effect of wealth taxes on wealth accumulation is small. We find the largest wealth effects when using the threshold change for couples, which is also the largest experiment (eliminating the wealth tax on the treated as opposed to cutting it to 1%). Still, as just described, even there the response is modest when interpreted as an elasticity with respect to the net-of-tax return on wealth.

References

Alstadsaeter, Annette, Johannesen, Niels, & Zucman, Gabriel. 2017. Tax Evasion and Inequality. NBER Working Paper.

Boserup, Simon Halphen, Kopczuk, Wojciech, & Kreiner, Claus Thustrup. 2014. Stability and persistence of intergenerational wealth formation: Evidence from Danish wealth records of three generations. Working Paper.

Brülhart, Marius, Gruber, Jonathan, Krapf, Matthias, & Schmidheiny, Kurt. 2016. The Elasticity of Taxable Wealth: Evidence from Switzerland. NBER Working Paper No. 22376.

Kleven, Henrik. 2016. Bunching. *Annual Review of Economics*, **8**, 435–464.

Leth-Petersen, Søren. 2010. Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit? *American Economic Review*, **100**(3), 1080–1103.

OECD. 1988. *Taxation of Net Wealth, Capital Transfers and Capital Gains of Individuals*. Tech. rept. Organization for Economic Cooperation and Development, Paris.

Piketty, Thomas. 2014. *Capital in the Twenty First Century*. Harvard University Press.

Saez, Emmanuel, & Stantcheva, Stefanie. 2016. A simpler Theory of Capital Taxation. NBER Working Paper No. 22664.

Saez, Emmanuel, & Zucman, Gabriel. 2016. Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data. *The Quarterly Journal Of Economics*, **131**.

Seim, David. 2015. Behavioral Responses to Wealth Taxes: Evidence from Sweden. *American Economic Journal: Economic Policy*, forthcomming.

Zoutman, Floris T. 2015. The Effect of Capital Taxation on Household Savings. Working Paper.

Figure 1:

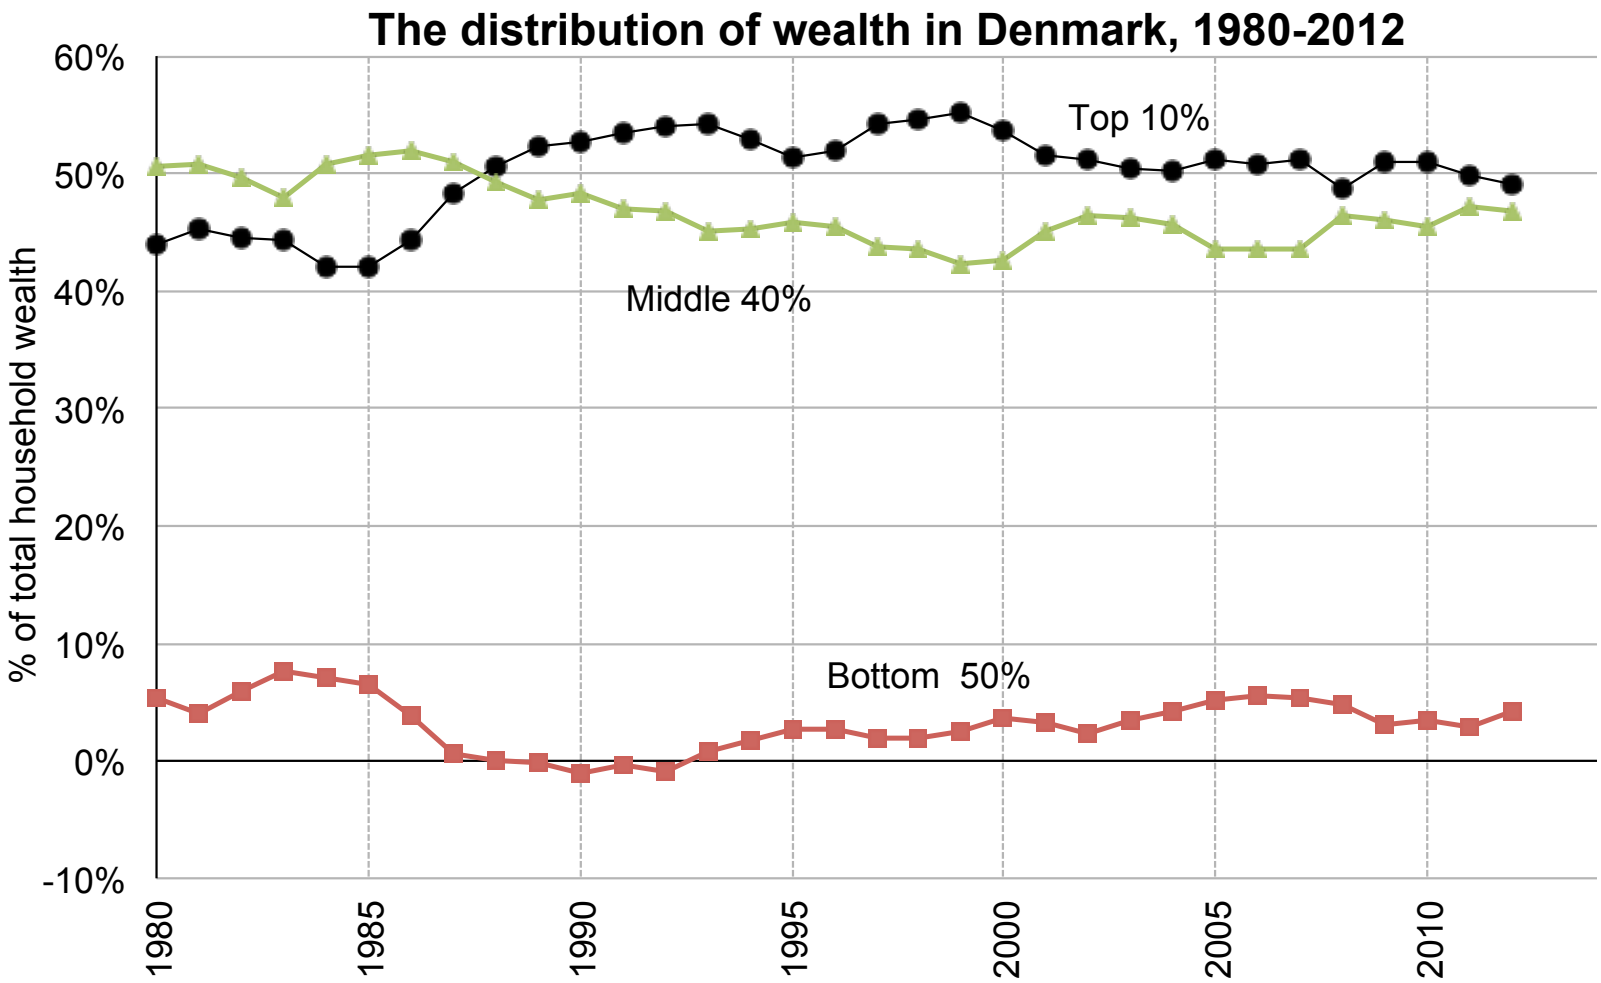

Notes: This figure shows the share of total household Danish net wealth owned by the bottom 50% of the distribution, the middle 40% (adults between the median and the 90th percentile), and the top 10%. The unit of observation is the adult individual (aged 20 or above). The wealth of married couples is split equally among spouses. Wealth includes all the non-financial and financial assets of Danish households net of debts, including funded pension assets and unlisted businesses. It matches the total amount of household wealth recorded in Denmark's household balance sheet. Source: Online Appendix.

Figure 2:

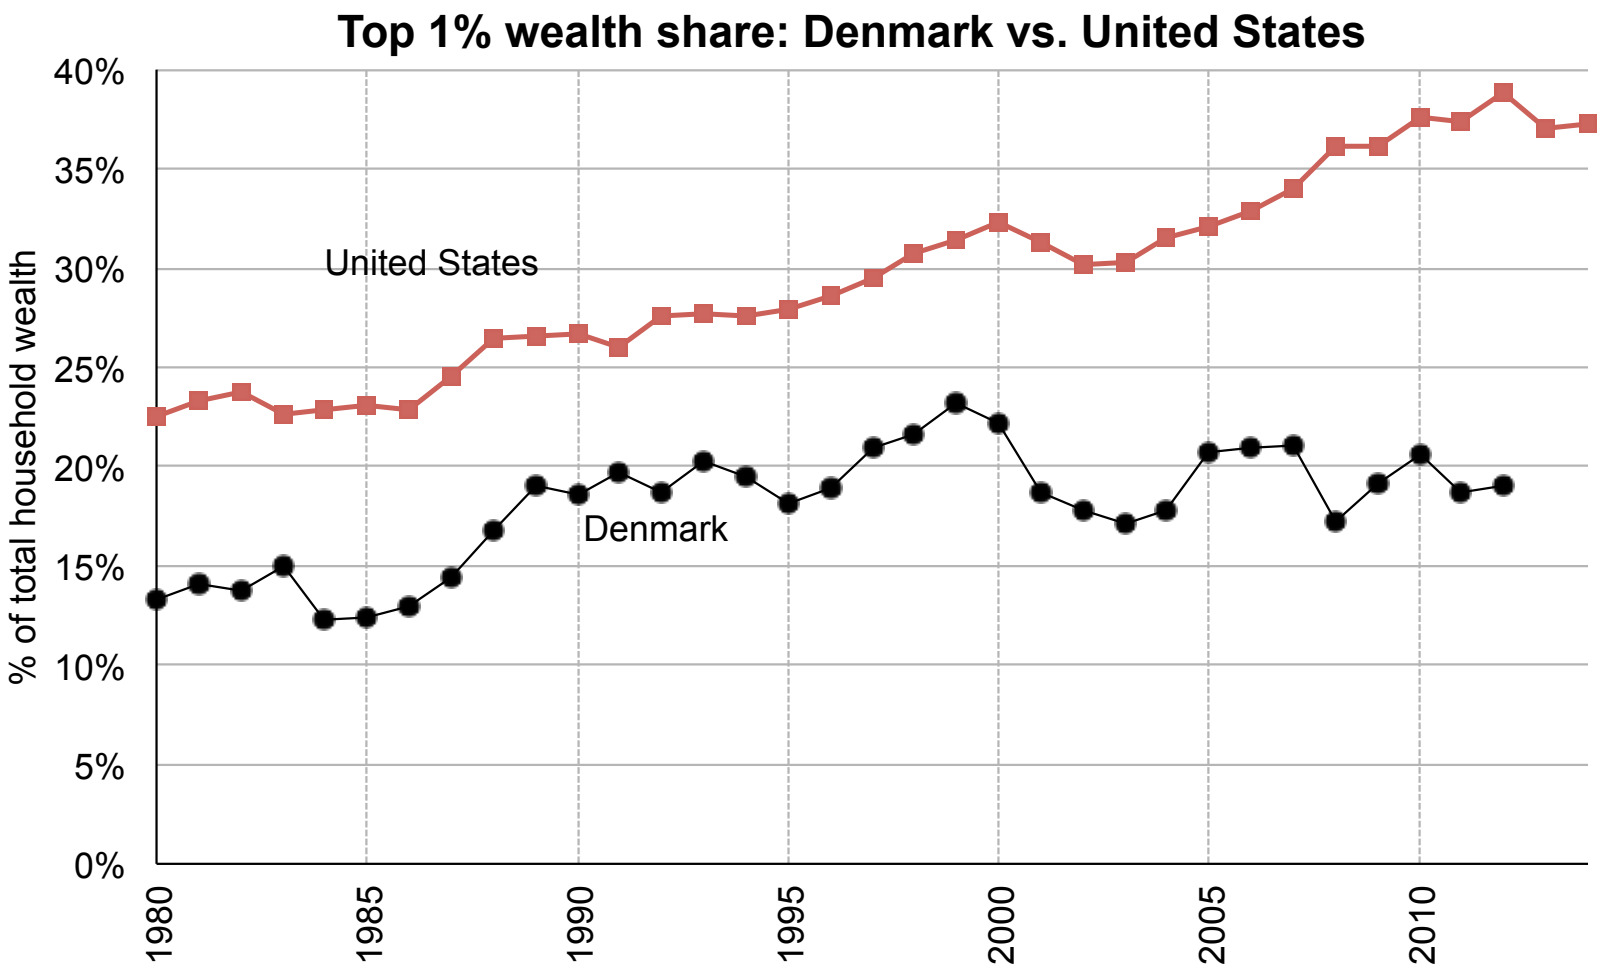

Notes: This figure shows the share of total household owned by the top 1% in Denmark and in the United States. In both countries, the unit of observation is the adult individual (aged 20 or above), and the wealth of married couples is split equally among spouses. Wealth includes all the non-financial and financial assets of households net of debts, including funded pension assets and unlisted businesses. It matches the total amount of household wealth recorded in Denmark's and the United States' household balance sheets. Source: Online Appendix, and Saez & Zucman (2016).

Figure 3: The 0.1% and Top 0.01% Wealth Shares in Denmark vs the United States

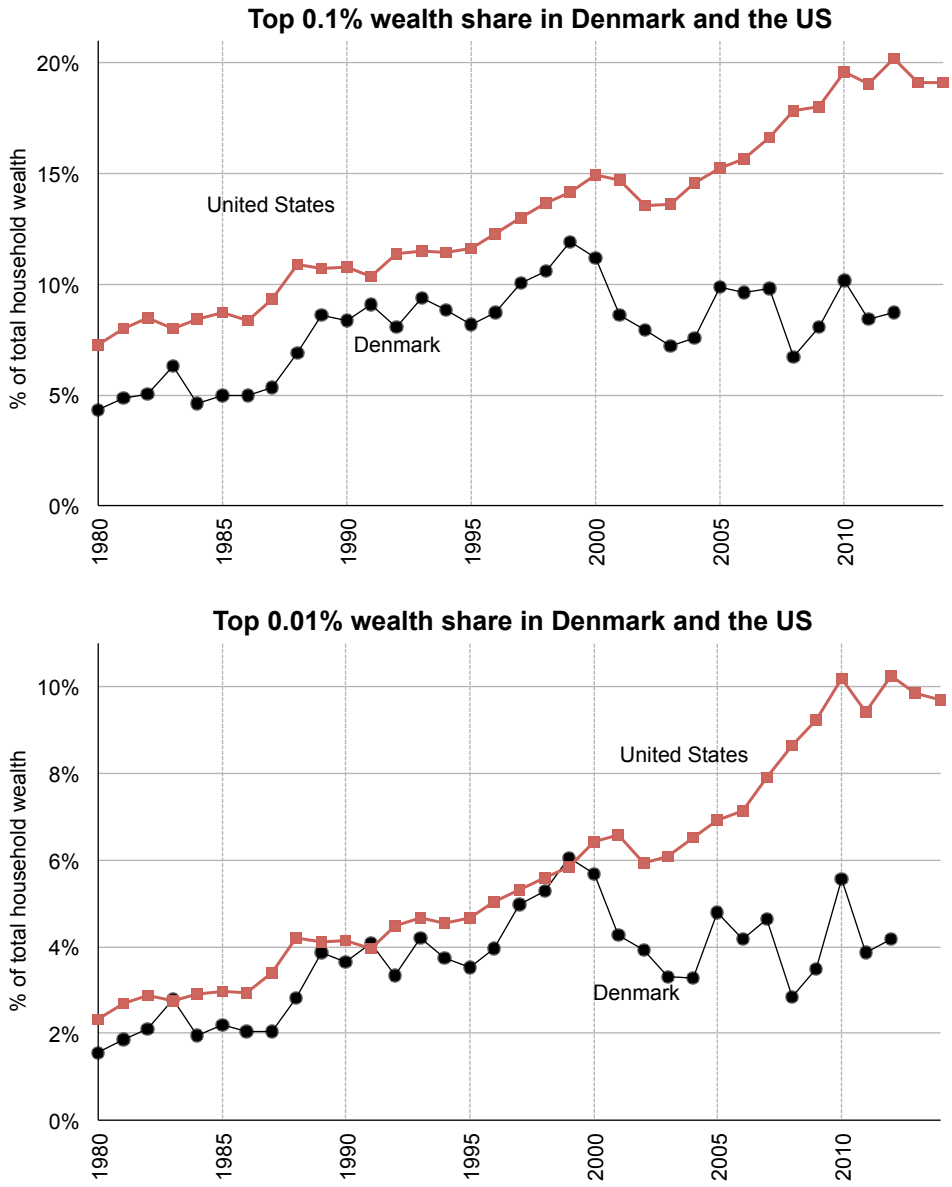

Notes: This figure shows the share of total household owned by the top 0.1% (top panel) and top 0.01% (bottom panel) in Denmark and in the United States. In both countries, the unit of observation is the adult individual (aged 20 or above), and the wealth of married couples is split equally among spouses. Wealth includes all the non-financial and financial assets of households net of debts, including funded pension assets and unlisted businesses. It matches the total amount of household wealth recorded in Denmark's and the United States' household balance sheets. Source: Online Appendix, and Saez & Zucman (2016).

Figure 4: Marginal Tax Rate and Exemption Thresholds in the Danish Wealth Tax

Panel A: Marginal Wealth Tax Rate Over Time

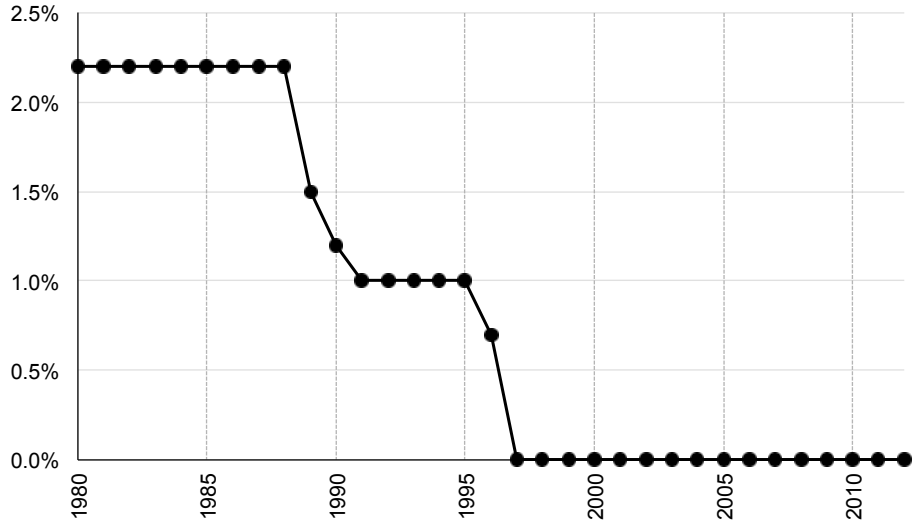

Panel B: Exemption Thresholds Over Time

Notes: This figure shows the marginal tax rate (top panel) and the exemption level of the Danish wealth tax (bottom panel). Source: authors' computations.

Figure 5: Market-Value Wealth of Households in Top 0.5% vs P95-P96.5

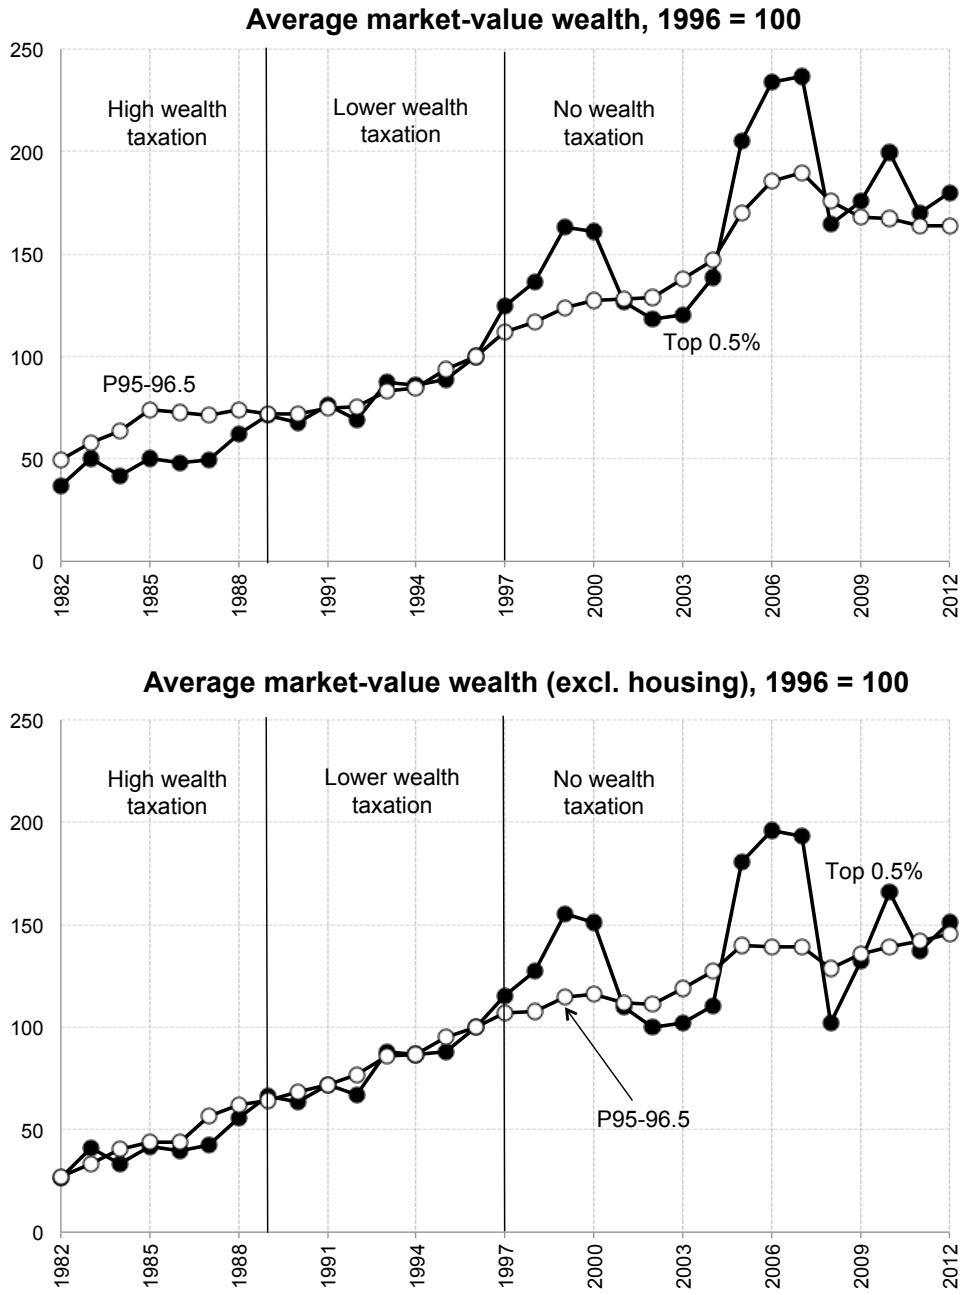

Notes: This figure shows the average wealth at market value (including tax-exempt assets) of households in the top 0.5% and between the 95th percentile (P95) and the 96.5th percentile (P96.5). Households are ranked by market-value wealth. Most households in the top 0.5% were taxable up to 1997, while most household in P95-P96.5 were below the exemption threshold. The top panel includes all wealth, while the bottom panel excludes housing wealth. The wealth of each group is normalized to 100 in 1996, and is deflated by the national income price index. Source: Online Appendix.

Figure 6: Difference-in-Differences Comparing Households Above and Below Threshold

Panel A: Raw Time Series of Treatments and Controls

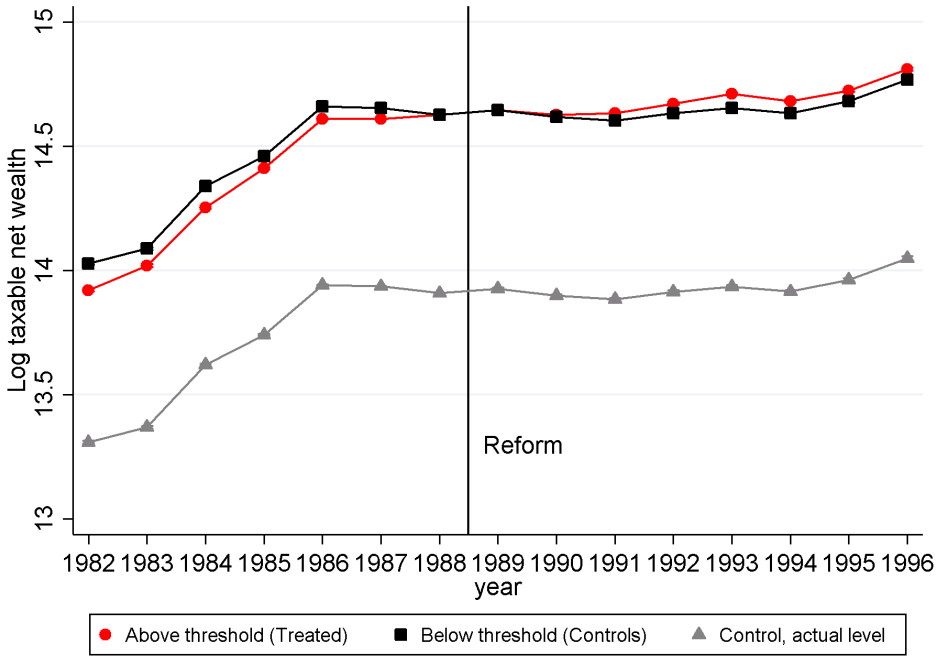

Panel B: Non-Parametric Controls for Pre-Reform Housing Shares

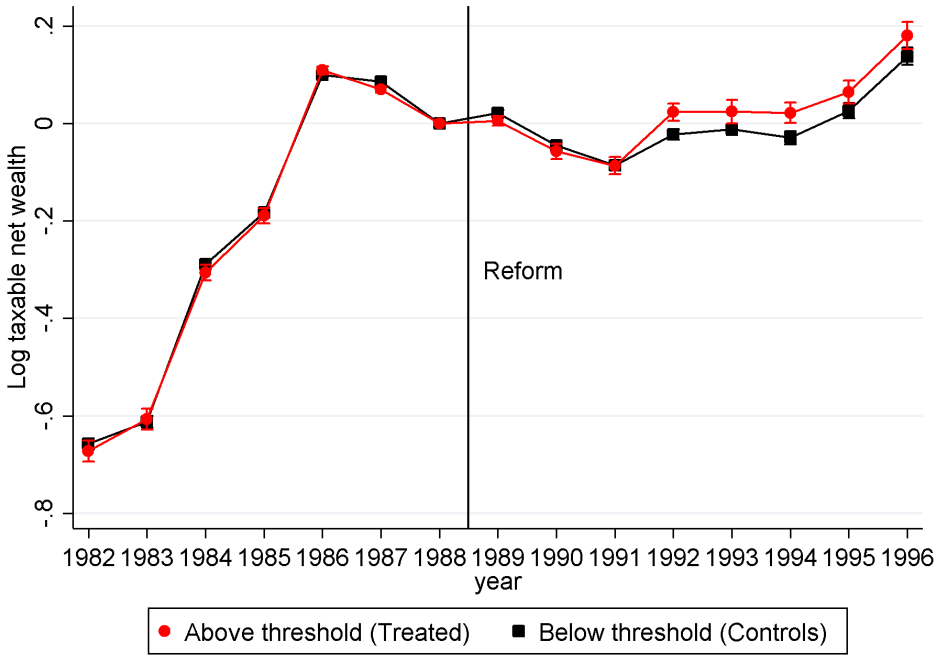

Notes: See text for details on the quasi-experiment and empirical specification. Source: authors' computations using DST data.

Figure 7: Difference-in-Differences Comparing Households Bound and Unbound By Tax Ceiling

Panel A: Raw Time Series of Treatments and Controls

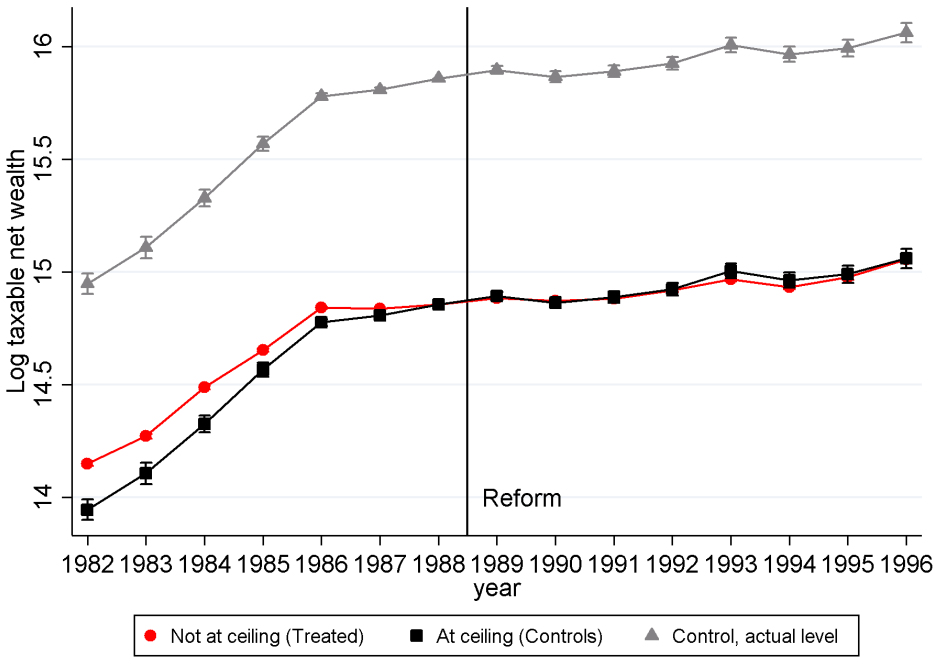

Panel B: Non-Parametric Controls for Pre-Reform Equity Shares

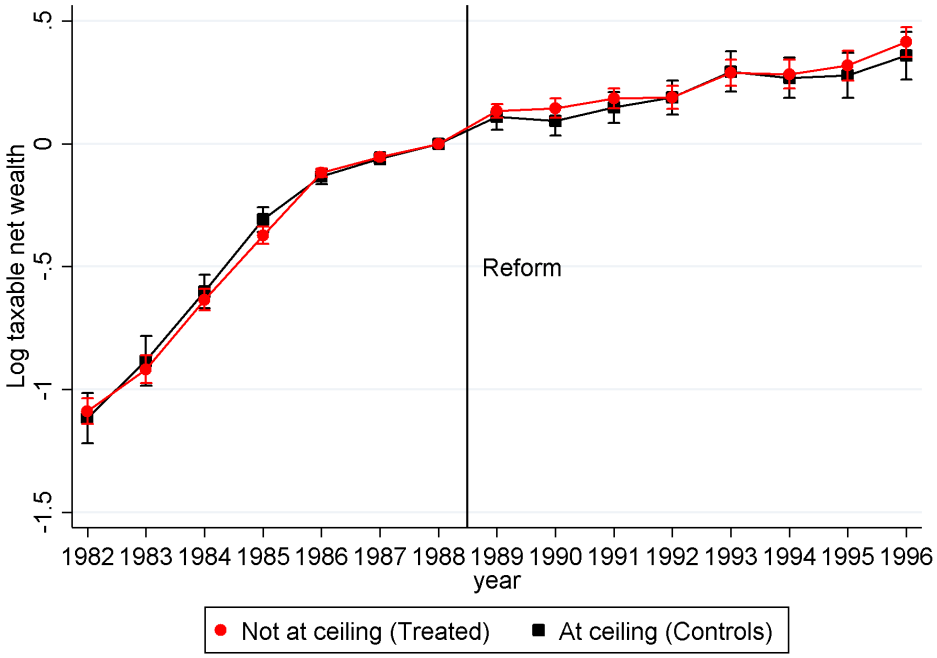

Notes: See text for details on the quasi-experiment and empirical specification. Source: authors' computations using DST data.

Figure 8: Difference-in-Differences Comparing Couples and Singles Between Threshold and 2 × Threshold

Panel A: Raw Time Series of Treatments and Controls

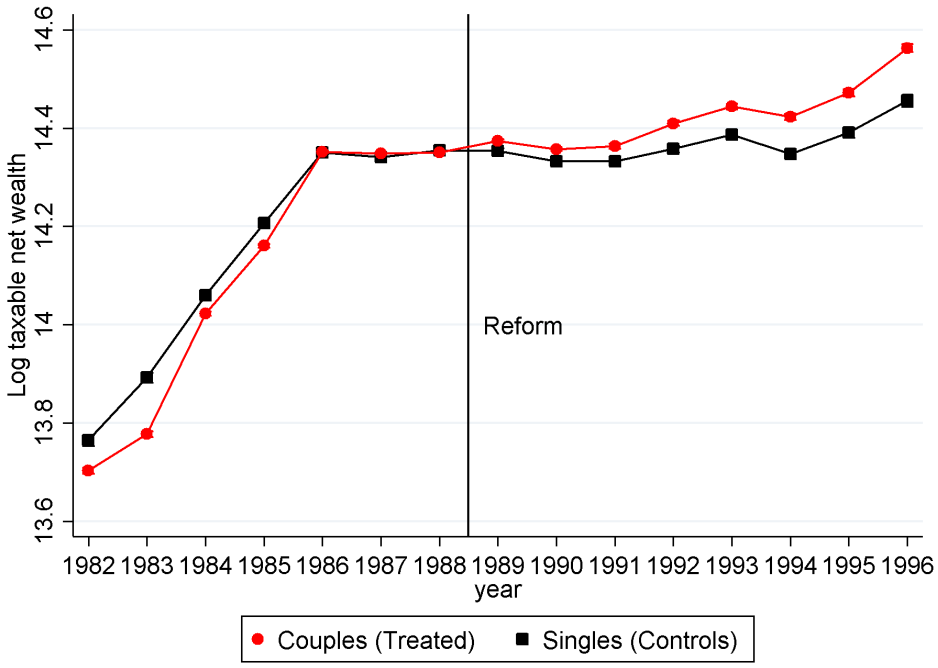

Panel B: Non-Parametric Controls for Pre-Reform Housing Shares

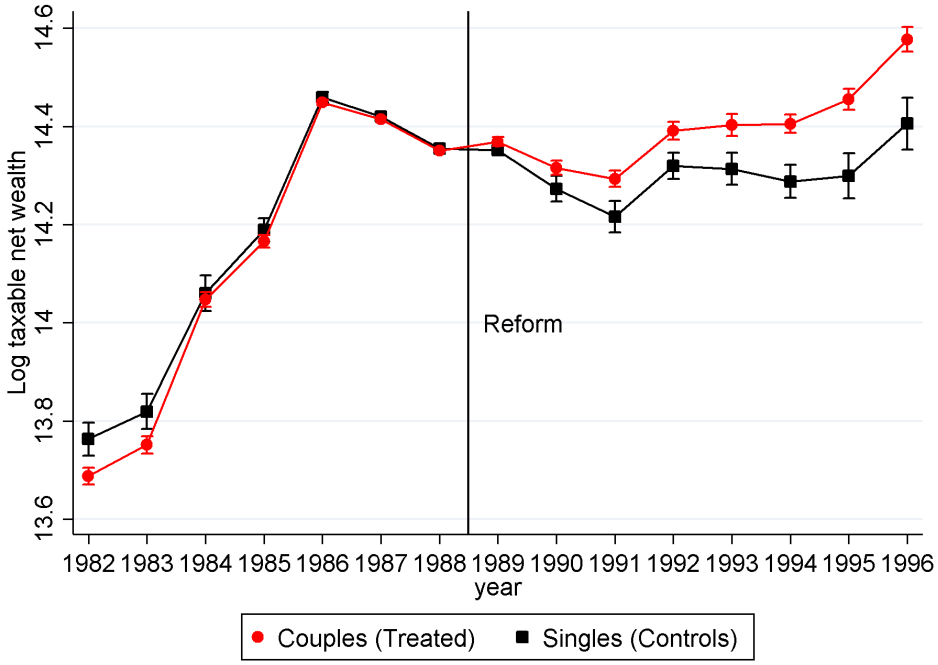

Notes: See text for details on the quasi-experiment and empirical specification. Source: authors' computations using DST data.

Figure A.1: Fraction of Households Affected by Tax Ceiling

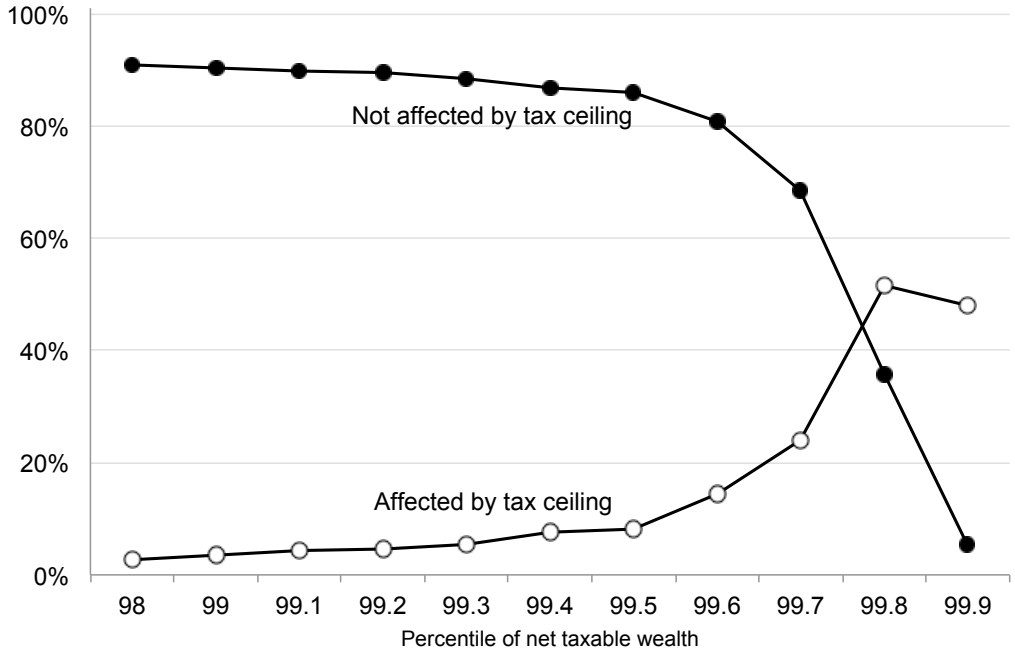

Notes: This figure shows the fraction of taxpayers affected by the tax ceiling in 1988, by percentile of net taxable wealth. Source: Online Appendix.