

WEALTH TAXATION AND WEALTH ACCUMULATION: THEORY AND EVIDENCE FROM DENMARK*

KATRINE JAKOBSEN
KRISTIAN JAKOBSEN
HENRIK KLEVEN
GABRIEL ZUCMAN

Using administrative wealth records from Denmark, we study the effects of wealth taxes on wealth accumulation. Denmark used to impose one of the world's highest marginal tax rates on wealth, but this tax was greatly reduced starting in 1989 and later abolished. Due to the specific design of the wealth tax, the 1989 reform provides a compelling quasi-experiment for understanding behavioral responses among the wealthiest segments of the population. We find clear reduced-form effects of wealth taxes in the short and medium run, with larger effects on the very wealthy than on the moderately wealthy. We develop a simple life cycle model with utility of residual wealth (bequests) allowing us to interpret the evidence in terms of structural primitives. We calibrate the model to the quasi-experimental moments and simulate the model forward to estimate the long-run effect of wealth taxes on wealth accumulation. Our simulations show that the long-run elasticity of taxable wealth with respect to the net-of-tax return is sizable at the top of the distribution. *JEL* Codes: H20, H31, E21, D31.

I. INTRODUCTION

What are the economic effects of taxing household wealth? Although an enormous literature estimates the elasticity of labor supply and taxable income, much 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 rising wealth inequality in the United States (Saez and Zucman 2016).

*We thank Raj Chetty, John Friedman, Lawrence Katz, Wojciech Kopczuk, Petra Persson, Emmanuel Saez, Kurt Schmidheiny, Michael Smart, Stefanie Stantcheva, Danny Yagan, and anonymous referees for helpful comments and discussions. We also thank Maxim Massenkoff, Yannick Schindler, and Shreya Tandon for excellent research assistance. We gratefully acknowledge support from the Center for Economic Behavior and Inequality (CEBI) at the University of Copenhagen, financed by grant #DNRF134 from the Danish National Research Foundation.

© The Author(s) 2019. Published by Oxford University Press on behalf of the President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2020), 329–388. doi:10.1093/qje/qjz032.
Advance Access publication on October 9, 2019.

How would wealth taxes affect the saving and consumption decisions of the rich? How would wealth taxes affect avoidance and evasion decisions? Would they reduce wealth inequality, and by how much?

Answering these questions is difficult because of several empirical challenges. First, while many countries collect data on labor supply and taxable income, very few collect individual data on wealth. Second, it has been difficult to find compelling variation in wealth taxation that allows for the estimation of causal effects. What is more, because wealth is always very concentrated—much more than labor income—it is crucial to estimate behavioral responses for the very wealthiest individuals. Sources of exogenous variation at the top of the wealth distribution have so far been elusive. Third, to assess the desirability of wealth taxes and of capital taxes more broadly, it is important to obtain estimates of long-run effects. While tax design always depends on long-run effects, this is a bigger challenge for capital taxes than for labor taxes due to the dynamic and slow-moving nature of wealth accumulation.

In this article, we break new ground on these questions. Our laboratory is Denmark, which offers data and quasi-experimental variation that allow us to overcome the challenges described above. 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 OECD countries levied similar taxes (OECD 1988), but the Danish wealth tax was the largest of its kind. The marginal tax rate on wealth equaled 2.2% up until the late 1980s, corresponding to a very high rate on the return to wealth.¹ The Danish government implemented large changes to the wealth tax starting in 1989—cutting the marginal rate to 1% and doubling the exemption threshold for married couples—before eventually abolishing the tax in 1997. These policy changes represent some of the largest natural experiments with wealth taxation ever conducted. In addition, a key advantage of the Danish setting is that the authorities have been collecting micro-level data on wealth for the entire population since 1980.

Our article makes three main contributions. The first is to provide quasi-experimental evidence on the effects of the 1989

1. For example, assuming a rate of return on wealth equal to 4.4%, a marginal wealth tax of 2.2% corresponds to a 50% tax on the flow of capital income.

reform on wealth accumulation. We consider two different empirical strategies and samples. One strategy exploits the doubling of the exemption threshold for couples (but not singles), which eliminated wealth taxes among couples located roughly between the 98th and 99th percentiles of the wealth distribution. This allows us to estimate impacts of wealth taxes on the moderately wealthy using a difference-in-differences design comparing couples in the exempted range to singles in the same range or to couples in other ranges. The other strategy exploits that, among the very wealthiest households, some face a 0 marginal tax rate on wealth due to a tax ceiling that limits the total average tax rate from personal taxes (income, social security, and wealth taxes). Therefore, the tax cuts had different impacts on those bound and unbound by the ceiling. This allows us to estimate impacts of wealth taxes on the very wealthy using a difference-in-differences design comparing bound and unbound taxpayers within the top 1%.²

The quasi-experimental analysis shows that wealth taxes have sizable effects on taxable wealth, with the effects being considerably larger at the extreme top of the distribution than further down. We view our evidence as compelling in the sense that in both of our approaches, the trends in taxable wealth for the treatment and control groups are parallel prior to the reform and then begin to diverge immediately after the reform.³ The effect on wealth builds up over time and is equal to about 19% after eight years for the moderately wealthy (couples DD) and 31% after eight years for the very wealthy (ceiling DD). These effects include both behavioral and mechanical effects: even if households did not change their behavior in response to wealth taxes, the increase in the after-tax rate of return would mechanically increase wealth over time. We show that the mechanical effect accounts for about one-tenth of the effect for the couples DD and about one-fifth of the effect for the ceiling DD.

2. The ceiling strategy represents a novel empirical approach in the large literature on behavioral responses to taxes. This approach offers a promising way to identify behavioral responses among the very wealthy that could be implemented in a number of countries with wealth taxes. This is because most countries with wealth taxes (including Norway, Sweden, France, Spain, and Germany) have such ceiling rules.

3. As we clarify later, the pretrends are not always parallel in the raw data, but they are parallel after adjusting for linear, group-specific pretrends.

Our second contribution is to develop a theoretical model allowing us to interpret the reduced-form impacts in terms of structural primitives. To keep the model relatively simple, we leave out aspects that are not central to our setting and sample. In particular, because wealthy people tend to be relatively old—most of those in the top 1% are above 50 years of age—we focus on the savings motives that are central to older, wealthy people. We argue that the life cycle motive and the bequest motive (or more broadly, utility of residual wealth) are important, while the precautionary motive is second order. Within such a model, we demonstrate how the reduced-form impact on wealth is driven by four conceptual effects: a substitution effect on consumption proportional to the elasticity of intertemporal substitution (EIS), a substitution effect on bequests proportional to a bequest elasticity, a wealth effect on the demand for consumption and bequests, and finally the mechanical effect discussed above. The importance of the bequest elasticity in determining the reduced-form impacts depends on the weight of the bequest motive in household preferences, and we show that this weight has to be large to rationalize the life cycle profile of wealth among very wealthy people. Therefore, the bequest elasticity is very important for understanding wealth responses at the top.

Our third contribution is to connect the theory and evidence to investigate the long-run effects of wealth taxes on wealth accumulation. We calibrate the parameters of the model to match the empirical life cycle profile of wealth at the top of the distribution as well as the quasi-experimental estimates of the short- to medium-term impacts of wealth taxes. When matching the model to the moderately wealthy in the couples DD—an empirical effect on taxable wealth of 19% after 8 years—and simulating the model forward, we obtain a 30-year effect of 30%. When matching the model to the very wealthy in the ceiling DD—an empirical effect on taxable wealth of 31% after 8 years—the long-run effect is considerably larger, 65% after 30 years. Although these effects may seem large, note that the underlying tax incentives driving them are also large. The implied long-run elasticity of taxable wealth with respect to the after-tax rate of return equals 0.77 for the moderately wealthy and 1.15 for the very wealthy.

Given that our estimates rely on tax records, they capture both real responses and evasion/avoidance responses. Most assets were third-party reported under the Danish wealth

tax—thus limiting the scope for evasion—but some were self-reported and therefore susceptible to evasion. If the amount of evasion responds to changes in marginal tax rates, this will be picked up by our elasticity estimates. Wealthy taxpayers may have access to relatively effective evasion vehicles, such as offshore accounts.⁴ It is worth noting, however, that repatriation of offshore wealth is unlikely to be part of our estimated responses. The wealth tax cuts did not come with an amnesty for previously unpaid taxes, implying that such repatriation would trigger back taxes and potential penalties. But our estimates may reflect other forms of evasion and avoidance, and it would be difficult to separate out those responses.⁵

Our article can be viewed in two ways. One view is that it contributes to a nascent literature studying the effects of wealth taxes on taxable wealth (Zoutman 2015; Brülhart et al. 2016; Seim 2017). Compared to this literature, we consider a larger natural experiment and we estimate behavioral responses at the very top of the wealth distribution. We provide clear graphical evidence on the short- to medium-term responses to wealth taxes. Unlike earlier work, our article provides a tractable dynamic framework to shed light on the theoretical mechanisms driving the reduced-form impacts, and it structurally estimates the model to explore the long-term consequences of wealth taxation.

Another view is that our article provides a first attempt to causally estimate the long-run elasticity of capital supply with respect to capital taxes. From this perspective, it is not crucial that we study wealth taxes per se, but that the Danish wealth tax allows us to estimate a key parameter for assessing the efficiency implications of capital taxes more broadly. Saez and Stantcheva (2018) show that the long-run elasticity of capital supply is a sufficient statistic for optimal capital taxation, but there is virtually

4. Using leaked data from HSBC Switzerland and Mossack Fonseca (“Panama Papers”), Alstadsæter, Johannesen, and Zucman (2019) show that essentially all of the wealth in offshore accounts belongs to the top 1% and that most of it belongs to the top 0.1%.

5. One possible strategy is to consider bunching at the kink point created by the exemption threshold. In the context of wealth taxation, bunching almost surely reflects evasion/avoidance responses rather than real responses. We show that there is very little bunching at the kink, consistent with modest evasion/avoidance responses. However, bunching at the kink may understate responsiveness within brackets, which prevents us from using the bunching evidence to rule out significant evasion and avoidance.

no evidence on what a reasonable value of this elasticity might be. Besides the empirical challenges discussed already, a reason for the lack of evidence may be that the seminal theoretical contributions guiding the debate focused on “corner solutions” that did not bring out the key role of the capital supply elasticity. In the Chamley-Judd framework (Chamley 1986; Judd 1985), the optimal capital tax is 0 in steady state because long-run capital supply is infinitely elastic. In the Atkinson-Stiglitz framework (Atkinson and Stiglitz 1976), the optimal capital tax is 0 because there is no heterogeneity in wealth, conditional on labor income. In other words, in one framework capital taxes are undesirable because they are too costly for efficiency, and in the other framework capital taxes are undesirable because they do not improve equity. But in general capital taxes do pose a trade-off between efficiency and equity, and it is governed by the long-run parameters we estimate here.⁶

In the process of producing the findings described above, we provide a number of bonus contributions. It is worth highlighting some of those here. First, our structural approach yields an estimate of the bequest elasticity with respect to the net-of-tax rate on capital.⁷ Although a large literature discusses the incentive effects of taxes on the size of bequests—typically focusing on estate and inheritance taxes—there is very little empirical evidence on the question. Piketty and Saez (2013) highlight that the bequest elasticity is a key parameter for optimal inheritance taxation. Reviews by Kopczuk (2009, 2013a, 2013b) summarize the few existing estimates of this parameter and discuss the challenges associated with interpreting them. We estimate the bequest elasticity based on a fundamentally different approach using variation in wealth taxes (rather than wealth transfer taxes) on wealthy, older people.

6. To be clear, our quasi-experimental estimates do not represent “all-inclusive” long-run elasticities of capital supply with respect to capital taxes. As we discuss later, our estimates capture the effect of wealth taxes on the already wealthy rather than the forward-looking effect on those who aspire to become wealthy. The potential aspiration effect of wealth taxes—including career decisions, entrepreneurial risk taking, and early-life saving decisions made in anticipation of future taxes—would be extremely difficult, if not impossible, to identify convincingly.

7. Or more precisely, we estimate the elasticity of end-of-life wealth. While we refer to this as a “bequest elasticity,” we are not able to distinguish between actual bequest motives and other motives driving utility of residual wealth (e.g., the “capitalist spirit” as discussed in Carroll 2002).

Our findings suggest that bequest elasticities are large at the top of the wealth distribution.⁸

Second, to calibrate our model, we carefully document the empirical life cycle profiles of wealth at the top of the distribution. Because we have access to full-population administrative wealth data over a long time horizon, we are able to provide particularly clean and striking evidence. We show that wealthy people tend to accumulate wealth through most of their lives; only after they reach 80 years of age do their wealth profiles flatten or fall slightly. As a result, people at the top of the wealth distribution tend to die close to their wealth peak. For example, among those who make it to the top 1% of the wealth distribution during their lifetime, the average person is still in the top 2% at age 90 and have almost 20 times the amount of per capita wealth. These findings show just how inaccurate the pure life cycle model is for wealthy individuals, and they would be difficult to rationalize without some form of bequest motive or utility of residual wealth. This part of the article contributes to an empirical literature documenting age-wealth profiles among the elderly (see, e.g., [Love, Palumbo, and Smith 2009](#); [Poterba, Venti, and Wise 2011, 2018](#)). Although these studies relied on small survey data sets, the precision and power of our data allow us to provide clean graphical evidence even for the extreme top of the wealth distribution. We also contribute to a literature trying to explain wealth concentration and the life cycle saving behavior of the rich (see, e.g., [Carroll 2002](#); [De Nardi 2004](#); [Kaplow 2011](#); [Benhabib and Bisin 2018](#)).

The article is organized as follows. [Section II](#) describes the data and documents the evolution of wealth inequality, [Section III](#) presents quasi-experimental evidence on the effects of wealth taxes, [Section IV](#) develops the theoretical model, [Section V](#) combines the model and quasi-experimental evidence to structurally estimate long-run effects of wealth taxes, and [Section VI](#) concludes. All appendix material is presented in the [Online Appendix](#).

8. These bequest elasticities (i.e., residual-wealth elasticities) will include any changes in inter vivos transfers and gifts. Such gift responses are most naturally interpreted as tax evasion. There is a tax exemption threshold for gifts, but it is very small relative to the wealth levels of our population of interest. Gifts above the exemption threshold are not desirable from a tax minimization perspective. However, unreported gifts (in cash or in kind) are difficult to detect for tax authorities and may respond to wealth taxes. The Danish data do not allow us to provide direct evidence on such responses, but they will be part of our estimates of taxable wealth responses.

II. DANISH HOUSEHOLD WEALTH: DATA AND DISTRIBUTION

II.A. *Wealth Data*

We base our analysis on the administrative wealth registry maintained by the Danish Statistical Agency. This registry includes annual wealth data for the entire Danish population since 1980. The Danish authorities initially collected these data to administer the wealth tax, but they continued to do so after the abolition of the wealth tax in 1997. The data are not censored or top coded, which is a key advantage given our focus on the top of the wealth distribution. We combine the wealth registry with other administrative registries containing data on income and socioeconomic characteristics.

The wealth registry includes detailed information on end-of-year financial assets, nonfinancial assets, and debts. As a rule, these assets are recorded in the registry at their prevailing market prices. Most assets and liabilities are reported by third parties to the Danish government, which makes the data very reliable (see [Leth-Petersen 2010](#); [Boserup, Kopczuk, and Kreiner 2014](#)). For instance, the value of bank deposits is reported by banks, the value of listed stocks and bonds is reported by financial institutions (banks, mutual funds, and insurance companies), and the value of mortgages is reported by mortgage lenders. Nonfinancial assets are recorded using land and real estate registries. Before the wealth tax was abolished in 1997, all assets other than those reported by third parties had to be self-reported by households. This included cash, large durables (such as cars, boats, and private planes), noncorporate business assets, unlisted securities (i.e., bearer bonds, unlisted equities, and shares of housing cooperatives), assets held abroad, and interpersonal debts.

The Danish wealth data are considered to be very high quality, and they have been used to study retirement savings ([Chetty et al. 2014](#)), intergenerational wealth mobility ([Boserup, Kopczuk, and Kreiner 2014](#)), and the accuracy of survey responses ([Kreiner, Lassen, and Leth-Petersen 2015](#)). The data do have two limitations, however. First, they exclude funded pension wealth before 2012, because such assets were not subject to wealth taxation. This is not a major issue for our purposes, because we are primarily interested in the effects of wealth taxation on taxable wealth. Moreover, because there are strict limits on the absolute amount that can be invested in tax-preferred pension accounts, pension wealth is always a small fraction of wealth at the top of the

distribution, the focus of our analysis. Second, there is a break in the wealth series in 1997, the year in which the wealth tax was abolished. After 1997, although the Danish administration continued to collect wealth data from third parties, it stopped asking households to self-report assets not reported by third parties.⁹ Because of this break in the data, our quasi-experimental analysis of behavioral responses to wealth taxation focuses on the large 1989 reform for which we have consistently measured taxable wealth before and after the reform.

II.B. Computing Wealth Inequality

To provide context, we start by documenting the evolution of wealth inequality in Denmark over the 1980–2012 period. We compute homogeneous series of wealth shares in which we match 100% of aggregate wealth at market value recorded in Denmark's household balance sheet. This implies that the wealth levels and wealth shares for Denmark are comparable to existing series for other countries, including those estimated for the United States by [Saez and Zucman \(2016\)](#).¹⁰ In keeping with standard national accounting concepts, our definition of wealth includes all financial and nonfinancial assets that belong to Danish residents, minus debts. In particular, it includes all funded pension wealth but excludes the present value of future government transfers as well as consumer durables and valuables. Average wealth per adult person was \$237,000 in 2012 (using the market exchange rate to convert Danish kroner to U.S. dollars), a level similar to that of the United States, where it is \$234,000.

The quality of the Danish data allows us to compute particularly reliable estimates of the wealth distribution. In most countries one has to rely solely on indirect methods to estimate wealth inequality, such as the capitalization method or the estate multiplier method (see [Zucman 2019](#) for a survey). In Denmark, by contrast, we directly observe the market value of most wealth components for the entire population in the administrative wealth registry. To capture 100% of the macroeconomic amount

9. Moreover, there were changes in the coverage of third-party wealth reporting in 1997, implying that even third-party reported wealth by itself suffers from a data break.

10. Similar wealth series are being produced for a growing number of countries, as published on the World Wealth and Income Database at <http://WID.world> ([Alvaredo et al. 2017](#)).

of household wealth, we supplement the wealth registry as follows. First, we impute funded pension wealth throughout the 1980–2012 period, using individual-level pension wealth that was added to the administrative data from 2012 onward.¹¹ Second, we impute assets not reported by third parties by capitalizing the respective income flows. Specifically, we compute noncorporate business assets by capitalizing business income (the capitalization rate equals the aggregate stock of business assets from the national accounts divided by the aggregate flow of business income from individual income tax returns), while we impute unlisted equities by capitalizing dividend income. Importantly, we only make these imputations when computing the distribution of wealth in this section. For our main analysis of behavioral responses to wealth taxes, we focus on reported taxable wealth (thus excluding pensions) because this is the most appropriate outcome for this purpose.

II.C. Trends in Wealth Concentration

Figure I shows wealth shares in three broad classes: the bottom 50%, the next 40%, and the top 10%. These wealth shares have been relatively stable in Denmark over the past three decades. Throughout the period, the bottom 50% of the distribution owns a tiny fraction of aggregate wealth: their assets are barely higher than their debts. Therefore, almost all wealth is owned by the richest half of the population, and it is shared about equally between the middle 40% and the top 10%. Although the wealth shares in the figure are overall stable, wealth inequality did increase somewhat from the mid-1980s to the early 1990s. During this time, the top 10% wealth share grew and the bottom 50% wealth share shrank. This evolution was driven by the dynamics of asset prices, in particular housing prices, which fell significantly during this period. Because the share of housing in asset portfolios tends to be decreasing in the level of wealth, housing slumps

11. The imputation is done as follows. In 2012, we observe that about 40% of pension wealth belongs to wage earners while 60% belongs to retirees. We assume that these shares were the same before 2012. We then allocate the pension wealth of workers proportionally to their wage incomes (winsorized at the 99th percentile) and the pension wealth of retirees proportionally to their pension benefits paid out of pension funds. We have checked that the distribution of imputed pension wealth for 2012 is very close to the observed distribution of pension wealth for that year. Saez and Zucman (2016) use a similar imputation procedure for the United States.




FIGURE I

Distribution of Wealth in Denmark, 1980–2012

This figure shows the share of total household wealth in Denmark 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), splitting household wealth in married couples equally among the spouses. Wealth includes all financial and nonfinancial assets, net of any debts. It matches the total amount of household wealth recorded in the Danish household balance sheet.

hurt the bottom more than the top, leading to a rise in wealth inequality.

Figure II zooms in on the top of the wealth distribution—the sample that is more relevant for our tax reform study—and contrasts Denmark with the United States. Several insights are worth noting. First, wealth inequality is markedly lower in Denmark than in the United States. In 2012, the top 1% accounts for about 20% of total wealth in Denmark, while it accounts for almost 40% in the United States. Average wealth in the population is similar in the two countries, but the top 1% are twice as wealthy in the United States as they are in Denmark.¹²

12. The average person in the top 1% of the U.S. distribution owned net wealth of \$9.3 million in 2012 (roughly 40 times average wealth), as opposed to \$4.8 million in Denmark (roughly 20 times average wealth).

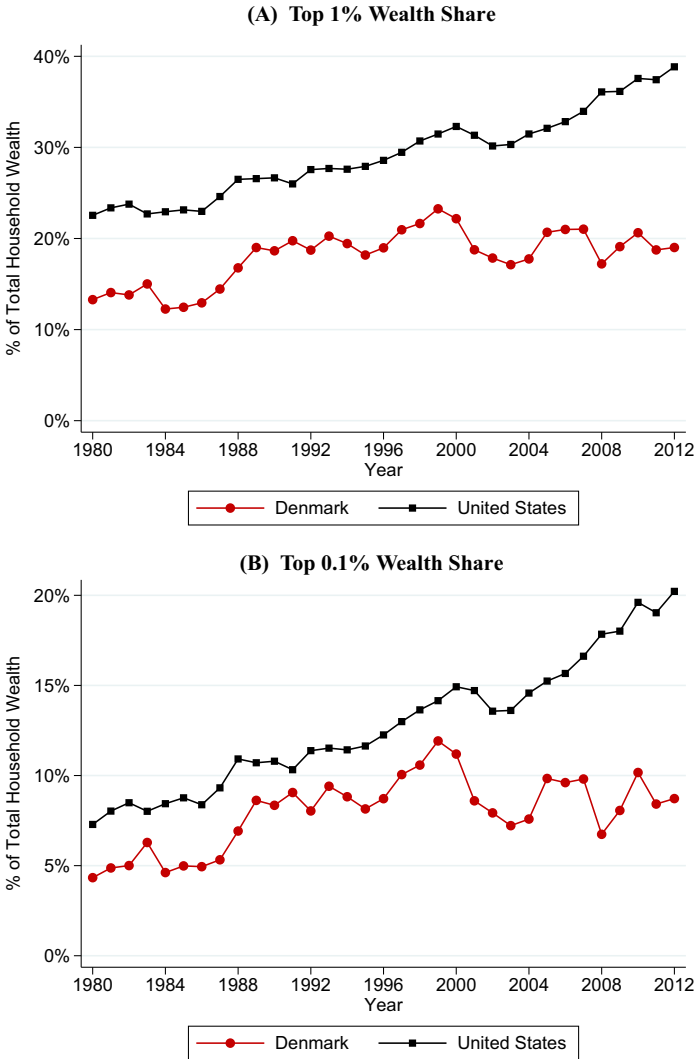


FIGURE II

Top 1% and Top 0.1% Wealth Shares in Denmark versus the United States

This figure shows the share of total household wealth owned by the top 1% (Panel A) and the top 0.1% (Panel B) in Denmark versus United States. The U.S. series is the one estimated by [Saez and Zucman \(2016\)](#). In both countries, the unit of observation is the adult individual (aged 20 or above), splitting household wealth in married couples equally among the spouses. Wealth includes all financial and nonfinancial assets, net of any debts, and it adds up to the total amount of household wealth recorded in the Danish and U.S. household balance sheets.

Second, the gap between the countries has widened over time. Top wealth shares were increasing in both countries until the late 1990s, but then they begin to diverge as wealth inequality stabilized in Denmark while it continued to increase in the United States. Third, the similarity between the two countries until the late 1990s and the subsequent divergence look more striking as we move into the extreme tail of the distribution. As shown in the bottom panel of [Figure II](#), the top 0.1% wealth share in Denmark was only 2–3 percentage points lower than in the United States around 2000 but then started to diverge very strongly. If we consider top 0.01% wealth shares (not shown), they were essentially the same in the two countries at the turn of the century and then diverged.

To conclude, despite the reduction and ultimate abolition of the wealth tax in Denmark in the 1990s, wealth accumulation at the top of the distribution (relative to the population as a whole) has not picked up speed in Denmark as compared with the United States. In other words, the aggregate patterns documented here do not provide a smoking gun for behavioral effects of wealth taxes. Of course, this does not imply that wealth taxes did not affect wealth accumulation and wealth inequality. It simply means that if the wealth tax cuts caused wealth to grow faster at the top, this unequalizing force must have been offset by confounding equalizing forces. In our analysis of the causal effect of wealth taxation, we do find that lower wealth taxes cause wealth to grow faster.¹³

III. THE EFFECT OF WEALTH TAXES: EVIDENCE

III.A. Tax Variation and Empirical Strategies

Denmark taxed wealth on an annual basis until 1997. Taxable wealth equaled the total net wealth of households, excluding pension wealth. Taxable wealth components thus included cash, deposits, bonds, equities, housing, large durables, and business assets, net of any debts. A number of these components were third-party reported by financial institutions, leaving little scope for tax

13. One important confounding reason wealth inequality has stabilized in Denmark (despite wealth tax cuts) is likely to be the sharp rise of pension wealth, from around 50% of national income in the late 1980s to 178% in 2014. Because pension wealth is relatively equally distributed, rising pension wealth tends to reduce inequality.

evasion. But some components were self-reported, namely, cash, durables, unlisted equities, noncorporate business assets, and assets held abroad.

Wealth was taxed at a flat rate above an exemption threshold. The exemption threshold varied over time (differentially for singles and couples) as we discuss later, but it was always above the 97th percentile of the household wealth distribution during the period we study. Wealth above the exemption threshold was taxed at 2.2% until two major reforms in the late 1980s and 1990s. 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 changes are illustrated in [Figure III](#).

This setting offers two sources of exogenous variation: the kink point at the exemption threshold and the tax reforms. Let us first consider the former. The kink point is very sharp as a tax rate jump of 2.2% on the stock of wealth translates into a very large tax rate jump on the return to wealth. As a result, taxpayers have strong incentives to bunch at the kink, allowing for a bunching approach to estimating taxable wealth responses.¹⁴ However, while bunching approaches are useful for uncovering evasion and avoidance responses to wealth taxes, they are not useful for uncovering real responses to such taxes. Taxable wealth depends not only on individual decisions but also on asset prices that are highly uncertain and move continuously through the tax year. Given such asset price movements, it would be virtually impossible for a taxpayer to bunch at the exemption threshold using real savings responses. Therefore, we do not pursue a bunching strategy as our main approach. We present a bunching analysis in Section A of the [Online Appendix](#), but we view this evidence primarily as informative of avoidance responses.¹⁵

14. [Kleven \(2016\)](#) provides a review of bunching approaches. In the context of wealth taxation, [Seim \(2017\)](#) presents bunching evidence using a kink point in the Swedish wealth tax. He finds clear bunching at the kink, but the implied elasticity of taxable wealth is small.

15. There is some bunching at the Danish wealth tax kink, but it is small—even smaller than in the Swedish context analyzed by [Seim \(2017\)](#)—and the implied elasticity of taxable wealth is tiny. The presence of bunching is useful for rejecting the null of no avoidance responses, but bunching is unlikely to capture the global responsiveness of avoidance to changes in tax rates. In a world with fixed avoidance costs and/or lumpy assets (“optimization frictions”), the amount of bunching understates the responsiveness of avoidance because taxpayers may overshoot or undershoot the threshold.

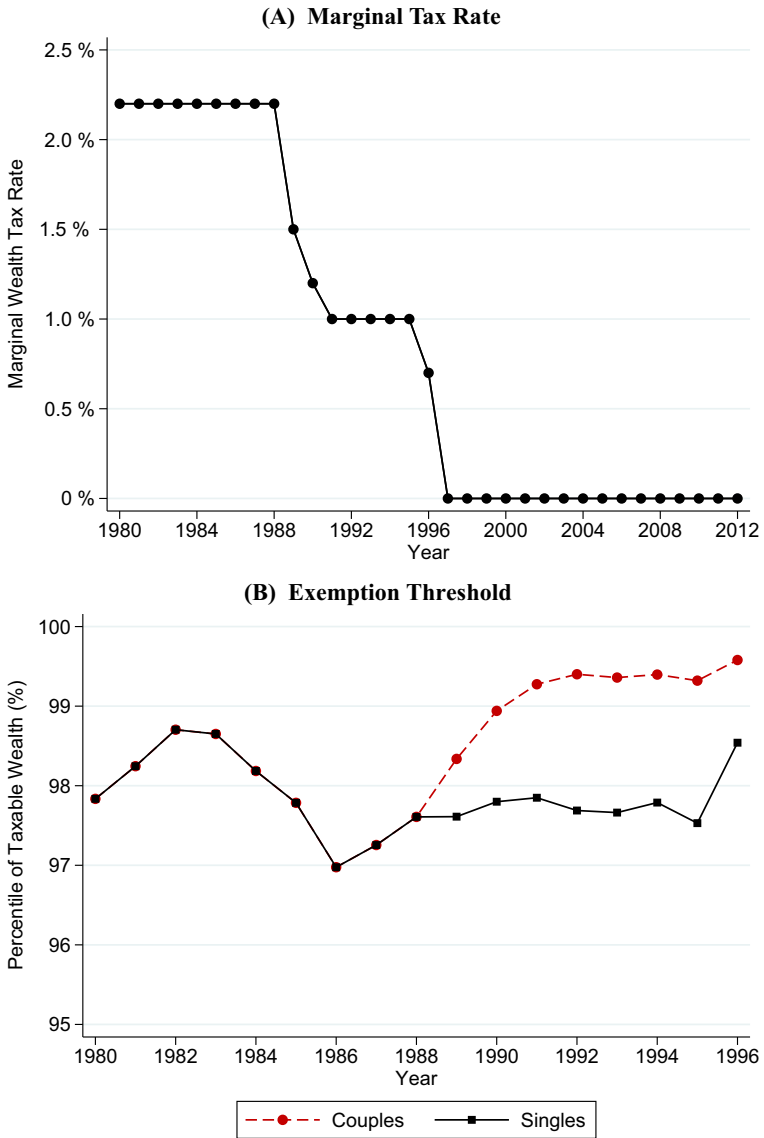


FIGURE III
Wealth Tax Variation

This figure shows the evolution of the marginal tax rate (Panel A) and the exemption threshold (Panel B) in the Danish wealth tax.

Given these considerations, our main analysis is based on the tax reform variation. In particular, we focus on the 1989 reform rather than the subsequent elimination of the tax, because of a data limitation discussed earlier: after abolishing the wealth tax in 1997, Statistics Denmark no longer records purely self-reported wealth. This break in the taxable wealth series makes it difficult to study the wealth tax abolishment, so we focus on the earlier tax cuts that do not have this limitation. To estimate behavioral responses to the 1989 reform, we consider difference-in-differences (DD) approaches in which we compare treatment and control groups in a balanced panel of taxpayers. We develop two DD approaches.¹⁶

The first approach uses the fact that the 1989 reform increased the exemption threshold for couples relative to singles. Before the reform, singles and couples faced the same nominal exemption threshold for wealth taxation. This is difficult to justify on equity grounds, because a couple is less wealthy in per capita terms than a single individual at the same level of household wealth. To rectify this issue, the exemption threshold for couples relative to singles was doubled between 1989 and 1992. These threshold changes are illustrated in [Figure III](#), Panel B. The implication of the reform is that couples in a certain range of the household wealth distribution—those between the 97.6th and the 99.3rd percentiles—became exempt from wealth taxation, allowing us to estimate responses by the moderately wealthy. We compare couples in the affected range (treatments) to singles in the same range or to couples below the range (controls). We refer to this strategy as the couples DD.

The second approach uses the fact that the 1989 reform reduced the tax rate from 2.2% to 1%. To define groups that were differentially affected by this tax cut, we exploit the existence of a ceiling on the total tax liability from all personal taxes (income taxes, social security taxes, and wealth taxes) as a fraction of taxable income. This ceiling—known as *Det Vandrette Skatte loft* (“horizontal tax ceiling”)—was in place to limit the total average

16. A compelling aspect of estimating wealth responses using the 1989 tax reform is that there was essentially no room for anticipation effects. The tax bill was first proposed on November 3, 1988, it was passed in parliament on December 21, 1988, and it took effect on January 1, 1989. Given that wealth accumulation is a forward-looking variable, the absence of anticipation is important for the validity of our empirical strategies.

tax rate on households with large wealth relative to income.¹⁷ The tax ceiling was set at 78% at the time of the 1989 reform. Whenever the total average tax rate exceeded this limit, tax liability would be reduced by the excess amount. For households bound by the ceiling, the marginal tax rate on wealth was equal to 0—before and after the reform—making them a natural control group. [Online Appendix Figure A.I](#) shows the fraction of taxpayers bound by the ceiling at different quantiles of the wealth distribution. The ceiling starts binding for a substantial fraction of households as we move into the extreme tail of the wealth distribution, allowing us to estimate behavioral responses by the very wealthy. We compare taxpayers unbound by the ceiling (treatments) to taxpayers bound by the ceiling (controls) within the top 1% of the distribution. We refer to this strategy as the ceiling DD.¹⁸

The empirical analysis is based on a standard DD event study specification, that is,

$$(1) \quad \log W_{it} = \sum_{j \neq 1988} \beta_j \cdot Year_{j=t} \cdot Treat_{it} + \gamma_i + \eta_t + v_{it},$$

where W_{it} denotes the wealth of household i in year t , $Year_{j=t}$ is a dummy equal to 1 when the year equals t , $Treat_{it}$ is a

17. The definition of taxable income used to assess the ceiling rule included labor income, pension income, and capital income (interest income, dividends, capital gains, etc.). The exact income definition was quite complicated and underwent some changes over time. See [Skatteministeriet \(2002\)](#) for details on the history of the personal income tax code in Denmark.

18. It is useful to define the treatment and control groups more formally. If we denote taxable income by z , taxable wealth by W , and the wealth exemption threshold by \bar{W} , then the tax ceiling binds if

$$\frac{t \cdot z + \tau \cdot (W - \bar{W})}{z} \geq 0.78 \quad \Leftrightarrow \quad \frac{W - \bar{W}}{z} \geq \frac{0.78 - t}{\tau},$$

where t is the average income tax rate (including social security taxes) and τ is the marginal wealth tax rate. With a top marginal income tax rate of 68% at the time of this reform (the upper bound on the average income tax rate) and a wealth tax rate of 2.2%, the ceiling starts binding for those with a taxable wealth-income ratio $\frac{W - \bar{W}}{z}$ of at least 5 and typically much higher. In the data, the average household bound by the ceiling has a wealth-income ratio well above 10. Since taxable income z includes capital income, the households affected by the ceiling tend to be those with large assets in a form that do not generate a correspondingly large flow income. Examples include valuable real estate or equity with unrealized capital gains. Such households are the controls in the ceiling DD.

dummy equal to 1 when household i is in the treatment group at time t , γ_i is a household fixed effect, η_t is a year fixed effect, and v_{it} is an error term. The DD coefficient β_t captures the effect of the tax reform in year t relative to the prereform year, 1988.

The assignment of treatment status depends on the empirical strategy, either being a couple in the exempted wealth range (couple DD) or being unbound by the tax ceiling (ceiling DD). A basic issue with specification (1) is that concurrent treatment status $Treat_{it}$ is endogenous to the outcome variable. To avoid bias from such endogeneity, we construct instruments based on prereform variables. Specifically, defining $Treat_i^{pre}$ as an indicator for being in the treatment group based on prereform behavior, we instrument $Year_{j=t} \cdot Treat_{it}$ in equation (1) using $Year_{j=t} \cdot Treat_i^{pre}$. For prereform treatment status to be a strong predictor of postreform treatment status, household behavior must be sufficiently persistent over time. To increase persistence, we focus on households with the same status in several consecutive prereform years. As a baseline, treatment status is assigned based on six prereform years (1982–88), but we show that results are robust to shorter and longer treatment windows.¹⁹

The IV coefficients $\hat{\beta}_t$ estimated based on specification (1) represent treatment-on-the-treated (TOT) effects. We also consider reduced-form specifications, that is, regressing log wealth directly on the instruments, $Year_{j=t} \cdot Treat_i^{pre}$. The reduced-form coefficients represent intention-to-treat (ITT) effects. The discrepancy between ITT and TOT effects depends on the persistence of treatment status. Even though we consider households with the same treatment status over several prereform years, their status is not perfectly persistent over the postreform period. This attenuates the ITT effects relative to the TOT effects.

As always, the difference-in-differences approach relies on the assumption of parallel trends. We assess the validity of this assumption by inspecting the pretrends of the comparison groups, and where these are not parallel, we adjust the series for differential pretrends. Specifically, using n prereform years to estimate the trend, we consider the following extension of the baseline event

19. Households that switch status during the specified treatment window (1982–88 in the baseline) are dropped from the estimation sample.

study specification

$$(2) \quad \log W_{it} = \sum_{j \notin (1988-n, 1988)} \beta_j \cdot Year_{j=t} \cdot Treat_{it} + \theta^T \cdot t \cdot Treat_i^{pre} + \gamma_i + \eta_t + v_{it},$$

where θ^T is a linear differential pretrend identified based on n prereform years (i.e., the omitted years in the first term on the right-hand side). In the implementation below, we estimate the pretrend using four prereform years.²⁰

These specifications give the effect of “treatment” on log wealth, where treatment results either from the exemption of couples in a certain wealth range or from the tax rate cut on households unbound by the ceiling. It is useful to convert these effects into elasticities with respect to the net-of-tax rate on wealth. This allows us to compare magnitudes across different specifications and compare the quasi-experimental estimations to the structural elasticity estimations presented later. To calculate elasticities, we relate the ITT effect on wealth to the expected change in the net-of-tax rate on wealth. The expected tax change accounts for the variation in tax treatment driven by the churn in and out of treatment and control groups after the reform. Specifically, we define

$$(3) \quad \varepsilon = \frac{E[\hat{\beta}_t^{ITT}]}{E[\Delta \log(1 - \tau_{it}) | T] - E[\Delta \log(1 - \tau_{it}) | C]},$$

where $E[\cdot]$ is the expectations operator, $\hat{\beta}_t^{ITT}$ is the ITT coefficient in year t , and τ_{it} is the wealth tax rate on household i at time t . The denominator represents the expected log-change in the net-of-tax rate for the treatment group T relative to the control group C . The elasticity defined in equation (3) is based on the average effect over the postreform window.

20. In general, the presence of pretrends in the raw data suggests that the unadjusted difference-in-differences estimates will be biased by nonreform confounders. Controlling for such confounders by extrapolating the pre-event trend is valid only under the assumption that the postevent behavior of the confound can be inferred from the pre-event trend (see, e.g., Freyaldenhoven, Hansen, and Shapiro 2019; Roth 2019). We discuss the validity of this assumption when presenting the results from the specifications where a pretrend adjustment is necessary.

It is useful to walk through the elasticity calculation by way of a concrete example. Consider the couples DD. Households in the treatment group—couples located in the exempted wealth range before the reform—may experience one of three treatment states after the reform: they may stay in the exempted range for couples (thus experiencing a tax cut of 2.2 percentage points), they may move above the new exemption threshold for couples or become singles above the threshold applying to singles (thus experiencing a tax cut of 1.2 percentage points), or their wealth may fall below the old exemption threshold (where the tax cut is 0). The expected tax change for the treatment group accounts for this churn and is therefore smaller than the mechanical tax change of 2.2 percentage points. Similarly, the expected change for the control group is different from 1.2 percentage points because they may fall below the exemption threshold or become couples in the exempted range. By scaling the ITT effect using the expected tax change, the resulting parameter represents a TOT elasticity. An alternative approach would be to divide the TOT effect by the mechanical, reform-induced tax change for the treatment group relative to the control group (i.e., a tax cut of 2.2 percentage points relative to 1.2 percentage points). These two approaches are not equivalent in our setting, but in practice they give almost the same result.²¹

The elasticity in [equation \(3\)](#) is defined with respect to the net-of-tax rate on wealth, $1 - \tau$. An alternative elasticity can be defined with respect to the net-of-tax-rate rate of return to wealth, that is, $(1 - \tau)R - 1$ where R is the gross rate of return on wealth. This is a more meaningful elasticity concept in the context of wealth taxation, but it requires us to make an assumption about the (unobserved) rate of return R . We consider both types of elasticity calculations.

III.B. Descriptive Statistics

Before investigating behavioral responses to the wealth tax, we present descriptive statistics in [Table I](#). The table shows means

21. The reason why these approaches are not equivalent is that the treatment category (a tax cut of 2.2 percentage points) and control category (a tax cut of 1.2 percentage points) are not mutually exclusive because some of the churn is driven by households in either group falling below the old exemption threshold and thus receiving zero treatment. These movements create bias in the elasticity calculated directly from the TOT effects. In practice, this has a small effect on the elasticity as the movements from the treatment and control groups create offsetting biases of similar magnitudes.

TABLE I
DESCRIPTIVE STATISTICS

	Full population		Couples DD		Ceiling DD	
	(1)	Couples within (treatment) (2)	Singles within (control) (3)	Couples below (control) (4)	Unbound (treatment) (5)	Bound (control) (6)
Taxable wealth	336,746	3,155,285	3,144,293	1,788,803	4,680,904	19,859,232
Market value wealth	744,734	3,937,661	3,686,328	2,410,357	6,054,538	18,719,788
Gross labor income	203,795	134,471	60,741	165,343	206,594	288,987
Gross total income	311,219	606,962	420,399	422,403	866,325	1,771,303
Asset share in equities	0.01	0.06	0.09	0.03	0.12	0.28
Asset share in housing	0.35	0.30	0.30	0.41	0.28	0.19
Fraction self-employed	0.14	0.67	0.45	0.48	0.65	0.79
Age	46.25	60.21	63.57	59.49	61.34	57.84
Fraction married	0.39	1.00	0.00	1.00	0.75	0.74
Years of schooling	11.03	10.80	11.16	10.27	11.85	12.18
Observations	18,965,710	32,354	10,815	42,924	35,966	2,947

Notes. This table shows means of wealth, income, and demographics for the full population (column (1)) and for our treatment and control groups (columns (2)–(6)). Columns (2)–(4) show the comparison groups in the couples DD: couples and singles located within or below the exempted wealth range (but always within the top 5%). Columns (5) and (6) show the comparison groups in the ceiling DD: households bound or unbound by the tax ceiling (within the top 1%). The assignment to treatment and control groups is based on prereform variables, restricting attention to households whose status stays constant during 1982–88. The table shows statistics at the household level (for adults aged 20 or above). Individual-specific variables (age, schooling) have been averaged across spouses. The statistics are based on pooled prereform data from 1982 to 1988. Monetary values are reported in Danish Kroner in 2017 terms (The \$/DKK exchange rate was 6.2 as of December 31, 2017). Market value wealth is computed as described in Section II.B. Gross total income includes all labor income, pensions, and capital income.

of wealth, income, and demographics for households in the full population (column (1)) and for households in our treatment and control groups (columns (2)–(6)). As discussed already, the assignment of treatment status is based on prereform variables and restricts attention to households whose status stays constant during 1982–88. The statistics in the table are based on pooled data between 1982 and 1988. The table reports both taxable wealth and total market value wealth, the latter computed as described in [Section II.B](#).

The following points are worth highlighting. First, our population of interest is very different from the general population. The treatment and control groups consist of households who are wealthier, older, and more self-employed than the average household in the population. They also hold a larger share of their wealth in equities and a somewhat smaller share in housing.²² Second, market value wealth is generally larger than taxable wealth, but less so in our estimation sample of wealthy taxpayers. This is primarily because pension wealth (which is not part of taxable wealth) weighs less heavily in the portfolio of the wealthy. Third, the difference between labor income and total income (including capital income) is relatively small in the full population, but large among the wealthy who receive most of their income in the form of asset returns.

Finally, there are some noticeable differences in prereform means for the treatment and control groups. This is to be expected given how these groups are defined. The couples DD—especially when we compare couples and singles within the exempted range—is much more balanced than the ceiling DD, where we compare households who are bound and unbound by the tax ceiling. Those bound by the ceiling are much wealthier, hold more of their wealth in equities and less in housing, and are more self-employed than the treated group of unbound taxpayers. This

22. Like a number of countries, Denmark subsidizes homeownership through a mortgage interest deduction. The value of this deduction used to be larger at higher incomes, but an income tax reform enacted in 1987 lowered the deduction and made it uniform across income tax brackets. [Gruber, Jensen, and Kleven \(2018\)](#) investigate behavioral responses to this change in the mortgage interest deduction and find zero effect at the extensive margin of housing investments. In any case, the mortgage interest deduction is not very important among the wealthy population studied here (as they have little or no mortgage debt) and the empirical strategies used in this article rely on different tax variation (within the group of wealthy people) than the variation created by the 1987 income tax reform.

lack of balance could be a concern for the ceiling DD approach, but only insofar as it affects the credibility of the parallel trends assumption.

III.C. Couples DD: Responses by the Moderately Wealthy

We first consider behavioral responses by the moderately wealthy using the couples DD strategy described in [Section III.A](#). This strategy exploits the fact that the 1989 reform doubled the exemption threshold for couples, thus eliminating wealth taxation for couples between the 97.6th and 99.3rd percentiles of the household wealth distribution. We compare these households to two alternative control groups: (i) singles located in the exempted range and (ii) couples located below the exempted range (within the top 5%). The advantage of the first specification is that the comparison groups have the same level of household wealth, but the disadvantage is that both groups are treated to some degree. Although couples in the exempted range have their tax rate cut to 0, singles in this range have their tax rate cut to 1%. The second specification is based on an untreated comparison group and therefore larger identifying variation, but it will require us to deal with differential trends in different parts of the wealth distribution. Under both specifications, the assignment of treatment status (from marital status and wealth bracket) is based on the observed values in prereform years (1982–88 in the baseline).

[Figure IV](#) provides evidence based on using singles as the comparison group. Panel A shows the time series of log taxable wealth for couples in the exempted wealth range (dots) and singles in the same range (squares) between 1980 and 1996, with both series normalized to 0 in the year before the 1989 reform. Panel B shows the differences between these two series. Two key insights emerge from the figure. First, the two groups are on similar trends prior to the reform. Although there are some differences in the early 1980s, the trends are almost perfectly parallel in the five years leading up to the reform. Second, the two series begin to diverge immediately after the reform. The difference in wealth levels between the groups is gradually increasing over time, consistent with a change in the savings rate. Overall, this figure provides clear evidence of behavioral responses to the reduction in wealth taxation.

The results in [Figure IV](#) correspond to reduced-form or ITT effects. The comparison groups are based on prereform treatment

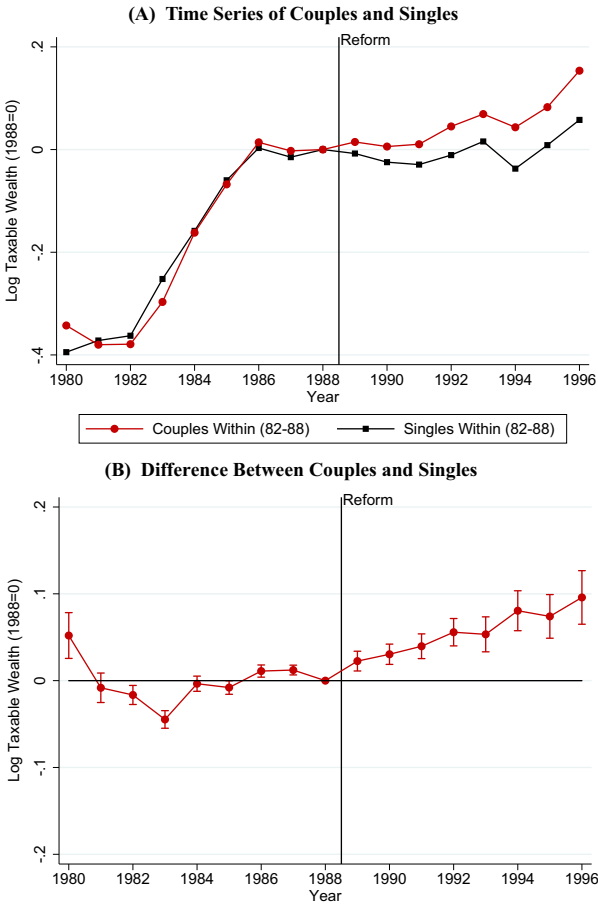


FIGURE IV

Difference-in-Differences Comparing Couples and Singles within Exempted Range

This figure shows the effects of the 1989 wealth tax reform on taxable wealth based on the couples DD in which we compare couples and singles within the exempted wealth range. The assignment to treatment and control groups is based on prereform variables, restricting attention to households whose status stays constant during 1982–88. The sample is a balanced panel of households observed in all years 1980–96. Panel A shows the evolution of log taxable wealth in the two comparison groups, normalized to 0 in the prereform year 1988. Panel B shows the differences between these two series, that is, our reduced-form or intention-to-treat (ITT) estimates. The 95% confidence intervals are based on robust standard errors clustered at the household level.

status, which is not perfectly persistent over time and this attenuates the observed effects.²³ Figure V investigates the persistence of treatment status and converts the ITT effects into TOT effects. Panel A documents the degree of persistence by showing the fraction treated in the two groups over time. By construction, couples within the exempted range are 100% treated in the six prereform years, while singles within this range are 0% treated in those years. After the reform, taxpayers may switch status due to changes in relationship status (through marriage, divorce, or widowhood) or changes in their wealth bracket. The figure shows that the “control group” (singles) is very persistent, reflecting the fact that it is unusual in this predominantly older sample to become married and at the same time stay within the same wealth bracket. The “treatment group” (couples) is less persistent because of wealth changes and spousal death or separation. Eight years after the reform, the difference in treatment intensity is about 50%. Panel B converts the ITT series into a TOT series by dividing the former with the differences in treatment intensity from Panel A (Wald estimator).²⁴ This implies that the dynamically growing effect on taxable wealth is enhanced, an implication of the gradual reduction in persistence. The TOT effect on log wealth is equal to 0.186 in the last postreform year (1996), that is, an increase of about 18% over eight years.

When considering the effects on wealth, it is important to keep in mind that these include both behavioral and mechanical effects. The tax reform raises the after-tax rate of return on wealth, which would increase wealth accumulation even if behavior were fixed. How much of the effects can be explained by such mechanical effects? This is not an entirely trivial question to answer due to two complications. The first complication is that the mechanical tax savings cannot be based on observed wealth because this includes any behavioral responses, but must be based on a measure of counterfactual wealth. Consistent with the DD design, we impute counterfactual wealth after the reform as observed wealth before the reform (in 1988) plus the growth rate in

23. In terms of the regression framework in Section III.A, the estimates in Panel B correspond to the reduced-form version of specification (1) where log wealth is regressed directly on the instruments, $Year_{j=t} \cdot Treat_i^{pre}$.

24. In regression terms, the TOT series in Panel B correspond to the IV estimates from equation (1) where the concurrent year-by-treatment dummies, $Year_{j=t} \cdot Treat_{it}$, are instrumented using $Year_{j=t} \cdot Treat_i^{pre}$.




FIGURE V

Couples DD: Treatment Effect on the Treated

This figure converts intention-to-treat (ITT) estimates into treatment-on-the-treated (TOT) estimates based on the couples DD in which we compare couples and singles within the exempted wealth range. The assignment to treatment and control groups is based on prereform variables, restricting attention to households whose status stays constant during 1982–88. Panel A shows the persistence in treatment status over time. By construction, couples in the exempted range have a treatment status of 100% before the reform, while singles in this range have a treatment status of 0% before the reform. After the reform, taxpayers may switch status due to changes in relationship status or changes in wealth bracket. Panel B compares the ITT and TOT series, where the TOT estimates are obtained by instrumenting treatment status in equation (1) using treatment status in the prereform years 1982–88.

wealth experienced by the control group. The second complication is that the tax savings earned in a given year will grow over time according to a rate of return that is not directly observed in the data. We will assume an annual rate of return equal to 5%. This falls within the range of existing estimates of wealth returns at the top of the distribution (see, e.g., [Fagereng et al. 2016](#)) and it corresponds to what we assume in the calibration exercise presented later. Based on these assumptions, we calculate the cumulative mechanical tax savings in each year due to the wealth tax cuts. The details of this calculation are provided in [Online Appendix, Section B](#) and the results are presented in [Figure A.IV](#).

[Figure A.IV](#) shows the series of total effects (squares) and behavioral effects (triangles), with the differences between the two being the mechanical effects. We see that the mechanical effects are small, an effect on log wealth of 0.02—or 11% of the total effect—after eight years. The reason the mechanical effects are modest has to do with the progressive nature of the wealth tax: taxes are saved only above the exemption threshold located around the 98th percentile of the wealth distribution. That is, while the behavioral responses are governed by the change in the marginal after-tax return (which is very large), the mechanical effect is governed by the change in the average after-tax return (which is more modest). This is a nice feature of the quasi-experiment we are analyzing. If we had considered similar rate changes in a proportional wealth tax, the mechanical effects would have been much larger.

In the [Online Appendix](#), we provide a number of robustness checks. First, [Figure A.V](#) investigates if our results are sensitive to defining comparison groups based on outcomes (marital status and wealth levels) in specific prereform years. The figure shows the evolution of log taxable wealth in the two comparison groups—couples and singles within the exempted wealth range—when using different prereform windows to define treatment status: 1980–88, 1982–88 (baseline), 1984–88, and 1986–88. The figure shows that the main implication of using a longer treatment window is to make the pretrends more parallel, especially in the early 1980s. Reassuringly, the results are similar across the alternative specifications. In all four panels, the wealth trends are almost parallel in the last five years before reform, and the divergence in wealth is about the same after reform. Based on this graph, we conclude that our results are not sensitive to the length of the treatment window.

Second, [Online Appendix](#) Figure A.VI provides a set of placebo tests assuming that the reform happened in earlier years: 1983, 1984, 1985, and 1986. The comparison groups are still couples and singles within the exempted wealth range, but the group assignment is based on outcomes prior to the placebo reform rather than the actual reform. When studying placebo reforms in the early 1980s, we have to shorten the window used to assign treatment status to three years (e.g., 1980–82 for the 1983 placebo reform). The figure shows ITT and TOT series for each of the four placebo reforms. The patterns lend further support to our interpretation of the data. In three out of four panels, there is a precisely estimated zero effect of the placebo reform in 1988, the last year before the actual reform starts affecting the patterns. Only the 1984 placebo reform appears to generate an effect. However, this is due to the fact that 1983 is an outlier year (see also [Figure IV](#)), and so normalizing the series to 0 in 1983 (as we do when the reform is assumed to happen in 1984) creates an illusory effect.

Third, in [Online Appendix](#) Figures A.VII–A.IX we consider the approach in which the comparison group consists of couples below the exempted wealth range. In this case, the comparison group is completely unaffected by the tax cuts, and the experiment is therefore larger. The figures are constructed in the same way as the corresponding figures for the previous strategy. Consider the raw series of log taxable wealth in [Figure A.VII](#), Panel A. The graph shows that couples in the exempted range are on a flatter trend than those below the exempted range in the years before the 1989 reform, while they are on the same and subsequently steeper trend in the years after the reform. This change in relative trends is consistent with an effect of the tax cuts on wealth accumulation. At the same time, we note that the timing of the trend break does not coincide exactly with the tax reform but happens a little too early. This points to the possibility of confounding shocks that have different effects on different parts of the wealth distribution. Such confounders may bias the estimated treatment effect and we therefore have less confidence in this specification.

These concerns notwithstanding, it is useful to turn the raw wealth series in [Online Appendix](#) [Figure A.VII](#), Panel A into DD estimates that are comparable to those obtained from the previous strategy. This requires us to adjust for the differential pretrends of the treatment and control groups. This is done in Panel B using specification (2). The dashed series show the raw differences between the treatment and control groups, and the solid series

show the pretrend-adjusted differences between the two groups. The next figure documents persistence and presents the series of TOT effects. We see that the TOT effect builds up gradually and is equal to 0.265 log points after eight years. This is much larger than the treatment effect obtained from the previous strategy, but recall that the underlying tax variation is also much larger here. In fact, as we show later, the estimates are similar in elasticity terms.

The [Online Appendix](#) presents results from one additional specification, a cross between the previous two. Instead of using singles in the exempted range or couples below the exempted range as controls, this strategy uses singles below the exempted range as controls. The rationale behind this strategy is that a couple with household wealth W has the same wealth per capita as a single person with wealth $\frac{W}{2}$. Therefore, this strategy compares couples in the exempted range (those between the singles' threshold and twice the singles' threshold) to singles in the same per capita range (those between half the singles' threshold and the singles' threshold). The results are presented in [Online Appendix](#) Figures A.X–A.XII and they look quite compelling. The raw pretrends are almost perfectly parallel in the eight years before the reform, followed by a clear divergence in the eight years after the reform. Again, we note that the treatment effect starts a little too early, raising concerns about confounders that are not present in our main strategy of comparing couples and singles in the same range of household wealth.

To conclude, we have presented findings from several DD specifications that take advantage of the doubling of the exemption threshold for married couples. We have compared these treated couples to different control groups, either singles or other couples within or below the exempted wealth range. Taken together, these specifications provide evidence of quite sizable taxable wealth responses to wealth taxation.²⁵

25. Although we focus on the impact of the wealth tax reform on wealth accumulation, it is worth noting that the reform also changed the marriage incentives. The fact that the nominal exemption threshold was the same for singles and couples prior to the 1989 reform created a significant marriage penalty for wealthy individuals. By doubling the exemption threshold for couples, the reform eliminated this penalty and strengthened the incentives to marry at the top of the distribution. [Online Appendix](#) Figure A.XIII provides evidence on the potential marriage responses. It shows the evolution of marriage rates in different wealth quantiles above and below the prereform threshold: the top 1% and top

III.D. Ceiling DD: Responses by the Very Wealthy

We now turn to the behavioral responses of the very wealthiest taxpayers using the ceiling DD. This strategy consists in comparing taxpayers who are unbound by the tax ceiling (treatments) to taxpayers who are bound by the tax ceiling (controls). The treatment group experienced a reduction in the marginal wealth tax rate from 2.2% to 1%, while the control group experienced no change in their marginal tax rate. Because the ceiling starts binding only at the very top of the wealth distribution (as shown in [Online Appendix Figure A.I](#)), we compare bound and unbound taxpayers within the top 1% of the wealth distribution. We assign taxpayers to treatment and control groups using six prereform years, thus dropping taxpayers who frequently switch ceiling status. To further increase the persistence of treatment status, we also drop observations who are only marginally bound by the tax ceiling. Specifically, the bound group includes those whose wealth tax liability would have to fall by at least 20% for them to become unbound, but the results are robust to alternative cuts.

[Figure VI](#) and [Online Appendix Figure A.XIV](#) are constructed in the same way as the preceding figures for the couples DD. In [Figure VI](#), Panel A, we see that the treatment group (unbound) is on a flatter trend than the control group (bound) during the prereform period. This pattern reverses just after reform, and the treatment group is on a considerably steeper trend during the entire postreform period. The switch from a flatter to a steeper trend around the reform provides strong evidence of behavioral responses to the reform. [Figure VI](#), Panel B shows the differenced series, with the raw differences in dashed and the pretrend adjusted differences in solid. The pretrend adjustment is based on

2.5–1% percentiles are above the threshold (where the incentive to marry becomes stronger after the reform), while the top 5–2.5% and top 10–5% percentiles are below the threshold (where the incentive to marry is unchanged). Interestingly, while the four groups are on parallel trends before the reform, the marriage rate increases in the higher percentiles relative to the lower percentiles after the reform. This is consistent with a behavioral response to the marriage penalty. However, the figure also highlights a potential caveat to this interpretation by showing that there is a general fanning out across the four groups. The top 1% increases relative to the top 2.5–1% (even though they are both treated), and the top 5–2.5% increases relative to the top 10–5% (even though they are both untreated). This suggests the presence of confounding effects on marriage. Therefore, while the patterns in [Online Appendix Figure A.XIII](#) are intriguing and consistent with a marriage response, the evidence is not conclusive.

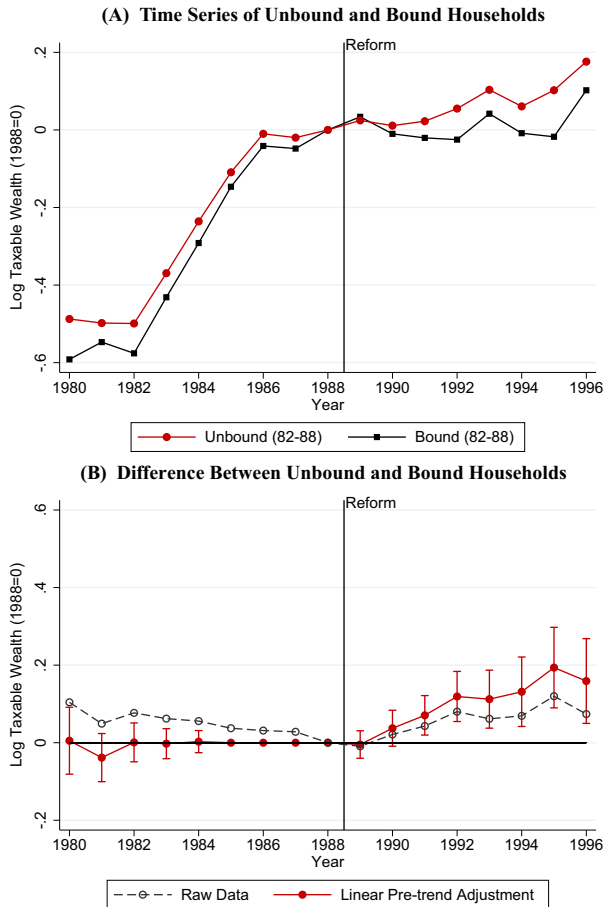


FIGURE VI

Difference-in-Differences Comparing Households Unbound and Bound by Tax Ceiling

This figure shows the effects of the 1989 wealth tax reform on taxable wealth based on the ceiling DD in which we compare households who are unbound by the tax ceiling (treatments) to those who are bound by the tax ceiling (controls). The assignment of treatment status is based on prereform variables, restricting attention to households whose status stays constant during 1982–88. The sample is a balanced panel of households observed in all years 1980–96 and located in the top 1% of the wealth distribution before the reform. Panel A shows the evolution of log taxable wealth in the two comparison groups, normalized to 0 in the prereform year 1988. Panel B shows the raw differences between these two series and the pretrend-adjusted differences (using [equation \(2\)](#)). The 95% confidence intervals are based on robust standard errors clustered at the household level.

four prereform years using specification (2). The adjusted DD series looks quite compelling: it features almost perfectly parallel pretrends in the decade leading up to the reform combined with a clear and growing divergence in the eight years following the reform.

Figure VII documents the persistence of ceiling status and converts the ITT estimates into TOT estimates. As shown in Panel A, the fraction treated in the treatment group is 100% during the prereform years (by construction) and falls only slightly after the reform, while the fraction treated in the control group starts from 0% and increases gradually after the reform. The control group is less persistent in this case, because it is more common for those bound by the ceiling to become unbound due to wealth and income shocks than it is for unbound taxpayers to become bound. In the last year of the postreform period, the difference in treatment intensity is a little more than 50%. When converting the ITT effects into TOT effects in Panel B, we estimate a treatment effect on log taxable wealth equal to 0.312, an increase in wealth of about 30%.

Online Appendix Figure A.XIV splits the total effect on wealth into behavioral and mechanical effects. The method is the same as for the couples DD: we calculate annual tax savings for the treatment group using a measure of counterfactual wealth and simulate cumulative tax savings assuming an annual return of 5%. The figure shows that the mechanical effects are larger for the ceiling DD than for the couples DD. The main explanation is that the ceiling approach captures responses by wealthier households, that is, households located farther from the exemption threshold. As a result, the reform-induced change in the average after-tax return is larger in this sample. We find that the mechanical effect on log taxable wealth equals 0.068 after eight years, corresponding to 22% of the total effect.

The Online Appendix provides robustness checks similar to those shown for the couples DD. Figure A.XV explores the implications of using different prereform windows to define treatment status. It is apparent that for the ceiling DD, specifying a longer treatment window ensures more parallel pretrends. Still, even though the pretrends differ across specifications, they all show clear evidence of behavioral responses: the treatment group is on a flatter trend before the reform and a steeper trend after the reform. The specifications with shorter treatment windows (in particular, the 1986–88 window in Panel D) imply larger behavioral

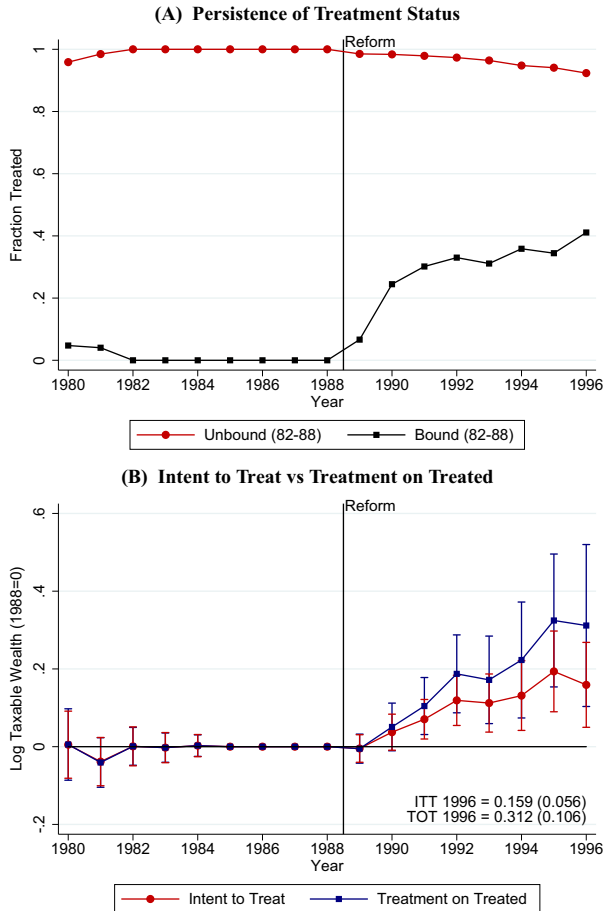


FIGURE VII

Ceiling DD: Treatment Effect on the Treated

This figure converts intention-to-treat (ITT) estimates into treatment-on-the-treated (TOT) estimates based on the ceiling DD in which we compare households bound and unbound by the tax ceiling. The assignment to treatment and control groups is based on prereform variables, restricting attention to households whose status stays constant during 1982–88. Panel A shows the persistence in treatment status over time. By construction, unbound households have a treatment status of 100% before the reform, while bound households have a treatment status of 0% before the reform. After the reform, households may switch status due to wealth or income shocks that change their tax ceiling status. Panel B compares the ITT and TOT series, where the TOT estimates are obtained by instrumenting treatment status in [equation \(2\)](#) using treatment status in the prereform years 1982–88.

responses than those reported above once we adjust for pretrends. We prefer the specification with a longer treatment window, because it ensures better pretrends in the raw data and is relatively conservative.

[Online Appendix](#) Figure A.XVI shows placebo tests based on assuming the reform happened in earlier years. The analysis is done in the same way as the corresponding analysis for the couples DD. Overall, the placebo tests look quite compelling. In all four panels, there is no significant effect of the placebo reform in 1988, the last year before the actual reform.

III.E. Summary of DD Estimates

[Table II](#) shows DD estimates of taxable wealth responses to the 1989 reform and converts these estimates into elasticities. The columns refer to the different quasi-experimental specifications: the couples DD and the ceiling DD, with and without adjusting for pretrends. The estimates without pretrend adjustment in columns (1), (3), and (5) are based on specification (1), whereas the estimates with pretrend adjustment in columns (2), (4), and (6) are based on specification (2). For each specification, we show both ITT and TOT effects. As described in [Section III.A](#), the ITT effects are obtained from a reduced-form specification in which log wealth is regressed directly on the instruments (constructed from prereform behavior), and the TOT effects are obtained from an IV specification. We report both the average effect over the postreform window (1989–96) and the effect in the last postreform year (1996). Although it is standard to show the average effect, the last-year effect is arguably more informative for a dynamic outcome like the stock of wealth. Still, the “last-year effect” shown here does not correspond to the long-run effect as we show in the structural analysis later. Finally, the table converts the average effects on log wealth into elasticities using the definition in [equation \(3\)](#). We show elasticities with respect to the net-of-tax rate, $1 - \tau$, and with respect to the net-of-tax rate of return, $(1 - \tau)R - 1$, assuming a gross rate of return of $R = 1.05$.

The absolute effects on log wealth vary considerably across the different strategies/samples, but so does the underlying tax variation driving them. For example, if we consider TOT effects adjusted for pretrends, the average effect equals 0.060 log points when comparing couples and singles within the exempted range, and it equals 0.135 log points when comparing couples within and

TABLE II
DIFFERENCE-IN-DIFFERENCES ESTIMATES OF TAXABLE WEALTH RESPONSES TO THE 1989 WEALTH TAX REFORM

	Couples DD			Ceiling DD		
	(1)	(2)	(3)	(4)	(5)	(6)
ITT effects on log wealth:						
Average effect	0.057 (0.008)	0.038 (0.009)	0.014 (0.005)	0.064 (0.006)	0.057 (0.021)	0.102 (0.033)
1986 effect	0.096 (0.016)	0.068 (0.017)	0.021 (0.011)	0.100 (0.012)	0.074 (0.034)	0.159 (0.056)
TOT effects on log wealth:						
Average effect	0.090 (0.012)	0.060 (0.014)	0.031 (0.011)	0.135 (0.013)	0.095 (0.034)	0.171 (0.052)
1986 effect	0.186 (0.029)	0.133 (0.033)	0.058 (0.029)	0.277 (0.032)	0.144 (0.066)	0.312 (0.106)
Elasticities:						
Elasticity wrt. $1 - \tau$	8.859 (1.189)	5.910 (0.813)	1.092 (0.288)	5.147 (0.311)	6.357 (0.910)	11.325 (1.138)
Elasticity wrt. $(1 - \tau)R - 1$	0.345 (0.048)	0.232 (0.033)	0.050 (0.010)	0.217 (0.010)	0.219 (0.030)	0.393 (0.038)
Observations	104,839	104,839	182,835	182,835	94,503	94,503
Household and year FE	X	X	X	X	X	X
Linear pretrends		X		X		X

Notes. This table shows the intent-to-treat (ITT) and treatment-on-the-treated (TOT) effects of the 1989 wealth tax cuts on log wealth and the implied elasticities of taxable wealth. The table displays the average effect over the postreform period (1989-96) and the last-year effect (1996). The columns refer to the different empirical strategies: the couples DD (either couples versus singles within the exempted range or couples within versus below the exempted range) and the ceiling DD (unbound versus bound taxpayers). In all specifications, the assignment to treatment and control groups is based on preretirement variables, restricting attention to households whose status stays constant during 1982-88. Columns (1), (3), and (5) show raw DD estimates without any pretrend adjustment (specification (1)), and columns (2), (4), and (6) show DD estimates adjusted for linear differential pretrends (specification (2)). For the couples DD using singles as controls, our preferred specification is the one without pretrend adjustment (column (1)). For the couples DD using couples below the exempted range as controls and for the ceiling DD, our preferred specifications are those with pretrend adjustment (columns (4) and (6)). The estimation sample is a balanced panel of households observed in all years 1980-96. Robust standard errors are clustered at the household level.

below the exempted range. The last-year effects equal 0.133 log points and 0.277 log points, respectively. However, because the tax variation in the second strategy is larger than in the first strategy, the implied elasticities of taxable wealth are about the same. In both cases, the elasticity with respect to the after-tax rate of return is just above 0.2. Turning to the ceiling DD, the effects are larger both in absolute terms and in elasticity terms, but here we are considering a different population of very wealthy taxpayers. The TOT effect on their wealth equals 0.171 log points on average and 0.312 log points in the last year. The elasticity with respect to the after-tax rate of return is about 0.4.²⁶

IV. THE EFFECT OF WEALTH TAXES: THEORY

IV.A. Life Cycle Model with Utility of Residual Wealth

In this section we develop a model for studying the effects of wealth taxation on the wealthy. Our goal is to construct a model that is sufficiently simple to derive analytical results, but at the same time rich enough to facilitate interpretation of the empirical results and allow for informative calibration exercises. To understand what the key features of such a model should be, we highlight two empirical facts regarding the wealthy. First, as mentioned earlier, wealthy people tend to be older people. Almost 80% of those in the top 1% of the wealth distribution are above age 50, as opposed to only 31% in the general population. Second, wealthy people continue to accumulate wealth into very old age and therefore die with large amounts of wealth. This is documented in detail in the next section.

To match these empirical facts, our model incorporates utility of residual wealth. This may be interpreted as capturing a bequest motive—and we refer to it as such—but it may also capture other utility-of-wealth motivations (see, [Saez and Stantcheva 2018](#) for a discussion of different mechanisms). The specific mechanism is

26. [Online Appendix](#) Tables A.I–A.II show heterogeneity in the estimated wealth responses by age (below and above the median age in each estimation sample). It would be natural to also study heterogeneity by the presence and number of children, especially considering that wealth responses by wealthy taxpayers may be partly motivated by bequest motives. Such an analysis is not feasible, however, because parents cannot be linked to children born before 1960 in the Danish data. Most people liable to pay wealth tax (i.e., older people) in the 1980s and 1990s would have had children before 1960.

not important for our purposes.²⁷ Although our model accounts for the bequest motive as well as the standard life cycle motive for saving, it abstracts from precautionary savings and uncertainty. The precautionary savings motive matters for the lower tail of the distribution, but it is second order for understanding savings behavior at the top of the wealth distribution (see, e.g., [Carroll 2002](#); [De Nardi 2004](#)).²⁸

Households live for T periods and their preferences are specified as follows

$$(4) \quad \frac{\sigma}{\sigma - 1} \sum_{t=0}^T \delta^t (c_t)^{\frac{\sigma-1}{\sigma}} + \delta^T V(W_{T+1}),$$

where c_t is consumption in period t , W_{T+1} is wealth at the end of life (bequests), σ is the elasticity of intertemporal substitution (EIS), and δ is the discount factor. To capture utility of bequests, we adopt the following parameterization

$$(5) \quad V(W_{T+1}) = A \frac{\alpha}{\alpha - 1} \left(\frac{W_{T+1}}{A} \right)^{\frac{\alpha-1}{\alpha}},$$

where A determines the strength of the bequest motive (under $A = 0$ the model corresponds to the pure life cycle model) and α is a bequest elasticity. This is a warm-glow bequest motive as

27. We use the model (together with the quasi-experimental moments) to estimate the long-run responsiveness of wealth accumulation to wealth taxation. This depends on the curvature of utility from wealth (which we estimate structurally), but it does not depend directly on the specific reason for utility from wealth. On the other hand, the specific reason may matter for normative policy analysis. For example, utility of wealth due to warm-glow of bequests will in general have different optimal tax implications than utility of wealth due to social status, because the former is associated with positive externalities (calling for Pigouvian subsidies) and the latter is associated with negative externalities (calling for Pigouvian taxes). However, in either case, the long-run elasticity of capital supply that we estimate is a key parameter, because it determines the fiscal externality against which we would trade off the potential benefits from redistribution and externalities.

28. We also abstract from labor supply responses to wealth taxes. Existing evidence from Denmark suggests that labor supply is relatively inelastic to labor taxes ([Kleven and Schultz 2014](#); [Kleven 2014](#)), suggesting that labor supply is also inelastic to capital taxes and perhaps especially for the population of older, wealthy people studied here. Furthermore, our explorations of the data shows no evidence of labor supply responses to wealth taxation when using the same empirical strategies (ceiling DD and couples DD) as those used for estimating savings responses.

introduced by Andreoni (1989, 1990) and used for studying estate taxation by for example Farhi and Werning (2010), Piketty and Saez (2013), and Kopczuk (2013a). For simplicity of exposition, we abstract from estate taxes (as our focus is on wealth taxes rather than on wealth transfer taxes) and model warm glow as a function of gross wealth.

In each period, there is a tax rate τ on household wealth above an exemption threshold \bar{W} . For someone with wealth above the exemption threshold in period t , the budget constraint is given by

$$(6) \quad \begin{aligned} c_t &= y_t + RW_t - \tau R(W_t - \bar{W}) - W_{t+1} \\ &= y_t + (1 - \tau)RW_t + \tau R\bar{W} - W_{t+1}, \end{aligned}$$

where y_t is (exogenous) labor income net of income tax, W_t is wealth at the beginning of the period, and R is the gross rate of return. We assume that R is time invariant, but this is straightforward to generalize and has no important implications for our results. The second line of the budget constraint (6) is a “virtual income” representation: it writes the budget as if the net-of-tax return equals $(1 - \tau)R$ on all units of wealth but provides a lump-sum income of $\tau R\bar{W}$ to compensate for the fact that the tax is not paid below the threshold. Combining all the per period budget constraints, we can express the lifetime budget constraint as

$$(7) \quad \begin{aligned} &\sum_{t=0}^T \frac{c_t}{((1 - \tau)R)^t} + \frac{W_{T+1}}{((1 - \tau)R)^T} \\ &= \sum_{t=0}^T \frac{y_t}{((1 - \tau)R)^t} + \sum_{t=1}^T \frac{\tau R\bar{W}}{((1 - \tau)R)^t} + W_0^n, \end{aligned}$$

where $W_0^n \equiv (1 - \tau)RW_0 + \tau R\bar{W}$ is initial (exogenous) wealth after tax.

Households maximize lifetime utility equations (4)–(5) subject to the lifetime budget constraint (7) with respect to consumption and bequests. The first-order conditions for c_t and c_{t+1} yield the standard Euler equation,

$$(8) \quad c_{t+1} = (\delta(1 - \tau)R)^\sigma c_t,$$

while the first-order conditions for W_{T+1} and c_T give

$$(9) \quad W_{T+1} = A c_T^{\frac{\alpha}{\sigma}}.$$

The solution to the model is described by the lifetime budget equation (7), the Euler equations (8) for all t , and the bequest condition (9). These conditions determine c_0, \dots, c_T and W_{T+1} . Wealth W_t in any given period can then be backed out using the per period budget constraints.

Using the Euler equations in each period, we can write consumption in period t and bequests in terms of consumption in period 0, that is,

$$(10) \quad c_t = (\delta(1-\tau)R)^{t\sigma} c_0,$$

$$(11) \quad W_{T+1} = A(\delta(1-\tau)R)^{T\alpha} c_0^{\frac{\alpha}{\sigma}}.$$

Inserting these conditions into constraint (7), we can express the choice of c_0 as

$$(12) \quad \sum_{t=0}^T q_t \cdot c_0 + q_b \cdot c_0^{\frac{\alpha}{\sigma}} = \sum_{t=0}^T \frac{y_t}{((1-\tau)R)^t} + \sum_{t=1}^T \frac{\tau R \bar{W}}{((1-\tau)R)^t} + W_0^n,$$

where $q_t \equiv \frac{(\delta(1-\tau)R)^{t\sigma}}{((1-\tau)R)^t}$ denotes present-value expenditures on consumption in period t relative to period 0, and $q_b \equiv \frac{A(\delta(1-\tau)R)^{T\alpha}}{((1-\tau)R)^T}$ denotes present-value expenditures on bequests relative to consumption in period 0. This expression is useful for characterizing the effects of wealth taxes.

IV.B. The Effect of Wealth Taxes

Consider a permanent change in the wealth tax rate, $d\tau$, holding the exemption threshold \bar{W} constant. The tax change is announced in period 0 and may affect wealth from the end of this period, W_1 . Initial after-tax wealth W_0^n is predetermined. We investigate the effect on households that are above the threshold \bar{W} (and stay above it over time), as opposed to the effect on households that are sometimes below and sometimes above the threshold over their lifetime. The former scenario is simpler to analyze, and it fits our quasi-experimental setting in which we estimate responses by households above the exemption threshold. The potential response by those who are below the exemption threshold

but expect to rise above it in the future is not captured by our empirical design and would be very hard to estimate in general.

We characterize analytically how the reduced-form effect of changing the wealth tax rate—what we have estimated empirically—relates to the structural parameters of the model. We start by deriving the effect of taxes on first-period wealth W_1 , and then show how the effect accumulates over time. The effect of taxes includes both substitution and wealth effects. To characterize the wealth effect, it is useful to define the amount of initial resources a household would have to receive to be able to afford an unchanged bundle of consumption and bequests when the net-of-tax return changes. This can be obtained by differentiating the lifetime budget constraint (7) with respect to $1 - \tau$, holding behavior $\{c_t\}_0^T$, W_{T+1} constant but allowing initial wealth to adjust. We denote this compensating change in initial wealth by dW_0^C .

We may state the following proposition:

PROPOSITION 1 (FIRST-PERIOD REDUCED-FORM EFFECT). Consider a permanent change in the wealth tax rate τ from period 0 onward. The reduced-form elasticity of first-period wealth W_1 with respect to the net-of-tax rate $1 - \tau$ can be expressed as

$$\begin{aligned}
 \frac{dW_1}{d(1-\tau)} \frac{1-\tau}{W_0} = & \sigma \cdot \left\{ \frac{\sum_{t=0}^T tq_t}{\sum_{t=0}^T qt + qb \frac{\alpha}{\sigma} c_0^{\frac{\alpha}{\sigma}-1} W_0} \frac{c_0}{W_0} \right\} \\
 & + \alpha \cdot \left\{ \frac{Tqb}{\sum_{t=0}^T qt + qb \frac{\alpha}{\sigma} c_0^{\frac{\alpha}{\sigma}-1} W_0} \frac{c_0^{\frac{\alpha}{\sigma}}}{W_0} \right\} \\
 (13) \quad & + \frac{dW_0^C}{d(1-\tau)} \frac{1-\tau}{W_0} \cdot \left\{ \frac{1}{\sum_{t=0}^T qt + qb \frac{\alpha}{\sigma} c_0^{\frac{\alpha}{\sigma}-1} W_0} \right\},
 \end{aligned}$$

where $q_t \equiv \frac{(\delta(1-\tau)R)^{t\sigma}}{((1-\tau)R)^t}$, $qb \equiv \frac{A(\delta(1-\tau)R)^{T\alpha}}{((1-\tau)R)^T}$, and $dW_0^C \leq 0$ is a compensating wealth change allowing the household to afford an unchanged bundle of consumption and bequests when $1 - \tau$ changes. The first term is a substitution effect on consumption (positive), the second term is a substitution effect on bequests (positive), and the third term is the wealth effect (negative).

Proof: See [Online Appendix C](#). \square

This result shows that the reduced-form elasticity of wealth (one period after the tax change) is an involved function of all the parameters of the model. There are three qualitative effects on wealth accumulation. First, there is a substitution effect on consumption. A larger net-of-tax return induces households to shift consumption to later in life, thereby increasing wealth accumulation. This effect is proportional to the EIS σ .²⁹ Second, there is a substitution effect on bequests. A larger net-of-tax return reduces the price on bequests, further increasing wealth accumulation. This effect is proportional to the bequest elasticity α , and it depends on the weight of the bequest motive A . The bequest effect vanishes when A (and therefore q_b) goes to 0, but can be important when A is large. We show later that A has to be very large to rationalize the life cycle profiles of wealthy people, which puts the bequest elasticity α at center stage. Finally, there is the wealth effect. A larger net-of-tax return increases lifetime resources, which leads to larger consumption and lower savings. The presence of the wealth effect implies that the total reduced-form effect is ambiguous in sign.

A simplifying special case is where bequests and consumption goods are equally elastic, that is, $\alpha = \sigma$. This is a natural benchmark assumption if bequests are viewed as future consumption (for the next generation). With $\alpha = \sigma$, we obtain the following result:

PROPOSITION 2 (FIRST-PERIOD REDUCED-FORM EFFECT SIMPLIFIED).

Assuming $\alpha = \sigma$ (bequest elasticity equals the EIS), the reduced-form elasticity of first-period wealth W_1 with respect to the net-of-tax rate $1 - \tau$ simplifies to

$$(14) \quad \frac{dW_1}{d(1-\tau)} \frac{1-\tau}{W_0} = \sigma \cdot \left\{ \frac{\sum_{t=0}^T tq_t + T Aq_T}{\sum_{t=0}^T q_t + Aq_T} \frac{c_0}{W_0} \right\} + \frac{dW_0^C}{d(1-\tau)} \frac{1-\tau}{W_0} \cdot \left\{ \frac{1}{\sum_{t=0}^T q_t + Aq_T} \right\},$$

where $q_t \equiv \frac{(\delta(1-\tau)R)^{t\sigma}}{((1-\tau)R)^t}$ and $dW_0^C \leq 0$ is the compensating wealth change allowing the household to afford an unchanged bundle of consumption and bequests when $1 - \tau$ changes. The

29. Our wording is somewhat loose here. The effect is only proportional in σ when taking the complicated term in braces (which itself depends on σ) as given.

first term is the substitution effect on both consumption and bequests (positive), and the second term is the wealth effect (negative).

Proof: Follows from setting $\alpha = \sigma$ (and thus $q_b = Aq_T$) in equation (13). \square

This result provides a simpler characterization of the reduced-form elasticity. There is only one substitution term (capturing both consumption and bequest responses) and it is proportional to the structural elasticity $\sigma = \alpha$. However, we do not rely on the assumption of $\sigma = \alpha$ in the structural estimation presented below but identify these two parameters separately.

The one-period effect derived above is helpful for establishing economic intuition, but our empirical analysis provides estimates over more than one year. We have estimates of the reduced-form elasticity $\frac{dW_t}{d(1-\tau)} \frac{1-\tau}{W_t}$ in each year t over eight years. How does the theoretical effect of wealth taxes accumulate over time? We consider the effect in any period t as a function of the effect in period 1 provided above.

From the per period budget constraint (6), we have

$$W_t = y_{t-1} + (1 - \tau)RW_{t-1} + \tau R\bar{W} - c_{t-1}.$$

We may substitute out W_{t-1} using the one-period lagged equation, and then substitute out W_{t-2} using the two-period lagged equation, and so on. This process allows us to write

$$(15) \quad W_t = [(1 - \tau)R]^{t-1} W_0^n + \sum_{j=0}^{t-1} (y_j - c_j) [(1 - \tau)R]^{t-1-j} + \sum_{j=1}^{t-1} \tau R\bar{W} [(1 - \tau)R]^{t-1-j}.$$

Thus, W_t equals the sum of initial net wealth W_0^n with returns, t periods of savings $y_j - c_j$ with returns, and the virtual income adjustment with returns. Using this expression, we may state the following proposition:

PROPOSITION 3 (PERIOD- t REDUCED-FORM EFFECT). Consider a permanent change in the wealth tax rate τ from period 0 onward.

The reduced-form elasticity of wealth in period t , W_t , with respect to the net-of-tax rate $1 - \tau$ can be expressed as

$$(16) \quad \frac{dW_t}{d(1-\tau)} \frac{(1-\tau)}{W_t} = dM_t + dB_t,$$

where dM_t is a mechanical effect given by

$$(17) \quad \begin{aligned} dM_t = & (t-1)[(1-\tau)R]^{t-1} \frac{W_0^n}{W_t} \\ & + \sum_{j=0}^{t-1} (t-1-j)[(1-\tau)R]^{t-1-j} \frac{y_j - c_j}{W_t} \\ & - \sum_{j=1}^{t-1} \left(1 - \frac{\tau(t-1-j)}{1-\tau}\right) [(1-\tau)R]^{t-j} \frac{\bar{W}}{W_t}, \end{aligned}$$

and dB_t is a behavioral effect given by

$$(18) \quad dB_t = \sum_{j=0}^{t-1} \frac{q_j}{[(1-\tau)R]^{1-t}} \left\{ \frac{dW_1}{d(1-\tau)} \frac{(1-\tau)W_0}{W_0} \frac{1}{W_t} - j\sigma \frac{c_0}{W_t} \right\}.$$

In [equation \(18\)](#), the term in braces, $\frac{dW_1}{d(1-\tau)} \frac{(1-\tau)W_0}{W_0}$, is characterized in [Proposition 1](#) for the general case and in [Proposition 2](#) for the special case of $\alpha = \sigma$.

Proof: See [Online Appendix C](#). \square

Although the first-period effect consists only of behavioral terms, the t -period effect consists of both behavioral and mechanical terms. The mechanical term reflects that as $1 - \tau$ increases, the individual earns larger net-of-tax returns on initial wealth and savings in each period. This increases wealth over time, holding consumption and savings behavior fixed. As a result, estimating significant reduced-form effects on wealth does not necessarily imply that people are elastic in their consumption or bequest behavior. However, as we saw in the previous section, the mechanical effect is a minor fraction of the total effect in our empirical setting due to the progressive design of the wealth tax. In addition to the mechanical effect, there will be behavioral effects as captured by the expression in [equation \(18\)](#). These consist of the one-period

effect derived earlier (substitution and wealth effect) accumulating over time, along with additional substitution effects on consumption from period 1 onward.

What have we learned from this theoretical exercise? First, the elasticity of wealth with respect to taxes is very far from being a structural parameter. It is endogenous to all the parameters of the dynamic setting, and its size depends mechanically on the time horizon of the estimate. Second, the magnitude of any reduced-form elasticity is most naturally assessed in terms of the structural elasticities of intertemporal substitution (σ) and bequests (α) implied by it. This requires a theoretical model and taking a stand on other parameters. It is in general possible to observe “large” effects under modest structural elasticities due to the mechanical effects. Third, the theoretical characterization allows us to calibrate the model to quasi-experimental moments obtained over the short to medium run and then use the calibrated model to assess long-run effects. Such long-run simulations would rely on parametric assumptions, but in a way that respects shorter-run nonparametric moments. Given the gradual, dynamic nature of wealth accumulation, it would be difficult (or impossible) to capture the long-run effects without a parametric model.

Finally, given our objective of estimating elasticities that can be used to assess optimal capital taxation, it is useful to relate our analysis to classic results calling for zero capital taxation. The result of [Chamley \(1986\)](#) and [Judd \(1985\)](#) that capital should be untaxed in steady state arises because in their models, the steady-state elasticity of capital supply is infinite (see, e.g., [Auerbach and Hines 2002](#); [Piketty and Saez 2013](#); [Saez and Stantcheva 2018](#)). Although we have characterized finite elasticities of capital supply, note that our model nests a version of the Chamley-Judd framework when assuming infinite horizons ($T \rightarrow \infty$) and no utility of residual wealth ($A \rightarrow 0$). The elasticity in steady state is then infinite and the optimal capital tax rate is 0. It would be next to impossible to evaluate the infinite-elasticity prediction in the data, because the steady-state equilibrium is not observed. Moreover, as shown by [Straub and Werning \(forthcoming\)](#), the steady-state prescription may not be very relevant as convergence to the steady state can be extremely slow in the Chamley-Judd framework, potentially taking centuries. We take a different approach, matching our model with a finite horizon and utility of residual wealth to estimated responses over the short to medium run to estimate a (finite) elasticity in the long run. This elasticity can be used to

assess optimal capital taxation within this general class of models (including, e.g., [Saez and Stantcheva 2018](#)) but does not by itself rule out the Chamley-Judd steady-state prediction.

V. CONNECTING THEORY AND EVIDENCE: LONG-RUN EFFECTS

V.A. Empirical Life Cycle Profiles of Wealth

To calibrate our model, we start by studying the empirical life cycle profiles of wealth at the top of the distribution. Besides informing our calibration exercise, these life cycle profiles will provide insights on the savings behavior of the rich and contribute to an area where we have relatively little evidence. Because we have household-level information on wealth for the full population over a long time period, we are able to present particularly clean and striking evidence.

[Figure VIII](#) shows life cycle profiles of wealth between the ages of 20 to 90 for the full population (Panel A) and for the top percentiles of the population (Panel B).³⁰ We highlight the following key points regarding the construction of the figure. First, the graphs show profiles of taxable wealth and therefore do not include pension wealth. Pension wealth is not very important at the top of the distribution (which is our main interest), but it is significant when considering the full population. Second, the graphs show profiles of normalized log wealth. Denoting log wealth for individual i at age a in year t by $\log(W_{iat})$, we define normalized log wealth as $\omega_{iat} \equiv \log(W_{iat}) - E[\log(W_{iat}) | t]$. That is, we normalize log wealth for each individual in each year by the average wealth in the population in that year, so that the life cycle profiles are not confounded by asset price inflation.³¹ Third, we consider unbalanced panels of individuals, because this allows us to show very wide age ranges. We consider balanced panels in narrower age ranges later. Finally, when showing individuals in the top percentiles of the distribution, we select individuals who are in the top $p\%$ at some point during their lives, but not necessarily in

30. Here we consider individuals (rather than households) as the units of analysis, but split household wealth of married couples equally between the spouses.

31. This normalization implies that the mean of ω_{iat} in each year t equals 0. The mean of ω_{iat} in each age bin is calculated as $E[E[\omega_{iat} | a, t] | a]$, that is, we first calculate means in bins of age a and year t , and then we calculate means in age bins over all years. This gives different calendar years the same weight in the calculated means.

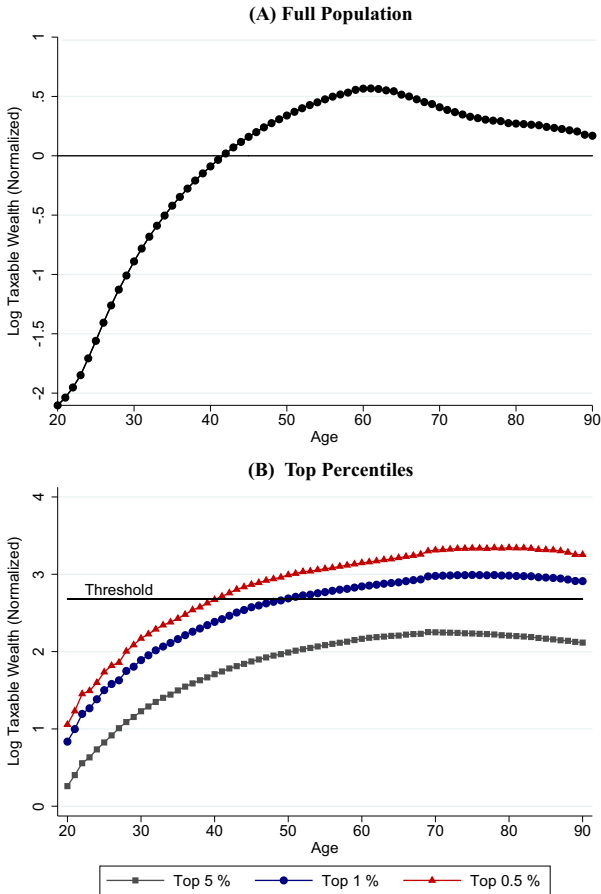


FIGURE VIII

Empirical Life Cycle Profiles of Wealth

This figure shows life cycle profiles of taxable wealth between the ages of 20 and 90 in an unbalanced panel of individuals over the period 1980–2012. To eliminate the confounding effects of inflation as people grow older, we normalize log wealth for each individual in each year by the average log wealth in the population in that year. The graphs show averages of this normalized wealth measure in different age bins. Panel A shows the life cycle profile in the full population of individuals with positive net wealth, and Panel B shows life cycle profiles in the top percentiles of the population. The top-percentile samples include individuals who are in the top $p\%$ for at least three years of their observed life span, keeping them in the data for their entire observed life span.

every year. That is, we do not condition the sample on being in the top $p\%$ at each age, but allow them to build up wealth gradually and draw down wealth at the end of life. This way of selecting the sample is more informative for understanding the life cycle savings behavior of the wealthy.³²

The following key findings emerge from [Figure VIII](#). The average person in the population (Panel A) accumulates wealth until just after age 60 and draws down wealth thereafter. In other words, the average person reaches her wealth peak around the age of retirement, consistent with the predictions of a pure life cycle model without any bequest motive or utility of wealth.³³ At the same time, wealth is still higher than average wealth in the population (the horizontal line at 0) at age 90, suggesting that the pure life cycle savings motive is not the only factor at play.

When turning to the wealthiest segments of the population (Panel B), the picture changes dramatically. Wealthy individuals tend to accumulate wealth through most of their lifetime. There is no draw-down of wealth until after the age of 80, and even then the draw-down is only marginal. For example, those who reach the top 1% of the wealth distribution during their lifetime surpass the exemption threshold for the wealth tax (demarcated by the horizontal line) around age 50 and stay well above it until age 90. [Online Appendix Figure A.XVII](#) shows the same type of graph but with wealth percentiles (instead of amounts) on the y-axis and including the extreme top of the distribution. There we see that those who reach the top 1% during their lifetime are, on average, located above the 98th percentile cutoff at the age of 90. Those who reach the top 0.5% are located at the 99th percentile cutoff at age 90, and those who reach the top 0.1% are located at the 99.8th percentile cutoff at age 90. There must be some form of utility of residual wealth (due to a bequest motive or another mechanism) to rationalize these empirical patterns.

Have these life cycle patterns changed over time? Specifically, have they changed from before to after the 1989 reform? [Online Appendix Figure A.XVIII](#) investigates this question by comparing the life cycle profiles of wealth before the reform (1980–88) and after the reform (1989–96). In this figure, we focus on individuals

32. However, we do condition the sample on being in the top $p\%$ for at least three years so as to reduce noise from transitory wealth shocks.

33. This pattern would be reinforced by including pension wealth, because in general people pay into such schemes during their working lives and draw benefits during retirement.

in the top 1% of the wealth distribution, selected in the same way as the top 1% in the previous graph. We focus on the age range 60–90 during which the top 1% is always above the wealth tax threshold on average. The series are otherwise constructed in the same way as in the previous figure. We see that there are no striking changes in the life cycle profile over time: both profiles stay quite flat into very old age and provide no indication that the wealth tax cuts increased wealth accumulation (or reduced draw-down) for people close to death. This time series comparison cannot be taken as causal evidence of the effect of the reform, but we do use one aspect of it in the calibration below. Specifically, in the calibration of the bequest motive, we use the fact that wealth accumulation in the very last years of life (after the age of 80) is unchanged from before to after the reform.

A potential concern with these graphs, especially in terms of interpreting what happens at very old ages, is that they are based on unbalanced samples. People drop out of the samples as they die, which may affect the wealth profiles due to selection on mortality. As we move out into the tail of the age distribution, we are increasingly considering people who live long, and such people may tend to be wealthier. If so, this would understate the within-person wealth draw-down at the end of life. To investigate this issue, [Online Appendix](#) Figure A.XIX shows life cycle profiles of wealth for the top 1% in a balanced sample. Panel A compares wealth profiles in balanced and unbalanced samples between 70–90 years of age, and Panel B compares wealth profiles before the reform (1980–88) and after the reform (1989–96) in the balanced sample. Panel A shows that the balanced sample does feature stronger wealth draw-down at the end of life, consistent with some selection on mortality. However, the extent of this draw-down is small. Moreover, it is driven largely by changes after the wealth tax abolishment in 1997, which can be seen from Panel B. This graph shows that both before and after the 1989 reform—but within the wealth tax period 1980–96—the wealth profile among the very old is almost completely flat even in the balanced sample. Overall, considering balanced samples does not change the fundamental insights from the unbalanced samples.

V.B. Calibration

To study the long-run effects of wealth taxes on the wealthy, we calibrate the model from [Section IV](#) to fit the empirical life

cycle profile of wealth at the top along with the quasi-experimental estimates of the impact of wealth tax reform. Because our theoretical model is a representative agent framework, we calibrate it to fit the average wealth profile in specific wealth ranges that vary by experiment. In particular, the quasi-experimental analysis considered two different empirical strategies and samples: the couples DD gave treatment effects roughly between the 98th and 99th percentile cutoffs, while the ceiling DD gave treatment effects above the 99th percentile cutoff. The life cycle profiles in [Online Appendix Figure A.XVII](#) show that the couples DD is captured well by the “top 1% sample” (i.e., the sample of those who reach the top 1% at some point in their lives) during the age range from 60 to 90. The figure also shows that the ceiling DD is captured well by the “top 0.3% sample” during the same age range. Therefore, we show calibrations for each of these experiments and samples over a 30-year life span, where age 61 corresponds to period $t = 0$ and age 90 corresponds to period $t = T$. Wealth at death, W_{T+1} , is thus defined as wealth at the end of the 90th year.

To fit the empirical life cycle profile of wealth in the baseline with $\tau = 0.022$, we set initial wealth W_0 equal to observed wealth at age 60 and calibrate the bequest parameter A so that end-of-life wealth W_{T+1} equals observed wealth at age 90. The calibration of A uses the optimal bequest condition (11) under $\tau = 0.022$. Thus, our calibrations ensure that the baseline wealth path matches the observed start and end points over the considered age range. To capture the shape of the wealth profile between periods 0 and T , we calibrate the discount factor δ given a reasonable value of the gross rate of return R . Specifically, when setting $R = 1.05$, the calibrated value of δ is around 0.97.³⁴ As we shall see, our parsimonious model is able to fit the empirical life cycle profile very well.

We calibrate the EIS σ and the bequest elasticity α to match the quasi-experimental moments from the analysis of the 1989 reform. The structural parameters are matched to the TOT effects, and we exploit the full dynamic pattern of those effects. We have estimates for each year between 1989 and 1996, giving us eight moments to identify two parameters. The reason the dynamic

34. Because we do not explicitly account for income taxes, R should be interpreted as the gross rate of return after the taxation of capital income. Even so, $R = 1.05$ falls well within the range of estimated returns at the top of the wealth distribution (see [Fagereng et al. 2016](#)).

pattern of the estimates is informative is that the two elasticities σ , α have different implications for the time path of behavioral responses: the EIS channel is relatively strong in the short run, while the bequest elasticity channel is relatively strong in the long run. Therefore, the two elasticities determine, not just the overall magnitude of the effects, but also the concavity or convexity of the effects over time. When σ is larger (smaller) relative to α , the time path of effects is more concave (convex).

Specifically, we calibrate σ and α by minimizing a standard quadratic loss function, that is,

$$(19) \quad \mathcal{L}(\sigma, \alpha) = \sum_{t=1}^8 [\text{TOT}_t - \Delta \log W_t(\sigma, \alpha)]^2,$$

where TOT_t is the estimated treatment effect in year t and $\Delta \log W_t(\sigma, \alpha)$ is the model-predicted effect in year t under the structural primitives σ , α . Recall that we have TOT estimates of both the total effect and the behavioral effect, where the latter excludes the mechanical effect. We will match on the total effect, letting the mechanical and behavioral effects be “free variables.” In general, our model does not exactly match the mechanical effect in the data, in part because our representative agent approach does not explicitly model the different tax parameters for singles and couples. We consider a unique threshold \bar{W} and a unique tax rate τ , translating the tax reform (part of which changed the exemption threshold for couples relative to singles) into an implied average change in the tax rate $\Delta \tau$. To put it differently, although our empirical estimates are based on experiments that change tax rates partly through household-specific threshold changes, our calibration exercises translate this into a simpler experiment that changes the tax rate on a representative household.

Although we could estimate σ and α based solely on [equation \(19\)](#), we bring in an additional empirical moment to discipline the calibration. As we saw in the previous section, the age profile of wealth (among the wealthy) tends to be flat toward the end of life. Specifically, we showed that the wealth profile is roughly flat after the age of 80, and that this is true both before and after the wealth tax reform. Calibrations that imply strongly increasing or decreasing wealth during the last years of life do not seem reasonable in light of these facts. As a result, we require that the average wealth growth during the last 10 years of life remain

the same after the reform as before the reform. The estimation of σ and α is therefore based on minimizing [equation \(19\)](#) subject to this requirement on the wealth path during the last years of life.

To summarize, our calibration procedure consists of two interconnected steps. In the first step we calibrate one set of parameters (such as the weight on bequests A and the discount factor δ) to match the wealth profile in the baseline with $\tau = 0.022$, taking the structural elasticities σ and α as given. In the second step we calibrate σ and α to match the quasi-experimental estimates of the effects of tax reform, taking the parameters from the first step as given. We loop back and forth between these two steps until we converge to a fixed point where the elasticities found in the second step correspond to those used as inputs in the first step.

V.C. Simulating the Long-Run Effects of Wealth Taxation

We first consider the effects of wealth taxation on the moderately wealthy. That is, we calibrate the model to the estimates from the couples DD in which we compare couples and singles in the exempted range of the wealth distribution (between the 98th and 99th percentiles). The results are presented in [Figure IX](#). Panel A shows three wealth paths: the observed wealth path (dotted line) and the simulated wealths path before and after reform (solid lines). The simulated wealth path before the reform is calibrated to fit the observed wealth path. The simulated wealth path after the reform is based on a reduction of the wealth tax rate by 1 percentage point, corresponding to the differential tax cut between the comparison groups in the couples DD. Our calibration ensures that the differences between the before-reform and after-reform wealth paths respect the quasi-experimental estimates during the first eight years. This can be seen in Panel B, in which we compare the simulated effects on log wealth (solid line) to the estimated effects (dotted line). This panel also splits the total effect into the underlying mechanical and behavioral effects (dashed lines).

The following insights are worth highlighting. First, the effect of the wealth tax reduction on the stock of wealth grows for about 25 years and then stabilizes. At the end of life, wealth is 30% higher than it would have been absent the reform. Second, while this effect seems large, the underlying tax incentive driving it is also large. Defining the net-of-tax return

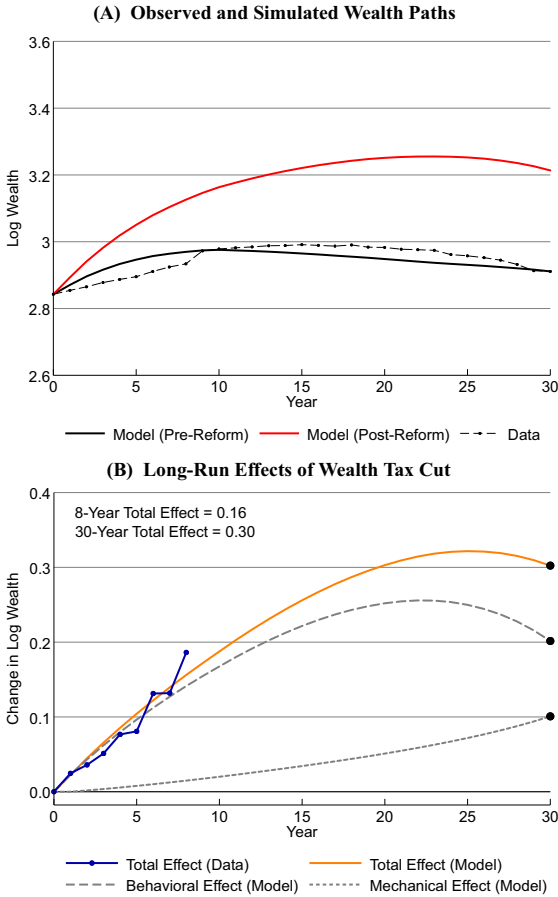


FIGURE IX

Long-Run Effects of Cutting Wealth Taxes: Couples DD (Moderately Wealthy)

This figure shows the long-run effects of wealth tax cuts when calibrating our model to the couples DD (comparing couples and singles in the exempted range). These are effects for the moderately wealthy (between the 98th and 99th percentile cutoffs). The reform experiment cuts the wealth tax rate by 1 percentage point, corresponding to the differential tax cut between the treatment and control groups. Panel A shows the observed life cycle profile of wealth, the simulated life cycle profile before the reform (calibrated to fit the empirical profile), and the simulated life cycle profile after the reform. Panel B illustrates the total effects, the mechanical effects, and the behavioral effects on taxable wealth over 30 years, demonstrating that the model matches the quasi-experimental estimates over the initial 8 years.

as $(1 - \tau)R - 1$, the percentage change in the return is equal to $-\frac{\Delta\tau \cdot R}{(1-\tau)R-1} = 39\%$. Therefore, the long-run elasticity of wealth with respect to the net-of-tax rate of return is 0.77. Third, the mechanical effect becomes increasingly important over time due to the compounding effects of a larger net-of-tax rate of return. After 30 years, the mechanical effect accounts for more than one-third of the total effect.

As a robustness check, [Online Appendix Figure A.XX](#) shows a calibration based on the couples DD in which we compare couples within the exempted range to those below. This comparison captures a larger experiment—a reduction of the wealth tax rate by 2.2 percentage points—and we therefore expect the effects to be larger. Indeed, the graph shows that the long-run effect on the stock of wealth equals 50%, much larger than in the baseline specification. Relating this effect to the change in the net-of-tax rate of return, $-\frac{\Delta\tau \cdot R}{(1-\tau)R-1} = 86\%$, the long-run elasticity of wealth is close to 0.6. It is reassuring that the two different couples DD strategies imply roughly similar long-run elasticities of wealth, despite being based on very different tax variation and reduced-form effects.

We now turn to the effects of wealth taxation on the very wealthy, calibrating the model to the ceiling DD and taxpayers within the top 1%. In this case, the experiment is a reduction in the marginal wealth tax by 1.45 percentage points, corresponding to the average reduction across treated taxpayers.³⁵ The results are presented in [Figure X](#). As before, the model does a good job of fitting the baseline wealth path to the observed path (Panel A) and the wealth responses to the quasi-experimental moments (Panel B). The effect of cutting wealth taxes on long-run wealth is equal to 65%. Given a percentage change in the net-of-tax return equal to $-\frac{\Delta\tau \cdot R}{(1-\tau)R-1} = 57\%$, the implied long-run elasticity of wealth equals 1.15. Slightly less than half of this effect is mechanical, and the rest is behavioral.

[Table III](#) summarizes the simulation results and shows the “structural primitives” underlying each calibration. The table also investigates the robustness of the results to the assumed rate of return R . The following points are worth noting. First, the total

35. While most treated taxpayers in the ceiling DD see their marginal tax rate reduced by 1.2 percentage points, some of them benefit from the increase in the exemption threshold for couples (to the 99.3rd wealth percentile) and therefore see their marginal tax rate reduced by 2.2 percentage points.

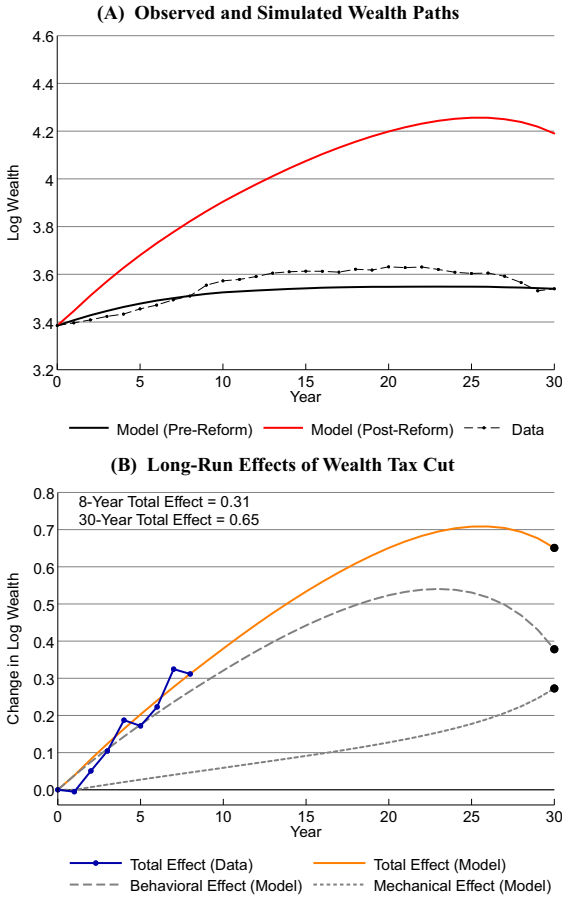


FIGURE X

Long-Run Effects of Cutting Wealth Taxes: Ceiling DD (Very Wealthy)

The figure shows the long-run effects of wealth tax cuts when calibrating our model to the ceiling DD. These are effects for the very wealthy (within the top 1%). The reform experiment cuts the wealth tax rate by 1.56 percentage points, corresponding to the tax cut for the average person in the treatment group. Panel A shows the observed life cycle profile of wealth, the simulated life cycle profile before the reform (calibrated to fit the empirical profile), and the simulated life cycle profile after the reform. Panel B illustrates the total effects, the mechanical effects, and the behavioral effects on taxable wealth over 30 years, demonstrating that the model matches the quasi-experimental estimates over the initial 8 years.

TABLE III
SIMULATION ANALYSIS OF LONG-RUN EFFECTS

	Couples DD				Ceiling DD	
	Couples vs singles		Couples within vs below		Unbound vs bound	High return
	Low return (1)	High return (2)	Low return (3)	High return (4)	Low return (5)	High return (6)
Panel A: 30-year effects on log wealth						
Total effect	0.302	0.313	0.499	0.509	0.651	0.639
Mechanical effect	0.101	0.142	0.240	0.354	0.273	0.419
Behavioral effect	0.202	0.170	0.259	0.155	0.378	0.219
Elasticity wrt. $(1 - \tau)R - 1$	0.774	1.359	0.581	1.005	1.150	1.913
Panel B: structural primitives						
R	1.05	1.07	1.05	1.07	1.05	1.07
δ	0.973	0.955	0.973	0.955	0.975	0.957
A	16.51	12.99	16.74	13.56	25.70	19.27
σ	2.62	2.25	2.20	1.83	6.65	4.36
α	2.14	1.86	1.50	1.31	3.81	2.83

Notes: This table summarizes the simulation results shown in Figures IX, X, and A.XX, and reports the structural primitives underlying each calibration. The table also investigates the sensitivity of our results to changing the assumed rate of return R : our baseline assumption of a 5% before-tax return is shown in the columns labeled "Low return," while the alternative assumption of a 7% before-tax return is shown in the columns labeled "High return."

effect on wealth is very robust to the assumed rate of return R , but it affects the decomposition into mechanical and behavioral effects. When the return is higher, more of the effect is mechanical and less is behavioral. Second, the rate of return has a big impact on the reduced-form elasticity with respect to $(1 - \tau)R - 1$, but this is a trivial implication of changing the denominator of the elasticity. In some sense, this highlights an issue with focusing on reduced-form elasticities in this context. Third, while the long-run elasticity of taxable wealth equals 0.58–0.77 for the moderately wealthy and 1.15 for the very wealthy (under $R = 1.05$), the underlying structural elasticities—the EIS σ and the bequest elasticity α —are larger in magnitude. For example, the EIS is around 2 for the moderately wealthy and even larger for the very wealthy. This is much larger than existing estimates of this parameter (see, e.g., [Best et al. forthcoming](#)).³⁶ When comparing the structural elasticities obtained here to those in the existing literature, two key aspects must be kept in mind. One is that we are considering behavioral responses by very wealthy households, something the existing literature has not been able to do. The other is that we are considering effects on taxable wealth, which will include both real savings responses and any evasion or avoidance responses. The larger elasticities among the very wealthy than among the moderately wealthy are consistent with larger avoidance at the top. Therefore, the elasticities we estimate represent upper bounds on real wealth accumulation responses.

VI. CONCLUSION

In this article we address one of the most important unanswered questions in public finance: what is the effect of capital

36. [Online Appendix](#) Figure A.XXI illustrates how the calibration of α and σ works. The two parameters are set to ensure that the model fits both the quasi-experimental moments during years 1,...,8 and the observed flatness of the wealth profile at the end of life (before and after the reform). While different combinations of the two parameters can provide a reasonable fit of the shorter-term effects (a lower α can be compensated for by a higher σ), the additional requirement on the end-of-life profile nails both parameters. As discussed in the previous section, the concavity or convexity of the quasi-experimental moments is in itself informative of the relative magnitudes of α and σ . In this particular case, however, our calibration is not very robust when using only those moments, because the time path of the estimates happens to be almost linear. If we had observed stronger concavity or convexity, the composition of the effect on α and σ would have been better identified.

taxation on capital supply? The answer to this question is critical for assessing the desirability of taxing capital income or wealth. There is an existing empirical literature studying different aspects of capital taxation—including a handful of papers on wealth taxation and many more on wealth transfer taxation—but there is no consensus on what might be a reasonable range for the elasticity of capital supply, particularly in the long run. [Saez and Stantcheva \(2018\)](#) show that this elasticity parameter provides a sufficient statistic for optimal capital taxation, but they do not cite any empirical evidence on its value. As discussed in the beginning, the lack of evidence in this area can be explained by a number of methodological difficulties. There are major empirical challenges related to both measurement and identification, as well as conceptual challenges related to modeling savings motives and the dynamic nature of wealth accumulation. Through a combination of quasi-experimental analysis based on administrative wealth records, theoretical modeling, and calibration, we have tried to make progress on this question.

For the bigger picture, it is worth discussing what our estimates may be missing. We highlight three potential limitations. First, our estimates capture the effect of wealth taxes on those who are already wealthy, as opposed to the forward-looking effect on those who aspire to become wealthy. The aspiration effect of wealth taxes, even if it is quantitatively important, is extremely difficult to identify empirically. One would have to compare the savings behavior of the potentially wealthy across economies that vary permanently in their levels of capital taxation, but such cross-country analyses are typically not persuasive.

Second, our estimates capture the effect of wealth taxes conditional on staying in Denmark. In other words, we do not consider the potential migration response to wealth taxes at the top. While there is growing evidence on migration responses to labor income taxes at the top (see [Kleven et al. forthcoming](#) for a review), there is virtually no evidence on migration responses to capital or wealth taxes. Such responses are difficult to study due to a lack of statistical power: we are studying the extreme tail of the wealth distribution, and given the low frequencies of international migration, very few individuals are moving in and out of the country from year to year. Furthermore, moving to another country is arguably not the most natural response to wealth taxes. Because capital tends to be more mobile than people, it would be more natural to move wealth across borders (which would be picked up by

our taxable wealth estimates) than to move the household across borders.

Finally, our quasi-experimental approach captures the effect of wealth taxes in partial equilibrium. The changes in wealth accumulation that we find may have implications for asset prices and wage rates, making the general equilibrium effect different from the partial equilibrium effect. Although this is important to keep in mind, it is a limitation of any quasi-experimental study (i.e., of any well-identified study). Our partial equilibrium estimates provide a set of moments that economists can target when calibrating general equilibrium models.

UNIVERSITY OF COPENHAGEN AND CENTER FOR ECONOMIC BEHAVIOR
AND INEQUALITY

THINK TANK DEA

PRINCETON UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH

UNIVERSITY OF CALIFORNIA, BERKELEY AND NATIONAL BUREAU OF
ECONOMIC RESEARCH

SUPPLEMENTARY MATERIAL

An [Online Appendix](#) for this article can be found at *The Quarterly Journal of Economics* online. Code replicating tables and figures in this article can be found in [Jakobsen et al. 2019](#), in the Harvard Dataverse, [doi:10.7910/DVN/PFQU4R](https://doi.org/10.7910/DVN/PFQU4R).

REFERENCES

- Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman, "Tax Evasion and Inequality," *American Economic Review*, 109 (2019), 2073–2103.
- Alvaredo, Facundo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman, *The World Wealth and Income Database (WID.world)* (2017).
- Andreoni, James, "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence," *Journal of Political Economy*, 97 (1989), 1447–1458.
- , "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving," *Economic Journal*, 100 (1990), 464–477.
- Atkinson, Anthony B., and Joseph E. Stiglitz, "The Design of Tax Structure: Direct versus Indirect Taxation," *Journal of Public Economics*, 6 (1976), 55–75.
- Auerbach, Alan J., and James R. Hines, "Taxation and Economic Efficiency," in *Handbook of Public Economics*, vol. 3, A. J. Auerbach and M. S. Feldstein, eds. (Amsterdam: Elsevier, 2002).
- Benhabib, Jess, and Alberto Bisin, "Skewed Wealth Distributions: Theory and Empirics," *Journal of Economic Literature*, 56 (2018), 1261–1291.
- Best, Michael C., James Cloyne, Ethan Ilzetzki, and Henrik J. Kleven, "Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches," *Review of Economic Studies* (forthcoming).

- Boserup, Simon H., Wojciech Kopczuk, and Claus T. Kreiner, "Stability and Persistence of Intergenerational Wealth Formation: Evidence from Danish Wealth Records of Three Generations," University of Copenhagen, Working Paper, 2014.
- Brülhart, Marius, Jonathan Gruber, Matthias Krapf, and Kurt Schmidheiny, "Taxing Wealth: Evidence from Switzerland," NBER Working Paper no. 22376, 2016.
- Carroll, Christopher D. "Why Do the Rich Save So Much?," in *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, J. B. Slemrod, ed., (Cambridge, MA: Harvard University Press, 2002).
- Chamley, Christophe, "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica*, 54 (1986), 607–622.
- Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben H. Nielsen, and Tore Olsen, "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark," *Quarterly Journal of Economics*, 129 (2014), 1141–1219.
- De Nardi, Mariacristina, "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71 (2004), 743–768.
- Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri "Heterogeneity and Persistence in Returns to Wealth," NBER Working Paper no. 22822, 2016.
- Farhi, Emmanuel, and Iván Werning, "Progressive Estate Taxation," *Quarterly Journal of Economics*, 125 (2010), 635–673.
- Freyaldenhoven, Simon, Christian Hansen, and Jesse M. Shapiro, "Pre-event Trends in the Panel Event-study Design," *American Economic Review*, 109 (2019), 3307–3338.
- Gruber, Jonathan, Amalie S. Jensen, and Henrik J. Kleven, "Do People Respond to the Mortgage Interest Deduction? Quasi-Experimental Evidence from Denmark," NBER Working Paper no. 23600, 2018.
- Jakobsen, Katrine, Kristian Jakobsen, Henrik Kleven, and Gabriel Zucman, "Replication Data for 'Wealth Taxation and Wealth Accumulation: Theory and Evidence from Denmark'," (2019), Harvard Dataverse, doi:10.7910/DVN/PFQU4R.
- Judd, Kenneth L., "Redistributive Taxation in a Simple Perfect Foresight Model," *Journal of Public Economics*, 28 (1985), 59–83.
- Kaplow, Louis, "Utility from Accumulation," *National Tax Association Proceedings* (2011), 102nd Annual Conference on Taxation, 189–194.
- Kleven, Henrik J. "How Can Scandinavians Tax So Much?," *Journal of Economic Perspectives*, 28 (2014), 77–98.
- , "Bunching," *Annual Review of Economics*, 8 (2016), 435–464.
- Kleven, Henrik J., Camille Landais, Mathilde Muñoz, and Stefanie Stantcheva, "Taxation and Migration: Evidence and Policy Implications," *Journal of Economic Perspectives* (forthcoming).
- Kleven, Henrik J., and Esben A. Schultz, "Estimating Taxable Income Responses Using Danish Tax Reforms," *American Economic Journal: Economic Policy*, 6 (2014), 271–301.
- Kopczuk, Wojciech, "Economics of Estate Taxation: A Brief Review of Theory and Evidence," *Tax Law Review*, 63 (2009), 139–157.
- , "Incentive Effects of Inheritances and Optimal Estate Taxation," *American Economic Review, Papers and Proceedings*, 103 (2013a), 472–477.
- , "Taxation of Intergenerational Transfers and Wealth," in *Handbook of Public Economics*, vol. 5, A. J. Auerbach, R. Chetty, M. S. Feldstein, and E. Saez, eds. (Amsterdam: Elsevier, 2013b).
- Kreiner, Claus T., David D. Lassen, and Søren Leth-Petersen "Measuring the Accuracy of Survey Responses Using Administrative Register Data: Evidence from Denmark," in *Improving the Measurement of Household Consumption Expenditures*, C. D. Carroll, T. F. Crossley, and J. Sabelhaus, eds. (Chicago: University of Chicago Press, 2015).

- Leth-Petersen, Søren, "Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit?," *American Economic Review*, 100 (2010), 1080–1103.
- Love, David A., Michael G. Palumbo, and Paul A. Smith, "The Trajectory of Wealth in Retirement," *Journal of Public Economics*, 93 (2009), 191–208.
- OECD, *Taxation of Net Wealth, Capital Transfers and Capital Gains of Individuals* (Paris: OECD, 1988).
- Piketty, Thomas, *Capital in the Twenty First Century* (Cambridge, MA: Harvard University Press, 2014).
- Piketty, Thomas, and Emmanuel Saez, "A Theory of Optimal Inheritance Taxation," *Econometrica*, 81 (2013), 1851–1886.
- Poterba, James, Steven Venti, and David Wise, "The Composition and Drawdown of Wealth in Retirement," *Journal of Economic Perspectives*, 25 (2011), 95–118.
- , "Longitudinal Determinants of End-of-Life Wealth Inequality," *Journal of Public Economics*, 162 (2018), 78–88.
- Roth, Jonathan, "Pre-test with Caution: Event-study Estimates After Testing for Parallel Trends." Harvard University, Working Paper, 2019.
- Saez, Emmanuel, and Stefanie Stantcheva, "A Simpler Theory of Optimal Capital Taxation," *Journal of Public Economics*, 162 (2018), 120–142.
- Saez, Emmanuel, and Gabriel Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *Quarterly Journal Of Economics*, 131 (2016), 519–578.
- Seim, David, "Behavioral Responses to Wealth Taxes: Evidence from Sweden," *American Economic Journal: Economic Policy*, 9 (2017), 395–421.
- Skatteministeriet. *Skatteberegningsreglerne for personer-før og nu* (Copenhagen: Danish Ministry of Taxation, 2002).
- Straub, Ludwig, and Iván Werning, "Positive Long-Run Capital Taxation: Chamley-Judd Revisited," *American Economic Review* (forthcoming).
- Zoutman, Floris T., "The Effect of Capital Taxation on Household Savings," Norwegian School of Economics, Department of Business and Management Science, Working Paper, 2015.
- Zucman, Gabriel, "Global Wealth Inequality," *Annual Review of Economics*, 11 (2019), 109–138.