ERASMUS UNIVERSITY ROTTERDAM

Erasmus School of Economics – MSc. Policy Economics

Behavioural responses to wealth taxes: Evidence from the Netherlands

Version date: 22-10-2023

Abstract

In this study, I use a reform of the Dutch wealth tax in 2017 to study the behavioural response to changes in the marginal tax rate on wealth. Using the standard difference-in-difference-in-differences approach in the tax elasticity literature, I estimate the elasticity of taxable wealth and investigate several mechanisms through which taxpayers respond. My baseline results indicate that taxpayers in the first bracket (up to €100,000) exhibit an elasticity of taxable wealth with respect to the net-of-tax rate between zero and 4.7. Graphical evidence suggest the assumption of constant trend differentials required to interpret this estimate as causal is plausible. In the top bracket (above €1 million), I estimate an elasticity of 43.3, but the constant trend differentials assumption is violated. More plausible estimates give an elasticity of 24.9 in the top bracket. Further analyses indicate that part of the response constitutes a shift of wealth towards closely held firms, which can be used as a tax shelter from the wealth tax. Furthermore, taxpayers decrease their taxable wealth through increased inter vivos giving and take on more deductible debt to reduce their net taxable wealth.

Keywords: Elasticity of taxable wealth; Wealth taxation; Difference-in-difference-in-differences; Tax avoidance and evasion

Name student: Thijs Busschots

Student number: 568280

Supervisor: Dinand Webbink

Second assessor: Aart Gerritsen

Contents

1. Introduction

A growing body of research documents an increasing wealth inequality in developed countries (Saez and Zucman, 2016; Piketty et al., 2018; Chancel et al., 2022). In response, a group of economists have proposed the implementation of a wealth tax to reduce wealth inequality and raise government revenue (Piketty, 2014; Saez and Zucman, 2019; Piketty et al., 2022; Guvenen et al., 2023).[1] However, not all economists agree. Opponents of a wealth tax argue that a tax on net wealth is not desirable from an economic perspective (Kopczuk and Mankiw, 2019; Pestieau and Broadway, 2019).[2]

To properly assess a wealth tax proposal, empirical evidence on the behavioural response to a wealth tax is needed (Auerbach and Hasset, 2015; McGrattan, 2015; Jakobsen et al., 2020; Advani and Tarrant, 2021). Such evidence is relevant, as the desirability and optimal rate of a wealth tax depend on the degree to which a wealth tax creates economic distortions (Saez and Stantcheva, 2018; Jakobsen et al., 2020). For instance, a tax on wealth will reduce the returns on saving, causing a welfare loss by distorting intertemporal prices (Chamley, 1986; Judd, 1985; Bernheim, 2002). If taxpayers reduce their saving, a wealth tax may reduce economic growth through a reduction in capital accumulation (Lucas, 1990; Seim, 2017; Scheuer and Slemrod, 2021), unless taxpayers are able to avoid the tax. If taxpayers are able to circumvent the wealth tax, the ability of the policy to reduce inequality and raise government revenue is weakened. As a summary statistic for the behavioural response, the theoretical and empirical literature focus on the elasticity of taxable wealth (Saez et al., 2012; Saez and Stantcheva, 2018).

While some evidence on the elasticity of taxable wealth exists, the number of studies is scant, and findings vary strongly. Papers using bunching techniques are notorious for their low estimates, estimating elasticities with respect to the marginal net-of-tax rate of 0.27 in Sweden (Seim, 2017), 0.56 in Norway (Ring, 2021), 0.8 in Switzerland (Brüllhart et al., 2022) 2 in Colombia (Londońo-Velez and Ávila-Mahecha, 2023) and no response in France (Garbinti et al., 2023). Difference-in-difference style estimations finds higher elasticities, from 11.3 in Denmark (Jakobsen et al., 2020), 13.8 in the Netherlands (Zoutman, 2018), 32.4 in Spain (Duran-Carbé et al., 2019) and up to 46.1 in Switzerland (Brüllhart et al., 2022). The limited number of studies is partly explained by the fact that few countries levy a wealth tax, making wealth tax reforms that can be used for causal analysis rare (OECD, 2018; Jakobsen et al., 2020).

In this study, I exploit one of the rare wealth tax reforms to provide empirical evidence on the behavioural response of Dutch taxpayers. A Dutch wealth tax reform in 2017 effectively introduces three tax brackets with progressive tax rates where the wealth tax was previously a flat tax at 1.2%.[3]

---

[1] A call for a tax on net wealth is not new. Kaldor (1956) advocated developing countries to implement a wealth tax and Allais (1977) proposed a tax of 2% on wealth in order to reduce other tax burdens in France.
[2] See, among others, OECD (2018), Adam and Miller (2021) and Scheuer and Slemrod (2021) for a review of the conditions under which a tax on net wealth is desirable and how it relates to capital income taxation.
[3] Legally, the Dutch wealth tax is a capital income tax. However, taxable capital income is assumed based on net taxable wealth, making the tax equivalent to a net wealth tax (Cnossen and Bovenberg, 2001; Brys, 2006; Zoutman, 2018).

Households with taxable wealth above around €1 million (top bracket) faced a significant increase in the effective marginal tax rate to around 1.6% following the reform, and taxpayers with wealth below around €100.000 (bottom bracket) faced a significant decrease to around 0.6%. Households in the middle bracket faced a relatively minor increase in the marginal tax rate (to around 1.35%), and hence function as a natural control group.

To identify the causal effect of a change in the marginal tax rate on the amount of taxable wealth taxpayers report, I closely follow the approach developed in Jakobsen and Søgaard (2022) which allows for a graphical representation of the standard difference-in-difference-in-differences method used in the elasticity of taxable income literature (Saez et al., 2012). The approach controls for initial wealth, leaving only time-variation in the marginal tax rate to estimate a causal effect. This method circumvents the usual endogeneity issues regarding mean reversion and differential growth trends arising from the difference in initial wealth between taxpayers treated in different brackets (Saez et al., 2012; Weber, 2014; Kumar and Liang, 2020; Jakobsen and Søgaard, 2022).

For taxpayers in the first bracket (up to around €100,000), I find estimates ranging from a zero elasticity of taxable wealth with respect to the net-of-tax rate up to 4.7. This implies that taxpayers with taxable wealth below €100,000 respond to an increase in the marginal tax rate with 1 percentage point by decreasing their taxable wealth with 4.7%.[4] Taxpayers in the control group did not show a change in the growth of their taxable wealth, indicating that the identifying assumption of constant trend differentials holds.

For taxpayers facing an increase in the marginal tax rate in the top bracket (above about €1,000,000), I find an elasticity of 43.4. However, a placebo test indicates the assumption of constant trend differentials may not hold, potentially introducing a bias. An analysis of the components of net taxable wealth indicate it is the investments asset class that causes the constant trend differentials assumption to fail. All other wealth components show constant trend differentials in the validation region when considered in isolation. Focussing only on savings, taxpayers in the top bracket show an elasticity of 24.9 with respect to the marginal net-of-tax rate.

In additional analyses, I explore the mechanisms through which taxable wealth is affected. First, taxpayers who own a closely held firm have a stronger response to a change in the marginal tax rate than taxpayers who do not own a closely held firm. In addition, the value of wealth held in the form of shares in closely held firms increases in response to an increase in the wealth tax rate. However, placebo tests show the constant trend differentials assumption is questionable here. These findings are in line with the use of closely held firms as tax shelters (Gordon and Slemrod, 2000; Alstadsæter et al., 2014), also confirming anecdotal evidence of such practices in the Netherlands (Ministerie van Financiën,

[4] Recall that the elasticity of taxable wealth with respect to the net-of-tax rate and the semi-elasticity with respect to a 1 percentage point change in in the wealth tax rate are equivalent by approximation as $\ln(1-\tau) \approx \ln(-\tau)$ when τ is small.

2023). Second, I find part of the response runs though inter-vivos gifts (for example, gifts to children). My estimates show that taxpayers in the top bracket increase their giving by 28.1% in response to a 1 percentage point increase in the marginal tax rate. In the first bracket, the effect is smaller, at 1.2%. Third, taxpayers in the top bracket appear to increase their amount of deductible debt in box 3. This highlights that also debt deductibility is a relevant policy design aspect in a net wealth tax.

My findings make several contributions to the literature. Firstly, I offer the first empirical evaluation of the behavioural responses to the introduction of effectively progressive brackets in the Dutch wealth tax in 2017. In doing so, I add new empirical evidence to the emerging elasticity of taxable wealth literature (Seim, 2017; Zoutman, 2018; Duran-Cabré et al., 2019; Agrawal et al., 2020; Jacobsen et al., 2020; Ring, 2021; Brülhart et al., 2022; Alstadsæter et al., 2022; Garbinti et al., 2023). Furthermore, exploiting a wealth tax reform that affects a broad range of the wealth distribution, I can quantify the elasticity for households with relatively little wealth (below €100,000), while most empirical studies focus only on the very top of the wealth distribution (Seim, 2017; Duran-Cabré et al., 2019; Agrawal et al., 2020; Jakobsen et al., 2020; Londońo-Velez and Ávila-Mahecha, 2023).

For the Netherlands in specific, only Zoutman (2018) provides an estimate of the elasticity of taxable wealth, but based on a tax reform in 2001. A new estimate on more recent data is useful, as factors that determine the elasticity of taxable wealth are not fixed over time (Perret, 2021; Garbinti et al., 2023). For example, the technology available to evade and avoid taxes changes over time, as well as the ability of the tax authority to detect such techniques: banking secrecy has been weakened, reducing the ability to hide wealth abroad (Leenders et al., 2023). Also, the tax moral (willingness to pay tax) may change over time. Such changes in context can affect the elasticity of taxable wealth. Furthermore, better data availability poses an improvement. Whereas Zoutman (2018) employs survey data on around 35.000 households, I use register data on the full population of Dutch households.

The estimates in this paper also offers a very practical use to Dutch policy makers. For instance, the estimated elasticities can be used to compute the implied revenue-maximising tax rate in the top wealth bracket. Using the elasticity of 25, tax revenue would be maximised at a top wealth tax rate of 2.1% (versus 1.6% currently). Translated to capital income taxes, the elasticity of 25 implies an elasticity of taxable wealth with respect to the net-of-tax rate of return of 1.12 when assuming a rate of return of 5%.[5] The resulting revenue-maximising marginal tax rate on capital income is 38% (versus 31% currently). Due to a lack of empirical evidence, the Dutch currently bases wealth tax revenue projections on a rule of thumb for the behavioural response to a change in the wealth tax rate. My results indicate the assumption of zero response for taxpayers with wealth below €100,000 (first bracket) is justified. However, the assumption that taxpayers with taxable wealth above €100,000 exhibit a response that erodes 20% of the mechanical increase in tax revenue is too conservative for taxpayers

[5] See Appendix 10.6 for the translation of taxable wealth elasticities to capital income elasticities and Appendix 10.7 for the calculations of revenue effects.

with taxable wealth above €1 million (for taxpayers with taxable wealth between €100,000 and €1,000,000, I cannot assess the validity of the rule of thumb). Based on my results, the behavioural effect for taxpayers with taxable wealth above €1 million is more likely to erode 75% of the extra tax revenue from a small change in the marginal tax rate.

Second, this paper contributes to the understanding of the mechanisms through which taxable wealth responds to a change in the marginal tax rate. As such, I also contribute to the empirical literature nascent to that on the elasticity of taxable wealth. In particular, I add empirical evidence to the literature on tax motivated (inter vivos) gifts (see Kopczuk (2013) for a review) and the tax evasion and avoidance literature (see Slemrod and Yitzhaki (2002) for a review), specifically on the use of closely held firms as tax shelters (Gordon and Slemrod, 2000; Alstadsæter et al., 2014). While Slemrod (1995; 1998; 2001), Chetty (2009) and Saez et al. (2012) underline how quantifying shifts to other tax bases is essential to determine tax revenue and welfare effects, little previous empirical work on the elasticity of taxable wealth has been able to identify these mechanisms explicitly (Duran-Cabré et al., 2019; Brülhart et al., 2022).

A third contribution of this paper is methodological. I offer the first application of the approach allowing for graphical validation of the identifying assumptions as proposed by Jakobsen and Søgaard (2022) in the elasticity of taxable wealth literature. The constant trend differentials assumption is central to any regression model where mean reversion and differential growth trends are tackled by controlling for initial income or wealth, but the plausibility of this assumption has remained obscure in most applications (Jakobsen and Søgaard, 2022). Furthermore, I derive and apply a formula to calculate a bias that may occur when there is no pure control group and the underlying elasticity in the treatment and control group differ. This type of bias is recognised in the literature (Saez, 2004; Saez et al., 2012), but ignored many empirical studies (Kumar and Liang, 2020). Additionally, I provide the first application of the novel synthetic-difference-in-difference methodology (Arkhangelsky et al. 2021) in the elasticity of taxable wealth literature as robustness check.

The outline of this paper is as follows. Section 2 provides a review of the relevant theoretical and empirical literature. A discussion on the used data is provided in Section 3. Section 4 briefly discusses the Dutch tax system, the reform of the wealth tax in 2017 and in what ways taxpayers can respond to the changed tax rate. The methodology is laid out in Section 5. Section 6 presents the results and robustness tests and finally Section 7 concludes and provides policy implications.

2. Related literature

In this section, I discuss the theoretical literature that forms the basis of this study and previous empirical tax elasticity literature. Section 2.1 describes the relevance of the elasticity of taxable wealth in the optimal tax literature and motivates the use of the elasticity of taxable wealth as summary statistic for the behavioural response to wealth taxes. Many of the cited papers are concerned with the elasticity of taxable income, but as is shown by Saez and Stantcheva (2018), these insights also translate to the taxation of wealth.[6] Section 2.2 gives an overview of empirical studies that have applied the insights from the elasticity of taxable income literature to wealth taxes.

2.1 The elasticity of taxable wealth in optimal tax theory

In the optimal taxation literature, the desirability of a tax mainly depends on two factors: equity and efficiency (Mirrlees, 1971; Saez, 2001; Saez and Stantcheva, 2018). A tax is more desirable if it is better able to redistribute based on given political preferences (equity). It becomes less desirable if it creates more economic distortions for a given tax revenue (efficiency). Greater economic distortions decrease welfare and are hence undesirable from the perspective of a social planner.

To quantify the economic distortions from taxation, the economic literature turns to the elasticity of the tax base with respect to the marginal tax rate (Saez et al., 2012). In the classical Mirrlees (1971) model for optimal labour income taxes, the distortionary effect from the tax is captured by the compensated elasticity of labour supply with respect to the marginal net-of-tax rate. However, focussing only on labour supply ignores other relevant dimensions of response, such as work effort, career decisions, income shifting and substitution towards non-monetary remuneration for work. Recognizing the need for a broader measure, Lindsey (1987) estimated the net effect of all the possible response margins by using taxable income as a summary statistic. In two seminal contributions, Feldstein (1995; 1999) shows that economists should indeed be interested in the elasticity of taxable income to predict effects of a change in the marginal tax rate for government revenues.

The behavioural response underlying the total elasticity can be categorized into real and reporting response (Slemrod, 1995; 1996; 2001). Real responses are changes in the economic decisions of the taxpayer caused by changes in relative prices, whereas reporting responses are tax avoidance and evasion efforts that reduce the tax liability without changing economic behaviour. A real response to a wealth tax is, for example, a reduced savings rate. In this case, the tax altered real economic choices as the consumption pattern of the taxpayer changed. A reporting response, for example, is to reallocate or hide personal savings so it faces no or lower taxation. In this case, the taxpayer may still hold the same amount of savings as without the tax. Only the reported/legal status of the savings has changed.

Under certain conditions the elasticity of the tax base is a sufficient statistic for welfare analysis, meaning the decomposition of the total effect into real and reporting responses is not necessary to derive welfare implications. The decomposition is not necessary if taxpayers are fully rational and equate the marginal cost of reducing the tax base through either real or reporting responses to the marginal benefit, which equals the marginal tax rate (Feldstein, 1999). In this simple setting, the social costs of reporting

[6] An assumption that is required for this extension, is that individuals derive utility from wealth. Saez and Stantcheva (2018) present several theoretical and empirical works justifying this assumption (Poterba and Rotemberg, 1987; Piazzesi et al., 2007; Stokey, 2009; Kiyotaki et al., 2011). Furthermore, utility in consumption must be linear, so there are no consumption smoothing motives, and the model can be analysed statically.

and real responses are both equal to the marginal tax rate, meaning the total elasticity and the marginal tax rate are sufficient statistics to compute welfare effects.

Three factors cause the assumption of equal social costs of real and reporting responses to fail. First, taxpayers may fail to perfectly equate the marginal costs and benefits of each response margin due to optimization errors (Chetty, 2009). Such errors may arise when taxpayers misestimate the costs of tax avoidance and evasion.[7] In this case, the marginal tax rate is no longer the relevant social cost for all margins of response and the total elasticity of taxable wealth is no longer a sufficient statistic to determine the deadweight loss. Second, a part of the response may include a shift to another tax base where the shifted wealth is still taxed, but at a different (lower) tax rate (Slemrod, 1995; 1996; 1998; Gordon and Slemrod, 2000; Saez et al., 2012). Intuitively, the tax revenue gain at this alternative tax base partly compensates the loss of tax at tax base where the wealth is shifted from. Third, reporting responses may have positive or negative externalities, causing their social costs to differ from real responses (Doerrenberg et al., 2017). In these cases, a division between real and reporting effects is required, as well as an estimate of the social costs associated with the reporting effect.

Motivated by the importance to identify mechanisms through which taxable wealth responds, I complement my estimates of the elasticity of taxable wealth with estimates on components of taxable wealth and wealth in other tax bases (particularly to closely held firms). These mechanisms can enlighten to what degree responses constitute real or reporting behaviour.

2.2 Previous estimates of the elasticity of taxable wealth

While there is an extensive literature on the behavioural responses to taxation of capital income (see Bernheim (2002) and Attanasio and Weber (2010) for a review), the literature on the elasticity of taxable wealth is small and has emerged only recently (Table 1).[8] To my knowledge, Table 1 includes all the estimates of elasticity of taxable wealth that are currently available in the literature.

Jakobsen et al. (2020) suggest several reasons for the rarity of empirical work on wealth taxes. Firstly, few countries have a net wealth tax (OECD, 2018; Perret, 2021), meaning there are few opportunities for natural experiments. Consequently, it is difficult to find exogeneous variation in wealth taxes that can be used for causal identification. Secondly, wealth taxes are often levied above a high threshold in the wealth distribution, while data on wealth at the top is of lesser availability/representativeness and finding a suitable control group for the very wealthy is difficult. Thirdly, wealth accumulation is slow, making it important to measure long-run effects. However, measuring long-run effects is especially challenging, as it requires an extended period of data with a stable tax schedule. Lastly, wealth taxes may have received little attention, as early literature ought the

[7] Indeed, empirical evidence shows, for example, that taxpayers overestimate the likelihood of being caught for tax avoidance and the size of the associated fines (Andreoni et al., 1998).
[8] Earlier work does study behavioural responses to one-off taxation of net wealth in the context of bequest and gift taxation (see Kopczuk (2013) for a review).

optimal capital taxes to be zero, following the seminal contributions by Atkinson and Stiglitz (1976), Chamley (1986) and Judd (1985). Modern academic works stands more positively towards taxing capital income and wealth.[9]

Table 1: Existing estimates of the elasticity of taxable wealth with respect to the net-of-tax rate

Paper	Country	Method	Time horizon	Elasticity
Seim (2017)	Sweden	Bunching	N/A	0.09-0.27
		DiD	2 years	0
Zoutman (2018)	Netherlands	DiD	1 year	11.6
		DiD	3 years	13.8
Duran-Cabré et al. (2019)	Spain	DiD	1 year	15.3
		DiD	4 years	32.4
Agrawal et al. (2020)	Spain	DiD	4 years	2.3-5.1
Jakobsen et al. (2020)	Denmark	DiD	8 years	8.9-11.3
Ring (2021)	Norway	RDD	N/A	4.67
		Bunching	N/A	0.56
Brüllhart et al. (2022)	Switzerland	DiD	At event	18.2
		DiD	1 year	12.6
		DiD	2 years	36.8
		DiD	3 years	42.5
		DiD	4 years	41.1
		DiD	5 years	43.2
		DiD	6+ years	46.1
		Bunching	N/A	0.7-0.8
Alstadsæter et al. (2022)	Norway	DiD	2 years	4.91
Garbinti et al. (2023)	France	Bunching	-	0
Londońo-Velez and Ávila-Mahecha (2023)	Colombia	Bunching	N/A	2

Notes: This table reports the main estimates presented in the existing literature on the elasticity of taxable wealth. As bunching is a cross-sectional method, there is no corresponding time horizon (N/A).

The empirical literature measuring the elasticity of a tax base can be categorized into two common methods: difference-in-difference (DiD) and bunching (He, et al., 2021; Aronsson et al., 2022). The DiD designs compare taxpayers who are treated by a change in the marginal tax rate to a control group of taxpayers whose marginal tax rate remains constant or changes less. Treatment status is often based on tax brackets, or regional tax code differences. The bunching estimates identify peaks in the density of taxpayers at tax-kinks in the wealth distribution that lead to kinks in the household budget curve (Saez, 2010; Seim, 2017; Londońo-Velez and Ávila-Mahecha, 2023). [10]

[9] In more recent theoretical and empirical literature, it follows that a zero capital taxes are only optimal under very specific conditions and that in realistic scenario's, a positive capital tax is warranted (Conesa et al., 2009; Diamond and Saez; 2010; Diamond and Spinnewijn, 2011; Saez, 2013; Saez and Stantcheva, 2018; Gerritsen et al., 2020; Straub and Werning, 2020, Piketty et al., 2022). Recently, researchers also advocate the benefits of a tax on net wealth opposed to (solely) taxing capital income (Seim; 2017; Saez and Zucman, 2019; Piketty et al., 2022; Guvenen et al., 2023).

[10] The bunching approach assumes taxable wealth follows a smooth density function across the population in absence of any taxes. Assuming strictly quasi-concave utility, continuously distributed preferences and abilities, and no uncertainty produces such a smooth distribution of the taxable base across the population (Saez, 2010; Chetty, 2011; Seim, 2017; Gelber et al., 2020).

The bunching technique has received criticism, because optimization frictions can bias the estimate of the elasticity of taxable income downward (Chetty et al., 2011; Chetty, 2012; Kleven and Wasseem, 2013; He, et al., 2021; Aronsson et al., 2022). Due to these optimization frictions, taxpayers may not bunch at kink points in the tax schedule like theory predicts (Saez, 2010; Seim, 2017). Such optimization frictions can explain why the response to a change in the tax schedule becomes larger over time. Taxpayers need time to adjust (Saez et al., 2012).

Seim (2017) studies the response to wealth taxation in Sweden, by examining whether taxpayers bunch at a cut-off in the tax schedule where the marginal tax rate goes from 0% to around 1.5%. The value of assets, is self-reported by the taxpayer, enabling taxpayers to underreport wealth and bunch at the exemption threshold. Indeed, Seim (2017) finds excess mass at the threshold, implying an elasticity in the range of 0.09-0.27. Most of the effect is driven by taxpayers underreporting the value of their car (a self-reported asset), indicating that the elasticity mainly consists of reporting effects rather than real effects. Additionally, Seim (2017) performs a reduced form[11] difference-in-differences analyses, exploiting a shift in the exemption threshold. The upward shift in the exemption threshold decreased the marginal tax rate from 1.5% to 0% for those within 100.000 DAK above the old threshold. Households with taxable wealth within 100.000 DAK below the old threshold are used as control group, as their marginal tax rate did not change. Despite the substantial change in the marginal tax rate, Seim does not find any response to the change in the exemption threshold.

Zoutman (2018) analyses a Dutch tax reform in 2001 that abolished the tax on capital income, leaving only a net wealth tax.[12] Because capital income varied across taxpayers with a similar level of net wealth, the abolishment of the capital income tax created changes in the marginal tax rate of difference size and directions within a given level of wealth. Zoutman uses this variation in a difference-in-difference design, comparing taxable wealth growth of taxpayers with similar net wealth, but facing different changes in the marginal tax rate. The results indicate an elasticity of 13.8 within a 3-year time window. In a heterogeneity analysis, Zoutman finds that households in the top 25% of the wealth distribution respond slightly stronger to the wealth tax and homeowners are slower in their response.

Certain limitations in the study by Zoutman (2018) can be identified. First, Zoutman investigates the proposed method to mitigate a bias from mean reversion from Weber (2014), by basing his instrument on a further lag. The analysis turns out to be very sensitive to the used base-year and the estimates fall strongly when the instrument is lagged further. In a specification where multiple instruments are used, the elasticity in a 3-year time window falls from 13.79 to 4.83 but remains statistically significant. Second, Zoutman controls for labour income in the reform period, which is

[11] Reduced form, because Seim (2017) does not scale estimated effect by the share of compliers: those households who are still in the treated wealth range during the treated years. The found effect should hence be interpreted as an intention to treat effect.
[12] The Netherlands is the first and only country in the world that effectively levies a tax on net wealth without also taxing some form of capital income (Zoutman, 2018).

possibly a bad control as it may be affected by the tax on wealth. Other studies have shown that labour income can be affected by capital taxation (Ring, 2021).

Duran-Cabré et al. (2019) find significantly higher elasticities of taxable wealth than other studies, when examining the response to the reintroduction of a net wealth tax in Spain. The authors have no data on the years before the wealth tax was reintroduced, but instead use the first year the tax was levied as control year because the tax was imposed unexpectedly.[13] As source of variation, the authors use the fact the differences in the composition of wealth led to differences in the average tax rate: households with a larger share of tax exempt assets face a lower tax rate as percentage of their total wealth. The new wealth tax featured many tax-exempt asset classes, which taxpayers used by shifting their wealth holdings towards these exempt types.[14] As a result, the authors find a relatively high elasticity of taxable wealth with respect to the net-of-tax-rate of 32.4.

A further explanation of the large shift in taxable wealth found in Duran-Cabré et al. (2019) is found in a paper by Agrawal et al. (2020), who examine taxpayer mobility in response to the same Spanish reform. Agrawal et al. use regional variation in the wealth tax rate, set independently by each region. All regions chose a positive tax rate, except for Madrid, which chose to levy no wealth taxes. As a result, wealthy Spanish taxpayers migrated to Madrid. The size of the mobility implies an elasticity of taxable wealth of 2.3 to 5.1. The authors mention that the mobility effect could be large, as residents are able to register residence in another region while effectively living in another reason. As such, the costs of avoiding/evading the wealth tax through mobility are low.

Ring (2021) studies the effect of geographical discontinuities in real estate taxes resulting from a model that is used to determine real estate values for tax purposes in Norway. The estimates imply an elasticity of 4.67 but is statistically insignificant. As there are few evasion possibilities for the wealth tax in Norway, the elasticity is likely close to estimating a pure real response. By separately instrumenting for changes in the marginal and average tax rates, Ring splits the effect into a substitution and income effect. Contrary to most studies, Ring finds a significant income effect which dominates the substitution effect. As a result, a higher tax on housing wealth resulted in an increase in savings.

Londońo-Velez and Ávila-Mahecha (2023) use kinks and notches in the Colombian wealth tax to estimate the elasticity of taxable wealth with a bunching technique. At the notches, the tax liability shows large discontinuous jumps. The bunching estimates show an elasticity of 2. In an additional paper using the same data, Londońo-Velez and Ávila-Mahecha (2021) find that taxpayers whose tax evasion

[13] The reinstatement of the wealth tax was only confirmed in November 2011 (going into effect by December 2012), giving taxpayers little time to respond (Duran-Cabré et al., 2019). However, given the public debate leading up to the reintroduction of the wealth tax, a risk of anticipation effects remains.
[14] Jakurti and Süssmuth (2023) corroborate the finding that taxpayers shifted their wealth to tax-exempt asset classes using survey data. The authors cannot observe the region of residence of each respondent, also meaning the tax rate faced by each respondent cannot be observed. As a consequence, they are unable to estimate an elasticity, which is why the study is excluded from Table 1.

is exposed by the Panama Papers leak strongly increase their reported wealth and that greater enforcement efforts by the Colombian tax authorities increased reported wealth.

In Switzerland, Brülhart et al. (2022) use variation in the tax rate on wealth between different cantons. Comparing the evolution of taxable wealth between a canton that increased its wealth tax to a canton that did not (or less), they estimate an elasticity of taxable wealth of 12.6 after 1 year and up to 46.1 after 6 years. Most of the effect (49%) is driven by a change in the value of self-reported financial assets, which points towards a large reporting-component. They find no significant income effects.

Jakobsen et al. (2020) quantify the elasticity of taxable wealth by applying two DiD approaches to a Danish tax reform. The first approach uses an increase in the tax-exempt wealth threshold for couples, while the threshold remained constant for singles. Couples that were shifted to the exempt range by the reform are compared to singles who remain taxed. The second approach compares taxpayers in the top 1% of the wealth distribution who faced a reduction in the marginal tax rate to taxpayers in the top 1% who had a marginal tax rate of zero due to a tax liability ceiling.[15] To account for different trends in taxable wealth, the authors estimate a linear treatment-group-specific pre-trend and project it forward in the treatment period. The estimates show significant negative elasticities, with higher elasticities in the second approach, where the sample is limited to the wealthiest taxpayers.

In a paper focussed on the effect of information sharing requirements on reported taxable wealth, Garbinti et al. (2023) additionally check for bunching at kinks in the wealth tax rate in France. They find no bunching, indicating that French taxpayers do not manipulate their taxable wealth to remain in a tax bracket with a lower marginal tax rate. The authors do find bunching at a wealth cut-off above which taxpayers have to report their wealth holdings in more detail. The result highlights the fact that taxpayers do not only care about marginal tax rates, but also reporting requirements.

Many studies outlined above are sensitive to a bias stemming from unequal baseline treatment status. If the treatment and control group do not face the same marginal tax rate in the pre-reform period, then a standard DiD estimator will only yield an unbiased average treatment effect if the treatment effect is immediate and constant over time (Tazhitdinova and Vazquez-Bare, 2023). If these conditions are not met, pre-trends may be polluted with treatment effects stemming from the unequal baseline, hence not reflecting parallel *potential* outcomes. These conditions are not likely to hold in the setting of the elasticity of taxable wealth, as most studies have shown heterogeneous effects over time (Table 1). Studies with an unequal baseline treatment status, which are hence possibly biased, are Zoutman (2018), the couples specification in Jakobsen et al. (2020), Ring (2021) and Brülhart et al. (2022). In this paper, all taxpayers in the treatment and control group face the same marginal tax rate in the pre-reform period, meaning there is no threat of bias due to unequal baseline treatment status.

---

[15] The tax liability ceilings put an upper bound on the total tax liability at 78% of the taxable income. A taxpayer bound by the ceiling will have an effective marginal tax rate on wealth of 0%, as an increase or decrease (within a degree to which the ceiling remains to apply), will not affect the tax liability.

3. Institutional setting

The Dutch personal income tax is divided into three 'boxes': labour income is taxed in box 1, dividend income and realised capital gains from substantial ownership in a firm[16] is taxed in box 2 (combined with corporate income taxes on profits), and finally net wealth (assets minus liabilities) is taxed in box 3.[17] Wealth that is eligible for taxation in box 3 includes cash, savings and deposits, investments and real estate that is not the primary residence of the taxpayer and is measured on the 1st of January of the tax year (from 2011 onwards). The value of owner-occupied housing is not taxed in box 3. Instead, imputed rents to owner-occupied housing are added to labour income and taxed in box 1. In conjunction, mortgage debt is included in box 1, with a portion of the mortgage interest rate being deductible from box 1 income.[18] All other personal debt is included in box 3. Financial wealth held in dedicated accounts for pension saving are exempt from wealth taxation under certain conditions. Lastly, durable goods, such as cars, jewellery and art are not included in taxable wealth, unless they are held with the purpose of being an investment.

The tax in box 3 is levied by applying a flat tax rate to an imputed rate of return on net wealth. Hence, legally, the tax in box 3 is a capital income tax. However, the imputed capital income only depends on the amount of net wealth and not on actual capital income, so the tax in box 3 is economically equivalent to a net wealth tax and can be analysed as a net wealth tax (Cnossen and Bovenberg, 2001; Brys, 2006; Zoutman, 2018). In Appendix 1, the assumed capital income and tax rate are shown, which together results in the effective marginal tax rate on net wealth.

It should be noted that net wealth also plays a role in means-testing for certain allowances. Individuals with wealth above the tax-exempt threshold are not eligible for rent allowance.[19] For the healthcare and extra child allowance a higher threshold is in place, at €127.582 for singles and €161.329 for couples (in 2017). The means-testing thresholds did not show a discontinuous change simultaneous with the introduction of tax brackets in 2017, meaning there is no risk for confoundedness in this respect. Moreover, Bosch et al. (2019) do not find any evidence for Dutch taxpayers manipulating their taxable wealth around these thresholds.

3.1 The 2017 wealth tax reform

In 2017, the Dutch wealth tax was reformed. Prior to 2017, all net wealth above the exempt threshold was taxed at an effective flat rate of 1.2%. From 2017 onwards, taxpayers with more taxable wealth were assumed to realise higher capital income, resulting in three brackets with progressive effective tax

---

[16] Substantial ownership means the taxpayer owns at least 5% of the shares in a firm. If the taxpayers owns a smaller share, then the value of the stocks is taxed in box 3.

[17] See Cnossen and Bovenberg (2001) and Brys (2006) for a detailed discussion on the Dutch tax system from 2001 onwards and the taxation of capital in particular.

[18] Typically, the mortgage interest deduction is higher than the imputed rents, so owner occupied housing is effectively subject to a fiscal subsidy.

[19] The threshold is multiplied by the number of persons living at the address for which rent allowance is requested.

rates (Figure 1).[20] The marginal tax rate on taxable wealth up to €75,000 above the tax-exempt amount (€100,000 when including the tax-exempt amount) was lowered to 0.86% (first bracket). Taxable wealth above €75.000 and up to €975,000 faced a relatively minor increase in the marginal tax rate, from 1.2% to 1.38% (second bracket). Finally, taxable wealth above €975,000 (€1,000,000 when including the tax-exempt amount) faced an increased marginal tax rate of 1.62% (third bracket). In 2018, the marginal tax rate in the first bracket fell further to 0.61%. In the years following, the marginal tax rates in each bracket remain relatively constant. In absolute terms, the reform decreased the wealth tax liability for a taxpayer with €100,000 from €907 in 2016, to €645 in 2017 and €458 in 2018. A taxpayer with €2,000,000 in taxable wealth would face an increase in wealth tax liability from €23,707 in 2016, to €29,265 in 2017.

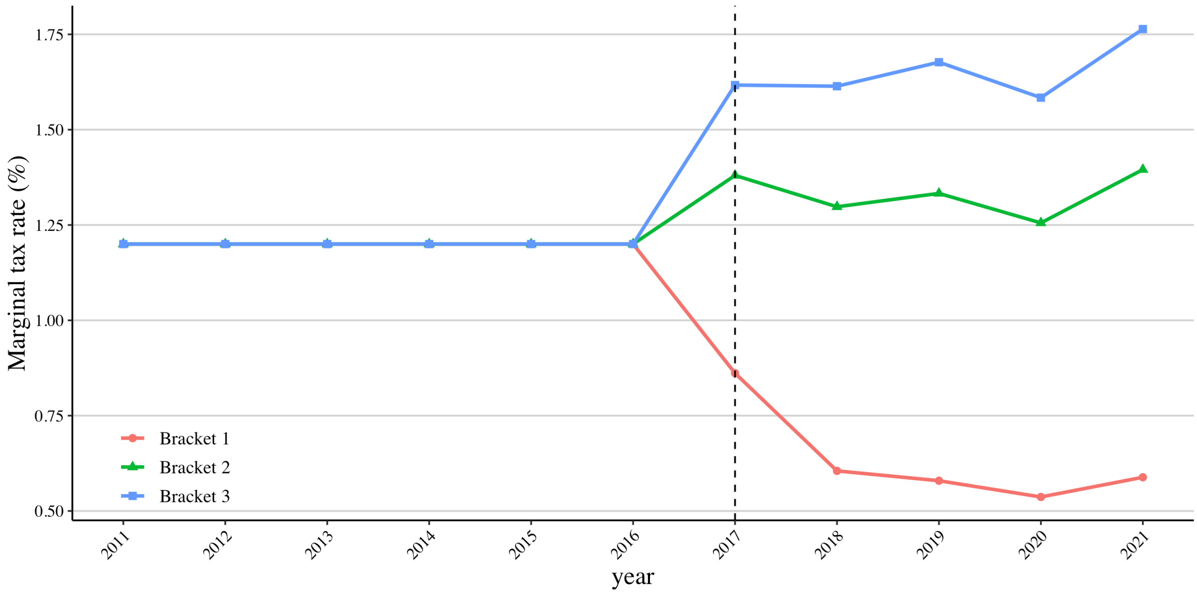

Figure 1: Marginal tax rate per wealth tax bracket. Notes: In the period before 2017, the marginal tax rate was equal for all wealth above the tax-exempt threshold. From 2017 onwards, differential marginal tax rates are introduced. The cut-off between the first and second bracket is at about €100,000 and the cut-off between the second and third bracket is at about €1,000,000.

The introduction of the progressive tax brackets has a significant impact on the incentive taxpayers face regarding their choice to hold and report taxable wealth. From 2018 onwards, taxpayers in the first bracket face a decrease of 52% in the benefit of taxes per euro of less reported taxable wealth.[21] Conversely, taxpayers in the top bracket see the tax benefit per reduced euro of reported taxable wealth increase by 38%. Taxpayers in the middle face a smaller increase in the incentive to report less taxable wealth, of 10%. Given the minor change for taxpayers in the middle bracket, this group will serve as the control group in the analysis.

---

[20] As noted earlier, box 3 taxes imputed capital income at a flat tax rate. The progressivity in the tax brackets introduced in 2017 stems from a change in imputed capital income over which the tax is levied. Assuming different levels of return on capital per bracket results in different effective marginal tax rates per bracket (Appendix 1).

[21] The percentage change in the marginal tax rate per bracket compares the flat rate of 1.2% to the average marginal tax rate per bracket in the years 2018-2021.

In the years around the reform, other taxes remained relatively stable. This feature of the 2017 wealth tax reform makes it attractive for an empirical analysis, as the risk of spill-over effects from other tax regime changes on the amount of taxable wealth are limited. A remaining potential confounding factor, however, is a change in the tax exemption for gifts spent on owner-occupied housing. Prior to 2017, people below the age of 40 were allowed to receive a tax-free gift of €53,016 under the condition that it is spent on purchasing or renovating owner-occupied housing. From 2017 onwards, this tax-exempt amount was increased to (slightly above) €100,000. In the empirical analysis, I assess for the sensitivity of my results to exclusion of all individuals who either received or gave a gift using this exemption. No other changes in gift and bequest taxation pose a risk to for confoundedness.

There is a small possibility for anticipation effects in the year 2016. The first moment taxpayers were able to learn about the wealth tax reform to be implemented by the 1st of January 2017 was in September 2015, when the Tax Plan for the year 2016 was published (Ministerie van Financiën, 2015). This leaves a period of four months for potential anticipation effect to accumulate (between September 2015 and January 2016). In the analysis, I find no evidence for such anticipation effects.

3.2 Avoiding the Dutch wealth tax

There are several ways in which taxpayers can reduce their taxable wealth. A real response would be to save less and/or consume more out of wealth. However, taxpayers may not resort to such real responses if there are possibilities at hand to reduce taxable wealth through reporting responses (Slemrod, 2001). In this study, I dedicate special attention to two of these possibilities (non-exhaustive).

First, taxpayers can reduce their taxable wealth (and remain eligible for allowances) by shifting wealth from box 3 to box 2 (Ministerie van Financiën, 2023). While box 3 taxes the stock of wealth at a rate of around 1.6% in the top bracket, box 2 taxes realised capital gains at around 25%. Assets with a low rate of return (such as savings deposits) will hence face a significantly lower tax burden when held in a closely held firm.[22] An additional benefit of holding wealth in box 2, is that the taxpayer has discretion over the timing of realised capital gains. Downside of holding wealth in box 2 are administrative costs and regulation limiting accessibility of wealth held in a closely held firm. Illustrative of the firm as a tax shelter, Alvaredo and Saez (2009) show that the introduction of a wealth tax in Spain caused taxpayers to shift wealth to closely held firms in order to avoid the new wealth tax and Alstadsæter and Kopczuk (2014) provide evidence that closely held firms are used as tax shelters in Norway. In the Netherlands, Lejour et al. (2022) report an increasing amount of closely held firms that are used for the purpose of holding wealth, possibly to avoid taxation in box 3. Especially in the

---

[22] A savings deposit of €100.000 that pays an interest of 2% will result in a tax liability of €1.617 when held in box 3 and taxed in the third bracket at a rate of 1.617% (2017). When the same deposit is held in box 2 (a closely held firm), it will result in a tax liability of €400 (€100.000 principal * 2% interest * 20% low corporate income tax rate). When the savings deposit pays an interest of 0%, the tax liability is €0 when held in a firm, but would still be taxed at €1.617 when held in box 3.

years 2016-2018, the amount of wealth held in this type of firms rises quickly, indicating that the tax reform in 2017 might have stimulated taxpayers to shelter more wealth in box 2.

Second, taxpayers can reduce their taxable wealth by giving it to others through gifts, such as their children. Such a response can be seen as a real response, but it does not have to be. In the Netherlands, it is possible to gift only on paper. In such a case, the gifted wealth is not transferred, but a debt arises where the giver owes the gifted amount to the receiver. The giver must pay interest over this debt. In the tax filing, the wealth that is gifted on paper is taxed at the receiver, hence lowering the taxable wealth of the giver.

A potential margin of response not studied here, is allocating more wealth to pension savings and investments accounts. Under certain conditions, wealth dedicated for pension is exempt from wealth taxes. Such pension wealth is not directly observed in the data and can hence not be studied. Earlier work by Chetty et al. (2014) using Danish data, shows that people barely respond to financial incentives that may increase pension savings.[23] Note that the absence of specific estimates on changes in pension savings does not mean this margin of response is excluded in the elasticity of taxable wealth. A reallocation of more wealth to pension accounts will still reduce the amount of taxable wealth.

4. Data

In Section 4.1, I describe the data sets used in the analysis and how my measure for taxable wealth is constructed. In Section 4.2, I describe what filters are applied to clean the data and summary statistics of the resulting dataset sample.

4.1 Household wealth data

Data on the wealth of Dutch households is drawn from the VEHTAB dataset of Statistics Netherlands. The dataset contains the full population of Dutch households. For each household, data on savings, financial assets, real estate, business wealth and other possessions are provided. These also include bank deposits abroad. Debt is subdivided in mortgage, student, and other debt. This decomposition of net wealth allows me to categorise net wealth into its fiscal subcategories: box 1 (owner occupied housing minus mortgage debt), box 2 (substantial ownership) and box 3 (all other net wealth).

Statistics Netherlands receives the data on wealth components from the Dutch Tax Authority in case the taxpayer has a tax liability in box 3. If the net wealth of the taxpayer is lower than the tax-exempt amount, filing taxes in box 3 is not obligatory. For these households, Statistics Netherlands draws the information from data that Dutch financial institution are obliged to report. Real estate values are from the housing valuation models used for municipal taxes (*Waardering onroerende zaken* , WOZ-value) and student debt is reported by the national authority responsible for providing and registering

[23] Each dollar of government subsidy is estimated to increase pension savings by 1 cent.

student debt (*Dienst Uitvoering Onderwijs*). For taxpayers with substantial ownership, information from the company's balance sheet is used to determine the business wealth. Real estate owned by businesses is valued at the WOZ-value. The value of the shares owned by the taxpayer with substantial ownership are based on the fiscal value of the firm, drawn from corporate tax data, divided across the owners based on their share in the firm. Ownership shares are observed in the corporate tax, or in the data of the Dutch Chambre of Commerce.

The taxable wealth per household is divided by two if the household is a couple.[24] Dividing wealth of couples by two ensures that the same tax bracket cut-offs can be applied to household taxable wealth of both singles and couples. As such, I assume that fiscal couples equally split their taxable wealth in their tax filing, which is the optimal choice under the presence of progressive tax brackets. Because an equal split is usually the optimal choice, the tax authority assumes an equal split when the fiscal couple does not state otherwise (Ministerie van Financiën, 2015).

While the household wealth data spans the period 2006-2021, I restrict the sample to the period 2011-2021. This restriction is necessary, as other datasets matched to the wealth data for filters and controls are only available from 2011 onwards. A second motivation for the period restriction is data quality. Student debt and "other" debt, such as consumption debt, stock financing, bonds, debt for real estate financing other than the residency of occupation and tax or allowance debt, is only fully observed starting from 2011. Since this debt is deductible in box 3, the box 3 tax liability will be overstated in the years before 2011. Healthcare debt is only fully observed starting from 2014, but only constitutes a small amount. Furthermore, a Dutch tax amnesty program caused a large inflow of foreign hidden wealth in the years 2009 and 2014 (Leenders et al., 2023). These inflows could distort wealth-growth for a part of the top 10% wealthiest households. To prevent these methodological breaks and noise to distort the results, I base my estimates on the years 2011-2021. The restriction to the period 2011-2021 is further motivated by the availability of other data, such as data on bequests and gifts, which is only available from 2011 onwards.

The dataset offers an advantage over other empirical studies that use smaller samples. The only estimate for the Netherlands is based on a sample of 10.000 households and relies on reweighting to make the sample representative for the Dutch population (Zoutman, 2018). Using the full population of tax return data gives more statistical power and, unlike survey data, does not suffer from underrepresentation of top wealth (Vermeulen, 2016). Data from Statistics Netherlands is not top coded.

A disadvantage of the wealth data at the household level, is that wealth taxable in box 3 is not perfectly observed in the data. To approximate taxable wealth, I add bank savings and deposits, investments, real estate other than owner-occupied housing and other wealth together and subtract student debt and debt other than mortgage debt. This measure for taxable wealth does not distinguish

[24] The status of couple is assigned if the variable INHSAMHH in the INHATAB dataset describes the composition of the household as being a couple, which is the case under the codes 21, 22, 31, 32, 33, 51, 52, 53, 54.

between taxable bank deposits and non-taxable “green” bank deposits (or investments). A shift from ordinary to green deposits as a response to a change in the marginal wealth tax cannot be observed. This response margin, however, is limited by the maximum amount of green investments that is tax exempt: €57.385 per person in 2017. Furthermore, this policy was in place both before and after the tax reform in 2017, meaning it does not create a discontinuous change in the tax policy, simultaneous to the change in marginal tax rates.

4.2 Sample restrictions

The final sample used in the analyses is subject to certain filters that are aimed at reducing noise and unrepresentative observations. First, I drop all taxpayers who never paid wealth taxes in box 3 during the period 2011-2021. As these taxpayers are not obliged to report their taxable wealth in their tax filings, their data is less reliable.

Second, I exclude individuals who received a bequest during the studied period. These bequests may result in substantial changes in wealth that are unrelated to the tax rate. They are, however, likely to be correlated with initial wealth. Previous studies have documented that individuals with higher wealth tend to receive larger bequests (Elinder et al., 2018), also specifically in the Netherlands (Groot et al., 2019). To prevent this factor to influence wealth growth, I drop households with individuals who received a bequest during the studied period. Information on bequests is drawn from the VRKTAB dataset, which documents bequests and bequest taxes. It is not obliged to report a bequest to the tax authority when the amount is below the tax-exempt amount.[25] These smaller bequests are not observed and hence cannot be filtered out.

Third, to prevent the change in the tax-exempt amount for gifts spent on owner-occupied housing to confound the results, I drop (households with) individuals who either give or receive a gift of this kind. These individuals can be identified using the SCHTAB dataset from Statistics Netherlands, which documents the type of tax exemption used for gifts reported in the tax filing. Taxpayers who give or receive other types of gifts are retained, as such behaviour can be an outcome of interested.

Fourth, to prevent imputation methods to affect the results, I drop households with individuals for which the data is not directly observed from the tax filing. If Statistics Netherlands does not observe a tax filing, taxable wealth is imputed based on a tax filing in the previous year, or the year before.

Finally, taxpayers who changed status between homeowner or renter during the studied period are dropped. Buying or selling owner-occupied housing is likely to coincide with a (large) change in wealth. Dropping these taxpayers will hence reduce noise in the data. I cannot identify taxpayers who are already a homeowner and switch to another residence. The buying of a new residence and selling

[25] The maximum tax-exempt bequest in 2017 was €638,089 for partners, €20,209 for children, €60,621 for sick or disabled children, €47,859 for parents and €2,129 for people with another relation to the deceased person.

the old residence can again cause (large) changes in wealth, but the status of the taxpayers will remain "homeowner" in the data.

Descriptive statistics for the resulting sample are presented in Table 1. The table shows taxpayers in the first bracket mainly hold wealth in the form of bank savings and only 6% owns shared in a closely held firm. In the second bracket, taxpayers more often hold other asset types, but still up to 25% of the taxpayers in the second bracket only own bank savings. Firm owners are more common in the second bracket, at 15% of the taxpayers. In the third bracket, savings is no longer the largest asset category on average. Instead, more wealth is held in investments. Within the third bracket, 41% of the taxpayers own a closely held firm and 96% own the house they live in. Age increases for higher brackets, but even in the first bracket, the 75% is older than 50. People who pay wealth taxes are hence relatively old.

Table 1: Descriptive statistics

	Bracket 1				**Bracket 2**				**Bracket 3**			
	P25	**P50**	**Mean**	**P75**	**P25**	**P50**	**Mean**	**P75**	**P25**	**P50**	**Mean**	**P75**
A: Taxable wealth in box 3												
Total net wealth	3	4	5	6	12	17	23	28	110	143	201	214
Savings	3	6	7	9	7	15	19	25	16	42	75	923
Investments	0	0	1	1	0	2	10	11	15	90	160	192
Real-estate	0	0	1	0	0	0	9	6	0	19	86	82
Other wealth	0	0	0	0	0	0	3	0	0	2	34	30
Student debt	0	0	0	0	0	0	0	0	0	0	0	0
Other debt	0	0	1	0	0	0	4	0	0	0	44	20
B: Other wealth												
Subst. ownership	1	12	52	48	3	30	115	109	13	106	421	365
Residence value	20	26	30	35	24	33	38	46	36	55	68	84
Mortgage debt	2	10	12	19	0	2	12	17	0	0	20	18
C: Demographics												
Single (%)	-	-	31	-	-	-	37	-	-	-	44	-
Firm owner (%)	-	-	6	-	-	-	15	-	-	-	41	-
Homeowner (%)	-	-	83	-	-	-	88	-	-	-	96	-
Minor child (%)	-	-	37	-	-	-	29	-	-	-	22	-
Age	50	60	60	70	54	64	64	73	59	68	68	76
Observations (household-year)		7,155,873				2,971,233				118,773		

Notes: This table presents descriptive statistics for the wealth and demographics of taxpayers in the sample, categorised by bracket over the period 2011-2021. The sample excludes taxpayers with imputed tax returns, who have received a bequest or gift for owner-occupied housing or have changes status between homeowner and renter. In case of couples, wealth is divided by two, to represent taxable wealth per taxpayers. Wealth is divided by €10,000 and rounded to round numbers.

5. Methodology

The main methodology in this paper follows the difference-in-difference-in-differences (DiDiD)[26] approach proposed by Jakobsen and Søgaard (2022). This approach is specifically designed to estimate

[26] I follow Jakobsen and Søgaard (2022) and refer to their approach as a DiDiD estimator. It should be noted that Jakobsen and Søgaard (2022) show that estimators described as difference-in-differences estimators in previous literature should in fact also be interpreted as DiDiD estimators, as they rely on the same estimation equation structure and identifying assumptions. Hence, despite a difference in terminology, the regression technique used

tax elasticities when variation in the marginal tax rate between brackets is created by a tax reform, as is the setting in this paper. In section 5.1, I first present the econometrical model after which I discuss how it addresses identification issues that are common in the tax literature. Section 5.2 dedicates attention to how the estimated local effects should be interpreted. Bunching estimates are not considered in detail, as the data shows no apparent bunching at the cut-offs between the brackets (Appendix 10.2).

5.1 Main estimation method

Following the standard literature (Auten and Carroll, 1999; Gruber and Saez, 2002; Saez et al., 2012; Kleven and Schultz; 2014; Jakobsen and Søgaard, 2022) the regression equation estimated in this paper is a two-stage-least-squares regression with the reduced form,

$$\Delta lnW_{it} = \beta_0 + \boldsymbol{\beta_1 W_{it-\mathrm{k}}} + \beta_2 D_{it}^{reform} + \beta_3 \Delta ln(1 - \tau_{it-k}^{p}) + v_{it}\,, \qquad (1)$$

where ΔlnW_{it} is the change in log taxable wealth for taxpayer *i* during period *t*. The term $\boldsymbol{W_{it-\mathrm{k}}}$ denotes a vector of dummies for bins of taxable wealth in the base-year. The dummy D_{it}^{reform} takes the value 1 in the period after the reform and 0 for the pre-reform period. Finally, $\Delta ln(1 - \tau_{it})$ captures the change in the log net-of-tax rate, making β_3 the variable of interest. The change in the log marginal tax rate is instrumented by the 'predicted' change (τ_{it-k}^{p}) that would occur if the taxable wealth of taxpayer i in pre-reform period *k* had remained constant during the reform period. The error term is denoted by v_{it}.[27]

The main coefficient of interest, β_3, has two interpretations. First, it can be interpreted as the elasticity of taxable wealth with respect to the net-of-tax rate thanks to the specification in logs. Second, it can be interpreted as a semi-elasticity with respect to a 1 percentage point cut in the wealth tax rate (Agrawal et al., 2020; Brülhart et al. 2022). As the marginal tax is very small, $ln(1 - \tau_{it})$ will be almost equal to $-\tau_{it}$. This means β_3 can also be interpreted as the percentage change in taxable wealth given a 1 percentage point cut in the wealth tax rate (semi-elasticity). For simplicity of interpretation, I follow Brülhart et al. (2022) and report $\beta_3/100$ in the main presentations of the results, as the percentage change in the taxable wealth over the period *t+s* given a 1 percentage point cut in the wealth tax rate.

The specification in logs would imply observations with negative or zero taxable wealth are treated as missing (the logarithm of zero or negative numbers is undefined). Negative net wealth in box 3 can occur when debt exceeds gross assets. To retain these observations, I set negative net wealth to zero and follow Gelber (2014) by adding 1 euro to the taxable wealth of each taxpayer. Given the magnitudes of wealth, a single euro will not distort the data. Preventing to treat observations of zero or negative taxable wealth as missing values avoids bias from non-random missing observations (Heckman, 1979). The occurrence of zero taxable wealth can be non-random, as an increase in the tax

---

in this study is equivalent to the "difference-in-difference" estimators in the standard tax elasticity literature (Auten and Carroll, 1999; Gruber and Saez, 2002; Saez et al., 2012; Kleven and Schultz; 2014).

[27] Note that I do not cluster standard errors. As neither the sampling nor the treatment assignment is based on clusters of taxpayers, clustering standard errors is unnecessary and would lead to unnecessarily large standard errors (Abadie et al., 2023).

rate may induce taxpayers to shelter taxable wealth to such a degree that the final net taxable wealth in box 3 becomes zero or negative.[28] Treating these taxpayers as missing, will lead to an underestimation of the response to the change in wealth tax. Except for Ring (2021)[29], none of the empirical papers on the elasticity of taxable wealth presented in Table 1, discuss the issue of non-random missing observations created by taking the log of taxable wealth. It is thus unclear whether those studies take this possible source of bias into account.

The specification of Equation (1) addresses a number of identification issues that have plagued the previous tax elasticity literature. Below I discuss these issues one by one, explaining how my specification addresses them. To aid the discussion, I start by setting out a simple model[30] following Feldstein (1999), Gruber and Saez (2002) and Saez et al. (2012) and build up Equation 1 from there. A household *i* reports taxable wealth W_{it} in year *t* and is subject to a marginal tax rate τ_{it}. The amount of taxable wealth the taxpayer reports, responds to the marginal net-of-tax rate with elasticity *e*. As such, reported taxable wealth equals $W_{it}^0 * (1-\tau_{it})^e$, where W_{it}^0 denotes the potential taxable wealth in the case of a zero marginal tax rate. Lastly, taking log differences gives,

$$\Delta lnW_{it} = e\Delta \ln(1-\tau_{it}) + ln\Delta W_{it}^0, \quad (2)$$

where the coefficient *e* captures the effect of a change in the marginal net-of-tax rate on taxable wealth. Differencing removes individual fixed effects (Gruber and Saez, 2002; Gelber, 2014) and removes serial correlation caused by the fact that wealth is a stock variable (Fagereng et al., 2020; Alstadsæter et al., 2022). Because the specification is in logs, *e* can be interpreted as an elasticity (Gruber and Saez, 2002; Saez et al., 2012). The change in potential taxable wealth ($ln\Delta W_{it}^0$) functions as an error term and is captured in v_{it} of Equation (1). Note that Equation (2) assumes no income effects, following the theoretical literature (Saez et al., 2012; Saez and Stantcheva, 2018) and most empirical applications (Seim, 2017; Zoutman, 2018; Duran-Cabré et al., 2019; Agrawal et al., 2020; Jacobsen et al., 2020; Brülhart et al., 2022; Alstadsæter et al., 2022; Garbinti et al., 2023).

5.1.1 Endogeneity of the marginal tax rate

The first issue to tackle, is the endogeneity of the marginal tax rate. A standard OLS estimation of Equation (2) will be biased because the change in the marginal net-of-tax rate is a function of taxable

---

[28] A memo from the Dutch Ministry of Finance (Ministerie van Financiën*,* 2023) indicates that wealth tax avoidance structures that reduce the wealth taxable in box 3 to zero, are not uncommon.

[29] Ring (2021) adds 100.000 NOK to the taxable wealth of each taxpayer.

[30] The original model concerns taxable income, assuming taxpayers derive utility from disposable income (consumption), but disutility from reported income, as generating income is costly and reported income is subject to taxation (Feldstein, 1999; Gruber and Saez, 2002; Saez et al., 2012). Extending the model to taxable wealth requires the assumption that taxpayers derives utility from "disposable" wealth and disutility from reported wealth. This is the same assumption as made in Saez and Stantcheva (2018) who motivate wealth in the utility function by a bequest motive (Cagetti and De Nardi, 2007; De Nardi, 2014), entrepreneurship (managing wealth) (Qaudrini, 1999; 2000) and service flows from wealth, such as a sense of security or usage of a capital good (Piazzesi et al., 2007; Stokey, 2009; Kiyotaki et al., 2011).

wealth and hence correlated with the error term by definition (Saez et al., 2012; Weber, 2014; Jakobsen and Søgaard, 2022). Furthermore, taxable wealth is only observed after the household responds to the change in the marginal tax rate. If the households respond by decreasing its wealth to such a degree that it ends up in a different tax bracket, then this household would erroneously be assigned to the control group while it received treatment.

To resolve the endogeneity issue, I follow the standard approach in the literature and determine the change in the marginal tax rate based on taxable wealth in a pre-reform year (Auten and Carroll, 1999; Gruber and Saez, 2002; Kleven and Schultz, 2014, Weber, 2014, Jakobsen et al., 2020; Jakobsen and Søgaard, 2022). Specifically, I instrument the change in the marginal tax rate by the change in the marginal tax rate a taxpayer would incur if its taxable wealth in the pre-reform period had remained constant during the reform period. This gives the reduced form equation,

$$\Delta lnW_{it} = e\Delta ln(1-\tau^{p}_{it-k}) + ln\Delta W^{0}_{it}. \quad (3)$$

where $\Delta ln(1-\tau^{p}_{it-k})$ is the 'predicted' change in the marginal tax rate had taxable wealth in pre-reform year *t-k* remained constant.

5.1.2 Validity of the instrument

The second issue to tackle, concerns the validity of the instrument. For *e* in Equation (2) to yield an unbiased intention-to-treat (ITT) effect in a reduced form estimation, the instrument must be uncorrelated with the potential outcome (independence). To yield a local average treatment effect (LATE) additionally, the assumption that the instrument only affects taxable wealth through the marginal tax rate (exclusion restriction) and a significant first stage must hold (Angrist et al., 1996; Angrist and Pischke, 2008). Formally, the independence assumption reads,

$$Cov(\Delta ln(1-\tau^{p}_{it-k}), \Delta lnW^{0}_{it}) = 0. \quad (4)$$

The independence assumption will fail if the level of wealth in the base year, upon which the instrument is based, correlates with potential wealth growth in the reform period.[31] Two common causes of such correlation have been identified in the tax elasticity literature (Saez et al., 2012; Weber, 2014; Jakobsen and Søgaard, 2022).

The first common cause for instrument endogeneity is mean reversion. Mean reversion occurs if taxable wealth growth contains a transitory element, resulting in negative serial correlation: a year with high wealth growth may be more likely to be followed by a period with low wealth growth and vice versa. Such negative serial correlation causes the independence assumption to fail, as a higher predicted marginal tax rate not only affects taxable wealth through the change in tax incentive, but also captures the negative serial correlation between initial wealth and subsequent wealth growth. Mean

[31] Note that the independence assumption is equivalent to the common trend assumption in a regular DiD setting. In absence of treatment, the change in the outcome variable (potential wealth growth) is assumed to be parallel between the treatment group and the control group.

reversion is shown to be a source of significant bias in empirical studies on the elasticity of taxable income (Weber, 2014).

The second common cause for instrument endogeneity is differential growth trends. Households with higher initial wealth might exhibit higher wealth growth than households with less initial wealth. In that case, the change in the marginal tax rate instrumented based on initial wealth will not be independent of potential wealth growth, as a higher level of initial wealth is correlated with wealth growth in the following period. Indeed, empirical studies have shown that those who have more wealth are able to generate higher returns on their wealth (Bach et al., 2020; Fagereng et al., 2020), possibly thanks to economies of scale, better access to investment opportunities or better financial management (Piketty, 2014).

To resolve the issues of mean reversion and differential growth trends, I follow the literature and control for initial wealth (Auten and Carrol, 1999; Gruber and Saez, 2002; Saez et al., 2012; Kleven and Schultz, 2014; Jakobsen and Søgaard, 2022). To see why controlling for initial wealth resolves these issues, note that the instrument τ^p_{it-k}, is a function of initial wealth, W_{it-k}, and the change in the marginal tax rate at a given taxable wealth level caused by the reform. It is the correlation between initial wealth and subsequent wealth growth that causes the mean reversion and differential growth trend issues. Controlling for initial wealth hence removes this part of the variation of the instrument. What remains, is (exogenous) time variation in the marginal tax rate caused by the reform within each level of taxable wealth. Because the identifying variation comes from variation over time, this method is only possible if panel data across multiple years are available.

When using only time variation in the marginal tax rate, the identifying assumption becomes one of constant wealth growth trend *differentials* instead of a common growth trend (Saez et al., 2012; Jakobsen and Søgaard, 2022).[32] In the absence of a change in the marginal tax rate, the differentials in taxable wealth growth across the taxable wealth distribution ought to remain constant. Formally,

$$Cov(\Delta ln(1-\tau^p_{it-k}), \Delta ln W^0_{it} | W_{it-k}) = 0. \tag{5}$$

For the conditional independence assumption (5) to hold, we no longer need parallel wealth growth trends between taxpayers with different levels of initial wealth. Instead, we need the differences in wealth growth trends between taxpayers with different levels of initial wealth to remain constant over time. Which holds if the dynamic process of wealth growth can be explained by,

$$E(\Delta\, ln\, W^0_{it} | W_{it-k}) = g(W_{it-k}) + \delta_t, \tag{6}$$

[32] As shown by Jakobsen and Søgaard (2022) the assumption of constant growth differential is at the basis of all previous studies that have aimed to resolve the issues of mean reversion and differential growth trends with initial wealth/income controls. Jakobsen and Søgaard (2022) are, however, the first to propose a method for graphical validation. Even though not referenced in Jakobsen and Søgaard (2022), it should be noted that Gelber (2014) already presents a graphical inspection that is similar to the approach suggested in Jakobsen and Søgaard (2022).

where $g(W_{it-k})$ is a function that explains potential wealth growth based on initial wealth, and δ_t is a common growth trend for each period. The equation allows potential wealth growth to differ each period through δ_t, but requires the difference in wealth growth between taxpayers with different base-year wealth levels to differ by a constant function over time. The function $g(W_{it-k})$ is specified non-parametrically by adding a set of initial wealth dummies, $\boldsymbol{W_{it-\mathbf{k}}}$, to the regression function.

5.1.3 Graphical validation and intuition

The treatment effect identified by the DiDiD set up and the plausibility of the constant trend differentials assumption can both be inspected graphically (Jakobsen and Søgaard, 2022). For illustration, and to get intuition into the method, I give an illustration of the DiDiD set up in Figure 2.

Figure 2 relates the changes in log taxable wealth over the pre-reform and reform period to initial taxable wealth at the start of each period. In both examples, the marginal tax rates are equal above and below the bracket cut-off (dashed line) and remain unchanged during the reform period. Example A considers a situation where the marginal tax rate in the bottom bracket falls in the reform period, while Example B considers a situation where the marginal tax rate in the top bracket rises. Under the assumption that the growth trend differentials would have remained constant in absence of a change in the marginal tax rate, the treatment effect is given by the difference in wealth growth between the reform period and pre-reform period for taxpayers in the treated brackets (identification region). The examples feature smaller effects closer to the cut-off, resembling smaller responses from taxpayer who are unaware of the exact location of the cut-off or optimization frictions (Chetty, 2012; Rees-Jones and Taubinsky, 2020).

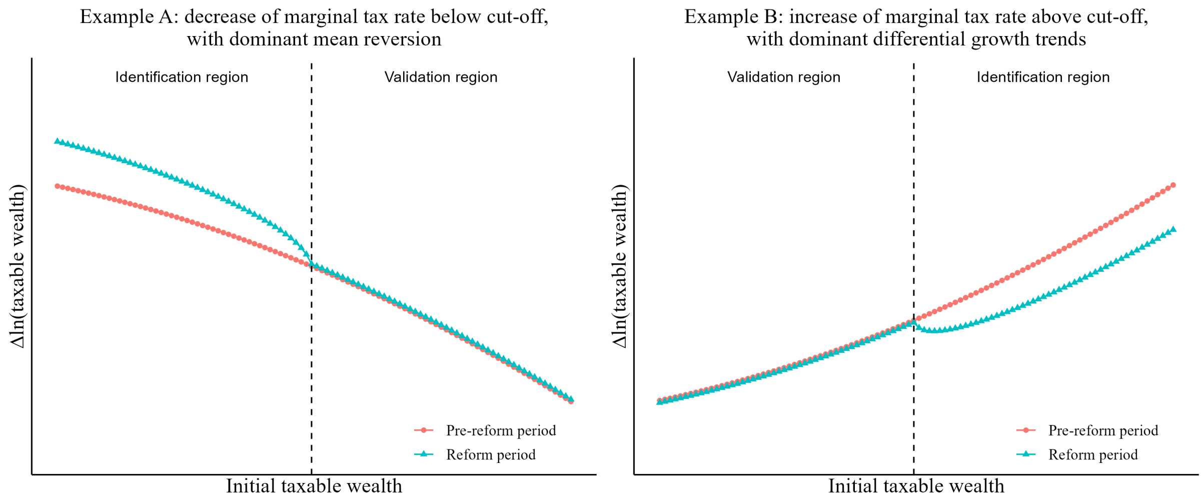

***Figure 2:* Illustration of identification strategy.** Notes: The figures relate the changes in log taxable wealth (ΔlnW_{it}) during the pre-reform and reform period to initial taxable wealth (W_{it-k}) at the start of each period. Under the assumption that the growth trend differentials would have remained constant in absence of a change in the marginal tax rate, the treatment effect is given by the difference in wealth growth between the reform period and pre-reform period for taxpayers in the treated brackets (identification region). To assess the plausibility of the identifying assumption, we can check if the trend differentials remain constant between the pre-reform and reform period among untreated taxpayers (validation region).

To assess the plausibility of the identifying assumption, we can check if the trend differentials remain constant between the pre-reform and reform period among untreated taxpayers (validation region). As such, the validation region acts as a placebo test for the identifying assumption (Jakobsen and Søgaard, 2022). In Examples A and B in Figure 2, there is a close match between the trend differentials in the validation region, providing evidence the identifying assumption is plausible.

The slope of the trend differentials indicates the presence of mean reversion or differential growth trends. If there is (dominant) mean reversion, the relation between initial wealth and subsequent wealth growth is expected to be negative, as illustrated in example A. If there are (dominant) differential growth trends, where people with more wealth can generate higher returns, the relation is expected to be positive, as illustrated in Example B. A flat relation could either mean there is no mean reversion or differential growth trends, or that both effects offset each other (Jakobsen and Søgaard, 2022).

5.1.4 Controlling for covariates

A possible way for the constant trend differentials assumption (6) to fail, is when the characteristics of taxpayers in a certain initial wealth bin change between the pre-reform period and reform period. Changes in the characteristics of taxpayers in certain parts of the taxable wealth distribution can bias the results if these characteristics influence wealth growth. To prevent such bias, I include control variables for taxpayer characteristics.

Control variables are incorporated through an inverse probability weighting (IPW) strategy following DiNardo et al. (1996) and Jakobsen and Søgaard (2022).[33] The inverse probability weighting method entails the reweighting of observations in the pre-reform period, to better match the observed characteristics of units in the reform-period. Specifically, I run a logit regression for each initial wealth bin separately, estimating the probability a household is in the reform period, $P(D_{it}^{reform} = 1|\boldsymbol{X}_{it}, W_{it-k})$, where D_{it}^{reform} is a dummy for being in the reform period, $\boldsymbol{X}$ is a vector of taxpayer characteristics and $W_{i,t-k}$ denotes set a of dummies per base year wealth bin. The observations in the pre-reform period are subsequently reweighted by the weights,

$$weight_i = \frac{P\left(D_{it}^{reform} = 1\middle|\boldsymbol{X}_{it}, W_{i,t-k}\right)}{P\left(D_{it}^{reform} = 0\middle|\boldsymbol{X}_{it}, W_{i,t-k}\right)}. \tag{7}$$

Intuitively, more weight is assigned to taxpayers in each wealth bin during the control period who are more similar in characteristics to taxpayers in the same wealth bin during the treatment period. This procedure controls for the change in the composition of taxpayers in each wealth bin between the pre-reform and reform period. The assumption of constant trend differential becomes an assumption of constant trend differentials conditional on the used covariates in the weighting procedure,

$$E(\Delta \ln W_{it}^{0}|W_{it-k}, \boldsymbol{X}_{it}) = g(W_{it-k}) + \delta_t. \tag{8}$$

[33] A similar reweighting technique is applied to estimate the elasticity of taxable income by Kawano et al. (2016).

In the covariates vector $\boldsymbol{X}$, I include dummies for being a single person household, a couple or other, age group[34], having children, a dummy if the household has substantial ownership in a firm and share of wealth held in financial investments, real estate, and other debt, as well as the share of other debt in gross wealth in the base year. For the control variables that might be affected by the change in the wealth tax, specifically the composition of wealth and substantial ownership, I only use their values in the base year (pre-reform) to ensure any response to the treatment is not included in the controls. Including controls variables that are affected by the treatment ('bad' controls), could otherwise bias my estimates (Rosenbaum, 1984; Angrist and Pischke, 2008).

The advantage of IPW compared to simply including covariates in the regression is twofold. First, reweighting preserves the simple graphical representation of wealth growth per bin, while with use of covariates in the regression, the figure would plot the less intuitive residual wealth growth (Jakobsen and Søgaard, 2022). Second, reweighting creates a counterfactual for the treatment group that is within support of the data, while including covariates can lead to extrapolation (King and Zeng, 2006; Abadie et al., 2010; 2015).

5.1.4 Heterogeneity across the taxable wealth distribution

In addition to estimating the aggregate elasticity of taxable wealth for taxpayers in the first and third bracket, I also explore heterogeneity in the elasticity across the wealth distribution within these brackets. To allow for heterogeneous effects across the wealth distribution, I follow Jakobsen and Søgaard (2022) and estimate local linear regressions. For several points h in the initial wealth distribution, I estimate a weighted local linear regression where the second stage equation reads,

$$\Delta lnW_{it+s} = \beta_0 + \boldsymbol{\beta_1 W_{it-1}} + \beta_2 D_{it}^{reform} + \beta_3 \Delta ln\left(1-\tau_{it-k}^{p}\right) + \beta_4 \Delta ln\left(1-\tau_{it-k}^{p}\right) * (W_{it-1} - h) + v_{it}\,. \quad (9)$$

The interaction term between the change in the marginal net-of-tax rate and the relative position to point h is included to allow the slope of the elasticity to very across different points in the taxable wealth distribution. Taxpayers in the validation region are assigned a weight of 1, but taxpayers in the identification region are assigned triangular weights within an interval around each point h. Within the third bracket, for example, I use an interval of €200,000, resulting in weights following the formula,

$$weight = \max(200{,}000 - | W_{it-k} - h|, 0)/200{,}000\,. \quad (10)$$

The width of the interval used in the local linear regressions is a discretionary choice. A smaller interval shows more granular heterogeneity but comes at the cost of fewer observations and hence statistical power. I choose tighter intervals in the first bracket (€10,000 in width) compared to the third bracket, as the first bracket has more observations and spans a smaller section of the taxable wealth distribution. In the analysis, I test for sensitivity to the width of the bins.

[34] The age groups are defined as <25, 25-34, 35-44, 45-54, 55-64. 65-74, 75-84 and >85.

5.1.6 Summary of the main estimation approach

In summary, the marginal tax rate is an endogenous variable because it is a function of taxable wealth. Therefore, I instrument the change in the marginal tax rate based on taxable wealth in a pre-reform year (initial wealth). This instrument is still correlated with the error term if initial wealth is correlated with subsequent wealth growth. Commonly, such correlation occurs due to mean reversion and growth trend differentials.

Identification issues arising from both mean reversion and differential growth trends originate from differences in initial wealth. Therefore, the issue can be resolved by controlling non-parametrically for initial wealth. Controlling for initial wealth leaves only time variation caused by the tax reform as identifying variation. As such, the effect of a change in the marginal tax rate is identified by examining the wealth growth differentials of taxpayers in a treated region of the taxable wealth distribution compared to taxpayers in the same region of the distribution in the period before the tax reform.

Any remaining concerns about a change in the characteristics of wealth groups between the pre-reform and reform period are addressed by reweighting the control group through an inverse probability weighting procedure that controls for such changes. Hence, the underlying identifying assumption is that the differentials in taxable wealth growth between initial wealth levels would have remained constant over time in absence of the tax reform, conditional on a set of used control variables.

Finally, heterogeneity in the estimated elasticities can be studied by running local linear regressions, where the control group is held constant, while a subsample of the treatment group used is shifted along the wealth distribution.

5.2 Interpretation caveats

When interpreting the results from the methodology laid out above, it is important to keep several caveats in mind. First, the two-stage-least-squares (2SLS) estimates of Equation (1) yield a local average treatment effect (LATE) for the group of compliers (Angrist et al., 1996; Angrist and Pischke, 2008). [35] A limitation of instrumenting the change in the marginal tax rate, is that taxpayers whose response causes them to switch brackets are not in the group of compliers and hence not in the LATE (Saez et al., 2012; Jakobsen et al., 2020). A reduced-form estimate would yield an ITT effect that does include these taxpayers, but also taxpayers who did not truly receive treatment. The ITT effect or the LATE hence offer their own advantages and disadvantages. The LATE can underestimate the result

---

[35] It is useful to point out the other types of populations in the instrumental variable framework within the context of this study. The group of always-takers consist out of taxpayers who always end up in bracket 3 during the reform period when either assigned to bracket 2 or 3 based on their pre-reform wealth. Conversely, never-takers always end up in bracket 2. Defiers are taxpayers who would switch to treatment when assigned to the control group and would switch to the treatment group when assigned to the control group. Among other assumption mentioned earlier, it is assumed there are no defiers present in the data so the estimates can be interpreted as a LATE (Angrist et al., 1996; Angrist and Pischke, 2008).

when a relatively large portion of the response to a change in the marginal tax rate consists of taxpayers who switch to a lower bracket as a result of their response. For example, a taxpayer with €1,500,000 who sets up a closely held firm and shelters €1,000,000 as response to an increase in the tax rate, will subsequently switch to the second bracket. As a result, the taxpayers is not a complier, and the response is not included in the LATE. In this respect, the ITT is attractive, as it does include the response of this fictitious taxpayer. However, the ITT also captures the response by taxpayers who switch to another tax bracket for other reasons and are subsequently not truly treated. Including these taxpayers in the estimated effect results in an underestimation of the true treatment effect. In the results discussed below, I focus on the LATE but also report the ITT effect for all estimates. In Section 6.4.1 perform a descriptive analysis of what portion of non-compliance may be caused by a response to treatment.

Second, taxpayers in the second bracket are not a pure control group, as they face a smaller change in the marginal tax rate instead of no change at all. This setting can introduce a bias in the estimates when the (potential) elasticity in the treatment and control group differ (Saez, 2004; Saez et al., 2012). In short, the DiDiD estimator will fail to difference out the effect from the (smaller) treatment in the control group when the potential response of the treatment and control group to treatment are not equal (See Appendix 10.8 for a formal discussion). The potential bias arising from heterogeneity in the elasticity of taxable wealth between the treatment and control group is recognised but receives little attention in the tax elasticity literature (Kumar and Liang, 2020). Jakobsen and Søgaard (2022, p.11) only mention the issue in a footnote, stating their finding of heterogenous elasticities implies their estimates contain a bias as they have no pure control group, but "this is less of an issue in our setting, where the individuals in the validation region are only treated by small tax changes". They fail, however, to quantitatively justify the expected bias is indeed small enough to ignore. In this paper, I do provide such calculations for my estimates in Appendix 10.8.

Third, when interpreting the elasticity found near the bracket cut-off, households close the cut-off may exhibit a smaller behavioural response due to optimisation frictions or unawareness of the exact cut-off between the brackets (Chetty, 2012; Rees-Jones and Taubinsky, 2020; Jakobsen and Søgaard, 2022). An example of optimisation frictions in the context of wealth is the fluctuations in value of investments in shares. Shares are included in taxable wealth, but stock market fluctuations are uncertain, so taxpayers cannot target a very precise amount of taxable wealth at a given date.

Fourth, assuming complete absence of spill-over effects is unrealistic. Spill-over effects occur when the marginal tax rate in one bracket affects behaviour in another bracket. Specifically, the prospect of higher taxation in the third bracket for taxpayers in the second bracket may reduce the incentive to accumulate taxable wealth for taxpayers in the second bracket. A response to an increased tax rate in the third bracket in the form of a decrease in wealth growth in the third bracket relative to the second bracket would hence be underestimated when wealth growth in the second brackets lowered by the reform too. Similarly, taxpayers in the middle bracket may be induced to reduce their taxable wealth to fall within the first bracket with a lower tax rate. An increase in wealth growth in the third bracket relative to the second bracket would hence be overestimated. The issue of spill-over effects between brackets is well known in the tax elasticity literature (Saez et al., 2012; Jakobsen et al., 2020), but the literature has not been able to address this issue econometrically. This aspect should hence be kept in mind when interpreting the results.

6. Results

In this section, I describe the results following the main estimation method described above. Section 6.1 presents the estimated elasticities in the first and third bracket and explores heterogeneity. Section 6.2 reports robustness checks. Section 6.3 discusses response margins and mechanisms.

6.1 Elasticity estimates

The baseline estimates compare the pre-reform period 2012-2016 to the reform period 2016-2020. I present results using a period of 4 years as baseline, as this is the maximum length useable in the data, before the useable range of taxable wealth to estimate effects in the first bracket is limited by an increase of the tax-exempt threshold in 2021. In the analysis, I also report estimates using other period lengths.

6.1.1 Response in the bottom bracket

As a first step, I inspect the identifying variation caused by the tax reform. Figure 3 Panel A shows the change in the marginal net-of-tax rate caused by the tax reform for taxpayers in each initial taxable wealth bin, as predicted by the instrument. [36] In the pre-reform period, taxpayers across the range of taxable wealth show no change in the marginal net-of-tax rate. In the reform period, taxpayers with taxable wealth below the cut-off between bracket 1 and 2 (at €100.000) face a significant increase in the marginal net-of-tax rate (identification region), whereas the net-of-tax rate falls slightly for taxpayers above the cut-off (validation region). The shaded area around the line for the reform period reflects the minimum and maximum change in the marginal net-of-tax rate during the reform period, as it fluctuates slightly year by year. In the regression estimates, I use the average change in the marginal net-of-tax rate over the reform period, indicated by the point averages in Figure 3 Panel A.

In the second step, I inspect the taxable wealth growth differentials across the initial taxable wealth distribution. For this purpose, I estimate the following regression for the pre-reform and reform period separately,

$$\Delta lnW_{it} = \beta_0 + \boldsymbol{\beta_1} \boldsymbol{W_{it-1}} + v_{it}, \qquad (11)$$

where $\boldsymbol{W_{it-1}}$ is a vector of dummies for initial wealth bins with a width of €10.000, based on taxable wealth in 2012 for the pre-reform period and 2016 for the reform period. The bin at €150.000 is

[36] For the estimates in the bottom bracket, I consider taxpayers with initial taxable wealth above €40.000 (to prevent interference of the tax-exempt threshold) and below €300.000 (the minimum of the range used in the analysis of the third bracket). The remaining taxpayers are binned together in bins with a width of €10.000.

excluded in the regression, so $\boldsymbol{\beta_1}$ captures the wealth growth of taxpayers at each point in the initial wealth distribution relative to taxpayers in the €150.000 bin.

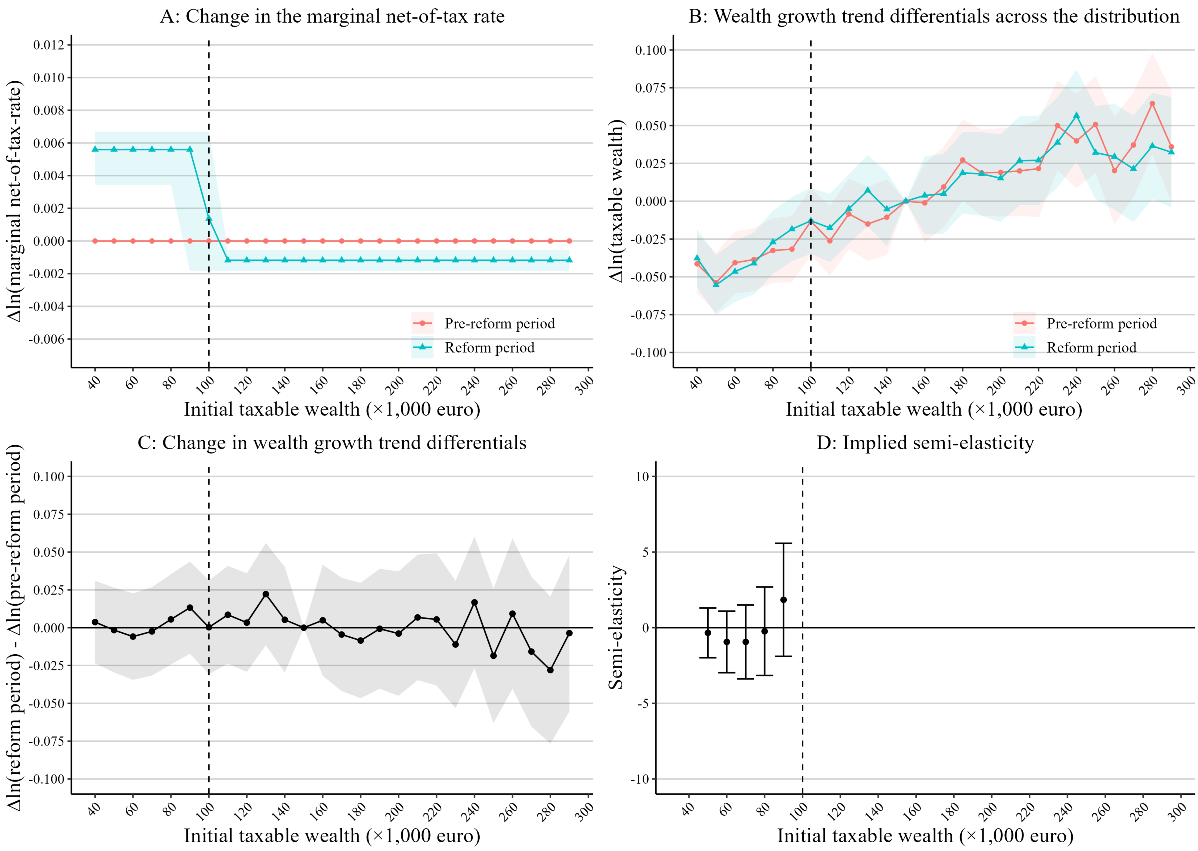

Figure 3: Response of taxable wealth to a change in the marginal tax rate in the bottom bracket. Notes: Panel A shows the change in the log net-of-tax rate as predicted by the instrument ($\Delta ln(1-\tau_{it-k}^{p})$) for the period 2012-2016 (pre-reform) and the period 2016-2020 (reform period). The points indicate the average change in the marginal net-of-tax rate during the period, with the shaded area indicating the fluctuations of the marginal tax rate during the period. Initial wealth is based on 2012 in the pre-reform period and 2016 in the reform period. Panel B shows the wealth growth trend differentials using Equation (11), where bin €150.000 is used as a reference. Panel C shows the change in the growth differentials, estimated following Equation (12). Panel D shows the implied semi-elasticities of taxable wealth with respect to a 1 percentage point cut in the wealth tax rate, estimated using Equation (9). The 95% confidence intervals are based on robust standard errors.

The plotted coefficients from regression (11) indicate that the constant trend differential assumption holds in the validation region (Figure 3, Panel B). In the validation region of the graph (above the cut-off), the differences in growth trends between each bin and the reference bin remain constant between the pre-reform and the reform period. Equivalently, taxpayers within the identification region show parallel growth differentials between the reform and pre-reform period, indicating taxpayers in the bottom bracket did not respond to a decrease in the marginal tax rate. The upward slopes of the lines indicates that differential growth trend is dominant over mean reversion: taxpayers with more initial wealth tend to accumulate more wealth subsequently.

In the third step, I quantify the change in the trend differentials per initial wealth bin between the pre-reform and reform period. Figure 3 Panel C shows the change in the taxable wealth growth differentials, estimated using,

$$\Delta lnW_{it+s} = \beta_0 + \boldsymbol{\beta_1 W_{it-1}} + \beta_2 D_{it}^{reform} + \boldsymbol{\beta_3 W_{it-1}} * D_{it}^{reform} + v_{it}, \quad (12)$$

with a dummy D_{it}^{reform} indicating the period and the interaction term $\boldsymbol{W_{it-1}} * D_{it}^{reform}$ capturing the change in the wealth growth differentials between the two periods per initial wealth bin.[37] The coefficients $\boldsymbol{\beta_3}$ plotted in Figure 3 Panel C reflect ITT effects. The interpretation of this figure is equivalent to a standard graphical representation of a DiD-estimate: when the coefficients in the validation region are statistically insignificant it provides confidence that the identifying assumption holds.[38] A formal test fails to reject the null-hypothesis that the coefficients in the validation region are jointly zero with a p-value of 0.607, providing formal proof the constant trend differentials assumption holds in the validation region. In the identification region, the change in the growth trend differentials remains flat, showing no significant ITT effect.

In the final step, I translate the effects found in Panel C to semi-elasticities while allowing for heterogeneity within the bracket, by estimating Equation (9) using 2SLS. At each initial wealth bin, I assign triangular weights with an interval width of €10.000 following Equation (10). Using this method, I find point-estimates that are statistically not distinguishable from zero and showing little heterogeneity (Figure 3 Panel D). The found confidence intervals range from a negative semi-elasticity of -3.2% to 5.6%. These estimates reflect the LATE for taxpayers who remain within the same bracket during the reform period (compliers).

6.1.2 Response in the top bracket

The steps presented above are repeated for the top bracket. Here, I limit the sample to taxpayers with initial taxable wealth above €300,000 and below €3,000,000, as above €3,000,000, the number of observations becomes too low for reliable estimates per bin. Households are binned into initial wealth bins with a width of €100,000. The bin-width is substantially larger than the €10.000 used for the bottom bracket, because the density of taxpayers is far smaller high in the taxable wealth distribution.[39]

As shown in Figure 4 Panel A, taxpayers with taxable wealth above the cut-off of €1,000,000 face a large decrease in the marginal net-of-tax rate (identification region), while taxpayers below the

[37] Note that Equation (12) is equivalent to the main estimation regression presented in the methodology section, but with the change in the marginal net-of-tax rate substituted by the interaction between the initial wealth dummy and the dummy for period.

[38] In DiD lingo, one would state that the insignificant coefficients in the pre-reform period indicate constant trends. Here, the insignificant coefficients in the validation region indicates constant trend differentials.

[39] Using bins with a width of €100,000 in a range up to €3,000,000, results in a minimum number of households in a given bin of 156. Restricting the upper limit of the range to €3,000,000, retains about 85% of the taxpayers in the top bracket (measured using wealth in 2016). In Section 6.2.3, I show choosing other upper limits has little effect on the estimated elasticities in Section 6.2.3.

cut-off face a smaller decrease in the marginal net-of-tax rate (validation region). In the pre-reform period, the marginal tax rate remained unchanged for all taxpayers.

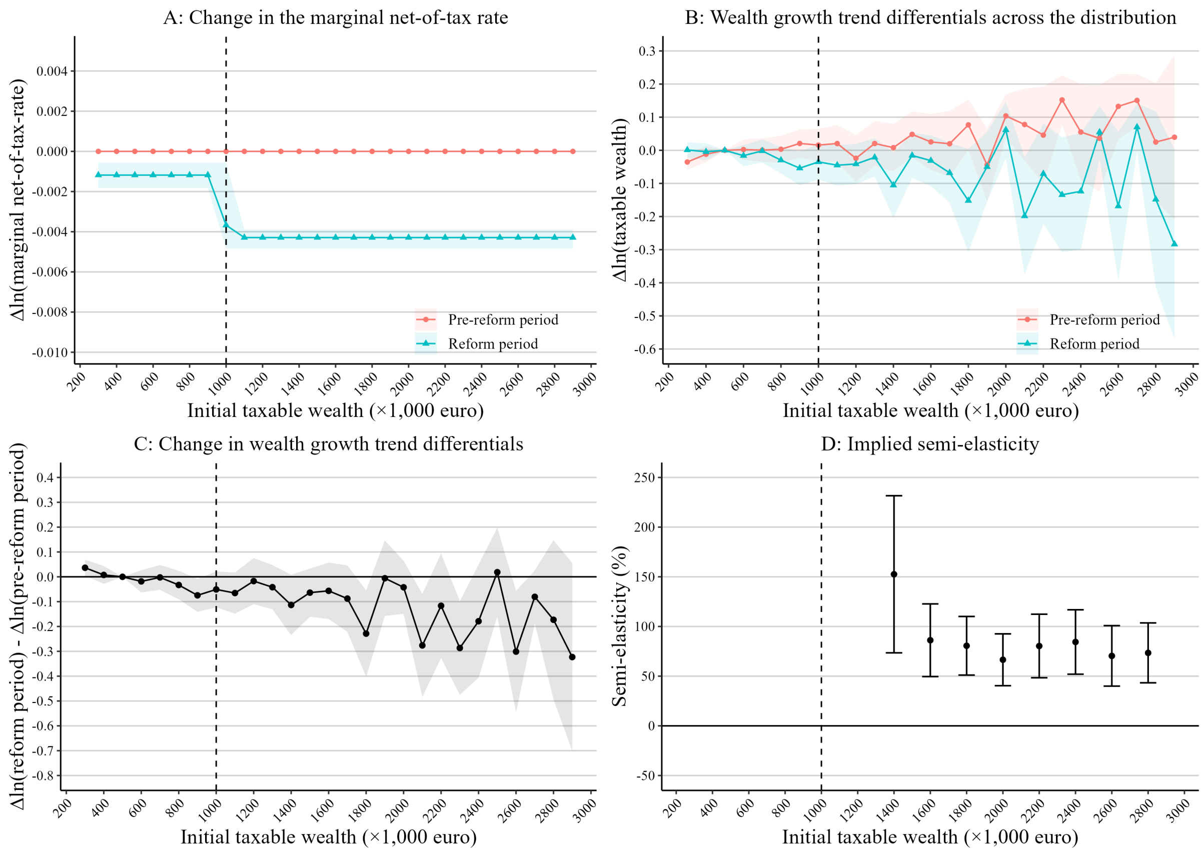

Figure 4: Response of taxable wealth to a change in the marginal tax rate in the top bracket. Notes: Panel A shows the change in the log net-of-tax rate as predicted by the instrument ($\Delta ln(1-\tau^{p}_{it-k})$) for the period 2012-2016 (pre-reform) and the period 2016-2020 (reform period). The points indicate the average change in the marginal net-of-tax rate during the period, with the shaded area indicating the fluctuations of the marginal tax rate during the period. Initial wealth is based on 2012 in the pre-reform period and 2016 in the reform period. Panel B shows the wealth growth trend differentials using Equation (11), where bin €500.000 is used as a reference. Panel C shows the change in the growth differentials, estimated following Equation (12). Panel D shows the implied semi-elasticities of taxable wealth with respect to a 1 percentage point cut in the wealth tax rate, estimated using Equation (9). The 95% confidence intervals are based on robust standard errors. For readability, the estimated elasticity for the local point around €1.200.000 is excluded (confidence interval ranges from -1002% to 52%).

The growth trend differentials in Figure 4 Panel B show noisy results, with large confidence intervals in the upper region of the taxable wealth distribution. In the validation region, the trend differentials are seemingly similar between the reform and pre-reform period, while the identification region exhibits lower taxable wealth growth during the reform period than in the pre-reform period, relative to the reference bin at €500,000. The slightly upward sloping line in the pre-reform period indicates there is some degree of differential growth trends: taxpayers with higher initial taxable wealth exhibit more subsequent wealth growth.

The ITT effect estimates in the identification region in Figure 4 Panel C are noisy but contain statistically significant decreases. These results indicate that the taxpayers in the top brackets responded to an increase in the marginal tax rate by decreasing their taxable wealth. However, a formal test of joint significance in the validation region rejects that all coefficients are insignificant (p-value = .000), indicating the identifying assumption to make this claim causal might be invalid for this region of the taxable wealth distribution.

Finally, Figure 4 Panel D shows estimated local elasticities using Equation (9) and an interval width of €200,000 for the triangular weights. I find significant semi-elasticities with very large confidence intervals in the lower region of bracket 3. From €1,600,000 upwards, the confidence intervals range from 40.5% to 122.7% and show little heterogeneity along the range of initial taxable wealth. The results imply that in response to a 1 percentage point increase in the marginal tax rate, taxpayers in the third bracket would decrease their taxable wealth with at least 40.5% over a period of 4 years. To put such large semi-elasticities into context, consider that a 1 percentage point increase in the wealth tax is a large increase. Given the tax rate of 1.2% in the pre-reform period, an increase of 1 percentage point is an 83.3% increase in the wealth tax rate. Therefore, a semi-elasticity of 40.5% implies an elasticity with respect to the marginal tax rate of (40.5/83.3 =) 0.486.

Note that Figure 4 Panel D excludes the local estimate for taxpayers around point €1.200.000 in the initial taxable wealth distribution. The confidence interval on this estimate swamps the other estimates (ranging from -1002% to 52%), making the figure less readable. A likely reason for the large confidence intervals in this region (and the point around €1,400,000), is the fact that the population of compliers is much smaller close to the cut-off, leading to less precise estimates. A relatively small decrease in taxable wealth during the reform period will cause taxpayers in this region to fall below the cut-off, losing their complier status (as discussed in Section 6.1.4 in more detail).

6.1.3 Aggregate elasticities and heterogeneity

In addition to the estimates for parts of the taxable wealth distribution, I estimate the aggregate elasticities in bracket 1 and bracket 3, using Equation (1). To explore heterogeneity, I vary the lengths of the estimation period and consider subsamples based on household characteristics. For each estimate, I report the p-value for a joint significance test on the coefficients in the validation region, where an insignificant p-value hence indicates the constant trend differential holds. Additionally, I report the F-statistic on the instrument in the first stage for each estimate. All F-statistics far exceed the rule-of-thumb threshold of 10 for a relevant first stage (Stock and Yogo, 2002), meaning a potential weak-instrument bias (Bound et al., 1995) is of no concern in the results presented below.

Using a period of four years for the pre-reform and reform period (baseline specification), I find an elasticity indistinguishable from zero in the first bracket and 43.7 in the third bracket (Table 3, Panel A). Both estimates reflect the LATE for taxpayers who remain within the same bracket during the reform period. The p-values on the joint significance test indicate the constant trend differentials assumption holds for the first bracket estimates. Conversely, the validation region for the third bracket fails the constant trend differentials assumption, indicating the elasticity of 43.7 might be biased.

Table 3: Effect of a change in the marginal tax rate on taxable wealth, varying period length

	Elasticity bracket 1			Elasticity bracket 3		
	F-stat first stage	Reduced form	2SLS	F-stat first stage	Reduced form	2SLS
A: Baseline result						
4 years (baseline)	740.8	0.7	0.8	135.9	36.0***	43.7***
		(0.7)	(0.7)		(5.2)	(6.3)
P-value assumption		.607			.000	
4 years (without IPW)	801.3	-0.1	-0.1		30.5***	36.9***
		(0.6)	(0.7)		(5.0)	(6.1)
P-value assumption		.935			.001	
4 years (without initial wealth bins)	726.1	-7.1***	-8.8***	135.9	16.7***	20.2***
		(0.5)	(0.6)		(3.7)	(4.6)
P-value assumption		N/A			N/A	
Observations		1,669,094			208,428	
B: Varying period length						
1 year	817.2	4.5***	4.7***	163.9	1.9	2.1
		(0.6)	(0.6)		(4.3)	(4.9)
P-value assumption		.225			.329	
Observations		2,089,339			273,851	
2 years	851.2	0.2	0.2	168.3	4.2	4.9
		(0.6)	(0.6)		(4.8)	(5.7)
P-value assumption		.988			.378	
Observations		1,920,577			248,035	
3 years	801.3	2.5***	2.4***	160.3	22.1***	26.4***
		(0.6)	(0.7)		(5.2)	(6.2)
P-value assumption		.654			.003	
Observations		1,794,369			227,035	
5 years	580.4	3.9***	4.7***	124.9	26.4***	31.4***
		(0.7)	(0.8)		(5.5)	(6.6)
P-value assumption		.005			.005	
Observations		1,157,408			159,132	

Notes: Regression estimates are results of variations in period length used for estimation of Equation (1). The sample used for estimations in bracket 1 contains taxpayers with initial wealth in the range €40.000-€300.000, where €100.000 is the cut-off between the identification and validation region. For the estimate using a period of 5 years, the bottom of the range for estimates in bracket 1 is raised to €50.000. The sample used for estimations in bracket 3 contains taxpayers with initial wealth in the range €300.000-€3.000.000, where €1.000.000 is the cut-off between the identification and validation region. Control variables included in the inverse probability weighting (IPW) procedure are dummies for being a single person household, a couple or other, age group, having children, a dummy if the household has substantial ownership in a firm and share of wealth held in financial investments, real estate, and other debt, as well as the share of other debt in gross wealth. The p-value for a joint significance test of the coefficients in the validation area is shown, where a statistically significant p-value indicates that the identifying assumption of constant trend differentials does not hold. Robust standard errors in parentheses.

To inspect sensitivity of the results to the inclusion of covariates, Table 3 Panel A presents the estimates without the IPW procedure. Excluding controls decreases the estimated elasticity in the first bracket to -0.1% and remains statistically insignificant, indicating the results in the first bracket are not sensitive to the control variables. In the third bracket, excluding the IPW procedure causes the elasticity to fall to 36.9%. This decrease underlines the value of the IPW procedure: without this correction some composition effects might be erroneously attributed to the effect of a change in the marginal tax rate.

Next, Panel A reports the estimate without controlling for initial wealth bins. This specification is not robust to mean reversion and differential growth trends. Recall that the upward slopes of the lines in Panel B of Figures 3 and 4 indicate taxpayers with higher initial wealth show higher subsequent wealth growth. Failing to control for initial wealth would bias the elasticity estimates downward: a decrease in the marginal tax rate appears to be coupled with higher taxable wealth growth, while the higher wealth growth is caused by higher initial wealth. Indeed, an estimate without initial wealth dummies gives a significantly lower estimates in both brackets. These findings highlight the importance of controlling for mean reversion and differential growth trends.

Table 3 Panel B shows the results are sensitive to the length of the period used for estimation. Using only one year before and after the tax reform, I estimate a significant semi-elasticity of 4.7% in the first bracket. Using two years, the effect becomes statistically indistinguishable from zero again. The 3-year estimate gives an elasticity of 2.4%. In each of these estimates, the constant trend differentials assumption holds in the validation region. For the 5-year estimate, I narrow the estimation range to €50.000 to €100.000, so the increase of the tax-exempt threshold to €50.000 in 2021 does not interfere with the results. The estimate gives a semi-elasticity of 3.7%, yet the assumption of constant trend differentials fails in the validation region.

The estimates with varying period lengths for bracket 3 seem to suggest that the statistically significant baseline result is the product of a bias. When using a period length of 1 or 2 years, the estimates are not statistically distinguishable from zero. The p-values on the joint significance tests provide evidence that the identifying assumption is plausible for these estimates. In the estimates using 3 to 5 years, the semi-elasticity becomes statistically significant, but test of the constant trend differentials assumption indicates it does not hold. These findings highlight the value of the procedure to visually inspect and formally test the underlying identifying assumption as proposed by Jakobsen and Søgaard (2022). Without checking and finding the constant trend differentials assumption fails, we could have erroneously interpreted the statistically significant elasticity estimates for the third bracket as a sound quantification of a behavioural response.

Lastly, I examine heterogeneity by applying the baseline specification to subsamples of taxpayers, varying household composition, age, homeownership, and closely held firm ownership (Table 4). The results show that taxpayers with a closely held have a stronger response to a change in the marginal tax rate. The semi-elasticity of taxable wealth is 20.6% for taxpayers with a closely held firm in the first bracket and 51.1% in the third bracket. Both subgroups exhibit constant trend differentials in the validation region, if one is willing to accept a 10% confidence level for the third-bracket estimate. A relatively high responsiveness from closely held firm owners is in line with the notion that a closely held firm can function as a tax shelter (Alvaredo and Saez, 2009; Alstadsæter and Kopczuk, 2014; Lejour et al., 2022). With regards to age, taxpayers in the pension age appear to be less responsive to a change in the marginal tax rate compared to taxpayers below the pension age. Possibly, shifting wealth is more costly or less convenient when one is more dependent on wealth for consumption (as is expected in old age). However, the higher point estimate for pensioners than non-pensioners found in the third bracket does not corroborate this theory. Homeowners appear to show a lower response to a change in the marginal tax rate than renters in the first bracket, but the estimates are within each other's confidence bounds. The same comparison cannot be made in the third bracket, as there are too few renters to estimate their elasticity. The differences between singles and couples is not statistically significant in either bracket.

Table 4: Effect of a change in the marginal tax rate on taxable wealth, varying household type

Period length	Elasticity bracket 1			Elasticity bracket 3		
	F-stat first stage	Reduced form	2SLS	F-stat first stage	Reduced form	2SLS
Total (baseline)	740.8	0.7	0.8	135.9	36.0***	43.7***
		(0.7)	(0.7)		(5.2)	(6.3)
P-value assumption		.607			.000	
Observations		1,669,094			208,428	
Singles	447.0	1.4	1.4	121.5	28.7***	33.2***
		(1.1)	(1.2)		(6.4)	(7.5)
P-value assumption		.018			.001	
Observations		562,230			88,571	
Couples	607.5	0.9	1.0	109.0	33.8***	41.7***
		(0.8)	(.0)		(7.5)	(9.3)
P-value assumption		.281			.001	
Observations		1,012,166			108,666	
Pension age	530.9	-0.6	-0.6	146.0	40.1***	47.2***
		(1.0)	(1.0)		(5.2)	(6.2)
P-value assumption		.987			.000	
Observations		685,724			115,896	
Non-pension age	459.5	3.4***	4.0***	64.3	28.6***	37.2***
		(1.1)	(1.3)		(12.9)	(16.7)
P-value assumption		.234			.171	
Observations		817,659			70,057	
Homeowners	667.4	0.5	0.6	151.4	35.1***	42.7***
		(0.7)	(0.8)		(5.2)	(6.4)
P-value assumption		.633			.000	
Observations		1,364,184			186,473	
Renters	320.6	2.4*	2.4*	-	-	-
		(1.4)	(1.4)		-	-
P-value assumption		.725			-	
Observations		288,599			19,855	
Firm owners	151.5	20.6***	25.5***	68.0	42.6***	52.1***
		(3.9)	(4.8)		(12.0)	(15.0)
P-value assumption		.231			.036	
Observations		129,359			48,243	
Non-firm owners	735.7	-1.1*	-1.2*	174.1	19.0***	21.6***
		(0.6)	(0.7)		(4.3)	(4.9)
P-value assumption		.724			.001	
Observations		1,494,034			147,995	

Notes: Regression estimates are results of variations in the specification of Equation (1). The sample used for estimations in bracket 1 contains taxpayers with initial wealth in the range €40.000-€300.000, where €100.000 is the cut-off between the identification and validation region. For the estimate using a period of 5 years, the bottom of the range for estimates in bracket 1 is raised to €50.000. The sample used for estimations in bracket 3 contains taxpayers with initial wealth in the range €300.000-€3.000.000, where €1.000.000 is the cut-off between the identification and validation region. The control variables included in the inverse probability weighting (IPW) procedure are dummies for being a single person household, a couple or other, age group, having children, a dummy if the household has substantial ownership in a firm and share of wealth held in financial investments, real estate, and other debt, as well as the share of other debt in gross wealth. The p-value for a joint significance test of the coefficients in the validation area is shown, where a statistically significant p-value indicates that the identifying assumption of constant trend differentials does not hold. Robust standard errors in parentheses.

6.1.4 Non-compliance

In the results discussed above, I have focussed on the LATE on the compliers: taxpayers who remain within the same bracket as predicted by their wealth in the base year, the instrument. As discussed above, the LATE hence excludes the response from taxpayers who react to such a degree that they end up in another bracket. To check if this occurs, I inspect the amount of people that switch between brackets per initial wealth bin and per period.

Figure 5 shows that mobility between brackets is substantial close to the bracket cut-offs. For taxpayers just above and below the cut-off between the first or third bracket, about 30% switches brackets during a four-year period. Further away from the cut-off, the share of taxpayers that switches brackets falls quickly. In the pre-reform period, the marginal tax rate did not change above or below the cut-off, so taxpayers had no incentive to exert effort to fall below the cut-off. Indeed, the distribution of the share of switchers around the cut-off appears to follow a smooth distribution around the cut-offs.

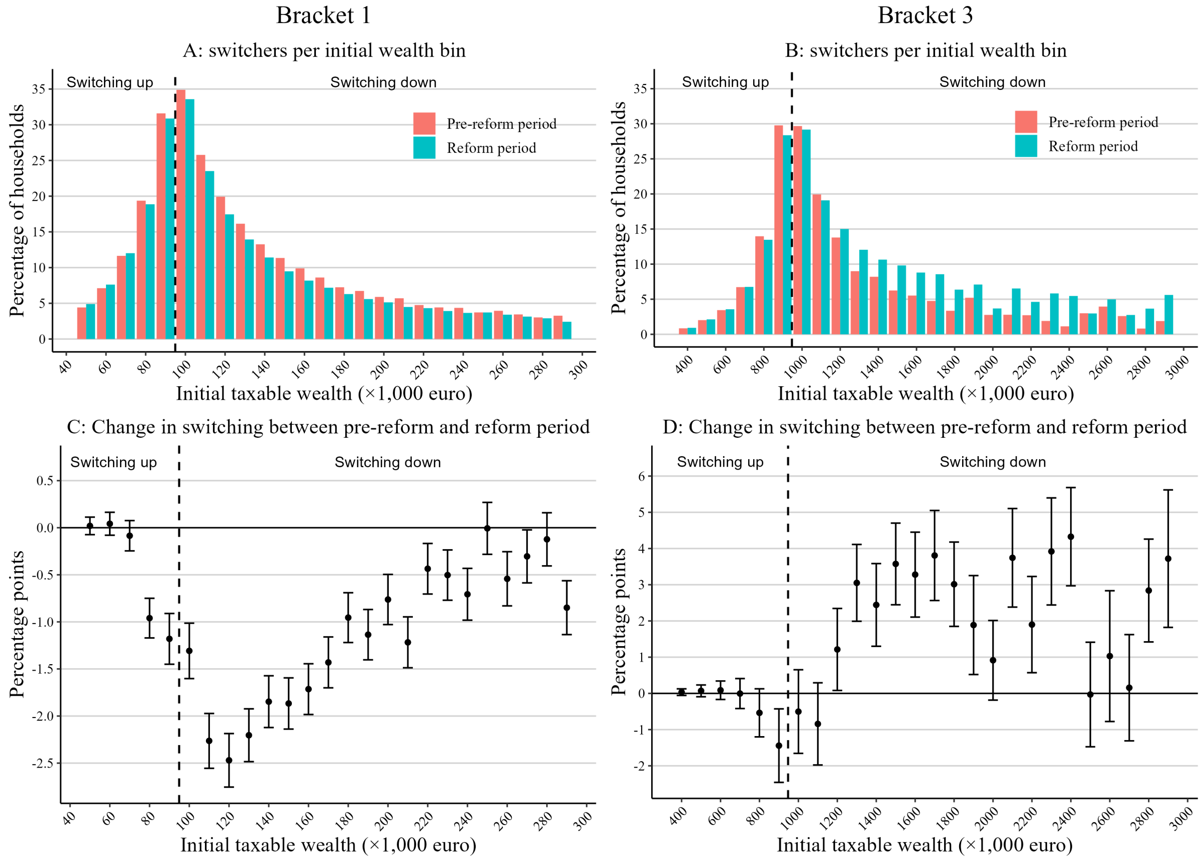

Figure 5: Share of taxpayers switching between brackets. Notes: this figure inspects the share of taxpayers that switch bracket during the pre-reform (2012-2016) and reform period (2016-2020) and hence do not comply with the instrument. In Panel A and B, bars to the left of the bracket cut-off (dotted line) show the share of taxpayers that switches up to a higher bracket, per initial wealth bin. The bars to the left of the cut-off shows the share of taxpayers that switch downward. The change in propensity to switch brackets is estimated using a simple regression of the form $Switch_{i,t} = \beta_0 + \boldsymbol{\beta_1 W_{it-1}} + \beta_2 D_{it}^{reform} + \boldsymbol{\beta_3 W_{it-1}} * D_{it}^{reform}$, where $Switch_{i,t}$ is dummy that takes the value of 1 for taxpayers with initial wealth below (above) the cut-off, but wealth at time *t* above (below) the cut-off. The coefficients $\boldsymbol{\beta_3}$ are plotted in Panel C and D. The 95% confidence intervals in panel C and D use robust standard errors.

In the reform period, the effective marginal rate increases at the cut-offs. As a result, taxpayers might want to switch to a lower bracket to face a lower marginal tax rate or prevent switching to a higher bracket. Indeed, taxpayers just below the cut-offs to bracket 2 show a statistically significant decrease in the propensity to switch to a higher bracket of about 1 percentage point (Figure 5 Panel C). Taxpayers with lower initial wealth do not show a significant change in the propensity to switch. On the right side of the €100,000 cut-off, we would expect an increase in the propensity to switch to a lower bracket, towards the lower marginal tax rate. However, the propensity to switch down from the second bracket falls during the reform period relative to the pre-reform period. In the third bracket, the change in shifting does follow the expected pattern (Figure 5 Panel D). Taxpayers just below the cut-off to the third bracket show a decrease in the propensity to switch upward during the reform period, possibly due to repelling effect of the higher marginal tax rate in the third bracket. Conversely, taxpayers with initial wealth in the third bracket switch downward more often during the reform period. Here, the effect is sizeable, with the propensity to switch downward increasing by about 4 percentage point. Even in high initial wealth bins, where the decrease in taxable wealth has to be substantial to cross the cut-off, there is a strong and significant increase in the share of taxpayers that switches to a lower bracket.

The findings from Figure 5 have several implications for my results. First, the relatively large amount of switching of taxpayers close to the cut-off implies the share of compliers is relatively low in these regions of the taxable wealth distribution. A lower share of compliers can explain why I find (much) larger confidence intervals for local estimates close to the cut-off in Figures 3 and 4. Second, the finding that taxpayers in the third bracket switch down to a lower bracket significantly more often following the reform, suggests that part of the response to a change in the marginal tax rate is indeed not captured by my LATE, implying the LATE will be smaller than the average treatment effect on the treated.

6.2 Robustness checks

To inspect the robustness of the estimates and assumptions, I perform several supplementary analyses.

6.2.1 Anticipation effects

If taxpayers anticipated the wealth tax reform, they could already adapt their taxable wealth in a pre-reform year. In that case, a part of the response to a change in the marginal tax rate will already be realised in the pre-reform period, causing a downward bias in the magnitude of my estimates. In my setting, a small window for such anticipation effects exists. Recall that the last recording of taxable wealth prior to the tax reform was on January 1st of 2016, while the reform was announced in September of 2015. During this four-month period, it is possible for taxpayers to learn about the reform and already alter their wealth holding per January 1st of 2016 in anticipation of the new tax scheme per January 1st of 2017. To check for such anticipation effects, I estimate the difference in taxable wealth growth between the periods 2014-2015 and 2015-2016, using the same method as applied above.

The results shown in Figure 6 indicate no strong sign for anticipation effects in either the first or third bracket. In the first bracket, the difference in taxable wealth growth between the period 2014-2015 and 2015-2016 is statistically indistinguishable from zero in all but the lowest initial wealth bin. The decrease in wealth growth at bin €40,000 falls just within the range of statistical significance, rendering a joint significance test for all coefficients in the first bracket significant at .003. In the third bracket, taxable wealth growth differentials also remains constant between the two period, except for a single initial wealth bin at €2,800,000. These results are consistent with taxpayers either not becoming aware of the tax reform, not being able to adjust their taxable wealth within such short notice, or not willing to adjust their taxable wealth in a tax year where the new tax schedule does not apply yet. Indeed, in a scenario where adjusting taxable wealth is costly, but flexible, taxpayers may prefer to delay adjusting their taxable wealth.

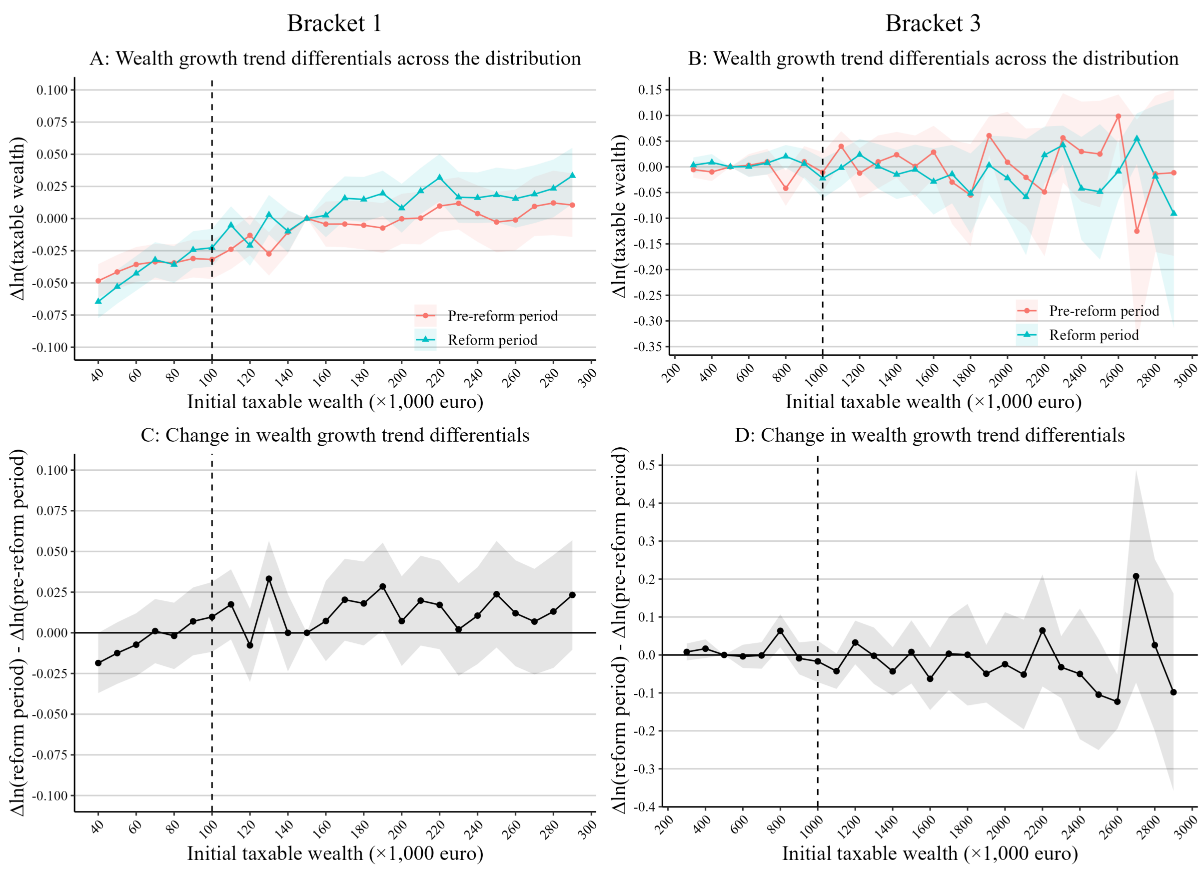

Figure 6: Inspection of anticipation effects. Notes: By comparing the period 2014-2015 and 2015-2016, this figure shows if taxpayers exhibit any response in anticipation of the tax reform. Panel A and B show the wealth growth trend differentials using Equation (11), where the bins €150.000 and €500.000 are used as reference for the estimates for bracket 1 and 3 respectively. The Panels C and D show the change in the growth differentials, estimated following Equation (12). The 95% confidence intervals are based on robust standard errors.

6.2.2 Placebo reforms

The same approach used to test for anticipation effects can be used to gain further insight into the constant trend differentials assumption. The p-values on tests of the identifying assumption which are presented alongside the elasticity estimates above, test for constant trend differentials in the same period as used in the estimates and are hence most relevant. However, showing that the constant trend differentials assumption hold more structurally would increase confidence in the identifying assumption. In all pre-periods pre-reform periods, both the validation and identification region did not experience a change in the marginal tax rate. Therefore, ideally, both the validation region and identification region should exhibit constant trend differentials when evaluating pre-reform years only. Appendix 10.3 reports these placebo reform estimates.

The identification region in bracket 3 performs well. In all the placebo runs, the trend differentials in the identification region of bracket 3 have remained constant. However, the constant trend differentials assumption fails in the validation region for most specifications. The identification region used for bracket 1 does not perform well. In most years, the trend differentials have not remained constant within the identification region. The validation region of bracket 1 performs somewhat better, but also rejects constant trend differentials in three of the six specifications.

A disadvantage of the placebo tests is the fact that I have to use shorter periods than used in the main estimates. As the number of pre-reform years in my sample is limited, I cannot perform placebo checks using a period length of 4 years before and after the reform. The fact that the placebo tests are estimated on shorter periods makes them more sensitive to noise and outliers, increasing the chance of joint significance test rejecting the null-hypothesis due to noise.

6.2.3 Varying discretionary choices

My main estimation approach requires certain discretionary choices. To alleviate the concern that my results are sensitive to such choices, I re-estimate my results using narrower bin widths, varying the upper cut-off in the range considered for the third bracket and without applying the filters used to clean the data. The results are presented in Appendix 10.4 and discussed verbally here.

Decreasing the size of the initial wealth bins does not affect the results. Using bins with a width of €5,000 instead of €10,000 in the estimates for the first bracket results in a semi-elasticity of 0.77% instead of 0.82%. Given the standard errors of about 0.7 percentage points, these estimates are neither economically nor statistically different. Similarly, using €50,000 bins instead of €100,000 bins for the estimates in the third bracket results in a semi-elasticity of 39.7% instead of 43.7%. While the magnitude of this difference can be economically relevant, standard errors of 6-7 percentage points imply the two estimates are not statistically different from each other. Smaller wealth bins also have no effect on the conclusions from the test for constant trend differentials in the validation region.

For the estimates in the top-bracket, I do not use all observations in my baseline estimates but drop taxpayers with initial wealth above €3,000,000. The fact that the estimated LATE does not include the response by taxpayers who switch brackets, implies such a sample restriction can affect the results. For example, when only considering taxpayers close to the cut-off between two brackets, solely taxpayers with little or no response will retain the complier status, as taxpayers with a larger response will switch brackets. If the range above (or below) a cut-off becomes larger, there is more room for taxpayers to exhibit a larger response and remain within the same bracket. To inspect the effect of this sample restriction on my estimates, I estimate the aggregate elasticities again using different caps on the range of initial taxable wealth. The results show higher point estimates as I increase the cap on initial taxable wealth, in line with the described mechanism, but these estimates are not statistically different from each other.

Lastly, I inspect the sensitivity of my results to the filters I use to clean the data.[40] When running the baseline specification regressions on my sample without filters, the estimates give a semi-elasticity of 1.2% in the first bracket and 42.8% in the third bracket. The differences from baseline estimates with filters is economically small and statistically insignificant. These results show my estimates are not sensitive to my filters.

6.2.4 Synthetic-difference-in-differences

As a complementary analysis and robustness check, I next estimate the elasticity of taxable wealth using the synthetic-difference-in-differences (SDiD) estimator as developed by Arkhangelsky et al. (2021).[41] The SDiD estimator is a useful complement to the DiDiD estimator as it offers a different approach to tackle the issues of mean reversion and differential growth trends. The DiDiD method controls for initial wealth and assumes the trend differentials between different levels of initial wealth remain constant over time in absence of treatment. The SDiD instead reweights units in the control group to match the trend in the treatment group by a fixed difference. Hence, an advantage of the SDiD is that it relies less strongly on a common trend to be present in the data, but instead uses reweighting to create it (Ferman, 2021; Arkhangelsky et al., 2021; Arkhangelsky and Imbens, 2022).

In Appendix 10.5 I further discuss the characteristics of the SDiD-estimator and present the results.[42] The SDiD estimates for the third bracket show a statistically insignificant result. The pre-trend of the synthetic control unit and the average in the third bracket line up well. The SDiD-estimate hence contradicts the finding of a large decrease when using the DiDiD method. It could be that the SDiD-estimator fails to adequately control for the differential growth trends and mean reversion, as it only

[40] Recall that the used filters drop taxpayers whose observations are imputed based on past tax filings, have received a tax-exempt gift for expenses on owner-occupied housing or a bequest or whose status changed from renter to homeowner or reversed during the studies period (as a proxy for purchase or sale of owner-occupied housing).

[41] For implementation of the method, I use the R package "synthdid" developed by Arkhangelsky et al. (2021), available at https://github.com/synth-inference/synthdid.

[42] I delegate most of the SDiD section to the Appendix, as practical limitations have hindered an analysis complete enough to warrant a presentation in the main text. Specifically, running the SDiD algorithm on a sample as large as the one used for the estimates in the first brackets takes a very long time (up to days). As discussed in Appendix 10.5, I have developed a method to reduce the time required for the algorithm but required changes in the estimation relatively late in the research project have made it infeasible to also present estimates for the first bracket using the year 2016. Despite these drawbacks, I still present the results that are available, to learn from their insights.

matches based on the period leading up to the base year, meaning it only matches on the increasing part of the mean reversion effect. The match might be improved by defined the instrument on an earlier base year, so the training data also contains the downward section of mean reversion. Using 2015 to define initial wealth, the pre-trend is again matched well and the outcome changes slightly. The effect now shows a small decrease in taxable wealth, which is in line with the sign of the effect found in the DiDiD analysis but remains statistically insignificant. Estimating the response in the first bracket using 2015 as a base year also shows a statistically insignificant result, with a positive point estimate of 0.005.

6.2.5 Non-pure control group

Lastly, I consider the threat of a bias stemming from the fact taxpayers in the second bracket experience a smaller change in the marginal tax rate, instead of no change at all (no pure control group). The DiDiD estimator will fail to difference out the effect from the (smaller) change in the marginal tax rate in the control group when the potential response of the treatment and control group to a change in the marginal tax rate are not equal (Saez, 2004; Saez et al., 2012; Kumar and Liang, 2020). Therefore, a bias may arise that increases in size when the underlying true elasticity in the treatment and control group lay further apart.

The true difference in the underlying elasticity in the treatment and control groups cannot be observed. The fact that I find no heterogeneity in the elasticity within the treatment groups (Figures 3 and 4), suggests a large difference is not likely between the underlying elasticity in the control regions in bracket 2 used for estimates on bracket 1 and 3 respectively. However, the empirical literature does find higher elasticities for taxpayers with more taxable wealth (Zoutman, 2018; Jakobsen et al., 2020). In Appendix 10.8, I formally derive the direction and size of the bias that would occur if the underlying elasticities do differ.

The main conclusion from the calculations in Appendix 10.8, is that the elasticities in brackets 1 and 3 are likely to be slightly overestimated. For illustration, assuming a true elasticity in the control group of half the true elasticity in the treatment group (large discrepancy) implies a bias by a factor of about 1.19 and hence a true elasticity of 36 underlying the estimate of 43.4 in the top bracket and 21 when using the estimate of 25. Similarly, assuming the elasticity in the control group for the bottom bracket is twice as large as the true elasticity in the bottom bracket gives a bias factor of 1.17, implying the estimate of 4.7% corresponds to a true elasticity of 4%. Under any smaller deviation between the true elasticity in the treatment and control group, the bias will be smaller.

6.3 Response margins and mechanisms

The total response captured by the aggregate elasticity does not provide information on the mechanisms through which taxpayers reduce their taxable wealth. Understanding these mechanisms is required to assess to what the degree the elasticity is driven by real or by reporting responses (Slemrod, 1996). In this section, I investigate the response per wealth category separately, estimate the increase in the value of substantial ownership (as indicator for an increase in the use of closely held firms as tax shelter), changes in residence value and mortgage debt, and the change in giving and receiving gifts.

6.2.1 Response per component of taxable wealth in box 3

Table 5 Panel A reports the semi-elasticities per component of box 3. First looking at bracket 1, the results show that looking at each component of box 3 separately, the elasticities point towards negative semi-elasticities. This would imply, that taxpayers in bracket 1 responded to a decrease in the marginal tax rate by decreasing their taxable wealth holdings. Do note, however, that the wealth of taxpayers in bracket 1 mainly consists of savings, with other wealth components only playing a role for a limit number of taxpayers (Table 1). The results may hence be driven by the response of a small group. Still, also the savings category shows a statistically significant negative semi-elasticity and for each wealth component the constant trend differentials assumption holds in the validation region.

Table 5: Effect of a change in the marginal tax rate on separate wealth components

	Elasticity bracket 1			**Elasticity bracket 3**		
Dependent variable	P-value assumption	Reduced form	2SLS	P-value assumption	Reduced form	2SLS
Total box 3 (baseline)	.607	0.7 (0.7)	0.8 (0.7)	.000	36.0*** (5.2)	43.7*** (6.3)
A: Box 3 components						
Savings	.750	-0.8* (0.5)	-0.9* (0.5)	.327	20.5*** (5.7)	24.9*** (6.9)
Investments	.532	-7.5*** (1.7)	-8.3*** (1.9)	.000	134.6*** (13.8)	163.4*** (16.8)
Real-estate	.152	-1.1 (1.7)	-1.3 (1.9)	.592	29.5 (19.2)	35.8 (23.5)
Other taxable wealth	.010	-8.8*** (1.8)	-9.8*** (2.0)	.569	-30.9 (23.8)	-37.6 (28.9)
Deductible debt	.617	9.3*** (1.9)	10.4*** (2.1)	.537	-107.2*** (23.3)	-130.1*** (28.4)
B: Other wealth components						
Substantial ownership	.403	-7.0*** (0.9)	-7.7*** (1.0)	.000	-101.9*** (16.2)	-123.6*** (19.8)
Mortgage debt	.304	0.22 (1.7)	0.25 (1.8)	.079	40.6*** (17.7)	49.2** (21.5)
Residence value	.026	-0.8*** (0.1)	-0.8*** (0.1)	.018	-5.2*** (1.4)	-6.4*** (1.7)
F-stat first stage		740.8			155.8	
Observations		1.669,069			208,401	

Notes: Regression estimates are results of variations in the outcome variable used in specification of Equation (1), with 2012-2016 being the pre-reform period and 2016-2020 the reform period. The sample used for estimations in bracket 1 contains taxpayers with initial wealth in the range €40.000-€300.000, where €100.000 is the cut-off between the identification and validation region. The sample used for estimations in bracket 3 contains taxpayers with initial wealth in the range €300.000-€3.000.000, where €1.000.000 is the cut-off between the identification and validation region. The control variables included in the inverse probability weighting (IPW) procedure are dummies for being a single person household, a couple or other, age group, having children, a dummy if the household has substantial ownership in a firm and share of wealth held in financial investments, real estate, and other debt, as well as the share of other debt in gross wealth. Below each estimate, the p-value for a joint significance test of the coefficients in the validation area is shown. A statistically significant p-value indicates that the identifying assumption of constant trend differentials does not hold in the validation region. Robust standard errors in parentheses.

The estimates on box 3 components in bracket 3, reveal that the issue of non-constant trend differentials is likely to be caused by the investment's component. The constant trend differentials assumption in the validation region is strongly rejected for the investments component but cannot be rejected when considering the other wealth components.[43] Moreover, the (likely biased) semi-elasticity estimates on investments is very large, at 163.4%, which may inflate the elasticity of total box 3 wealth.[44] Looking only at the savings component in box 3, I find an elasticity in the third bracket of 24.9%. This estimate is considerably lower than estimate for total box 3 but remains economically significant. Real estate does not respond to a change in the marginal tax rate, which is in line with the illiquid nature of real estate. The elasticity on debt deductible in box 3 is -130.1% and the constant trend differentials assumption holds in the validation region for this box 3 component. This result implies that taxpayers may attempt to reduce their net taxable wealth by increasing their deductible debt.

6.2.2 Response of other wealth components

Besides wealth that is taxable in box 3, a change in the marginal tax rate may influence wealth that is (un)taxed in other tax regimes. Specifically, taxpayers may shift wealth from box 3 to substitutes outside box 3 in response to a change in the wealth tax rate. Table 5 Panel B reports estimates of elasticities on such wealth components outside box 3.

Firstly, taxpayer may shelter wealth in closely held firms when the marginal tax rate increase or retrieve wealth from closely held firms when the marginal tax rate falls. Indeed, the results show an increase in the tax rate on wealth in box 3 seems to be associated with an increase in the value of shares held in closely held firms (substantial ownership). For taxpayers in the bottom bracket, the estimate implies that a 1 percentage point decrease in the marginal tax rate in box 3 leads to an increase in substantial ownership wealth of 7.7% over a period of 4 years. A joint significance test cannot reject the hypothesis that the constant trend differentials assumption holds in the validation region. In the top bracket, the estimate indicates a semi-elasticity of -123.6%. However, the constant trend differentials assumption is rejected.

Second, taxpayers may choose to pay off extra mortgage debt. As such, they will reduce their wealth in box 3 and increase the net-value of their owner-occupied housing, which is untaxed wealth. A downside of this approach is that paying off more debt will reduce the benefit of subtracting mortgage interest payments from box 1 income. The results show that the mortgage debt of taxpayers in bracket 1 is not responsive to the decrease in the marginal tax rate. Both in the identification and validation region, the growth trends of mortgage debt remain constant. In the top bracket, mortgage debt does

---

[43] This finding implies it would be fruitful to estimate the elasticity for taxpayers in bracket 3 who do not have investments. However, almost all taxpayers with at least €1,000,000 own some investments, and filtering them out shrinks the sample by too much to run the analysis. Moreover, the generalisability of an estimate based on such a select subgroup would be limited, making it of little relevance for policy.

[44] A semi-elasticity of 163.4% is very large, but still within the range of Brülhart et al. (2022), who find estimates up to 187% when focussing on reforms causing the largest change in the marginal tax rate.

respond. An increase in the marginal tax rate with 1 percentage point leads to a decrease in mortgage debt of 49.2% over a period of 4 years. The assumption of constant trend differentials cannot be rejected at the 5% confidence level, but does het rejected at the 10% confidence level. In both brackets, the point estimates indicate a small increase in the residence value (owner-occupied housing) in response to an increase in the marginal tax rate. Such an effect can be rationalised by owners investing in their homes to spend excess wealth. For the estimate on residence value too, however, also the validation region shows a significant change, meaning the constant trends differential does not hold.

6.2.3 Gifts

Taxpayers may alter their amount of taxable wealth through receiving and giving financial gifts. To inspect this mechanism, I use the same methodology as above, but with receiving and giving financial gifts as outcome variables.[45]

Table 6: Effect of a change in the marginal net-of-tax rate on financial gifts

	Semi-elasticity bracket 1			**Semi-elasticity bracket 3**		
	P-value assumption	Reduced form	2SLS	P-value assumption	Reduced form	2SLS
Giving	.632	-1.0*** (0.4)	-1.2*** (0.4)	.812	-23.1* (12.7)	-28.1* (15.5)
Receiving	.386	-0.3 (0.2)	-0.3 (0.2)	.486	2.8 (3.5)	3.3 (4.2)
F-stat first stage		740.8			155.8	
Observations		1.669,069			208,401	

Notes: Regression estimates are results of using the amount of gifts given and received as outcome variables in specification of Equation (1), with 2012-2016 being the pre-reform period and 2016-2020 the reform period. The sample used for estimations in bracket 1 contains taxpayers with initial wealth in the range €40.000-€300.000, where €100.000 is the cut-off between the identification and validation region. The sample used for estimations in bracket 3 contains taxpayers with initial wealth in the range €300.000-€3.000.000, where €1.000.000 is the cut-off between the identification and validation region. The control variables included in the inverse probability weighting (IPW) procedure are dummies for being a single person household, a couple or other, age group, having children, a dummy if the household has substantial ownership in a firm and share of wealth held in financial investments, real estate, and other debt, as well as the share of other debt in gross wealth. Below each estimate, the p-value for a joint significance test of the coefficients in the validation area is shown. A statistically significant p-value indicates that the identifying assumption of constant trend differentials does not hold in the validation region. Robust standard errors in parentheses.

The results reported in Table 6 show that taxpayers increase the amount of wealth they give away in response to an increase of the marginal wealth tax rate. In the first bracket, an increase in the marginal tax rate of 1 percentage point increases the amount of giving by taxpayers in the first bracket by 1.2%. The effect is significantly larger in the third bracket, where giving increases by 23.1% given the same change in the marginal tax rate. However, the effect in the third bracket is only statistically significant at the 10% confidence level. With regards to receiving gifts, no significant effect is found in either bracket. Speculatively, the sign on the estimates do suggest that receiving increases in the first

[45] Recall that these estimates exclude financial gifts used for expenses on owner-occupied housing to prevent confoundment with a change in the exemption for these gifts.

bracket and decreases in the third bracket, which is in line with wealth being shifted away from taxpayers with a high marginal tax rate towards taxpayers with a lower marginal tax rate. Finally, in all regressions, a joint significance test cannot reject the null-hypothesis that the constant trend differentials assumption fails in the validation region.

7. Discussion and conclusion

In this article I study the behavioural response of taxpayers to a reform in the Dutch wealth tax in 2017. Following the theoretical and empirical literature, I focus on the elasticity of taxable wealth with respect to the net-of-tax rate as summary statistic capturing all response margins to a change in the marginal tax rate (Saez et al., 2012). To estimate the elasticity, I use the difference-in-difference-in-differences approach standard in the literature (Auten and Carroll, 1999; Gruber and Saez, 2002; Saez et al., 2012; Kleven and Schultz; 2014). Following Jakobsen and Søgaard (2022), I am able to provide graphical and formal evidence on the plausibility of constant trend differentials assumption required for unbiased estimates. Such proof for the constant trend differentials assumption has not been presented in the elasticity of taxable wealth literature before.

It is inappropriate to select a single point-estimate for the elasticity of taxable wealth based on the evidence in this study. The estimated elasticities are shown to be sensitive to the time-horizon considered and vary significantly when examining subcomponents of taxable wealth. The evidence does provide a useful indication for the plausible range of the elasticity of taxable wealth in the Netherlands.

For taxpayers in the first bracket (up to around €100,000), I find estimates ranging from a zero elasticity of taxable wealth with respect to the net-of-tax rate up to 4.7 (both short and long term). The latter elasticity implies that taxpayers with taxable wealth below €100,000 respond to an increase in the marginal tax rate with 1 percentage point by decreasing their taxable wealth with 4.7%.[46] Taxpayers in the control group do not show a change in the growth of their taxable wealth, indicating that the identifying assumption of constant trend differentials holds.

For taxpayers facing an increase in the marginal tax rate in the top bracket (above about €1,000,000), I find an elasticity of 43.4 with respect to the marginal net-of-tax rate over a 4-year period. However, a placebo test indicates the assumption of constant trend differentials does not hold, indicating the estimate may contain a bias. An analysis of the components of net taxable wealth indicates it is the investments asset class that causes the constant trend differentials assumption to fail. All other wealth components show constant trend differentials in the validation region when considered in isolation. Focussing only on savings, taxpayers in the top bracket show an elasticity of 24.9.

Compared to the existing literature, my estimates for the elasticity in the top bracket are similar to the results in Duran-Cabré et al. (2019) and Brülhart et al., (2022). The estimates the bottom bracket

[46] Recall that the elasticity of taxable wealth with respect to the net-of-tax rate and the semi-elasticity with respect to a 1 percentage point change in in the wealth tax rate are equivalent by approximation as $\ln(1-\tau) \approx \ln(-\tau)$ when τ is small.

are lower and most similar to estimates in Norway (Ring, 2021; Alstadsæter et al., 2022). However, the estimates in the bottom bracket are more difficult to compare to existing literature, as there is little evidence on the response to wealth taxes by taxpayers with relatively little wealth. The sample in Zoutman (2018) also contains taxpayers with wealth below €100,000 but he does not present a heterogeneity analysis where the effect is estimated separately for taxpayers in the bottom region of the wealth distribution.

The heterogeneity analysis shows that taxpayers who own shares in a closely held firm respond far stronger to a change in the wealth tax rate compared to taxpayers without a closely held firm. In the first bracket, the elasticity increase to 25.1 when limiting the sample to business owners. In the third bracket, the elasticity increases to 52.1, but the constant trend differentials assumption fails in the validation region. Additionally, I find an increase in the wealth tax rate increases the value of substantial ownership (proxy for wealth held in closely held firms). These findings are in line with the notion that a closely held firm can function as a tax shelter (Alvaredo and Saez, 2009; Alstadsæter and Kopczuk, 2014; Lejour et al., 2022) and corroborate anecdotal evidence from the Dutch Ministry of Finance concluding that evading wealth tax in box 3 by transferring wealth to a closely held firm in common practice (Ministerie van Financiën, 2023).

Two other mechanisms which stand out are an increase in financial (inter vivos) gifts and debt deductible in box 3. My estimates show that taxpayers in the top bracket increase their giving by 28.1% in response to a 1 percentage point increase in the marginal tax rate. In the first bracket, the effect is smaller, at 1.2%. I do not find a change in the amount in gifts taxpayers receive. Deductible debt rises sharply in the top bracket, indicating a 1 percentage point increase in the marginal tax rate increases deductible debt by 130%. This result highlights that also debt deductibility is a relevant policy design aspect in a net wealth tax.

To illustrate the policy implications of these elasticities, I compute expected revenue effects from a change in the marginal tax rate in the top wealth bracket (Appendix 10.7). Starting from a marginal tax rate in the top bracket of 1.6%, an elasticity of 25 implies that 75.2% of the mechanical increase in tax revenue from a further increase in the top wealth tax rate is lost through erosion of the tax base.[47] Using the elasticity of 43, the behavioural effect rises to 129%, implying the behavioural effect dominates the mechanical effect. Under the latter elasticity, the current tax rate is above the Laffer curve (Laffer, 2004), meaning more tax revenue can be raised with a lower marginal tax rate. However, recall that this high estimate likely contains a bias. Using the elasticity of 25, there is room to increase tax revenue with a higher marginal tax rate. Specifically, tax revenue would be maximised at a top wealth tax rate of 2.1%. Using the translation to capital income taxes (Appendix 10.6), the elasticity of taxable wealth with respect to the net-of-tax rate of return becomes 1.12, using a rate of return of 5%, implying the revenue-maximising marginal tax rate on capital income is 39%. Do note, I assume no

[47] In the calculations, I use a Pareto parameter of 1.85 (see Appendix 10.7 for the calculations).

fiscal externalities when computing the revenue maximising tax rate. A shift to closely held firms and taxable gifts does imply not all taxable wealth that is shifted away from box 3 escapes taxation completely. To the degree these mechanisms create a positive fiscal externality, the revenue-maximising tax rate will be higher.

Finally, I discuss several limitations of this study. With respect to data, I observe taxable wealth with some error. The data on taxable wealth used in this study is directly observed from the tax filings, but tax-exempt green deposits or investments cannot be separated from taxable deposits and investments. The potential measurement error is limited by the cap on the amount of wealth taxpayers are allowed to own tax-free in these green deposits and investments. Any response to a change in the marginal tax rate that consists of taxable wealth portfolio rebalancing towards these tax-exempt asset classes is not included in my estimates.

A further data-related limitation is the short-term nature of my results. The longest time horizon over which I base my estimate is 5 years. Especially for taxation of wealth, responses in the long term are essential, as wealth is an accumulating stock and people lower in the wealth distribution might also save less in the long term (Jakobsen et al., 2020). Do note that the elasticities in Brülhart et al. (2022) are increasing over time, but level off after about 3 years. In the degree to which Dutch taxpayers respond similarly to the Swiss taxpayers in Brülhart et al. (2022), my time horizon is long enough not to underestimate the elasticity due to an overly short estimation period.

A limitation of my estimation is that I do not use a pure control group. Instead of facing no change in the marginal tax rate, my control group experiences a smaller change in the marginal tax rate than the two treatment groups. If the (potential) elasticities in the treatment and control group differ, a non-pure control can lead to a bias (Saez, 2004). Within the brackets, my heterogeneity analysis shows very little differences in the elasticity across the wealth distribution. Hence, it is plausible, that the control group used for estimates in the first and third bracket show similar (potential) elasticities. To the degree they might not, I analyse and quantify the potential bias that arises in Appendix 10.8.

Lastly, certain response margins have remained outside the scope of this study. Further research may study potential response margins such as immigration, the use of tax-deductible gifts to charity and the effect of a higher wealth tax on wealth allocation towards tax exempt asset classes such as green deposits and investment funds (if data limitations can be overcome). Another aspect that is not considered in this study, is the response of pensions savings to a change in the marginal tax rate. In the Netherlands, wealth saved in dedicated pension accounts is not taxed, meaning a change in the wealth tax rate could induces changes in pension savings. Because data on such pension savings is limited, I have not been able to investigate this margin of response due to data limitations. Lastly, my estimates excluded responses in the form of migration. Wealthy taxpayers may be persuaded to migrate out of the Netherlands in response to the increased tax on wealth. A migration response is unlikely given the small change in the tax rate, but other studies have shown such responses are possible and can be substantial (Agrawal et al., 2020).

Another avenue for further research is methodological advancement. In this study, I provide graphical evidence supporting the constant trend differentials assumption and a SDiD-estimate for the first time in elasticity of taxable wealth literature. However, as the elasticity of taxable wealth literature is developing fast, new methodological insights are already available. Garbinti et al. (2023) propose a 'dynamic bunching' technique which improves the standard bunching approach. It is beyond the scope of this paper to apply a dynamic bunching approach to the Dutch wealth tax reform of 2017, but further research may reap the benefits of this new tool in the tax elasticity literature.

8. Acknowledgements

I am grateful to Dinand Webbink for his coaching, the CPB Netherlands Bureau for Policy Analysis for the use of their facilities and René Schulenberg and Arjan Lejour for their supervision during my internship. Furthermore, I would like to thank Céline van Essen, Wouter Leenders, Simon Rabaté, Maarten van 't Riet and Francis Weyzig for their valuable comments and discussions. Any errors are my own.

9. References

Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. M. (2023). When should you adjust standard errors for clustering?. *The Quarterly Journal of Economics*, *138*(1), 1-35.

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American statistical Association*, *105*(490), 493-505.

Abadie, A., Diamond, A., & Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, *59*(2), 495-510.

Abadie, A., & L'hour, J. (2021). A penalized synthetic control estimator for disaggregated data. *Journal of the American Statistical Association*, *116*(536), 1817-1834.

Abadie, A., & Vives-i-Bastida, J. (2022). Synthetic controls in action. *arXiv preprint arXiv:2203.06279*.

Andreoni, J., Erard, B., & Feinstein, J. (1998). Tax compliance. *Journal of economic literature*, *36*(2), 818-860.

Advani, A., & Tarrant, H. (2021). Behavioural responses to a wealth tax. *Fiscal Studies*, *42*(3-4), 509-537.

Agrawal, D. R., Foremny, D., & Martínez-Toledano, C. (2020). Wealth tax mobility and tax coordination. *Available at SSRN 3676031*.

Allais, M. (1977). L'impôt sur le capital et la réforme monétaire. Hermann.

Alstadsæter, A., Bjørneby, M., Kopczuk, W., Markussen, S., & Røed, K. (2022). Saving effects of a real-life imperfectly implemented wealth tax: Evidence from Norwegian micro data. In *AEA*

Papers and Proceedings (Vol. 112, pp. 63-67). 2014 Broadway, Suite 305, Nashville, TN 37203: American Economic Association.

Alstadsæter, A., Kopczuk, W., & Telle, K. (2014). Are Closely Held Firms Tax Shelters?. *Tax Policy and the Economy*, *28*(1), 1-32.

Alvaredo, F., & Saez, E. (2009). Income and wealth concentration in Spain from a historical and fiscal perspective. *Journal of the European Economic Association*, *7*(5), 1140-1167.

Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, *91*(434), 444-455.

Angrist, J. D., & Pischke, J. S. (2008). Making regression make sense. *Mostly harmless econometrics: An empiricist's companion*, 27-112.

Atkinson, A. B., Piketty, T., & Saez, E. (2011). Top incomes in the long run of history. *Journal of economic literature*, *49*(1), 3-71.

Atkinson, A. B., & Stiglitz, J. E. (1976). The design of tax structure: direct versus indirect taxation. *Journal of public Economics*, *6*(1-2), 55-75.

Attanasio, O. P., & Weber, G. (2010). Consumption and saving: models of intertemporal allocation and their implications for public policy. *Journal of Economic literature*, *48*(3), 693-751.

Auerbach, A. J., & Hassett, K. (2015). Capital taxation in the twenty-first century. *American Economic Review*, *105*(5), 38-42.

Auten, G., & Carroll, R. (1999). The effect of income taxes on household income. *Review of economics and statistics*, *81*(4), 681-693.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, *111*(12), 4088-4118.

Arkhangelsky, D., & Imbens, G. W. (2022). Doubly robust identification for causal panel data models. *The Econometrics Journal*, *25*(3), 649-674.

Aronsson, T., Jenderny, K., & Lanot, G. (2022). The quality of the estimators of the ETI. *Journal of Public Economics*, *212*, 104679.

Bach, L., Calvet, L. E., & Sodini, P. (2020). Rich pickings? Risk, return, and skill in household wealth. *American Economic Review*, *110*(9), 2703-2747.

Bernheim, B. D. (2002). Taxation and saving. In *Handbook of public economics* (Vol. 3, pp. 1173-1249). Elsevier.

Bosch, N., Jongen, E., Leenders, W., & Mölmann, J. (2019). Non-bunching at kinks and notches in cash transfers in the Netherlands. *International Tax and Public Finance*, *26*, 1329-1352.

Bound, J., Jaeger, D. A., & Baker, R. M. (1995). Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American statistical association*, *90*(430), 443-450.

Brülhart, M., Gruber, J., Krapf, M., & Schmidheiny, K. (2022). Behavioral responses to wealth taxes: Evidence from Switzerland. *American Economic Journal: Economic Policy*, *14*(4), 111-150.

Brys, B. (2006). The Box system in the Netherlands: an alternative?. *Reflets et perspectives de la vie économique*, *45*(3), 39-52.

Chancel, L., Piketty, T., Saez, E., & Zucman, G. (Eds.). (2022). *World inequality report 2022*. Harvard University Press.

Chamley, C. (1986). Optimal taxation of capital income in general equilibrium with infinite lives. *Econometrica: Journal of the Econometric Society*, 607-622.

Chetty, R. (2009). Is the taxable income elasticity sufficient to calculate deadweight loss? The implications of evasion and avoidance. *American Economic Journal: Economic Policy*, *1*(2), 31-52.

Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*, *80*(3), 969-1018.

Chetty, R., Friedman, J. N., Leth-Petersen, S., Nielsen, T. H., & Olsen, T. (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark. *The Quarterly Journal of Economics*, *129*(3), 1141-1219.

Cnossen, S., & Bovenberg, L. (2001). Fundamental tax reform in the Netherlands. *International Tax and Public Finance*, *8*, 471-484.

Conesa, J. C., Kitao, S., & Krueger, D. (2009). Taxing capital? Not a bad idea after all!. *American Economic Review*, *99*(1), 25-48.

Diamond, P., & Saez, E. (2011). The case for a progressive tax: From basic research to policy recommendation. *Journal of Economic Perspectives*, *25*(4), 165-190.

Diamond, P., & Spinnewijn, J. (2011). Capital income taxes with heterogeneous discount rates. *American Economic Journal: Economic Policy*, *3*(4), 52-76.

DiNardo, J., Fortin, N., & Lemieux, T. (1995). Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, 64(5), 1001-1044.

Doerrenberg, P., Peichl, A., & Siegloch, S. (2017). The elasticity of taxable income in the presence of deduction possibilities. *Journal of Public Economics*, *151*, 41-55.

Doudchenko, N., & Imbens, G. W. (2016). Balancing, regression, difference-in-differences and synthetic control methods: A synthesis (No. w22791). *National Bureau of Economic Research*.

Duran-Cabré, J. M., Esteller-Moré, A., & Mas-Montserrat, M. (2019). Behavioural responses to the (re) introduction of wealth taxes. Evidence from Spain.

Elinder, M., Erixson, O., & Waldenström, D. (2018). Inheritance and wealth inequality: Evidence from population registers. *Journal of Public Economics*, *165*, 17-30.

Efron, B., & Gong, G. (1983). A leisurely look at the bootstrap, the jackknife, and cross-validation. *The American Statistician*, *37*(1), 36-48.

Fagereng, A., Guiso, L., Malacrino, D., & Pistaferri, L. (2020). Heterogeneity and persistence in returns to wealth. *Econometrica*, *88*(1), 115-170.

Feldstein, M. (1995). The effect of marginal tax rates on taxable income: a panel study of the 1986 Tax Reform Act. *Journal of Political Economy*, *103*(3), 551-572.

Feldstein, M. (1999). Tax avoidance and the deadweight loss of the income tax. *Review of Economics and Statistics*, *81*(4), 674-680.

Ferman, B. (2021). On the properties of the synthetic control estimator with many periods and many controls. *Journal of the American Statistical Association*, *116*(536), 1764-1772.

Garbinti, B., Goupille-Lebret, J., Muñoz, M., Stantcheva, S., & Zucman, G. (2023). Tax Design, Information, and Elasticities: Evidence From the French Wealth Tax (No. w31333). *National Bureau of Economic Research.*

Gelber, A. M. (2014). Taxation and the earnings of husbands and wives: evidence from Sweden. *Review of Economics and Statistics*, *96*(2), 287-305.

Gerritsen, A., Jacobs, B., Rusu, A. V., & Spiritus, K. (2020). Optimal taxation of capital income with heterogeneous rates of return. *Available at SSRN 3637931*.

Gordon, R. H., & Slemrod, J. (2000). Are" real" responses to taxes simply income shifting between corporate and personal tax bases?. In *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, edited by Joel Slemrod. New York: Harvard University Press and Russel Sage Foundation.

Groot, S., Lever, M., & Möhlmann, J. (2019). Effect van erfenissen en schenkingen op vermogensongelijkheid en de rol van belastingen. *CPB Notitie.*

Gruber, J., & Saez, E. (2002). The elasticity of taxable income: evidence and implications. *Journal of public Economics*, *84*(1), 1-32.

Guvenen, F., Kambourov, G., Kuruscu, B., Ocampo, S., & Chen, D. (2023). Use It or Lose It: Efficiency and Redistributional Effects of Wealth Taxation. *The Quarterly Journal of Economics*.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, 153-161.

He, D., Peng, L., & Wang, X. (2021). Understanding the elasticity of taxable income: A tale of two approaches. *Journal of Public Economics*, 197, 104375.

Jakobsen, K., Jakobsen, K., Kleven, H., & Zucman, G. (2020). Wealth taxation and wealth accumulation: Theory and evidence from Denmark. *The Quarterly Journal of Economics*, *135*(1), 329-388.

Jakobsen, K. M., & Søgaard, J. E. (2022). Identifying behavioral responses to tax reforms: New insights and a new approach. *Journal of Public Economics*, 212, 104691.

Jakurti, E., & Süssmuth, B. (2023). Behavioral responses to wealth taxes: Evidence from the Spanish Survey of Household Finances. *Economics Letters*, *223*, 110976.

Jongen, E. L., & Stoel, M. (2019). The elasticity of taxable labour income in the Netherlands. *De Economist*, *167*(4), 359-386.

Judd, K. L. (1985). Redistributive taxation in a simple perfect foresight model. *Journal of public Economics*, *28*(1), 59-83.

Kaldor, N. (1956). *Indian tax reform: report of a survey*. Ministry of Finance, Department of Economic Affairs.

Kawano, L., Weber, C., & Whitten, A. (2016). Estimating the elasticity of broad income for high-income taxpayers. *Available at SSRN 2852048*.

King, G., & Zeng, L. (2006). The dangers of extreme counterfactuals. *Political analysis*, *14*(2), 131-159.

Kiyotaki, N., Michaelides, A., & Nikolov, K. (2011). Winners and losers in housing markets. *Journal of Money, Credit and Banking*, *43*(2-3), 255-296.

Kleven, H. J., & Schultz, E. A. (2014). Estimating taxable income responses using Danish tax reforms. *American Economic Journal: Economic Policy*, *6*(4), 271-301.

Kleven, H. J., & Waseem, M. (2013). Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan. *The Quarterly Journal of Economics*, *128*(2), 669-723.

Kopczuk, W. (2005). Tax bases, tax rates and the elasticity of reported income. *Journal of Public Economics*, *89*(11-12), 2093-2119.

Kopczuk, W. (2013). Taxation of intergenerational transfers and wealth. In *Handbook of public economics* (Vol. 5, pp. 329-390). Elsevier.

Kopczuk, W., & Mankiw, N. G. (2019). Comments and discussion. *Brookings papers on economic activity*, 512-533.

Kumar, A., & Liang, C. Y. (2020). Estimating taxable income responses with elasticity heterogeneity. *Journal of Public Economics*, *188*, 104209.

Laffer, A. B. (2004). The Laffer curve: Past, present, and future. *Backgrounder*, *1765*(1), 1-16.

Leenders, W., Lejour, A., Rabaté, S., & Van't Riet, M. (2023). Offshore tax evasion and wealth inequality: Evidence from a tax amnesty in the Netherlands. *Journal of Public Economics*, *217*, 104785.

Lejour, A., Deinum, I., Mohlmann, J., & Vinke, R. (2022). De economische activiteit van besloten vennootschappen. CPB Publication.

Lindsey, L. B. (1987). Individual taxpayer response to tax cuts: 1982–1984: with implications for the revenue maximizing tax rate. *Journal of Public Economics*, *33*(2), 173-206.

Londoño-Vélez, J., & Ávila-Mahecha, J. (2021). Enforcing wealth taxes in the developing world: Quasi-experimental evidence from Colombia. *American Economic Review: Insights*, *3*(2), 131-48.

Londoño-Vélez, J., & Ávila-Mahecha, J. (2023). Behavioral Responses to Wealth Taxation: Evidence from Colombia. Unpublished.

Lucas Jr, R. E. (1990). Supply-side economics: An analytical review. *Oxford economic papers*, *42*(2), 293-316.

McGrattan, E. R. (2015). *Taxing Wealth* (No. 15-4). Federal Reserve Bank of Minneapolis.

Ministerie van Financiën. (2015). Wijziging van enkele belastingwetten en enige andere wetten (Belastingplan 2016). Retrieved from www.rijksfinancien.nl/belastingplan-memorie-van-toelichting/2015.

Ministerie van Financiën. (2023, April 14). Besluit op woo-verzoek over memo inventarisatie Vermogensverkeer Tussen box 2 en box 3. Retrieved from www.rijksoverheid.nl/documenten/woo-besluiten/2023/04/14/besluit-op-woo-verzoek-over-memo-inventarisatie-vermogensverkeer-tussen-box-2-en-box-3.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The review of economic studies*, *38*(2), 175-208.

OECD (2018). The Role and Design of Net Wealth Taxes in the OECD. *Organisation for Economic Co-operation and Development publishing*.

Perret, S. (2021). Why were most wealth taxes abandoned and is this time different?. *Fiscal Studies*, *42*(3-4), 539-563.

Pestieau, P., & Broadway, R. (2018). The Dubious Case for Annual Wealth Taxation. *CESifo DICE Report*, *16*.

Piazzesi, M., Schneider, M., & Tuzel, S. (2007). Housing, consumption and asset pricing. *Journal of Financial Economics*, *83*(3), 531-569.

Piketty, T., Saez, E., & Zucman, G. (2018). Distributional national accounts: methods and estimates for the United States. *The Quarterly Journal of Economics*, *133*(2), 553-609.

Piketty, T., Saez, E., & Zucman, G. (2022). Rethinking capital and wealth taxation. HAL Open Science Id: fhalshs-04104410f.

Poterba, J., & Rotemberg, J. (1987). Money in the utility function: An empirical implementation, new approaches to monetary economics: Proceedings of the second international symposium in economic theory and econometrics. 22.

Rauh, J., & Shyu, R. J. (2019). Behavioral responses to state income taxation of high earners: evidence from California (No. w26349). National Bureau of Economic Research.

Rees-Jones, A., & Taubinsky, D. (2020). Measuring “schmeduling”. *The Review of Economic Studies*, *87*(5), 2399-2438.

Ring, M. A. K. (2021). Wealth taxation and household saving: Evidence from assessment discontinuities in Norway. *Available at SSRN 3716257*.

Rosenbaum, P. R. (1984). The consequences of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society Series A: Statistics in Society*, *147*(5), 656-666.

Saez, E. (2001). Using elasticities to derive optimal income tax rates. *The review of economic studies*, *68*(1), 205-229.

Saez, E. (2004). Reported incomes and marginal tax rates, 1960-2000: evidence and policy implications. *Tax policy and the economy*, *18*, 117-173.

Saez, E. (2010). Do taxpayers bunch at kink points?. *American economic Journal: economic policy*, *2*(3), 180-212.

Saez, E. (2013). Optimal progressive capital income taxes in the infinite horizon model. *Journal of Public Economics*, *97*, 61-74.

Saez, E., Slemrod, J., & Giertz, S. H. (2012). The elasticity of taxable income with respect to marginal tax rates: A critical review. *Journal of economic literature*, *50*(1), 3-50.

Saez, E., & Stantcheva, S. (2018). A simpler theory of optimal capital taxation. *Journal of Public Economics*, *162*, 120-142.

Saez, E., & Zucman, G. (2016). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, *131*(2), 519-578.

Saez, E., & Zucman, G. (2019). Progressive wealth taxation. *Brookings Papers on Economic Activity*, *2019*(2), 437-533.

Scheuer, F., & Slemrod, J. (2021). Taxing our wealth. *Journal of Economic Perspectives*, *35*(1), 207-30.

Seim, D. (2017). Behavioral responses to wealth taxes: Evidence from Sweden. *American Economic Journal: Economic Policy*, *9*(4), 395-421.

Slemrod, J. (1995). Income creation or income shifting? Behavioral responses to the Tax Reform Act of 1986. *The American Economic Review*, *85*(2), 175-180.

Slemrod, J. (1996). High-income families and the tax changes of the 1980s: the anatomy of behavioral response. In *Empirical foundations of household taxation* (pp. 169-192). University of Chicago Press.

Slemrod, J. (1998). Methodological issues in measuring and interpreting taxable income elasticities. *National Tax Journal*, *51*(4), 773-788.

Slemrod, J. (2001). A general model of the behavioral response to taxation. *International Tax and Public Finance*, *8*, 119-128.

Slemrod, J., & Yitzhaki, S. (2002). Tax avoidance, evasion, and administration. In *Handbook of public economics* (Vol. 3, pp. 1423-1470). Elsevier.

Stock, J. H., & Yogo, M. (2002). Testing for weak instruments in linear IV regression (No. T0284). *National Bureau of Economic Research.*

Stokey, N. L. (2009). Moving costs, nondurable consumption and portfolio choice. *Journal of Economic Theory*, *144*(6), 2419-2439.

Straub, L., & Werning, I. (2020). Positive long-run capital taxation: Chamley-Judd revisited. *American Economic Review*, *110*(1), 86-119.

Stock, J. H., & Yogo, M. (2002). Testing for weak instruments in linear IV regression (No. w0284). *National Bureau of Economic Research.*

Vermeulen, P. (2016). Estimating the top tail of the wealth distribution. *American Economic Review*, *106*(5), 646-650.

Weber, C. E. (2014). Toward obtaining a consistent estimate of the elasticity of taxable income using difference-in-differences. *Journal of Public Economics*, *117*, 90-103.

Zoutman, F.T. (2018). The elasticity of taxable wealth: Evidence from the Netherlands. *Manuscript, November*.

10. Appendix

Appendix 10.1 Box 3 rules

Table A1: Overview of exemptions, bracket thresholds and effective marginal tax rates in box 3

Year	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
A: Tax base exemptions											
Tax exempt wealth (€)	20,785	21,139	21,139	21,139	21,330	24,437	25,000	30,000	30,360	30,846	50,000
Extra exemption for individuals with minor children (€)	2,779	-	-	-	-	-	-	-	-	-	-
Extra exemption for green investments	55,476	56,420	56,420	56,420	56,928	57,213	57,385	57846	58,540	59,477	60,429
Extra exemption for individuals in pension age with taxable box 1 income below around €14,300 and a maximum taxable wealth of around €280,000 (€)	27,516	27,984	27,984	27,984	28,236	-	-	-	-	-	-
Extra exemption for individuals in pension age with taxable box 1 income between around €14,300 and €19,900 and a maximum taxable wealth of around €280,000 (€)	13,785	13,992	13,992	13,992	14,118	-	-	-	-	-	-
Non-deductible debt threshold (€)	2,900	2,900	2,900	2,900	2,900	3,000	3,000	3,000	3,000	3,100	3,200
B: Cut-off points											
Between bracket 1 and 2 (€)	-	-	-	-	-	-	75,000	78,000	71,651	72,798	50,001
Between bracket 2 and 3 (€)	-	-	-	-	-	-	975,000	978,000	989,737	1,005,573	950,001
C: Return assumed by tax authority											
On wealth in bracket 1 (%)	4	4	4	4	4	4	2.87	2.02	1.93	1.79	1.90
On wealth in bracket 2 (%)	4	4	4	4	4	4	4.60	4.33	4.44	4.19	4.50
On wealth in bracket 3 (%)	4	4	4	4	4	4	5.39	5.38	5.59	5.28	5.69
D: Tax rate in box 3											
Tax rate (%)	30	30	30	30	30	30	30	30	30	30	31
E: Effective marginal tax rate											
Bracket 1 (%)	1.20	1.20	1.20	1.20	1.20	1.20	0.86	0.61	0.58	0.54	0.59
Bracket 2 (%)	1.20	1.20	1.20	1.20	1.20	1.20	1.38	1.30	1.33	1.26	1.40
Bracket 3 (%)	1.20	1.20	1.20	1.20	1.20	1.20	1.62	1.61	1.68	1.58	1.76

Notes: panel A shows the exemptions for wealth taxable in box 3. Panel B shows the cut-off points between the brackets introduced in 2017 in taxable wealth net of exemptions. Panel C shows the assumed rate of return per bracket. Panel D gives the tax levied over the assumed return following panel C. The rates in panel C and D are multiplied to give the marginal rax rate per bracket as reported in Panel E.

Appendix 10.2 Bunching results

The bracket cut-offs are defined around round numbers (€100,000 and €1,000,000). Consequently, one Has to be cautious not to interpret round-number-bunching as bunching due to the cut-off. However, as shown in Figure A1, no apparent bunching occurs. Just as Bosch et al. (2019) find Dutch taxpayers do not bunch at cut-offs for wealth in means-testing for allowances, I find they also do not bunch at box 3 tax cut-offs. The finding of no bunching as marginal tax rate thresholds is in line with the findings in France by Garbinti et al., (2023), but goes against the findings of significant bunching in Norway by Seim (2017) and Columbia by Londońo-Velez and Ávila-Mahecha (2023). These discrepancies can be rationalised. First, the Norwegian tax system features more self-reporting of asset values than the Dutch wealth tax. A higher degree of self-reporting allows taxpayers to manipulate their declared taxable wealth more easily and more precisely. For Dutch taxpayers, it is more difficult to manipulate their taxable wealth in a precise manner. Second, the Columbian wealth tax system creates discontinuous jumps in the tax liability at cut-offs, whereas the Dutch tax system only increases the effective marginal tax rate on wealth above the cut-off. As such, surpassing the threshold only slightly can increase the tax liability considerably in Columbia, while surpassing the thresholds slightly in the Netherlands does not result in a jump of the total tax liability. As such, the incentive for precise bunching is lower in the Netherlands than in Columbia.

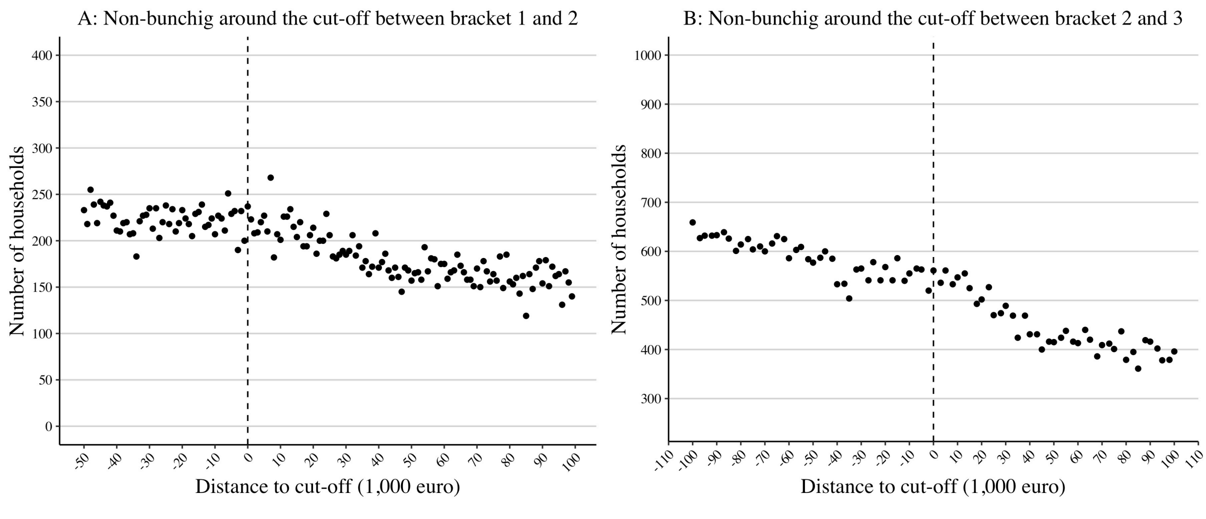

Figure A1: Density of households around bracket cut-offs. Notes: The figure shows the number of households binned by distance from the cut-off between the effective tax brackets. The years 2017-2021 are stacked. The bins have a width of €1,000 in Panel A and €2,500 in Panel B. Taxable wealth is divided by two for couples.

Appendix 10.3 Placebo reforms

Table A2: Tests for constant trend differentials in pre-reform periods

	Bracket 1		Bracket 3	
Pre-reform periods	**Validation**	**Identification**	**Validation**	**Identification**
2012 +/- 1 year	.010	.000	.126	.790
2013 +/- 1 year	.567	.023	.451	.733
2013 +/- 2 years	.803	.016	.047	.194
2014 +/- 1 year	.126	.002	.013	.599
2014 +/- 2 years	.387	.000	.058	.341
2015 +/- 1 year	.123	.003	.063	.122

Notes: Regression estimates are results of Equation (12), where the sample is restricted to pre-reform years only. For each regression, the p-values for a joint significance test of coefficients in the validation and identification region are reported. The sample used for estimations in bracket 1 contains taxpayers with initial wealth in the range €40.000-€300.000, where €100.000 is the cut-off between the identification and validation region. The sample used for estimations in bracket 3 contains taxpayers with initial wealth in the range €300.000-€3.000.000, where €1.000.000 is the cut-off between the identification and validation region. The control variables included in the inverse probability weighting (IPW) procedure are dummies for being a single person household, a couple or other, age group, having children, a dummy if the household has substantial ownership in a firm and share of wealth held in financial investments, real estate, and other debt, as well as the share of other debt in gross wealth. P-values are based on robust standard errors.

Appendix 10.4 Robustness to discretionary choices

Table A3: Effect of discretionary choices on the elasticity estimates

	Elasticity bracket 1			**Elasticity bracket 3**		
	F-stat first stage	Reduced form	2SLS	F-stat first stage	Reduced form	2SLS
Baseline	740.8	0.7	0.8	135.9	36.0***	43.7***
		(0.7)	(0.7)		(5.2)	(6.3)
P-value assumption		.607			.000	
Observations		1,669,094			208,428	
Half bin size	817.2	4.5***	4.7***	155.8	32.6***	39.7***
		(0.6)	(0.6)		(6.1)	(7.3)
P-value assumption		.225			.000	
Observations		2,089,339			208,428	
Cap at €4M				161.8	37.9***	45.5***
					(5.2)	(6.3)
P-value assumption					.000	
Observations					210,252	
Cap at €5M				164.9	39.1***	46.8***
					(5.2)	(6.2)
P-value assumption					.000	
Observations					211,033	
No filters	744.1	1.0	1.2	152.9	35.0***	42.8***
		(0.7)	(0.8)		(5.5)	(6.8)
P-value assumption		.602			.010	
Observations		1,941,557			264,525	

Notes: Regression estimates are results of variations in the specification of Equation (1). The baseline estimates restrict the sample for estimations on bracket 1 to taxpayers with initial wealth in the range €40.000-€300.000, where €100.000 is the cut-off between the identification and validation region. The sample used for estimations in bracket 3 contains taxpayers with initial wealth in the range €300.000-€3.000.000, where €1.000.000 is the cut-off between the identification and validation region. Control variables included in the inverse probability weighting (IPW) procedure are dummies for being a single person household, a couple or other, age group, having children, a dummy if the household has substantial ownership in a firm and share of wealth held in financial investments, real estate, and other debt, as well as the share of other debt in gross wealth. The p-value for a joint significance test of the coefficients in the validation area is shown, where a statistically significant p-value indicates that the identifying assumption of constant trend differentials does not hold. Robust standard errors in parentheses.

Appendix 10.5 Synthetic-difference-in-differences results

The SDiD estimator combines attractive features of the synthetic control (SC) and DiD estimators (Arkhangelsky et al., 2021). The SC method constructs a synthetic control units based on a weighted average of control units available in the data. The weights are chosen to minimise the error between the outcome variable of the treatment unit and synthetic control unit in the pre-treatment period. If a weighted average of control units can match the outcomes of the treated unit in the pre-treatment period, then a discrepancy between the outcome the treated and synthetic control unit during the treatment period can be attributed to the causal effect of the treatment.[48]

The SDiD estimator improves the SC method in two main ways (Arkhangelsky et al., 2021). First it allows for the outcomes of the treated unit and the synthetic control unit to differ by a constant (the DiD element). This modification allows the outcomes of the treated unit to lie outside the convex hull of outcomes in the donor pool, as the synthetic control unit is only required to match the *trend* of the outcome variable, instead of matching the *level* during the pre-treatment period. Second, the SDiD method additionally reweights the pre-treatment years in a way that makes the average outcome of the control units in the pre-treatment and treatment period differ by a constant.

The combination of reweighting and the DiD-component makes the SDiD a doubly robust estimator. Double, because when the structural model identifies the causal effect (DiD-style identification), a misspecification in the weights will not pose any issues, while under a miss-specified structural model, successful reweighting (SC-style identification) will still yield a causal effect (Ferman, 2021; Arkhangelsky et al., 2021; Arkhangelsky and Imbens, 2022). Hence, a major advantage of the SDiD method is that it relies less strongly on the common trend assumption.

An application of the SDiD method to a setting with many potential control units, as is the case for the application in this paper, offers both advantages and disadvantages.[49] An advantage is that the SC weighting procedure has many degrees of freedom, increasing the chance a good fit can be found (Doudchenko and Imbens, 2016; Ferman, 2021). A disadvantage is that many degrees of freedom can also lead to overfitting. The risk of over-fitting becomes larger when the variation in the outcome variable due to unobserved factors is larger relative to the variation caused by the observed factors, when the number of pre-treatment periods is smaller and when the donor pool is larger (Abadie and Vives-i-Bastida, 2022). However, constraining the weights to be non-negative and adding up to one acts as a regularisation method that prevents such overfitting (Ferman; 2021; Abadie and L'Hour, 2021). In the SDiD algorithm, Arkhangelsky et al. (2021) additionally employ a regularisation method based on the variance of the outcome variable of the control units to further prevent overfitting.

---

[48] In the setting of this paper, applying the SC method is infeasible. The treatment and control group are separated based on the level of the outcome variable in the pre-treatment period, so the average outcome of the treatment group lies outside any possible weighted combination of control units, as long as the weights are limited to be nonnegative and sum to 1.

[49] An application of the SDID method to estimate the elasticity of taxable income is provided in Rauh and Shyu (2019).

Still, the size of the sample poses computational difficulties. In the first bracket estimates, where the number of observations is a multitude higher than for the third bracket estimates, to runtime of the SDiD-algorithm becomes infeasibly long when using the full sample.[50] Therefore, I precede the SDiD-algorithm by filtering out taxpayers in the pool of control units who are less comparable to the treatment group. Specifically, I estimate the propensity to be treated based on taxpayer characteristics following the same IPW-procedure as detailed in Section 5.1.4 and only keep the top 25% potential control units with the highest propensity to be treated. Selection based on similarity between the treatment and control group to reduce the risk of over-fitting is similar in spirit to the suggestion of Abadie and L'Hour (2021) to penalise discrepancy between the treatment group and synthetic control unit in characteristics other than the outcome variable. My method hence narrows down the pool of potential control units based on comparability to the treatment group, further limiting the risk of over-fitting.

Finally, following the advice in Arkhangelsky et al. (2021), I compute standard errors using a Jackknife procedure. The Jackknife procedure repeats the estimation, while iteratively leaving out one additional unit from the sample (Efron and Gong, 1982). The sample standard error from the resulting distribution of estimates is used as standard error for the estimate using the full sample. Jackknife standard errors are preferred when applying the SDiD to a large panel, as they are more conservative than bootstrap standard errors and less computationally intensive than placebo standard errors (Arkhangelsky et al., 2021).

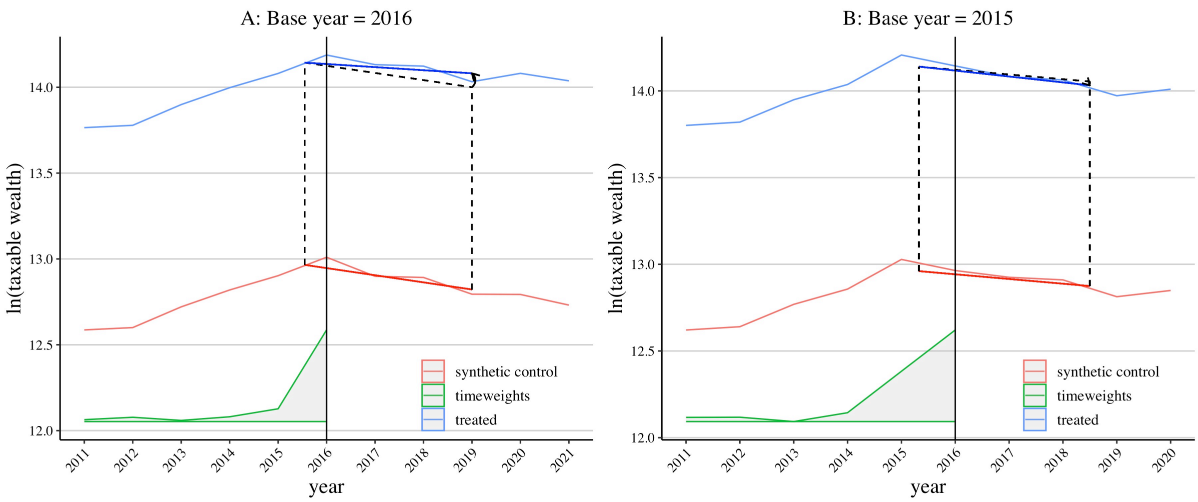

Figure A2: Synthetic-difference-in-differences results for response in third bracket. Notes: In Panel A, the treatment group is defined as taxpayers with an amount of taxable wealth in 2016 that would be taxed in the third bracket in the reform period. The pool of control units consists of taxpayers who would be taxed in bracket 2, with taxable wealth down to €300,000. In Panel B, the treatment status is determined based on taxable wealth in 2015. The arrow indicates the ITT effect as the difference between the mean observed change in log taxable wealth and the change in the synthetic control group.

[50] For illustration of the necessity of the IPW-procedure to make the SDiD estimator computationally feasible: even after shrinking the pool of potential control units by dropping the 75% taxpayers with the least resemblance to the treatment group, the SDiD algorithm still required a runtime of 15 hours.

When using 2016 as base year, the synthetic control unit for the third bracket matches the pre-trend of the third bracket well (Figure A2). In the reform period, the observed average decrease in taxable wealth is slightly smaller than the decrease in the synthetic control group. This result indicates the increase in the marginal tax rate resulting in a slight increase of taxable wealth. The SDiD-estimate hence contradicts the finding of a large decrease when using the DiDiD method. It could be that the SDiD-estimator fails to adequately control for the differential growth trends and mean reversion, as it only matches based on the period leading up to the base year, meaning it only matches on the increasing part of the mean reversion effect. The match might be improved by defined the instrument on an earlier base year, so the training data also contains the downward section of mean reversion. Panel B shows the results using 2015 as a base year. The pre-trend is again matched well and the outcome changes slightly. The effect now shows a small decrease in taxable wealth, which is in line with the sign of the effect found in the DiDiD analysis in terms of sign but remains statistically insignificant. Estimating the response in the first bracket using 2015 as a base year also shows a statistically insignificant result, with a positive point estimate of 0.005.

Appendix 10.6 Translating wealth elasticities to capital income elasticities

In the elasticity of taxable wealth literature, it is common to translate the estimates into an elasticity of taxable wealth with respect to the net-of-tax rate of return (Zoutman, 2018; Duran-Cabré et al., 2019; Jakobsen et al., 2020; Brülhart et al., 2022; Londońo-Velez and Ávila-Mahecha, 2023). By expressing the elasticity in terms of net-of-tax rate of return, the elasticity can be interpreted as the elasticity of taxable wealth with respect taxable capital income taxes, which plays a pivotal role in the model for optimal capital taxation derived in Saez and Stantcheva (2018). Furthermore, it allows the elasticity to be compared more directly to the elasticity of taxable (labour) income (Brülhart et al., 2022; Londońo-Velez and Ávila-Mahecha, 2023).

The elasticity of taxable wealth with respect to the net-of-tax-rate can be translated into an elasticity with respect to the net-of-tax rate of return by scaling the semi-elasticity by the corresponding change in the net-of-tax rate of return. Following Jakobsen et al., (2020), the change in the net-of-tax rate of return is given by,

$$\frac{\Delta\tau * R}{(1-\tau)R-1}, \tag{13}$$

where $\Delta\tau$ is the change in the marginal tax rate, R is the gross rate of return and τ is the before-change tax rate. The semi-elasticity of taxable wealth is expressed relative to a 1 percentage point change in the marginal tax rate ($\Delta\tau = 0.01$) and the tax rate prior to the reform was 1.2% for all taxpayers ($\tau = 0.012$). The gross rate of return cannot be observed in the data, so I consider a range between 4% and

6% ($R \in [1.04, 1.05, 1.06]$).[51] Using a rate of return of 1.04, 1.05 and 1.06, a 1 percentage point change in the marginal wealth tax rate results in a change in the marginal net-of-tax rate of return of respectively, 37.8%, 28.1% and 22.4%. Hence, a semi-elasticity in the top bracket of 43.3% implies an elasticity with respect to the net-of-tax rate of return of 1.15 when assuming a return of 4%, 1.54 for a return of 5% and up to 1.93 when assuming a return of 6%. Using the semi-elasticity of 24.9%, the elasticity with respect to the net-of-tax rate of return becomes 0.66, 0.89 and 1.11 using the same assumed returns respectively. Finally, the elasticity in the first bracket of 4.7% translates into an elasticity with respect to the net-of-tax rate of return of 0.12, 0.17 and 0.21 respectively.

Compared to the elasticity of taxable labour income of 0.24 in the Netherlands (Jongen and Stoel, 2019), the elasticity of taxable wealth with respect to the net-of-tax rate of return is slightly lower, but similar. The elasticity in the top bracket is much larger than the elasticity of taxable income, larger than in Kleven and Schultz (2014), who find an elasticity of capital income up to three times the size as the elasticity of labour income in Denmark. This implies capital income is less responsive to taxes than labour income taxes for taxpayers in the first bracket, but more responsive for taxpayers in the top bracket.

Appendix 10.7 Revenue effects

In this section, I apply the elasticities estimated in this paper to a theoretical framework for revenue effects and derive practical policy implications regarding tax base erosion in box 3 and the revenue maximising tax rate. As such, I can assess whether my estimates provide evidence for the assumptions the Dutch Government makes when projecting the effect of behavioural responses to a change in the wealth tax rate on wealth tax revenue. Due to a lack of empirical evidence, the Dutch Ministry of Finance currently makes assumptions on the size of the behavioural response to wealth taxes. No behavioural response is assumed for taxpayers with taxable wealth below €100.000 (the first bracket) and a small response for taxpayers with taxable wealth above €100.000. Specifically, the Ministry of Finance presumes 20% of the mechanical increase in wealth tax revenue from an increase in the tax rate will be lost due to behavioural responses eroding the tax base.

10.7.1 The model

For the analysis, I use the theoretical framework for tax revenue effects and optimal tax rates as laid out in Saez (2001; 2004) and Saez et al. (2012).[52] The model decomposes the change in tax revenue (dR)

---

[51] A rate of return of 4% corresponds to the rate of return the wealth tax law assumed taxpayers would gain on their taxable wealth in the years before 2017, which rose to closer to 6% in the years following. In equivalent calculations, Brülhart et al. (2022) use a return of 4.5%, and Londońo-Velez and Ávila-Mahecha (2018) and Jakobsen et al. (2020) use a return of 5%.

[52] The model is based on the optimal taxation framework originally constructed by Mirrlees (1971).

given a change in the marginal tax rate into the mechanical change (dM) and the behavioural response (dB),

$$dR = dM + dB. \tag{14}$$

The mechanical effect (dM) captures the change in tax revenue in absence of any behavioural response. It consists of the change in the marginal tax rate ($d\tau$) times the mean taxable wealth (W_m) above threshold $\overline{W}$ where the marginal tax rate applies, times the number of taxpayers above the threshold N,

$$dM = Nd\tau(W_m - \overline{W}). \tag{15}$$

The behavioural effect (dB) captures the change in tax revenue due to the effect of the behavioural response on the tax base following a change in the marginal tax rate. It is given by the change in reported taxable wealth due to the behavioural response ($dW_m = -e * W_m * d\tau/(1-\tau)$). The average behavioural response times the number of people and marginal tax rate gives the effect on total tax revenue, $dB = -NdW_m\tau$, or,

$$dB = -N * e * W_m \frac{d\tau}{1-\tau} * \tau.[53] \tag{16}$$

Combining the mechanical and behavioural effect gives,

$$\begin{aligned} dR = dM + dB &= Nd\tau(W_m - \overline{W}) - NeW\frac{d\tau}{1-\tau}\tau \\ &= Nd\tau(W_m - \overline{W})\left(1 - e\frac{W_m}{W_m - \overline{W}} * \frac{\tau}{1-\tau}\right). \end{aligned} \tag{17}$$

In a final step. let the Pareto parameter a denote the fraction $\frac{W_m}{W_m - \overline{W}}$. It is common for the distribution of top income and wealth to follow a Pareto distribution (Atkinson et al., 2011), in which a is a constant for any given $\overline{W}$,

$$dR = Nd\tau(W_m - \overline{W})\left(1 - e * a * \frac{\tau}{1-\tau}\right).[54] \tag{18}$$

The share of tax revenue that is lost due to behavioural responses is captured by the term between brackets and increases in the net-of-tax rate elasticity e, and the Pareto parameter a. The lower a, the thicker the tail of the taxable wealth distribution, indicating wealth is more concentrated at the top of the distribution (Atkinson et al., 2011). Because tax rate and the Pareto parameter are straight forward to obtain, the main empirical challenge lies in estimating the elasticity (Saez et al., 2012).

---

[53] Note that this formula ignores the effect a change in the marginal tax rate might have on the amount of taxpayers (N) in the tax bracket (Saez et al., 2012). This aligns with the fact that my estimates for the elasticity are local for taxpayers who do not switch between brackets. Both the estimates and the revenue effects should hence be interpreted as given for the population of taxpayers who remain within the same bracket and are hence underestimated to the degree a change in the marginal tax rate does induce taxpayers to switch tax brackets.

[54] Equation (18) is not only central to tax revenue, but also welfare analysis. Under strict conditions, the term between brackets in Equation (18) equals the welfare loss due to the distortions the tax creates, the marginal deadweight loss (Saez et al., 2012).

The formula's derived above assume that the entire loss of taxable wealth due to the behavioural response is no longer taxed. However, part of the tax base may be shifted to other forms of wealth or income that are still taxed. Such effects can be incorporated in the formula by including two extra parameters. First, *s* captures the fraction of taxable wealth that is shifted to another tax base (such as a closely held firm), leaving 1 – *s* as the fraction of taxable wealth that does not show up at an alternative tax base (such as a shift towards unreported cash or tax-exempt assets). Second, *t* captures the effective tax rate in the alternative tax base. The loss in tax revenue from the behavioural response now equals $d\tau * (\tau - s * t)$. Incorporating the fiscal externality into the equation for tax revenue gives,

$$dR = Nd\tau(W_m - \overline{W})\left(1 - e * a * \frac{\tau - s * t}{1 - \tau}\right). \tag{19}$$

Intuitively, the larger the fiscal externality ($s * t$), the smaller the revenue loss from the behavioural response.

10.7.2 Revenue effects

Having established the model, I next parameterise the model for taxpayers in the top bracket assuming no fiscal externalities using the estimates from this paper. For the elasticity with respect to the marginal net-of-tax rate in the top bracket, I consider 25 (estimated on savings) as lower bound and 43 as upper bound (the highest estimate on the full sample). Saez et al. (2012) note that the elasticity used in the formula above should be weighted by individual income, or in this case, taxable wealth. The reason is that the response of taxpayers with more taxable wealth has a larger effect on tax revenue. In these calculations, I use unweighted elasticities, motivated by the fact that I do not find heterogeneity in elasticities within tax brackets.

The Pareto parameter for wealth above €1,000,000 is 1.85 and stable across the years 2011-2020.[55] This Pareto parameter is larger than the Pareto parameter for capital income in the United States of 1.38 (Saez and Stantcheva, 2018) and substantially lower than the Pareto parameter for the distribution of taxable income in the Netherlands, at 3.16 (Jacobs et al., 2017). A lower Pareto parameter for the taxable wealth distribution than the taxable income distribution shows that in the Netherlands, wealth is significantly more concentrated at the top than income, which corroborates the findings of Chancel et al. (2022).

Using the elasticity of 25 and the current marginal tax rate in the top bracket of 1.6%, I find the behavioural effect excluding fiscal externalities equals 75.2% ($e * a * \frac{\tau}{1-\tau}$). This effect implies that 75.2% of the mechanical increase in tax revenue from an increase in the top wealth tax rate is lost through base erosion in box 3. Using the elasticity of 43, the behavioural effect rises to 129.3%, implying the behavioural effect dominates the mechanical effect. In the bottom bracket, an elasticity of 4.7 implies a behavioural effect of 3.5%, using a Pareto parameter of 1.25 and the marginal tax rate of

[55] Using taxable wealth in 2021, the Pareto parameter increases to 2.

0.6%. Using the zero-elasticity result in the bottom bracket would obviously result in a behavioural effect of 0%.

The results corroborate the assumption used by the Dutch Ministry of Finance that taxpayers with taxable wealth up to €100,000 do not respond (or negligibly) to a change in the marginal tax rate. I can only test the assumption of a behavioural effect of 20% for taxpayers with more than €100,000 partially. The calculations above show this assumption is likely to be invalid for taxpayers with taxable wealth above €1,000,000, where the behavioural effect is significantly larger. However, it remains possible that 20% is a valid rule of thumb for taxpayers in (the lower region of) bracket 2.

10.7.3 The revenue maximising top rate

The revenue maximising marginal tax rate in the top bracket can be calculated using just the Pareto parameter and the elasticity of taxable wealth with respect to the net-of-tax rate. The revenue maximising tax rate is set such that an extra euro of tax revenue is exactly offset by the behavioural response.

$$\tau^* = \frac{1}{1 + a * e}. \tag{20}$$

With a lower tax rate, there is still room to increase revenue by setting a higher tax rate, but above τ^*, the behavioural response dominates the mechanical effect, meaning a higher tax rate would only decrease tax revenue.

The revenue-maximising top rate equals 2.12% when using the elasticity with respect to the net-of-tax rate of 25. With the current marginal tax rate standing at 1.6%, the government can significantly increase wealth tax revenue by increasing the top rate. Using the high elasticity of 43, the revenue-maximising tax rate falls to 1.24%, which is close to the effective tax rate prior to the 2017 wealth tax reform. In the latter case, the marginal tax rate would be above the Laffer curve, implying the government can increase tax revenue by decreasing the tax rate (Laffer, 2004). Also here, I have assumed no fiscal externalities.[56]

Using the translation of the elasticity of taxable wealth with respect to the net-of-tax rate to the elasticity with respect to the net-of-tax rate of return in Appendix 10.4, I can use Equation (20) to similarly estimate the implied revenue-maximising *capital income* tax rate (Saez and Stantcheva, 2018). Using the elasticity of 25 in the top bracket and an assumed rate of return of 5%, the implied elasticity with respect to the net-of-tax rate of return is 1.12. The resulting revenue-maximising capital income tax rate is 37.8%.[57] When interpreting the tax in box 3 as a capital income tax (it's legal definition), the capital income tax in 2021 stands at 31%. An elasticity of 25 hence implies the current tax rate is below

[56] fiscal externalities can be incorporated in the revenue-maximising top rate as, $\tau^* = \frac{1+s*t*a*e}{1+a*e}$.
[57] Implicitly, I assume capital income follows a parallel distribution to taxable wealth, which is valid when assuming taxpayers above €1,000,000 face the same rate of return.

the revenue-maximising tax rate for the top bracket. Using the elasticity of 43 and an assumed return of 5% as input, however, the revenue-maximising tax rate falls to 26.1%.

Note that the formula for the revenue-maximising tax rate in the top bracket does not necessarily reflect the optimal tax rate from the perspective of social welfare. To derive the socially optimal tax rate, an assumption is required on the social welfare weights the social planner assigns to taxpayers in the top bracket (Mirrlees, 1971; Saez, 2004). As preferences for redistribution are a political question, I abstain from deriving the socially optimal marginal tax rate in the top bracket.

Appendix 10.8 Quantifying the potential bias from non-pure control group

The potential bias arising from heterogeneity in the elasticity of taxable wealth between the treatment and control group is recognised (Saez, 2004; Saez et al., 2012), but receives little attention in the tax elasticity literature (Kumar and Liang, 2020). The DiDiD-estimator estimates the elasticity as,

$$\hat{e} = \frac{E(\Delta \log(W)|T, W_{t-k}) - E(\Delta \log(W)|C, W_{t-k})}{E(\Delta \log(1-\tau)|T, W_{t-k}) - E(\Delta \log(1-\tau)|C, W_{t-k})} \quad (21)$$

where the numerator captures the difference in expected growth of log taxable wealth between the treatment (T) and control group (C), given initial wealth W_{t-k}. The denominator captures the differences in expected change in the marginal-net-of-tax rate between the treatment and control group. Recall that the identifying assumption states the numerator is zero in absence of treatment (constant trend differentials). For easy of readability, I drop the W_{t-k} controls for base year wealth from the expression,

$$\hat{e} = \frac{E(\Delta \log(W)|T) - E(\Delta \log(W)|C)}{E(\Delta \log(1-\tau)|T) - E(\Delta \log(1-\tau)|C)}. \quad (22)$$

To see how a bias may occur from heterogeneity in the underlying elasticities in the treatment and control group, assume the true elasticity in the control group is zero ($E(\Delta \log(W)|C) = 0$), while the true elasticity in the treatment group is e ($E(\Delta \log(W)|T) = e * E(\Delta \log(1-\tau)|T)$). If the control group is a pure control group (no change in the marginal tax rate, hence $E(\Delta \log(1-\tau)|C) = 0$), the estimated elasticity equals,

$$\hat{e} = \frac{e * E(\Delta \log(1-\tau)|T) - 0}{E(\Delta \log(1-\tau)|T) - 0}$$

$$\hat{e} = e * \frac{E(\Delta \log(1-\tau)|T)}{E(\Delta \log(1-\tau)|T)}$$

$$\hat{e} = e * 1 = e. \quad (23)$$

As such, the true elasticity in the treatment group is identified without a bias, even when the treatment and control group have different underlying elasticities. A bias does arise when the control group faces a change in the marginal tax rate too. If, for example, the control group faces a change in the marginal tax rate half the size of the change in the treatment group ($E(\Delta \log(1-\tau)|C) = 0.5 * E(\Delta \log(1-\tau)|T)$). The estimated elasticity becomes,

$$\hat{e} = \frac{e * E(\Delta \log(1-\tau)|T) - 0}{E(\Delta \log(1-\tau)|T) - 0.5 * E(\Delta \log(1-\tau)|T)}$$
$$\hat{e} = e * \frac{E(\Delta \log(1-\tau)|T)}{(1-0.5) * E(\Delta \log(1-\tau)|T)}$$
$$\hat{e} = \frac{e}{0.5} * \frac{E(\Delta \log(1-\tau)|T)}{E(\Delta \log(1-\tau)|T)}$$
$$\hat{e} = \frac{e}{0.5}$$
$$2\hat{e} = e. \qquad (24)$$

This example shows that the estimated elasticity is biased, overestimated by a factor 2, when the underlying elasticities treatment and control group are different, *e* versus 0 respectively, and the control group faces a change in the marginal tax rate half the size of the change in the treatment group.

To analyse the potential bias in my estimates, I extend the short discussion in Saez (2004) and Saez et al. (2012), by formally deriving a formula capturing the size and direction of the bias. Two factors describe the bias:

$$A = \frac{E(\Delta \log(1-\tau)|C)}{E(\Delta \log(1-\tau)|T)}, \qquad (25)$$
$$B = \frac{e_C}{e_T}, \qquad (26)$$

where A is the treatment intensity in the control group, expressed as a fraction of the treatment intensity in the treatment group. By definition, the treatment intensity in the control group is lower than in the treatment group but can take a different sign ($|A| < 1$). When $A = 0$, the control group is a pure control group. The factor B captures the degree to which the true elasticity in the control group differs from the true elasticity in the treatment group, expressed as a fraction of the true elasticity in the treatment group. Note that B cannot be observed. When $B = 1$, both elasticities are equal. Substituting these definitions into Equation (21) gives,

$$\widehat{e_T} = \frac{e_T * E(\Delta \log(1-\tau)|T) - e_T * AB * E(\Delta \log(1-\tau)|T)}{E(\Delta \log(1-\tau)|T) - A * E(\Delta \log(1-\tau)|T)}$$
$$\widehat{e_T} = \frac{(e_T - e_T * AB) * E(\Delta \log(1-\tau)|T)}{(1-A) * E(\Delta \log(1-\tau)|T)}$$
$$\widehat{e_T} = \frac{e_T(1-AB)}{1-A} * \frac{E(\Delta \log(1-\tau)|T)}{E(\Delta \log(1-\tau)|T)}$$
$$\widehat{e_T} = e_T * \frac{1-AB}{1-A}. \qquad (27)$$

This simple formula allows me to express the bias as a factor ($\frac{1-AB}{1-A}$) described by the degree to which the fundamental elasticity differs between the treatment and control group (*B*) and the degree to which the control group faces a change in the marginal tax rate too (A). If A = 0, there is not bias, as there is a pure control group. If A = 1, the control group exactly matches the treatment of the treatment group, essentially indicating there is no control group. The closer B is to 1, the smaller the bias.

The interpretation of the bias formula is simple. If the control group is a pure control group ($A = 0$), the estimate is always unbiased, as the bias factor equals 1 for any value of B. If the control group also experiences some change in the marginal tax rate ($A \neq 0$), a bias arises when the true elasticity in the treatment and control group differ ($B \neq 0$). If B equals 1, there is no bias, as there is no heterogeneity in the underlying elasticity. The more B deviates from 1 and A is larger than 0, the larger the potential bias becomes. Both are intuitive, the bias becomes larger if the treatment and control group are less similar in their potential elasticity and this bias becomes larger if the control group faces a larger change in the marginal tax rate, causing their response to become larger. To facilitate intuition into the magnitude of the bias given the size of factor A and B, Table A4 reports the bias factor per combination of both factors.

Table A4: Bias from non-pure control group and heterogeneous underlying elasticities

		Factor B										
		0	**0.2**	**0.4**	**0.6**	**0.8**	**1**	**1.2**	**1.4**	**1.6**	**1.8**	**2**
Factor A	**-0.3**	0.77	0.82	0.86	0.91	0.95	*1.00*	1.05	1.09	1.14	1.18	1.23
	-0.2	0.83	0.87	0.90	0.93	0.97	*1.00*	1.03	1.07	1.10	1.13	1.17
	-0.1	0.91	0.93	0.95	0.96	0.98	*1.00*	1.02	1.04	1.05	1.07	1.09
	0	*1.00*	*1.00*	*1.00*	*1.00*	*1.00*	*1.00*	*1.00*	*1.00*	*1.00*	*1.00*	*1.00*
	0.1	1.11	1.09	1.07	1.04	1.02	*1.00*	0.98	0.96	0.93	0.91	0.89
	0.2	1.25	1.20	1.15	1.10	1.05	*1.00*	0.95	0.90	0.85	0.80	0.75
	0.3	1.43	1.34	1.26	1.17	1.09	*1.00*	0.91	0.83	0.74	0.66	0.57
	0.4	1.67	1.53	1.40	1.27	1.13	*1.00*	0.87	0.73	0.60	0.47	0.33
	0.5	2.00	1.80	1.60	1.40	1.20	*1.00*	0.80	0.60	0.40	0.20	0.00
	0.6	2.50	2.20	1.90	1.60	1.30	*1.00*	0.70	0.40	0.10	-0.20	-0.50

Notes: This table shows the bias resulting from a non-pure control group combined with heterogeneous underlying elasticities in the treatment and control group. Factor A equals the treatment intensity (change in the marginal tax rate) in the control group expressed as fraction of the treatment intensity in the treatment group. Factor B equals the true elasticity in the control group expressed as fraction of the true elasticity in the treatment group. The true elasticity equals the estimated elasticity divided by the bias factor presented in the table.

Finally, I apply formula (27) to provide an indication of the potential bias in my estimates. The change of the marginal tax rate in the control group as percentage of the change in the treatment group (factor A) can be observed and equals -0.21 for the bracket 1 estimates and 0.28 for the bracket 3 estimates. Factor B cannot be observed, so has to be assumed. In the extreme case where taxpayers in the second bracket are not responsive at all (B=0), the estimated elasticity is biased by a factor of 1.39 in the top bracket, meaning the estimate of 43.4 corresponds to a true elasticity with respect to the marginal net-of-tax rate of 31.2. Similarly, the estimate of 24.9 would imply a true elasticity of 17.9.

Considering less extreme scenario's: a true elasticity in the control group of half the true elasticity in the treatment group implies a bias by a factor of about 1.19 and hence a true elasticity of 36.5 underlying the estimate of 43.4 in the top bracket. Finally, let's say the elasticity in the control group for the bottom bracket is twice as large as the true elasticity in the bottom bracket. The bias factor would be 1.17, implying the estimate of 4.7% corresponds to a true elasticity of 4%. Under any smaller deviation between the true elasticity in the treatment and control group, the bias will be smaller.