Behavioral Responses to the Norwegian Wealth Tax

ØYVIND BRUER-SKARSBØ

THESIS FOR MASTER OF PHILOSOPHY IN ECONOMICS

DEPARTMENT OF ECONOMICS

UNIVERSITY OF OSLO

MAY 2015

PREFACE

This thesis was written as a completion of the Master of Philosophy in Economics at the University of Oslo. I would like to thank everyone who has helped and supported me through my time as a student. The thesis was written under the excellent supervision of Elin Halvorsen at Statistics Norway. Without her vast knowledge about the Norwegian tax system and econometric modelling, this thesis would have never seen the light of day. I would also like to thank Statistics Norway for providing me with an inspiring working environment and access to their valuable data. I am also grateful to Oslo Fiscal Studies who provided me with a scholarship and an office space to write this thesis. A big thank you is also in order to my study group, consisting of geniuses like Ingrid Hjort, Trond Vigtel, Kristine Wika Haraldsen, and Camilla Karto. Last but not least, I would have never been able to do this without the love and support of my wife Marte and son August.

2015

Behavioral Responses to the Norwegian Wealth Tax

Øyvind Bruer-Skarsbø

http://www.duo.uio.no/

Printing: Reprosentralen, Universitetet i Oslo

Abstract

The effects of the wealth tax have been an endless source of debate in Norway. Those in favor of it claim that it has important distributional properties, while those who want to abolish it usually make their claim on efficiency grounds. A frequently used argument against the taxation of wealth has been that it discourages private savings. This thesis uses two empirical strategies to investigate this claim. A regression discontinuity design is applied to the tax bracket cutoff in 2011, and a differences-in-differences approach is applied to the change in allowance in 2010. While the results from the two strategies point in somewhat opposite directions, it does not support the claim that the Norwegian wealth tax discourages savings.

CONTENT

1 INTRODUCTION

The wealth tax is a much debated topic in many countries. This was recently made apparent by the overwhelming popularity of French economist Thomas Pikettys latest book, Capital in the Twenty-First Century. At the core of the debate is the inevitable trade-off between efficiency and equity that arises when designing any tax system. All governments levy taxes to offer some level of public service. Many also want to redistribute wealth or income from those who have a lot to those who have less. In order to redistribute resources, governments need to levy taxes, and by doing so they distort peoples choices. Resources are relocated, but at the expense of efficiency. Since all taxes more or less create this efficiency loss, the question is how to design a redistributive tax system that minimizes this cost. Much of the debate around the wealth tax concerns the question of whether the wealth tax is the best means to this end.

If we distinguish between the wealth tax and the property tax, the trend in OECD countries have been to remove the wealth tax and strengthen the property tax. Norway, France and Switzerland are among the few comparable countries that have had a tax on wealth for longer period of time. In Norway, there is a constant pressure for its removal, and the current government has stated that they want to gradually fade it out. There are two main arguments against the wealth tax in the public debate. Some claim that it forces business-owners to take out capital from the businesses to pay for the tax. In other words, the wealth tax makes some businesses capital constrained. Empirical studies has investigated this claim and found that this problem is minimal in Norway (Edson, 2012). The other argument is that the wealth tax discourages savings. This claim has not been tested empirically on data from Norway, and motivates the research question in this thesis. Using registry data from Norway, I will try to shed some light on the following claim:

I. The Norwegian wealth tax decreases the level of savings in Norway.

Savings are taxed via the capital income tax and the wealth tax. A key difference between these taxes is that the wealth tax is independent of how large the returns to the assets are. Since Norway has both types of taxes, the effective tax rate can be very high, sometimes exceeding 100%. Simple intuition could lead one to think that a decrease in the tax on savings would increase the level of saving. But this view does not take into account that people plan for the future. A decrease in the tax rate means that for the same amount of saving you keep a larger share of your returns. This might result in people consuming more today, while consuming the same later in life. Thus, it is not clear how the wealth tax affects savings.

This master thesis is organized as follows. Section 2 presents the features of the Norwegian wealth tax. In the recent years there have been several changes in the tax rates, allowances and valuations of different assets. The general trend has been that fewer people are affected by the tax, but those who are affected face a higher tax rate. These changes make up the foundation for my empirical approaches. Section 3 reviews the theoretical framework for the taxation of wealth. The classical literature on optimal tax theory has mostly found it optimal to have a zero tax on capital (Chamley, 1986; Judd, 1985), making no distinction between stocks, flows and transfers. More recent work has concluded that the taxation of wealth might be warranted from a distributional perspective (Diamond and Saez, 2011; Piketty and Saez, 2013). Section 4 presents the data used to estimate the effects of the wealth tax. Statistics Norway has provided me with accurate panel data on people's wealth from 2008 to 2011. I also operationalize the dependent variable to be active savings; the part of savings that involves actively placing part of your income in an asset. Section 5 discusses the two empirical strategies used to answer my research question. The first is a regression discontinuity design around the 2011 tax bracket cutoff. The second is a differences-in-differences approach that takes advantage of the panel structure of the data. The results presented in section 6 points in somewhat opposite directions. The regression discontinuity estimates indicate that there is a small negative effect of the wealth tax on active savings for those who are not married, whilst the effect is positive for those are married. But none of these results are significant using the optimal bandwidth. The differences-in-differences estimates are all significant and show that not having to pay the tax is associated with a 3-3.5 percent decrease in savings. In section 7 I discuss possible explanations for these findings, and conclude that the findings in this thesis do not support the claim that the wealth tax depresses savings.

2 THE NORWEGIAN WEALTH TAX

"*Net wealth tax increases re-distribution through the tax system, and a moderate, uniform net wealth tax carries modest economic costs. Net wealth tax primarily weakens incentives to save, and is expected to have little effect on the scale investment in Norway*." (NOU 2014: 13, p.20).

The above quote is taken from the report of the government appointed Scheel-commission from 2014: Capital Taxation in an International Economy. It reveals both the motivation for, and the concerns about, the Norwegian wealth tax. While it generates some revenue for the government, the main motivation is to insure progressivity and re-distribution in the Norwegian tax system. But the concern is that the current system of not treating assets uniformly, by exempting some assets and valuating assets differently, might make the social cost of having a wealth tax greater than its benefits.

Norway has had a wealth tax for many decades. The main arguments in favor of having the tax have been to generate revenue and to achieve an equitable distribution of income. There seems to be a consensus between academia and politicians that the labor income tax is the most efficient tool for re-distribution. This is also the case in Norway, were they have a progressive labor income tax designed in such a way that those with high income pay a larger share of their income in tax than do those with low income. But the problem with re-distribution via the labor income tax is that those with the highest income usually have a much larger capital income than labor income. Often they own a large share of the country's wealth, making them more exposed to the capital tax design then the labor income design. Since capital income tax is flat and set at a lower rate than the labor income tax, a tax system without the wealth tax would not be progressive for the wealthiest people in Norway.

During the last decade the Norwegian government has changed the allowances, rules of valuations and tax rates several times, seeking to increase both the horizontal and vertical equity of the tax system. Horizontal equity means that individuals who are similar in relevant characteristics should be treated equally by the tax system. Vertical equity refers to the fact that people with a high ability to pay taxes should pay a higher share of their income in tax than those with a low ability (Stiglitz, 2000, p.468).

Table 1 shows how the wealth tax has changed between 2006 and 2015. First of all, the system differentiates between couples and individuals. From 2009 and onward there is only one bracket, and married couples jointly receive twice the allowance that they get as individuals. For example, in 2014, an individual pays a 1% wealth tax on the net taxable wealth above 1 000 000. But if the individual is married, the couple only pays tax on net taxable wealth above 2 000 000. The tax is partly paid to the municipal level and partly to the state level. The tax rates reported in figure 1 are the total tax rates for both levels. Before 2009 there was a two bracket system with an ascending tax rate. For example,

Table 1: Changes in the Norwegian wealth tax.

YEAR	ALLOWANCE INDIVIDUALS	ALLOWANCE MARRIED COUPLES	TAXRATE
2006	200 000/540 000	400 000/1 080 000	0.9%/1.1%
2007	220 000/540 000	440 000/1 080 000	0.9%/1.1%
2008	350 000/540 000	700 000/1 080 000	0.2%/1.1%
2009	470 000	940 000	1.1%
2010	700 000	1 400 000	1.1%
2011	700 000	1 400 000	1.1%
2012	750 000	1 500 000	1.1%
2013	870 000	1 740 000	1.1%
2014	1 000 000	2 000 000	1 %
2015	1 200 000	2 400 000	0.85%

Source: The Norwegian Tax Administration (www.skatteetaten.no)

in 2006 you had to pay 0.9% in total tax rate if you were an individual with a net taxable wealth above 200 000 and below 540 000. Wealth above 540 000 was taxed at 1.1%. The allowance was doubled for married couples. The trend in these changes has been towards higher allowances and a decrease in the rates the last two years. Increasing the allowance has made the wealth tax better targeted at the wealthiest.

The revenues generated by the wealth tax are moderate but significant. About 14% of the Norwegian people are estimated to pay the tax in 2014, amounting to about 1.8 % of total tax revenue (Norwegian Ministry of Finance, 2014). More importantly are distributional effects of the tax. Figure 1 depicts how changes in the wealth tax have affected the share of people in wealth tax position and the average wealth tax paid. We see that the share of people has declined since 2001, and the average payment has gone up about 4 times from the 1997-level. This means that the changes have made the tax system more progressive.

Figure 1: Share of people paying wealth tax and average wealth tax.

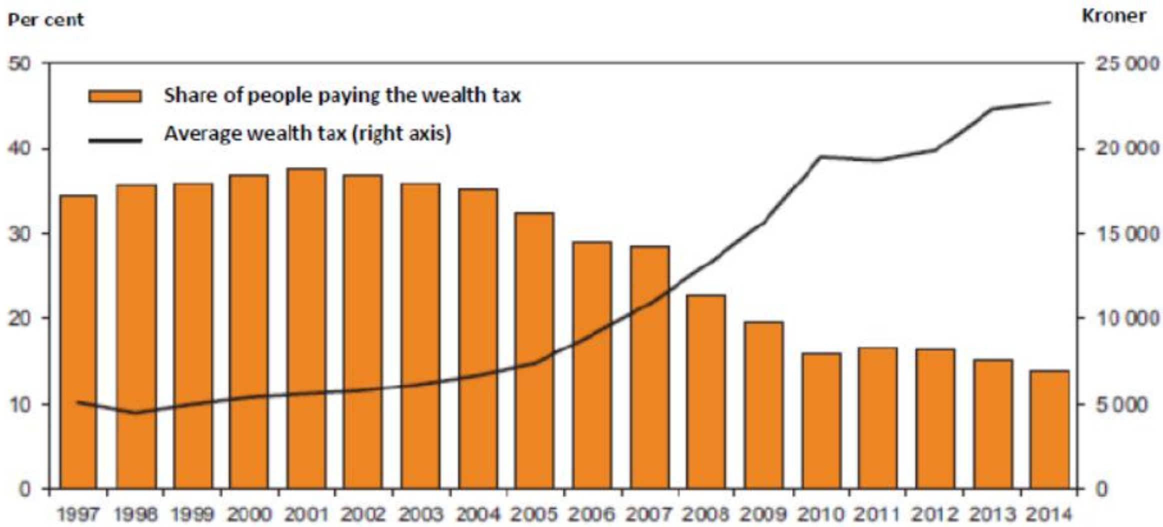

Source: The Norwegian Ministry of Finance (2015), p. 52.

Figure 2: Tax as a share of gross income, 2012.

A. 1% wealthiest people by decile

B. All people by income. In thousands (NOK)

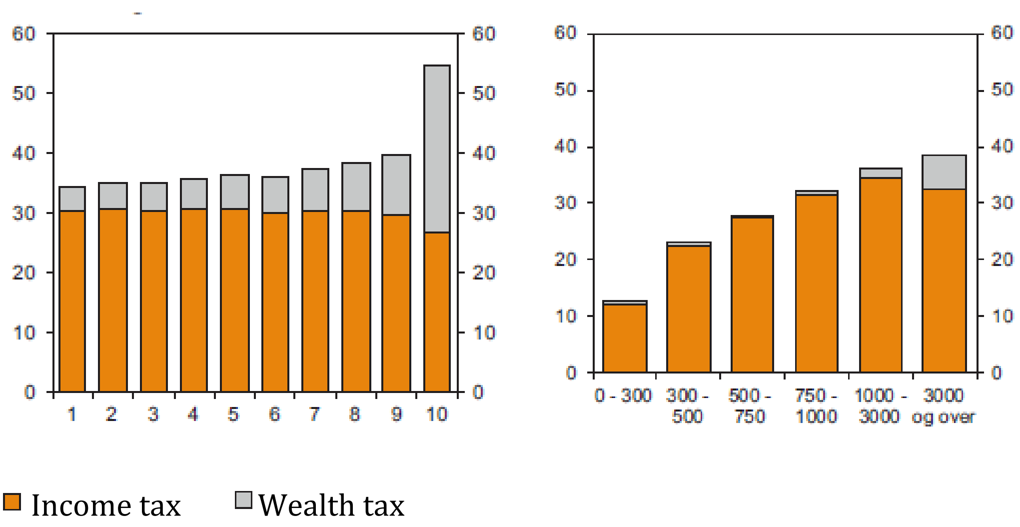

Source: The Norwegian Ministry of Finance (2015), p. 51.

Figure 2 gives further evidence that the Norwegian tax system actually is progressive, but that it would not have been without the wealth tax. Figure A divides the wealthiest one percent into 10 deciles and shows the share of gross income paid in tax for each group. Figure B shows the same but divides the whole population into six intervals. The two figures show that without the wealth tax the Norwegian tax system would in fact be regressive.

In the wealth tax scheme valuation varies across asset classes. Tax advantages are given to housing wealth, especially owner-occupied housing, and to pension wealth. Previously, favorable tax treatment was also given to equity in order to stimulate investment in so-called "productive" assets, but these allowances were removed in 2008 as part of the attempt to improve the redistributive effects of the wealth tax.

The favorable treatment of housing and pension is an example of how one may use tax instruments to obtain welfare gains in other areas, but at the expense of inefficiencies in the overall tax system. At the outset, the wealth tax contributes to very high effective tax rates on asset income, and with the favorable tax treatment of housing the effective tax rates vary a great deal between asset classes.

Table 2: Effective tax rates on the real income from different assets.

	Without wealth tax	With wealth tax
Interest-bearing accounts	56%	113%
Shares	56%	113%
Owner-occupied housing	0%	14%
Rental housing	56%	79%

Note: The calculations are done for the tax rates in 2012, a nominal rate of return of 4%, and an inflation rate of 2%, which correspond closely to the nominal rate of return to government bonds and consumer price inflation in Norway since 2000. The effective tax rates (ETRs) apply to an extra NOK of saving by a Norwegian resident investing in a Norwegian asset. The ETRs for shares are based on nominal depreciation rates which are a reasonable approximation to how the Norwegian tax system functions as tax depreciation depends on the cost price (not the repurchase price) and the expected life span of the asset. The ETRs for owner-occupied housing and rental housing are independent of the degree of debt finance versus self-finance when assuming that mortgage interest rates equal saving interest rates. Source: OECD 2012.

Table 2 presents effective tax rates on real income from different assets, based on calculations of the 2012-system made by the Ministry of Finance. It illustrates both the difference between housing and other assets in the income tax (“without wealth tax”), and the contribution of the wealth tax to the effective tax rates. Not only were the effective tax rates on financial assets very high in 2012 (they are somewhat lower now), but they also create a major distortion in the savings incentives away from financial savings and into unproductive real wealth. The high effective rates also make it more “profitable” to avoid or evade the taxes. In the end it might have severe negative consequences for the economic growth in Norway.

The aim of this thesis is to investigate the effect of the wealth tax on savings in assets that have a high ETR (mostly financial assets). But when interpreting results from such study, we should keep in mind that people might not be fully rational when deciding which assets to invest in. People might have very different motives for placing their money in shares, bank accounts or some other financial assets. The key implication being that, if people don’t perceive the ETR as being as high as it is, this might make them less responsive to the wealth tax, and thus the efficiency cost of taxation might be lower as result. It might even be the case that different people react differently.

3 THEORETICAL BACKGROUND

One of the primary tasks for any government is to provide a minimum level of public services such as police, defense and roads. Many governments provide a much more comprehensive menu of services than the ones mentioned above. Modern welfare states supply health care, education, social benefits and much more. These services are mostly paid for through taxes levied on individuals and firms. The problem with these taxes is that they influence people choices, altering their behavior, and ultimately affecting the economy through reducing how much people work and/or save.

3.1 THE TAX SYSTEM IN GENERAL

When discussing the tax system, economist usually start out by showing that, under certain conditions, the free market leads to the most efficient allocation of societies resources. This is captured in the First Fundamental Theorem of Welfare Economics (FWT), and states that in the absence of market power, asymmetric information and externalities, the private market allocates societies resources in a Pareto efficient way (Rosen & Gayer, 2008, p. 40). A Pareto efficient allocation is such that it is not possible to make one agent better off without making someone else worse off. While the assumptions underlying the FWT rarely are met in the real world, the theorem serves as a helpful (theoretical) tool when thinking about the cost of taxation. It provides a benchmark to which we can compare the incidence, effectiveness and optimality of different tax schedules.

The first welfare theorem is consistent with a great number of different allocations. But it provides no guidance for the distribution of wealth among individuals. The Second Fundamental Theorem of Welfare Economics (SWT) states that any distribution among individuals using lump-sum transfers is possible and will lead to a Pareto efficient allocation of resources (Rosen & Gayer, p. 45). A lump-sum tax is a tax based on some exogenous feature of individual (e.g. height or ability), making it impossible for the agent to change the amount of taxes paid by changing behavior. Although the tax has an income effect, it makes the agent poorer, it does not have a substitution effect, since there is no incentive for the agents to alter their behavior to avoid tax liability.

If we assume that the government has some revenue requirement, the FWT and SWT states that this revenue can be collected in a Pareto-efficient way with the use of lump-sum taxes. But in practice, lump-sum taxation is seldom used. The main problem is that governments do not have full information about individuals. Without full information, the exogenous features of individuals cannot be perfectly identified. If a government were to levy a high tax on those with high ability and low tax on those with low ability (in order to redistribute from high to low income people), with imperfect information, high ability people would have an incentive to pretend to be low ability in order to pay less in tax. This would in turn create an efficiency loss since high ability people would work less than they would with a lump-sum tax, and the result is a suboptimal allocation of resources. Since lump-sum taxes are unfeasible in practice, we are left with the option of using taxes that creates efficiency losses. The literature surrounding taxation is often referred to as the theory of the second best. Thus, the fundamental question for optimal tax theory concerns how the government should set its distortive taxes in order to minimize distortions and achieve its desired level of re-distribution?

The answer to the above question is inherently of a normative character, and has been answered in numerous ways via the optimal tax theory. The building blocks of this theory are the positive analysis of tax incidence and economic efficiency. The first is concerned with who bears the burden of a tax. The main theoretical insight from this analysis is that in a competitive market, it does not matter if you tax the consumer or supplier. The burden of a tax is determined by the elasticities of demand and supply (Stiglitz, 2000, p.518). If supply is fully elastic or demand is fully inelastic, the burden of the tax falls on consumers. If demand is fully elastic or supply is fully inelastic, the burden of the tax falls on producers. These results hold in the case of competitive partial equilibrium analysis. In a general equilibrium analysis, or if in the case of market failures, the results may differ substantially. The other building block of optimal tax theory concerns the question of economic efficiency. Stiglitz (2000) defines the efficiency loss from distortionary taxes in the following way: “for the same effect on individual welfare, how much extra revenue would a lump-sum tax have raised” (p. 522). This is often referred to as the excess burden or deadweight loss of taxation. While the theory of tax incidence is concerned with describing how the “pie” get split across people in the economy, the theory of economic efficiency concerns the size of the “pie”. Two results stand out as especially important from theory of economic efficiency. It is optimal to tax inelastic goods more, and the deadweight loss increases with the square of the tax. These results are depicted in figure 3. Point A is competitive equilibrium without any taxes in the market for some good. Here the marginal cost of producing another unit equals the marginal willingness to pay for that extra unit. When a tax t_1 is levied on the producers, the supply curve shifts inwards from S_1 to S_2. The new equilibrium point is B. At this point the tax creates a wedge between consumers’ willingness to pay for the good and the producers’ marginal cost of producing it, resulting in a deadweight loss equal to the triangle ABD. From this situation it is clear that the more elastic is demand (flatter demand curve), the larger is the deadweight loss. This has the policy implication that government could tax inelastic goods relatively more than goods that are elastic in demand, to minimize deadweight loss. If government raised the tax incrementally to t_2, this will shift the supply curve further inwards to S_3, generating a new equilibrium at point C.

Figure 3: The effect of a tax on producers.

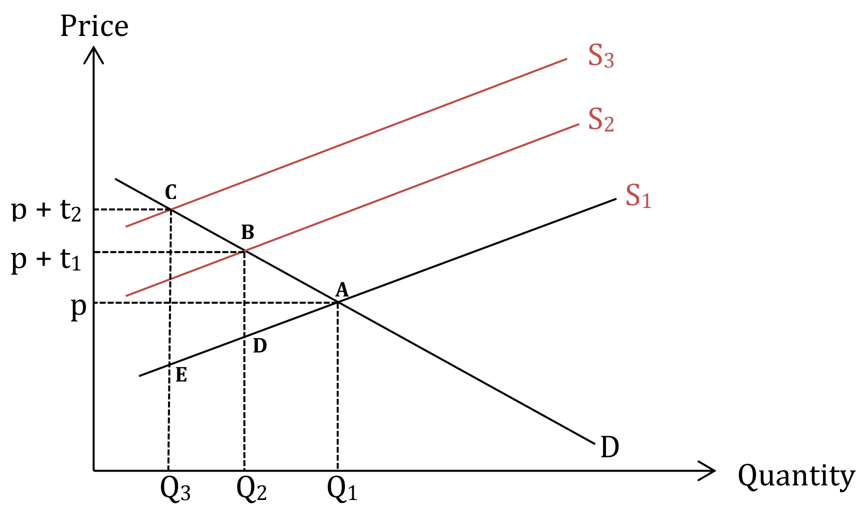

The extra deadweight loss due to this second increase in the tax rate is equal to the trapezoid BCDE. From figure 3 it is easy to see that the area of the trapezoid is larger than the area of the triangle ABD. This illustrates the earlier statement that the deadweight loss increases with the square of the tax. The policy implication being that a smaller tax rate across a broad base creates a lower efficiency loss than a large tax rate on a narrow base, for a given revenue requirement for the government.

The two main results presented in the last section are somewhat pulling in opposite directions. One states that governments should seek out goods that are inelastic in demand and tax them at a high rate, while the other implies a low rate across all goods. In addition, taxing inelastic good more can hardly be seen as optimal in a distributional perspective, since this would affect low income people the most. Balancing these considerations is one of the mains tasks of optimal tax theory.

The standard approach to optimal tax theory is that some benevolent social planner maximizes a social welfare function subject to some revenue requirement (and possibly other constraints, e.g. constraints on information). Thus, the choice of the welfare function has important implications for which policy recommendations are made on the basis of the model in question. Usually the utilitarian social function is used, which maximizes the sum of individual utilities (but other functions such as the Rawlsian are also used). The choice of function will be decisive for how the social planner values re-distribution among individuals. At the heart of optimal tax theory lies the normative question of how to achieve distributional goals while keeping efficiency costs at a minimum. Many important and robust results have been generated by this literature the last half of a century. Since this thesis is concerned with the wealth tax, the following section gives an overview of the literature on the taxation of wealth.

3.2 TAXATION OF WEALTH

There are many different tax schedules that could be implemented to meet government requirements for revenue and desire for re-distribution. We could tax corporate income, personal income, goods and services, property, et cetera. In general all tax bases can be classified as either stemming from consumption, capital or labor, or as some combination of these. When we speak of the income tax, the tax base is both capital and labor. While labor income is quite straight forward to define, capital income is more complex. The reason for this is that capital income is only one type of capital we could tax. Capital can be taxed as an income, a stock of assets or as transferred stocks. Income flows are the annual return to a capital stock. It may consist of assets such as bank deposits, stocks, mutual funds, property, et cetera. For example having money in a bank account generates an annual interest rate, which is a taxable income flow, usually taxed as capital income. The part of the after-tax return which is not consumed is then added to the individual's capital stock. This tax base can then be subject to a wealth tax, property tax or capital transaction tax. The third tax base is on transferred stocks of assets. Most often inheritance and gifts are subject to a capital tax. In Norway, wealth is taxed when it is earned and per se. Up until 2014 it was also taxed when transferred with an inheritance tax. This is a system we share with countries such as France and Switzerland. In addition, most of these countries also have property taxation, which is levied at the municipal level.

While the framework for optimal tax theory with respect to labor income taxation has been well established, the framework for analyzing taxation of wealth is less robust. This is partly due to complex nature of the tax base. Firstly, it consists of several tax bases. We could tax wealth when it is earned, when it is transferred and per se. In most models this is simplified by only looking at the effect of a tax on capital income. Another complicating factor is that it is difficult to distinguish assets owned personally from assets owned by the person's own corporation. Finally, different assets might have different rates of return. In the theory it is often assumed that all assets have the same return, making the implications of the models less useful for real world policies.

The classical optimal tax literature has often analyzed capital taxes together with a non-linear tax on wages. Most conclude that a capital tax of zero is optimal. The model of Atkinson and Stiglitz (1976) states that: when preferences are separable between labor supply and goods, governments should abstain from taxing capital if nonlinear income taxation is an option. Another theoretical argument for zero tax on capital was that the distortions the tax introduced would accumulate over time and eventually grow so large that it would be unacceptable in the eyes of the social planner (Chamley, 1986; Judd, 1985). Lately, there has been renewed interest in the theoretical foundations for wealth taxation, see Diamond and Saez (2011) and Piketty and Saez (2013). Their argument goes as follows. Taxes should be levied on the basis of the ability to pay. Income alone is not a sufficient gauge of well-being or of taxable capacity. The possession of wealth adds to this capacity, over and above the income it yields. Furthermore, since trends indicate an increased concentration of wealth in the upper part of the distribution, while at the same time capital income tax rates are declining, then the wealth tax may be used progressively instead.

The main argument against a tax on wealth (earned, transferred or in itself) is based on efficiency reasons, thus that it will inhibit savings and reduce capital investment. One direct implication of wealth taxation is that the after-tax return to capital will be lowered. Facing lower rates, potential investors may shift resources from savings to consumption since the price of consumption today has become relatively lower.

The simplest way to show the distortionary effects on savings is by considering a two-period model were a single agent chooses consumption between the periods. If the agent can freely borrow and lend in a fully competitive market, the no-tax optimal allocation is where the agent's indifference curve is tangent with the intertemporal budget constraint. The point of tangency is the Pareto-efficient allocation of consumption between the two periods. If we introduce a tax on savings, we essentially make consumption this period relatively cheaper than consumption tomorrow. The effect can be decomposed into an income and substitution effect. The income effect makes the agent consume less both periods, while the substitution effect makes consumption this period relatively cheaper to next period, making the agent consume more this period. The net effect is that consumption next period is reduced, but the effect on current consumption is uncertain. It depends on the magnitude of the effects. Since savings is the part of income that is not consumed, the effect on savings is uncertain. Thus, savings is reduced only if the substitution effect dominates the income effect.

Yet, there has been little evidence that wealth taxes substantially lowers savings. The change in household savings with respect to a shift in the after-tax rate of return, and the current or future consumer prices, is theoretically undetermined. Early empirical studies of the interest rate elasticity of savings conclude likewise (Boskin, 1978; Bernheim, 2002). Later empirical work have mainly pointed out that econometric issues concerning endogeneity and multi-collinearity make it almost impossible to correctly estimate the intertemporal elasticity of substitution (see for instance Carroll, 2001). The empirical question of the distortionary effect of the after tax-return on savings remains unresolved.

Another important aspect of the effect of a capital tax on savings can be illustrated using figure 3. In this case individuals are the suppliers of savings and firms demand investments. On the horizontal axis is the interest rate, which can be interpreted as the price of money. If a tax on wealth depresses savings then the supply curve will shift inwards, causing a deadweight loss equal to that described in part 3.1. This describes the effects in a closed economy, were investment equals savings in equilibrium. This need not be the case in an open economy. Considering the case of an infinitely elastic supply of funds from abroad (perhaps a fair assumption in the case of a small, open economy such as Norway). In figure 3 this would leave the figure unchanged, except for a horizontal line depicting the supply from abroad. An inwards shift in the domestic supply curve due to a capital tax would only change the amount of funds the firms get from domestic savers, not the total amount. The amount firms get is equal to where the demand curve intercepts with the horizontal foreign supply curve. The fraction of this which is domestic is given by the intercept of demand and domestic supply. Thus, in an open economy a reduction in domestic savings need not be equivalent with a reduction in investments.

The focus of this paper is on the effects of a wealth tax. As mentioned above, this is a tax on a stock of assets. But the implementation may vary a lot between countries. Firstly, a wealth tax may levied on an individual, firm or other legal entities. In most countries the wealth tax is levied on individuals, while property taxes are levied on individuals and firms. Secondly, the assets that make up the tax base for a wealth tax may vary between countries. We can categorize assets into tangible and intangible asset. The former consist of physical entities such as cars, boats and land. Intangible assets consist of non-physical entities such as pension rights, equities and human capital. In most countries that have a wealth tax, assets included in the tax base are ownership of land, property, investments, firms and cash equivalent banking assets (Mintz, 1991). Thirdly, rules of valuations or reliefs are critical for the effects of a wealth tax. Some assets might be valued at full market value, while for others it might be unfeasible for the authorities to calculate a market value, and thus they have to use other methods for setting the value. This might influence the saving decision of the individuals.

If there is asymmetric treatment of assets types in the wealth tax base and this causes a lower rate of return to financial assets than to real assets, then a wealth tax may induce families to invest in unproductive non-financial assets, for example residential housing. The end result being that the tax causes resources to be allocated in a way that is not optimal for the society as a whole.

3.3 OTHER CONSIDERATIONS

For lawmakers, designing a good tax system also involves striving for simplicity, such that administration and compliance cost are kept at a minimum. An argument against the taxation of wealth tax is that it generates tax avoidance, and this violates the principle of horizontal equity. When some people can avoid paying tax by moving their assets abroad or by reclassifying their assets into tax favored items, then the wealth tax' re-distributional property is weakened and the tax causes inefficient evasion behavior. A French study (Pichet, 2007) has found that tax flight of the wealthy in France is substantial. In an analysis of the Swedish wealth tax, Seim (2013) found evidence of an increase in net wealth as a response to a lower wealth tax rate, but that all could be explained by adjustments in the self-reported part of taxable wealth, while there was no change in third party reported wealth. He concludes that the wealth tax seems to have little effect on actual saving, but a significant effect on the degree of tax avoidance.

4 DATA

4.1 SOURCES OF DATA

The data used is provided by Statistics Norway (SSB). It consists mostly of register data from administrative records of income and wealth, used by the tax authorities. It is an unbalanced panel of all Norwegian tax payers from 2008 till 2011.

Norway has had an annual wealth tax for many years. Every year banks, employers, insurance companies, brokers and other intermediaries, report people's wealth and income to the tax authorities. They also report to the individual. Then the individual receives a pre-filled tax form, which is either accepted or revised. This system makes tax evasion difficult, and provides reliable data for the research design. Information on the value of stocks comes from The Norwegian Central Securities Depository (VPS). The data provides exact observations of individual asset composition for all stocks traded at the Oslo Stock Exchange, and their value at the end of each year. More problematic is the valuation of non-listed stocks. Since they are not traded in market, valuations are based on accounting principles, and they are most likely valued too low by the tax authorities.

The wealth tax gives a joint allowance for married couples. Consequently, it is interesting to look at changes in saving at the household level. From the administrative tax records each individual has a personal identification number. In the dataset this information has been linked to the National Registry (Folkeregisteret), which enables us to identify individuals who belong to the same household. We also get information on the level and type of education in each household.

4.2 VARIABLES

The aim of this thesis is to investigate empirically the effect of the wealth tax on savings. A critical point in the analysis is how one should define savings, both theoretically and empirically. The following section explains the choices made in this thesis, and is largely based on definitions from a working paper by Fagereng and Halvorsen (2015).

There are two theoretical definitions of savings that should yield the same result. The first defines saving as the part of income that is not consumed. The second defines saving as the first difference in wealth between two subsequent years. In standard lifecycle theory, agents maximize their lifetime utility each period given the following constraint

(1) $$a_t = (1+r)a_{t-1} + yl_t - c_t$$

where a_t is the net wealth in period t, yl_t is the after-tax labor income, c is consumption and r is the rate of return on savings. The budget constraint can be re-arranged to give the definition of saving (s) in terms of changes in wealth between two years

(2) $$s_t = a_t - a_{t-1} = ra_{t-1} + yl_t - c_t$$

Both definitions have their advantages and disadvantages. Given the dataset available for this thesis, and the research question posed, I will use the second definition, i.e. saving as the first difference in wealth. The challenge of using wealth to calculate saving is to uncover the correct valuations of different assets, since not all assets are valued at market value in the administrative records. The following sections will explain how I have operationalized the term savings, and the method used to calculate their true value.

When we observe changes in the portfolio of individuals each year, we can distinguish between passive and active savings. For example, having a certain amount of money placed in a bank account will yield some rate of return that, unless it is consumed, adds to the wealth of the individual. Given the above definition of saving, this will be defined as an increase in saving. In this paper I define this as passive savings. It includes revaluations of assets and their normal rate of return (explained in more detail below). In contrast, if the individual in addition to the passive return actively places some of his income to the account, this is a case of active saving. Given the research question in this thesis, the interest is in behavioral responses with respect to active saving.

There are many types of assets that people can choose to save in. Broadly speaking, these can be divided into financial, real and foreign assets. Financial assets consist of bank deposits, stocks (listed and non-listed), mutual funds, debt, and bonds. In this analysis I have left out Individual pension savings (IPS) and cash value of life insurance due to their small size and insignificance, and because of inconsistent treatment in the tax scheme. Real assets consist of houses (owner occupied, secondary housing and holiday houses), motorized vehicles and other physical assets (forest, property and farm). In the analysis the owners of production capital and transactions of holiday houses are not included. Houses are also left out. Since houses are given preferential treatment in the Norwegian tax system, there is no easy way of calculating their market value. The value of owner occupied housing has been given a substantial relief in the Norwegian tax system for the whole period this analysis covers. Until 2010 the valuation method used was very loosely linked to actual market value. In 2010 these rules changed, making them better linked to the market value, but only at around 25%. Finally, the foreign assets consist of bank deposits and ownership of property in foreign countries.

The first step in preparing the savings data is to scale up the values of the assets that have been subject to reliefs. Assets such as stocks (listed and unlisted), mutual funds and bonds were given different reliefs between 2004 and 2007, ranging from 65%-85%. From 2008 they were valued at a 100%. As mentioned before, the market value of listed stocks is well defined, but this is not the case for unlisted stocks (relief was given on the basis of their accounting value), and they are probably measured with error. The second step is to deflate the data by the Consumer Price Index (CPI). In order to single out the part of saving that is active, I take out the passive part of saving, defined as the normal rate of return on that type of asset. More specifically, this has been calculated for stocks by using the historical return of Oslo Stock Exchange. The normal return to mutual funds is calculated using historical return on a weighted average of Oslo Stock Exchange (30%) and MSCI World Index (70%). Lastly, the Treasury Bill is used to calculate normal returns to bonds. Formally, this can be written in the following way

$$\Delta a_{it} = a_{it} - R_{it} a_{it-1} \tag{3}$$

The change in the active saving in an asset type is the difference in wealth between two subsequent years (t and t-1) for person i, and then subtracting the normal rate of return. All other financial assets (debt, outstanding claims and other financial assets are represented by deflated differences only, and can be written in the following way

(4) $$\Delta a_{it} = a_{it} - a_{it-1}$$

The value of motorized vehicles is set to market value in the year of purchase. In the subsequent years they are valued at standardized rates, usually at 75% the year after the purchase, and then declining by 10 percentage points each year.

Summing up, active savings in a given year i, can be written as

(5) $$Active\, savings_{it} = \Delta a_{it}^{F} + \Delta a_{it}^{R} + \Delta a_{it}^{FA}$$

Active savings for person i in period t is the sum of changes in financial assets (F), real assets (R) and foreign assets (FA), using the definitions in (3) and (4). For a married couple saving is the change in their joint asset holdings. This is the dependent variable in the analyses.

The independent variable is easier to operationalize, since this entails just using the valuations in the tax records directly. The analysis aims at investigating whether being in a wealth tax position or not alters how much people save. I use the data to generate a variable for net wealth. Since the reliefs are already in the data, this amounts to adding up all assets that make up the tax base, and subtracting all liabilities.

The key variables in this analysis are net taxable wealth and savings. The distributions of net wealth and savings are both skewed (net wealth is right skewed and savings is left skewed), and both have very long tails with large extreme values. Log transformation of a variable is often used to reduce the influence of large outliers and to obtain a more normally distributed variable. When working with a variable such as income, which never take on negative values or the value zero, the log transformation has proved very useful. The transformation also enables the researcher to interpret the estimated effects as elasticity's (changes in terms of percentages). When variables can take on both negative and positive values, including zero, the logarithmic transformation is not defined. A solution to this problem is to transform the data using the inverse hyperbolic sine function (IHS) (Pence, 2006).

(6) $$g(x) = \ln\left(x + \sqrt{x^2 + 1}\right)$$

The advantages of using the IHS transformation in (6) is that it structures the data in a preferable way, bringing extreme values more to the middle. In this thesis both the dependent and the dependent variable are transformed before analysis. A more detailed description of IHS can be found in an article by Burbidge, Magee and Robb (1988).

4.3 SUMMARY STATISTICS

The key variables of this analysis are savings and net taxable wealth. Before starting the analysis it seems prudent to investigate how these key variables change throughout the lifecycle of an individual. Looking at all individuals, the trend would be that the stock of wealth is negative early in life (e.g. borrowing money to buy a house), and positive after the age of fifty (when the mortgage has been paid back). Since the focus of this thesis are the behavioral responses of those with a positive net wealth, figure 4 gives a graphical representation of the relationship between age and net taxable wealth for non-married with a wealth between 0 and 1 400 000 NOK. The figure shows that for this group wealth increases through life, peaking just before the age of 80, and decreasing after that.

Figure 4: Relationship between age and net taxable wealth for non-married, 2011.

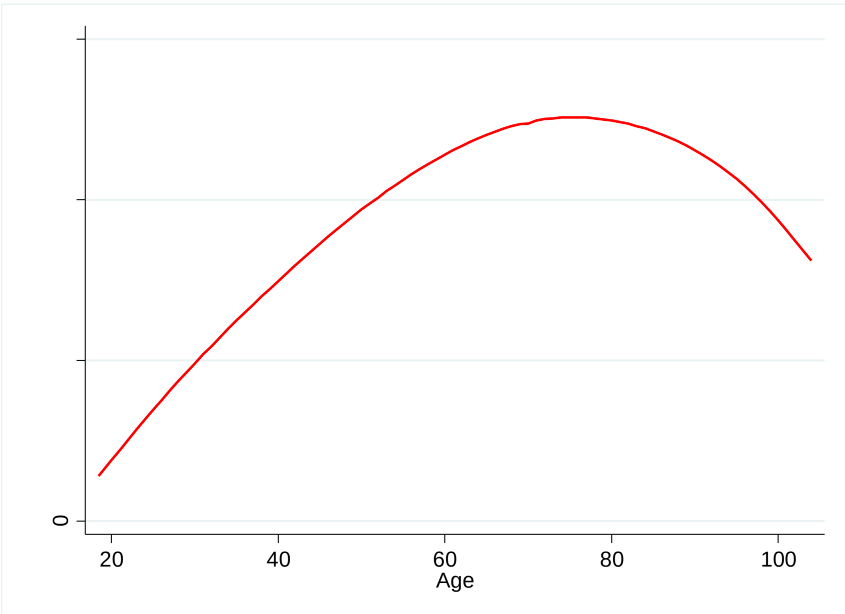

Note: The line has been smoothed using a local regression (LOWESS).

While the relationship between age and wealth is quite intuitive, it seems less clear how the relationship between active savings and age looks like. Using the plain definition of savings (passive plus active), figure 4 would suggest that savings increase with age since a higher stock of assets will bring more passive returns. Figure 5 depicts the relationship between active savings and age. As in figure 4, the sample has been restricted to non-married with a net taxable wealth between 0 and 1 400 000 NOK. Figure 5 shows that the active savings increase up until the age of 35, then decreases slowly through life.

Figure 5: Relationship between age and active savings for non-married, 2011.

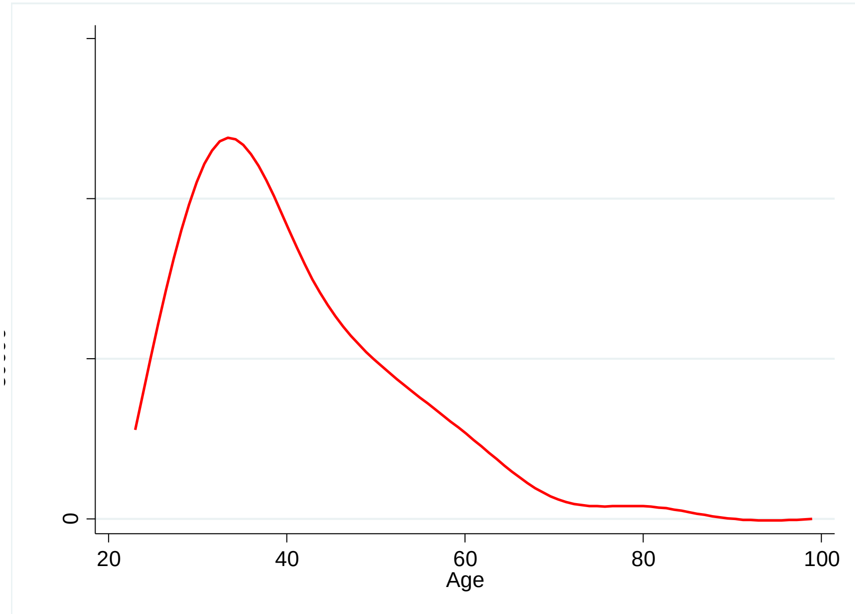

Note: The line has been smoothed using a median spline.

In the analysis I use data in the neighborhood of the thresholds for the net taxable wealth tax in 2011. They were respectively 700,000 NOK for unmarried persons, and 1,400,000 NOK for married couples. In order to get an overview I chose at the outset to sample a rather wide interval around these values, i.e. net taxable wealth from 0 to 1,400,000 for unmarried persons, and net taxable wealth from 0 to 2,800,000 NOK for married couples. Table 2 and 3 show means and standard deviations for selected variables in these two samples. Table 2 depicts how non-married people are distributed with respect to these variables. On average people with a wealth above the cutoff are older, have higher education, higher income and save more than those below the cutoff.

Table 3: Summary statistics for non-married with NTW between 0 - 1 400 000 NOK, 2011.

	All		NTW<700 000		NTW>700 000	
	Mean	**SD**	**Mean**	**SD**	**Mean**	**SD**
Age	46.29	22.96	43.01	22.33	63.76	17.75
Children	0.39	0.78	0.43	0.80	0.22	0.62
Male (%)	0.43	0.50	0.44	0.50	0.40	0.49
Higher Education	0.18	0.39	0.17	0.37	0.25	0.43
Labor Income	257 230	191 142	241 370	181 728	341 660	216 142
NTW	314 480	351 385	188 340	196 615	986 454	197 541
Wealth Tax Paid	488	1438	0	0	3077	2212
Active Savings	41 112	371 909	34 209	317 169	74 199	562 959
Observations	983 199		827 794		155 370	

Notes: This table presents means and standard deviations for different variables of different subsamples. NTW stands for Net Taxable Wealth. The first column named "All" consists of all non-married people in Norway in the year 2011 with a NTW between 0 and 1 400 000 NOK. The second and third column consists of those with a NTW below and above the tax cutoff. The variable children shows the average number of children below 18 years of age. Higher education are the average number of people with a lower college degree or more. Labor income consist of gross income from labor and pensions.

Table 4 is equivalent to table 3, except that now the unit of interest is the family (defined as those who are married and their families). The trend is much the same as we see with individuals. Couples with more than 1 400 000 NOK in wealth are on average older, better educated, higher income and save more than those with wealth below the cutoff. Comparing the two groups married and non-married, we see that the difference between those who are above and below the cutoff are very similar when it comes to education and income, but very different when it comes to age and active savings. For the non-married, those below the cutoff save on average 39 990 NOK less than those above. For married, this difference is 61 955 NOK. The difference in age for the non-married are 20,75 years, while it is 6.62 years for the married. Some of these differences reflect the fact that married are two people instead of one, but it also reflects the relative numbers across the cutoff. The non-married above the cutoff make up about 16 percent of the sample, while in the married sample they make up about 27 percent of the sample. These factors form the basis on which a causal effect is estimated, and it is of great importance to choose an identification strategy that takes these factors into account.

Table 4: Summary statistics for the married with NTW 0-2 800 000 NOK, 2011.

	All		**NTW<1 400 000**		**NTW>1 400 000**	
	Mean	**SD**	**Mean**	**SD**	**Mean**	**SD**
Age	58.41	14.77	56.65	15.23	63.27	12.18
Children	0.43	0.89	0.48	0.93	0.30	0.75
Male (%)	0.52	0.50	0.53	0.50	0.50	0.50
Higher Education	0.26	0.44	0.24	0.43	0.33	0.47
Labor Income	664 665	399 868	634 793	384 613	744 532	427 931
NTW	942 814	747 156	561 890	420 979	1 961 232	389 610
Wealth Tax Paid	1 681	3 634	0	0	6 176	4 555
Active Savings	39 134	684 267	21 818	602 894	83 773	857 566
Observations	434 051		318 498		115 553	

Notes: This table presents the same numbers as table 3 except at the family level (those who are married and their family). The "All" column consists of people with Net Taxable Wealth (NTW) between 0 and 2 800 000 NOK, while the two other columns consists of those over and below the cutoff. Higher education is based on the male of the couple, while the other variables are calculated on the basis of aggregated sizes across families.

The general insight from the descriptive statistics above are that any research design seeking to uncover the behavioral effects of the wealth tax with respect to saving, will have to take into account factors such as age, size of household, education and income. Indeed, the two research designs presented in the next chapter have quite different strategies for dealing with this problem. While both exploit the quasi-experimental design, the differences-in-differences approach deals with omitted variables bias by using the panel structure to control for unobservable characteristics, and include relevant observable variables as controls. The other approach, regression discontinuity design, utilizes the randomness that occurs around a tax bracket cutoff, ensuring that the distribution of characteristics is similar around a small interval on both sides of the cutoff.

5 METHOD

5.1 THE QUASI-EXPERIMENTAL DESIGN

This thesis will take the quasi-experimental approach in order to uncover the effect of a wealth tax on savings. The design is best understood in the context of an ideal experiment. If we imagine a situation without a wealth tax, we could randomly select people into a treatment and a control group. Those in the treatment group are subjected to wealth tax, and those in the control group are not. Since people have been selected randomly to each group, they do not differ in characteristics. This enables us to use the outcome of the control group as a counterfactual outcome for those who are treated had they not been treated. The same applies the other way around. The outcome of the treated serves as a counterfactual for the control group had they received treatment. If we define $Y_i(0)$ as the outcome for individual i when in the control group and $Y_i(1)$ the outcome for individual i when in the treatment group, then the fundamental problem for causal inference is that only one of the outcomes can be observed for a single individual. But if the two groups do not differ in characteristics, the average causal effect can be calculated by using the average value for one group as the counterfactual for the other group.

Ideal experiments are for obvious reasons not feasible to implement in most research designs. They are often costly and/or pose serious ethical challenges. But the quasi-experimental research approach is a helpful way to uncover causal effects, using statistical insights and methods taken from ideal experiments (Stock and Watson, 2012). The idea is that some event in the past, such as policy change and program implementation, are as-if random in assigning people to a treatment and control group. In order to answer the research question posed in this thesis, changes in the implementation of the Norwegian wealth tax is used a quasi-experiment. Since many factors might change how much people save, I use two approaches to try to uncover the effect caused by the wealth tax. Firstly, regression discontinuity design uses the fact that the probability of receiving treatment changes around some threshold, and this generates randomness in the assignment of treatment. In other words, we calculate the average treatment effect for those at the threshold, and use this as an approximation for effect on the whole population (Stock and Watson, 2012, p. 536). Secondly, by applying differences-in-differences approach to the whole sample, controlling for other factors that might influence savings.

5.2 REGRESSION DISCONTINUITY DESIGN

Of all the quasi-experimental strategies for identifying treatment effects (i.e. differences-in-differences and IV), regression discontinuity (RD) design is often thought of as the closest cousin of a randomized experiment (Lee and Lemieux, 2010). The basic idea is that we can use some observable variable (often called a running or assignment variable) to determine whether or not an individual receives treatment (Lee and Lemieux, 2010). The running variable has some cutoff point, and having a value below or above this point determines whether an individual is in the treatment or control group. If the probability of receiving treatment jumps from 0 to 1 at the cutoff, the RD design is called sharp (SRD). If the cutoff only partially determines the assignment of treatment, the appropriate RD design is called fuzzy (FRD). In the context of the Norwegian wealth tax, it seems appropriate to use a SRD design, given the precise data about wealth used by the tax authorities. In other words: making the plausible assumption that whether or not someone pays a wealth tax is a deterministic function of their wealth. This can be written in the following way

$$T_i = 1\left[X_i \geq c\right] \tag{7}$$

were treatment status is denoted T_i, X_i is the assignment variable and c is the cutoff value. The key identifying assumption is that the expected value of $Y_i(0)$ and $Y_i(1)$, given a value of X_i, is continuous in x (strictly speaking it is only required at $x = c$). This enables the use of average outcomes just below the cutoff as a counterfactual for those just above the cutoff. The average causal effect is then

$$\tau_{SRD} = E\left[Y_i(1) - Y_i(0) \middle| X = c\right] \tag{8}$$

Figure 5 shows this effect in a more intuitive way. On the vertical axis is the outcome variable, and on the horizontal axis is the assignment variable. If we plot the observed data on either side of the cutoff ($X = c$), and fit a flexible regression line, figure 5a depicts the case of a non-linear relationship between Y and X on both sides of the cutoff. We observe the expected value of $Y_i(1)$ for those at the right of the cutoff and the expected value of $Y_i(0)$ for those at the left of the cutoff. The continuity assumption allows us to use the average of those who are observed to the left of the cutoff, as counterfactual for those who at the right of the cutoff. Consequently, the average treatment effect is the difference between were these two lines intercept with the cutoff value ($X = c$).

Figure 6: The treatment effect in RD.

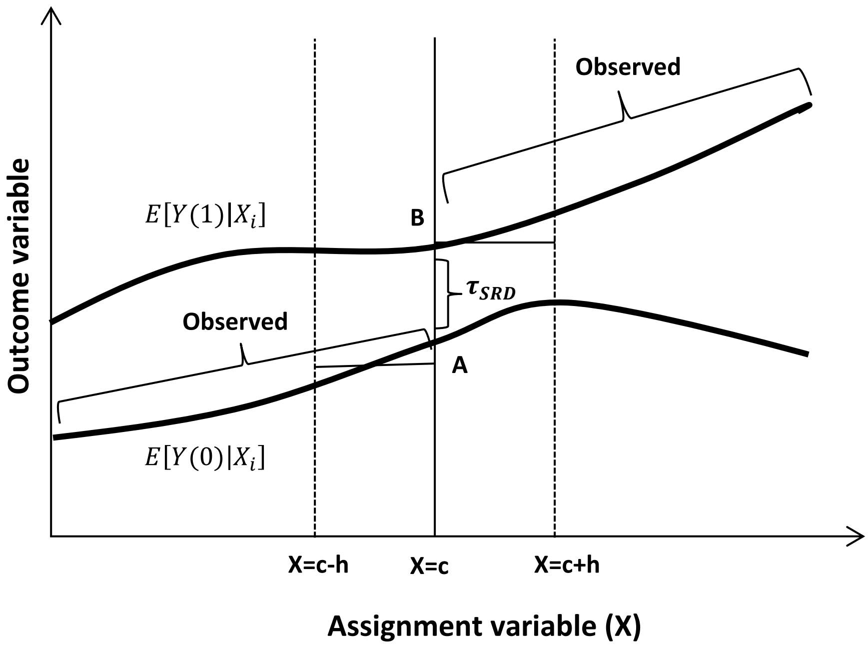

Imbens and Lemieux (2008) suggest the following procedure when applying SRD design. The first step is to graph the data to get a visual impression of the jump at the cutoff. This is done by first dividing the assignment variable into a number of bins, calculate the midpoint in each bin, then plot it against the average value of the outcome variable within each bin. To make the visual inspection easier, it is useful to plot a flexible regression line on each side of the cutoff. The binwidth should be big enough to make the plot look smooth, but small enough to make a jump visually clear. This plot is in a way the most essential part of an RD analysis. As Imbens and Lemieux puts it, the more formal part of the analysis is “essentially just sophisticated versions of this” (2008, p. 12).

The next step is to estimate the treatment effect. There are two ways of doing this, either using parametric or nonparametric regressions (Lee and Lemieux, 2010). The parametric approach uses all the values of the running variable and is therefore a global estimate of the treatment effect. This is problematic since identification in RD design lies in the jump at X=c. If the functional form assumption is not correct in parametric regression, it might lead to a bias in the treatment effect. One way to do this is to try very flexible specifications of X, and test how well different polynomials fit the data. The other approach to estimating the treatment effect is nonparametric and uses kernel regressions. Although kernel regression is fundamentally a local approach, and should be well suited for the RD design, it turns out that it performs badly in this setting due to boundary problems (Lee and Lemieux, 2010). This point is shown in figure 6. The true effect is equal to τ. If we use a rectangular kernel, as recommended by Imbens and Lemieux (2008), and the bandwidth (c-h, c+h), the kernel estimate is B-A. In other words, it systematically overstates the effect when the regression line is upwards sloping (Lee and Lemieux, 2010). Hahn, Todd and van der Klaauw (2001) has proposed to use a local linear regression to mitigate the bias. In the context of a rectangular kernel, this is equivalent to running a standard linear regression within the bin around the cutoff. From figure 6 we see that estimating a standard linear regression in the interval (c-h, c+h) will give a better estimate than the B-A estimate of the kernel regression. Due to the large register dataset available for this thesis, I will use a local linear regression to estimate the treatment effect. This entails estimating the following regression in a given bandwidth

(9) $$Y_i = \beta_0 + \beta_1 T_i + \beta_2 r_i + \beta_3 r_i T_i + e_i$$

Y_i is active savings for individual i, T_i is a binary treatment variable equal to 1 if in net taxable wealth is above the cutoff and zero otherwise, and r_i is the assignment variable (net taxable wealth). To make interpretation easier, it is standard to center the assignment variable as shown in (10), enabling a shift at the cutoff to be interpreted as a shift at the intercept.

(10) $$r_{i,\, centered} = r_i - cutoffvalue$$

Identification relies on estimates based on a small sample around the tax bracket cutoff. Consequently, any non-linearity in the specification will make the estimate sensitive to changes in the bandwidth. Thus, the choice of bandwidth is a crucial factor in regression discontinuity analysis. This thesis uses an optimal, data-driven bandwidth rule, as presented by Imbens and Kalyanaraman (2009).

The third step is to check the robustness of the results to employing different specification tests. There are two main conceptual concerns to be aware of when applying RD design (Imbens and Lemieux, 2010). The first concern is about other changes at the same cutoff value of the assignment variable. In the context of the Norwegian wealth tax, this may be the case if other variables affect how much people save around the threshold. The second concern is that people might be able to precisely manipulate the assignment variable. People might try to be on the “good” side of the cutoff, thus invalidating the RD approach.

In the analysis presented in the next chapter, the assignment variable is net taxable wealth. Looking at the year 2011, individuals with a net wealth above 700 000 NOK had to pay a wealth tax. The cutoff is therefore at 700 000 NOK. Treatment is defined has having to pay the tax, so the treated are those with a net wealth above 700 000 NOK. The control group consists of those with a net a net wealth of 700 000 NOK or less. Since married couples get a joint allowance that is double that of individuals, I have performed a separate analysis for this group. The outcome variable is active savings.

5.3 DIFFERENCES-IN-DIFFERENCES ESTIMATION

The second approach to uncovering the causal effect of a wealth tax on the level savings is the differences-in-differences (DiD) estimator. Once again, this method uses the potential outcome framework described in chapter 5.1. The general idea is that a tax reform “as-if” randomly assigns people to a treatment and control group, after controlling for factors that are determinants of dependent variable and are correlated with the independent variable. The choice of reform is the change in allowance in 2010 from 470 000 to 700 000 NOK for the non-married, and from 940 000 to 1 400 000 NOK for the married. Treatment is defined as having to pay the tax before 2010, and not having to pay it after 2010. The control group consists of those who had to pay the tax before 2010 and still have to pay it after the reform. The outcome of interest is active savings.

Again, it useful to think of the DiD approach in the context of an ideal experiment. If we randomly assign people to a treatment and control group, the randomization will enables us to use the control group as a counterfactual for the treatment group had they not been treated. The average effect of the treatment is difference in the mean outcomes of the two groups. In the context of quasi-experimental design, treatment is "as-if" randomly assigned if we control for all relevant characteristics. The panel structure of the data enables us to control for unobservable characteristics that are time-invariant and effects that do not differ across groups. Formally, we define this in the following way. If we consider two time periods, denoted with subscript t, were t =1 after treatment and t =0 before treatment, and two groups, denoted with subscript g = 1 if in treatment group and g =0 if in control group. The outcome for individual i in group g in period t is $Y_{igt}(1)$ if treated and $Y_{igt}(0)$ if not treated. In the case of no treatment we assume that there is only time-invariant (α_g) and constant group (λ_t) effects like stated in (11)

(11) $$\mathrm{E}\left[Y_{igt}(0)\right] = \alpha_g + \lambda_t$$

Define X_{gt} as a treatment indicator equal to 1 if treatment in the second period and 0 otherwise. Then (12) defines the expected outcome

(12) $$\mathrm{E}\left[Y_{igt}\right] = \beta X_{gt} + \alpha_g + \lambda_t$$

The average causal effect being

(13) $$\mathrm{E}\left[Y_{igt}(1) - Y_{igt}(0)\right] = \beta$$

The DiD estimator is found by subtracting the differences between the expected outcome for the treatment group with the differences in the expected outcome of the control group between the two periods.

(14) $$\left(E\left[Y_{i,\,treatment,\,1}\right] - \mathrm{E}\left[Y_{i,\,treatment,\,0}\right]\right) - \left(\mathrm{E}\left[Y_{i,\,\mathrm{control},\,1}\right] - \mathrm{E}\left[Y_{i,\,\mathrm{control},\,0}\right]\right) = \beta$$

From (14) we see that the time-invariant and constant group effects disappear and we are only left with β **.** This shows that the DiD approach enables us to control for characteristics that are not observable, thus making it a very powerful tool for estimating causal effects.

The most important assumption made in the DiD setup, is that in the absence of treatment, the two group would have experienced the same trend (Stock and Watson, 2012). This is shown formally in equation (11). In the context of this research design, this means that if the change in the wealth tax allowance in 2010 had not occurred, the only factors that would affected how much people save, are time invariant and entity fixed effects. This is a strong assumption, and represent a challenge to any practical implementation of DiD. A way to mitigate this problem is to include relevant and observable control variables. The analysis in section 6.2 includes the control for income, inheritance/gifts and whether the household has increased in size.

Two specifications are used to estimate the average treatment effect. The first is an OLS in levels were the average value of active savings is regressed for each individual before (2008-2009) and after (2010-2011) the reform, leading to the following equation

$$\overline{Y}_{it} = \beta_0 + \beta_1 Treatment_i + \beta_2 Time_t + \beta_3 \left(Treatment * Time\right)_{it} + \overrightarrow{X_{it}} + e_{it} \qquad (15)$$

$\overline{Y}_{it}$ is the average active savings for person i in period t, treatment and time are dummy variables, and $\overrightarrow{X_{it}}$ is the vector of control variables. The coefficient β_3 is the treatment effect.

The second specification is in first differences and can be written as in (16).

$$\Delta\overline{Y_{ig}} = \beta_0 + \beta_1 (Treatment{*}Time) + \Delta e_{ig} \qquad (16)$$

$\Delta\overline{Y_{ig}}$ is the difference in average active savings before and after the reform for person i in group g, and equivalently for the error term Δe_{ig}. This regression is run with and without the control variables mentioned above. Running 3 different regressions will be a check of whether the effect is stable across different specifications.

6 RESULTS

6.1 REGRESSION DISCONTINUITY

A nice feature of the RDD analysis is that it relies heavily on graphical analysis. This makes it easier for readers to feel confident in the identification strategy. A natural way to start the analysis is by plotting the outcome variable (active savings) against the assignment variable (net taxable wealth). Starting with individuals that are not married, I restrict the sample to only include people with a net wealth between 0 and 1 400 000 NOK and people with active savings between the absolute value of 500 000 NOK (removes noise from large outliers) in the year 2011. The sample used consists of 849 559 people.

Figure 7: Basic plot of net taxable wealth and active saving. Data on non-married, 2011.

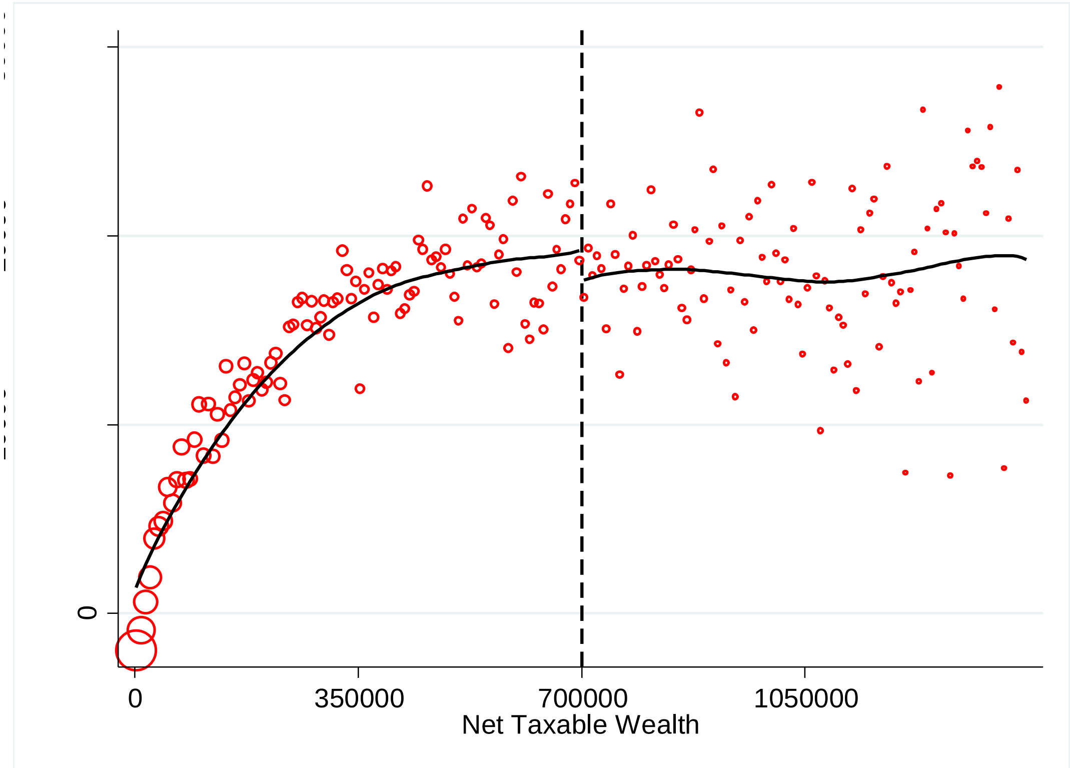

In figure 7 the sample have been divided into 7000 NOK bins by taxable net wealth, and then the median value of the assignment variable have been plotted against the mean value of the outcome variable for each bin. The size of each dot reflects the amount of observations in that bin. A flexible regression line has been fitted to the data on each side of the cutoff.

The tax bracket cutoff for the non-married is at 700 000 NOK, and indicated by a vertical line in the figure.

At first glance there are three aspects of the graph that is worth mentioning. Firstly, the amounts of observations in each bin change as we move up the horizontal axis. As discussed in section 4.3, the number of people in each bin is decreasing in wealth. But around the cutoff they are about the same. Secondly, the spread around the fitted regression line also increasing in wealth, making it harder to see a clear pattern in the regression line to the right of the cutoff. Thirdly, the functional form is clearly nonlinear.

Figure 7 depicts a negative jump at the cutoff, indicating that an estimate of the treatment effect should be negative. But the jump is not fully convincing. As discussed earlier, such a jump could be a consequence of nonlinearity around the cutoff. Another important visual check to do on this basic plot is that of jumps at other values of the assignment variable. This could indicate that other variables effect the treatment. Reassuringly, there does not seem to be any such jumps in figure 7.

Due to the large sample accessible for this analysis, the nonlinearity of the regression, and the large differences in characteristics when moving away from the cutoff, it seems sensible to use local linear regression to estimate the causal effect. Table 5 shows the estimates from a local linear regression with active savings as the outcome variable and net taxable wealth as the assignment variable. Before running the regression, the variables have been transformed using the inverse hyperbolic sine transformation (IHS). As described in section 4.2, this transformation is well suited to handle values that are both positive, negative and zero, as opposed to standard log transformation normally used when values are only positive. The estimates reported in table 5 are the average treatment effect for the treated, calculated at different bandwidths. In the first row the Local Average Treatment Effect (LATE) estimate is reported using the optimal bandwidth (OBW), as described by Imbens and Kalyanaraman (2009). The second row shows the estimated effect when using a twice as large bandwidth (indicated by the 200 in parenthesis) as the OBW, the third row shows the estimate using four times the OBW (indicated by 400), and so on. All estimates show a negative effect on savings caused by the wealth tax, as depicted in figure 8. The effect increases as the bandwidth increases. The estimates are significant at 1 percent level for the largest bandwidths.

Table 5: Estimates of average treatment effect using local linear regression. Non-married in 2011.

	Active Savings
LATE	-0.0515
	(-0.71)
LATE (200)	-0.1000
	(-1.82)
LATE (400)	-0.235***
	(-5.20)
LATE (600)	-0.452***
	(-10.71)
LATE (800)	-0.617***
	(-15.02)
N	2013709
OBW	0.945

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Since married couples receive double the allowance as the non-married, it seems paramount to investigate if effects differ between these groups. In 2011, married couples received a joint allowance of 1 400 000 NOK. Thus, this is the cutoff value which will be used in the following analysis. I restrict the sample to only include married people with a combined net wealth ranging from 700 000 to 2 100 000 NOK, and active savings between $\pm$ 1 000 000 NOK. This constitutes a sample of 426 196 couples. In figure 8 the data have been plotted in a similar way as in figure 7, except the grouping has been done over net taxable wealth in intervals of 14 000 NOK.

Similar to figure 7, it is evident that the spread around the fitted regression line increases with wealth. In contrast to figure 8, the plot in figure 9 seems to be more linear, but as wealth increases, and as the noise in the data gets larger, it exhibits a more nonlinear shape.

From the size of the plotted dots it is clear that the number of observations in each bin decreases with wealth, but not by as much as in figure 7. Similar to the effect found for the non-married, the jump in figure 8 is negative. Judging from the graph, being affected by the wealth tax decreases active savings.

Figure 8: Basic plot of net taxable wealth and active saving. Married in 2011.

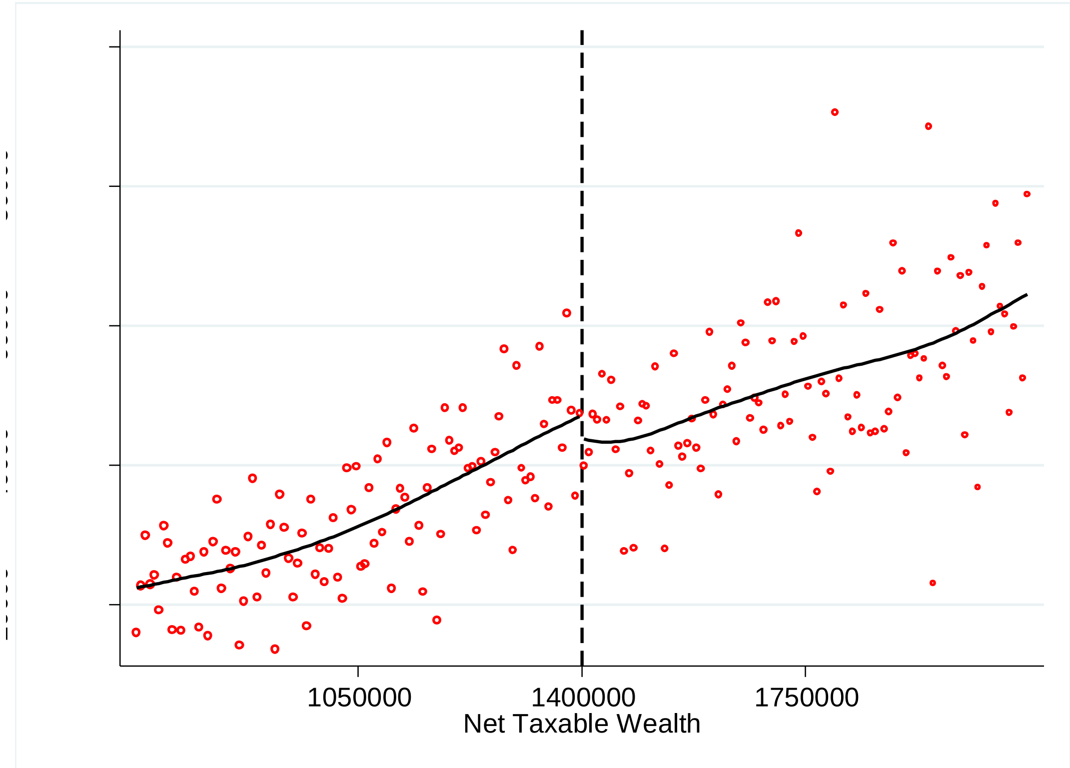

Table 6 shows the LATE estimates using different bandwidths. As in table 5, the dependent and independent variable has been transformed using the IHS. The estimates show that receiving treatment is associated with an increase in savings. The estimates increase when the bandwidth is doubled, but decreases after that. The estimate with the OBW is not statistically significant, but 200 and 400 are. The fact that the local linear regression estimates are positive, but the observed jump in figure 8 is negative, seems puzzling. One possible explanation is that the noisy data make the estimate especially sensitive to the choice of bandwidth.

To sum up, both plots indicate that having to pay the wealth tax decreases the level of active savings. But the estimates move in separate directions. The non-married reduce their active savings, while the married increase it. However, the results are not significant using the optimal bandwidths, but some are significant for larger windows around the cutoff.

Table 6: Estimates of average treatment effect using local linear regression. Married in 2011.

	(1) Active Savings
LATE	0.122
	(1.70)
LATE(200)	0.244^{***}
	(4.10)
LATE(400)	0.176^{**}
	(3.26)
LATE(600)	0.0659
	(1.24)
LATE(800)	0.0490
	(0.93)
N	811356
OBW	1.576

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

6.2 DIFFERENCES-IN-DIFFERENCES

The RD approach indicated that the wealth tax might have a negative effect on active savings for non-married, but the results were more indecisive when it comes to the married. It seems reasonable to try a different approach to see if we can shed some more light on this relationship.

The following analysis utilizes panel data on saving and wealth for the period 2008 to 2011. The reform of interest is the 2010 change in allowance, going from 470 000 to 700 000 NOK for non-married, and from 940 000 to 1 400 000 NOK for married. As described in section 5.3, estimates are provided with and without control variables using the first differences method and OLS in levels.

Table 7 shows the results from the different regressions. The sample has been restricted to people with a net taxable wealth between 0 and 1 400 000 NOK. Column (1) and (2) shows the results from the first difference approach with and without control variables. Column (3) shows result from the OLS regression using levels. The control variables included are income

Table 7: Results from DiD for non-married, 2008-2011.

Dependent Variable	(1) First Differences in Saving	(2) First Differences in Saving	(3) OLS Savinglevel
ATE[1]	-3.543***	-3.515***	-3.524***
	(0.089)	(0.089)	(0.089)
Income[1]		0.272***	0.145***
		(0.045)	(0.014)
Inheritance/gifts[1]		0.403***	0.271***
		(0.016)	(0.011)
Familysize[1]		0.295*	0.320***
		(0.115)	(0.031)
Treatment dummy			-0.834***
			(0.066)
Time dummy			-0.917***
			(0.054)
Constant	-0.922***	-0.921***	0.036
	(0.054)	(0.054)	(0.190)
N	118239	118239	236478

Robust standard errors in parentheses. Clustered at the individual level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
[1]First differences in specification (1) and (2), levels in (3).

(defined as the sum of labor and pension income), receiving gifts or inheritance, and whether there has been a change in the number of people in the household. All variables except change in household have been transformed using IHS. All individuals with missing values for one of the years have been excluded. The control group consists of those who had to pay the wealth tax before and after the reform, while the treatment group paid the tax before the reform, but not after.

The estimates of the average treatment effect (ATE) are stable across all three estimations. Not having to pay the tax after the reform is associated with a 3.5 percent decrease in savings. The result is significant at a 1 percent level, and does not change much specifications are changed and controls are included.

Table 8 shows the average treatment effect for those with a joint allowance. Being in the treatment group is associated with 3 percent decrease in active savings. The results are stable over all the estimations and significant at a 1 percent level. Just as with non-married, not having to pay the tax reduces savings. The effect is a bit smaller for the married. In chapter 7 I discuss various possible explanations for these findings.

Table 8: Results from DiD for married, 2008-2011.

Dependent Variable	(1) First Differences in Saving	(2) First Differences in Saving	(3) OLS Savinglevel
ATE[1]	-3.028***	-2.994***	-3.008***
	(0.102)	(0.101)	(0.101)
Income[1]		0.590***	0.200***
		(0.062)	(0.020)
Inheritance/gifts[1]		0.362***	0.253***
		(0.014)	(0.009)
Familysize[1]		0.121	-0.153***
		(0.147)	(0.024)
Treatment dummy			-1.275***
			(0.072)
Time dummy			-1.384***
			(0.052)
Constant	-1.371***	-1.400***	1.640***
	(0.052)	(0.053)	(0.292)
N	131745	131745	263490

Robust standard errors in parentheses. Clustered at the household level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
[1]First differences in specification (1) and (2), levels in (3).

6.3 VALIDITY

The internal validity of a statistical analysis concerns the question of whether estimated causal effects in a sample are representative for the population under study. The external validity concerns the ability to use these results to conclude about other populations (Stock & Watson, 2012, p. 354).

How informative are the research designs above about the behavioral responses of the Norwegian people to the wealth tax? First of all, the population under study is the people with a positive net taxable wealth. Also, the choice of sample has in various ways excluded those with a very high wealth. Thus, the results are only valid for those with a positive, but moderate stock of wealth.

Figure 9: Histogram of net tax wealth for non-married in 2011, binwidth = 5000.

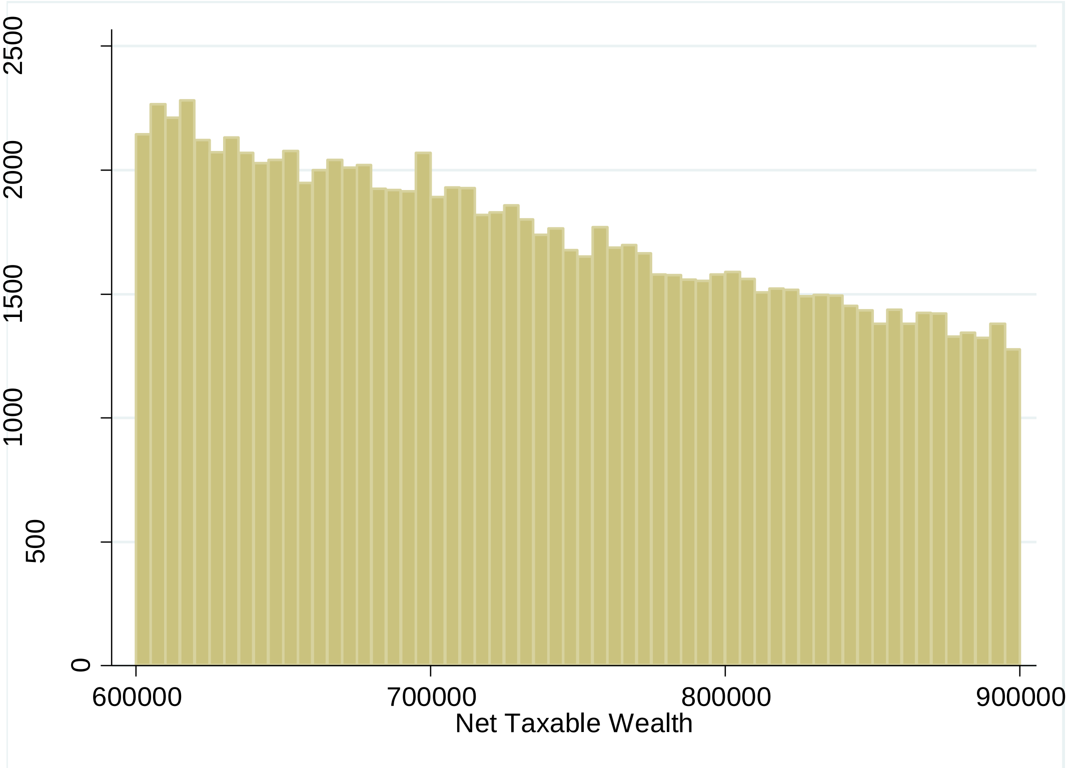

Figure 10: Histogram of net taxable wealth for the married in 2011, binwidth = 7692.

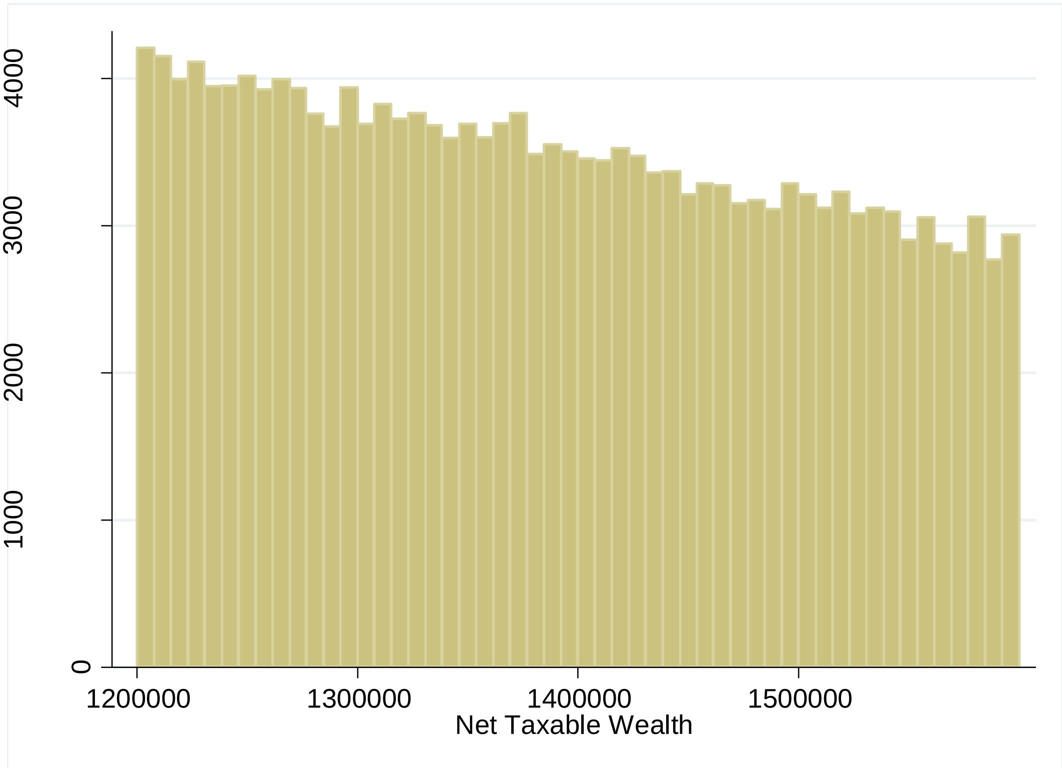

The internal validity of the RD design is very strong when applied to a correct setting, whilst its external validity is more questionable. The fundamental assumption in this analysis is that the people just below the cutoff are representative for those just above the cutoff. If people were able to perfectly manipulate their net taxable wealth, this should lead to an excess mass of people just below the cutoff, and the two groups would differ in characteristics.

Figure 9 and 10 are histograms of non-married people around the cutoff. Although there is a small spike just before the cutoff in figure 9, it can hardly be said to be perfectly manipulated. In figure 10 there is no spike at all. It does not appear to be any sorting around the threshold that could invalidate the research design. Of greater concern is the lack of a clear jump in when plotting the assignment and outcome variable. The noise in the data makes it very hard to see a clear jump in the raw data. It also makes the estimates very sensitive to the bandwidth, since the local linear regression uses very few observations around the cutoff to calculate the estimates. This poses a threat to the internal validity of the RD-estimates, and might explain why the estimates are positive for the married, whilst the graph shows a negative relationship. Thus, we should be careful about drawing strong conclusions from these estimates.

The crucial assumption in the DiD design is that the treatment is as-if randomly assigned to the groups. The assignment is as-if random if we control for all relevant variables that are correlated with treatment and is a determinant of active savings. The advantage of using panel data structure is that we can control for unobserved characteristics that differ across people but are constant over time, and factors that constant across people but change over time. The estimates will be biased if there are variables that vary both across time and people that are not included in the regression. In this analysis income, change in the size of household, receiving gift or inheritance are included in the regression without affecting the estimates of the treatment. But we cannot rule out the fact that there are relevant and omitted factors that influence the result.

Possible extension of the DiD approach used in this paper could be to use matching methods to insure comparability between control and treatment group. Unfortunately, this was beyond the reach of this master thesis.

7 SUMMARY

A visual inspection of figures 8 and 9 indicate a small negative effect of the wealth tax on active savings. This is confirmed by local linear regression estimates for non-married households but not for married couples, as the estimates for married couples are either positive or insignificant. The differences-in-differences analysis of the change in savings following a wealth tax relief show that those who did not have to pay wealth tax after 2010 actually reduced their active savings by 3-3.5 percent. This result is rather puzzling.

One possibility is that a tax on wealth may induce households to save more if they have a predetermined savings goal, or a specified amount of wealth that they wish to accumulate over time. The removal of a wealth tax may therefore decrease savings since less is needed in order to meet the after-tax wealth objective. However, inspection of the data reveals that the non-married group affected by the tax change paid on average 1700 NOK in wealth tax before 2010, while their average reduction in savings was about 50 000 NOK. This is far more than the gain from the tax elimination.

This illustrates a general point, which is that changes around the threshold for the wealth tax are not substantial in terms of the actual tax paid. It is reasonable to believe that individuals has to face a monetary change of some size in order to change their behavior, especially given that the households studied here are already quite well off in terms of both income and net wealth. Any response around the cutoff would therefore hinge on the assumption that knowing that one has to pay the tax has an effect in itself, or knowledge of how the wealth tax affects the effective tax rates of financial saving. The last point is often referred to as the salience of taxation. Recent literature have shown that people's perception of how the tax system work may play a significant part in determining the behavioral responses to changes in it (see Chetty, Looney and Kroft, 2007). For example, people have been known to confuse the terms average and marginal tax rates (Liebman and Zeckhauser, 2004). If people do not comprehend how large the reduction in effective tax rates were as a consequence of the 2010 change in allowance, that might influence their choices.

One could also argue the opposite; the fact that households have imperfect information about the actual monetary effect could lead to an inflated assessment of the tax relief and subsequently an increase in consumption that is larger than their increase in after-tax wealth. This could explain why the effect of a change in the tax schedule has a larger effect than crossing the tax threshold. While the former is a policy change, the latter is more a consequence of a slow and gradual wealth increase over time.

Given the diversity of the results presented here it is difficult to draw any clear conclusions about the effect of the wealth tax on household saving, except that these results do not support the claim that a wealth tax depresses financial saving. On the other hand, the discussion surrounding the efficiency of the wealth tax revolves mostly around groups of people who have a substantial amount of wealth and even investment opportunities. Until the tax is completely abolished or there are significant changes in the tax schedule at the top of the wealth distribution, we cannot answer whether the results presented in this analysis are also valid for more wealthy tax payers.

REFERENCES

Atkinson A. and J. Stiglitz (1976): The design of tax structure: Direct versus indirect taxation, Journal of Public Economics, 6, pp. 55-75.

Bernheim, B.D. (2002): Taxation and saving, in: A. J. Auerbach & M. Feldstein (ed.), Handbook of Public Economics, edition 1, volume 3, chapter 18, pp. 1173-1249. Amsterdam: North Holland.

Boskin, M. (1978): Taxation, Saving and the Rate of Interest, Journal of Political Economy, 86, pp. 3-27.

Burbidge, J.B., L. Magee, and A.L. Robb (1988): Alternative Transformations to Handle Extreme Values of the Dependent Variable, Journal of the American Statistical Association 83, pp. 123-127.

Carroll C. D. (2001): Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation), The B.E. Journal of Macroeconomics, vol. 1(1), pp. 1-38.

Chamley, C. (1986): Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives, Econometrica, 54(3), pp. 607-622.

Chetty, R., A. Looney and K. Kroft (2007): Salience and Taxation: Theory and Evidence, NBER Working Paper 13330.

Diamond, P. A. and E. Saez (2011): The Case for a Progressive Tax: From Basic Research to Policy Recommendations, Journal of Economic Perspectives, 25(4), p. 177.

Edson, C. (2012): The capital constraining effects of the Norwegian wealth tax, Discussion paper 724, Statistics Norway.

Fagereng, A. and E. Halvorsen (2015): Imputing Consumption from Norwegian Income and Wealth information, working paper, Statistics Norway.

Hahn, J., P. Todd, W. van der Klaauw (2001): Identification and Estimation of Treatment Effects with a regression-Discontinuity Design, Econometrica, 69(1), pp. 201-209.

Imbens, G.W. and K. Kalyanaraman (2009): Optimal Bandwidth Choice for the Regression Discontinuity Estimator, National Bureau of Economic Research Working Paper 14726.

Imbens, G.W. and T. Lemieux (2008): Regression Discontinuity Designs: A Guide to Practice, Journal of Econometrics, 142, pages 615-635.

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

Lee, D. S., T. Lemieux (2010): Regression Discontinuity Designs in Economics. Journal of Economic Literature 48, pages 281-355.

Liebman, J.B. and R.J. Zeckhauser (2004): Schmeduling, Mimeo, Harvard University.

Mintz, J.M., J. Seade (1991): Cash Flow or Income? The choice of Base for Company Taxation, The World Bank Research Observer, Vol. 6, No. 2, pp. 177-190.

Norwegian Ministry of Finance (2015): Prop. 1 LS (2014-2015), Skatter, avgifter og toll 2015 (In English: Taxes, fees and customs duty 2015), https://www.regjeringen.no/contentassets/4b5801220aa442e5b592dc9cecd0cce2/no/pdfs/prp201420150001ls0dddpdfs.pdf, 04.04.2015.

Norwegian Ministry of Finance (2014): Direkte skatter (In English: Direct Taxation), https://www.regjeringen.no/nb/tema/okonomi-og-budsjett/skatter-og-avgifter/direkte-skatter/id2353512/, 28.03.2015.

OECD (2012): OECD Economic Surveys: Norway 2012, OECD Publishing, http://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=ECO/WKP(2012)27&docLanguage=En, 24.03.2015.

Official Norwegian Report 13 (NOU) 2014: 13. Capital taxation in an international economy. https://www.regjeringen.no/contentassets/bbd29ff81485402681c6e6ea46655fae/no/pdfs/nou201420140013000dddpdfs.pdf, 28.03.2015.

Pence, K. M. (2006): The Role of Wealth Transformations: An Application to Estimating the Effect of Tax Incentives on Saving, Contributions to Economic Analysis & Policy: Vol. 5: Iss. 1, Article 20.

Pichet, E. (2007): The economic consequences of the French wealth tax, La Revue de Droit Fiscal, Vol 14, p. 5.

Piketty, T. (2014): Capital in the Twenty-First Century, Cambridge, Harvard University Press.

Piketty, T. and E. Saez (2013): A theory of optimal inheritance taxation, Econometrica, vol.81, n.5, pp. 1851-1886.

Rosen, H.S., T. Gayer (2008): Public Finance, Eight Edition, New York, McGraw-Hill.

Seim, D. (2013): Real or Evasion Responses to the Wealth Tax? Theory and Evidence from Sweden, working paper

Stiglitz, J. (2000): Economics of the Public Sector, Third edition, New York: W.W. Norton & Company.

Stock, J.H., M.M. Watson (2012): Introduction to Econometrics (3rd edition), Pearson Education Limited.