Universität Hohenheim
Fakultät Wirtschafts- und Sozialwissenschaften

EFFICIENCY OF SELECTED FISCAL POLICY INSTRUMENTS

Cand. Dr. oec. Vincent Dekker, M.Sc.

Institut für VWL
Lehrstuhl für VWL insb. Finanzwissenschaft (520D)

Inauguraldissertation zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften (Dr. oec.)

2017

Datum der Verteidigung: 07. August 2017
Dekan: Prof. Dr. Dirk Hachmeister
1. Gutachter: Prof. Dr. Nadja Dwenger
2. Gutachter: Prof. Dr. Nadine Riedel

Acknowledgements

I would like to thank Nadja Dwenger for supervising my thesis. I have benefitted a great deal from thorough discussions with her, on and off topic, and I am grateful that she encouraged me to visit both summer schools and conferences that have developed me as an economist. Many thanks also to Nadine Riedel, to whom I owe my entry into the academic world and who has, through several workshops and Skype sessions, contributed significantly to the development of this thesis. I am more than privileged to have both Nadja and Nadine as my advisors.

Many thanks to all my (ex-) colleagues in Hohenheim, especially to Martina Baumann, Patrick Gauß, Patricia Hofmann, Bodo Knoll, Davud Rostam-Afschar, Robert Stüber and Lukas Treber, who have all contributed to a very pleasant working environment that eased the pain of writing this thesis. Special thanks has to go to Konstantin Kuck, with whom I have shared countless coffee breaks and pub nights, most of them dedicated to research. His encouragement and thoughts helped me a great deal in writing this thesis and coming to a conclusion. I am furthermore indebted to Kristina Strohmaier, who had the misfortune to be stuck with me for three years of research and two projects. This thesis is based partially on our joint work and the discussions with her as well as some night shift sessions have all contributed a great deal to finishing my dissertation. I would also like to thank Datenlabor Hohenheim for providing access to several data sources such as Datastream, Orbis and PATSTAT.

Last but not least I owe a great deal to my parents, Ronald and Bettina, for encouraging and supporting me not only during my academic career, but throughout my whole life. Although she hasn't been part of my life for long, Hannah has made the last year of my doctoral studies a treat with her loving kindness and support. Without her, I wouldn't be finished today.

Declaration of Co-Authorship

Chapter 2: A Data-Driven Procedure to Determine the Bunching Window

The text of this chapter comes from joint work with Kristina Strohmaier (Ruhr-University of Bochum and University of Hohenheim) and Nicole Bosch (CPB, The Hague, The Netherlands). Work on the whole project was evenly shared with Kristina Strohmaier, Nicole Bosch aided in the subsection "Anatomy of Response" as well as with the institutional setting.

Chapter 3: The Effect of Transfer Pricing Regulations on Intra-Industry Trade

The text of this chapter comes from joint work with Kristina Strohmaier (Ruhr-University of Bochum and University of Hohenheim). The idea and the estimation strategy were developed through joint discussions. The Data, Hypotheses, Estimation Strategy and Results sections have been my own work, whilst Kristina Strohmaier contributed the Motivation and the Conclusion sections.

Chapter 4: Opening Pandora's Box - Do Intellectual Property Boxes Foster Innovation?

The text in this chapter is my own work.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Taxation is at the heart of modern welfare states. Governments collect taxes to be able to provide public goods to the taxpayer.[1] In this sense, the taxpayer does not forfeit the money but benefits from public infrastructure, redistributive systems and other services provided by the state. Yet, many people perceive taxation to be unfair or too high (Blaufus et al., 2015) and some engage in avoidance and evasion activities to minimise their individual tax burden (Mason and Calvin, 1984).[2] Others show real responses, such as the reduction of labour supply (individuals) or productivity (firms). These behavioural responses induce inefficiencies into the tax system that are economically undesirable. Efficient tax policies should therefore aim at minimising the behavioural response to distortions created by the tax system.

Since the first theoretical explorations into optimal taxation (Ramsey, 1927; Mirrlees, 1971; Diamond and Mirrlees, 1971a,b), both theoretical economists and empiricists have thoroughly analysed the topic and its inherent equity-efficiency

[1] It is acknowledged that this is a modern view on taxation. Historically, taxes (or equivalent policies such as the tithe) were also used for example as compensation, to finance wars or for a kings daily living expenses.

[2] This can also be extended to firms.

trade-off. From a theoretical point of view, extreme inequality can be a desirable, efficient solution to the taxation problem (e.g. lump sum taxation), but this does not satisfy the demands for equality expressed by parts of the population.[3] The optimal commodity taxation proposed by Ramsey (1927), for example, turns out to be regressive, thus having a starker negative effect for poorer households, which is not desirable from an equity point of view. Although this has been mitigated somewhat in the analyses by Mirrlees (1971); Diamond and Mirrlees (1971a,b), the general consensus in fiscal policy is therefore to accept a certain level of inefficiency in favour of equity. So conditional on the level of equity that is desirable and its inevitable level of inefficiency, the policymakers need to decide on the tax policy that induces a minimum of additional inefficiency. The aim is thus to find a second best solution to the taxation problem. To achieve this goal, it is critical to understand the (behavioural) responses to various policy instruments.

The thesis at hand intends to contribute to this understanding by dedicating each chapter to the analysis of a different fiscal policy instrument. **Chapter 2** focusses on the individual tax system in the Netherlands that exhibits tax brackets, as opposed to a smooth progressive tax system. The aim is to uncover the extent of behavioural responses to the kinks in the budget set that are created by the non-linear increases in the marginal tax rates at the tax brackets cutoff points. **Chapter 3** discusses the implications of the introduction of transfer pricing regulations (TPR) on intermediate goods trade. The chapter thus analyses an anti-avoidance measure implemented by many governments in recent years and evaluates the consequences for allocative and distributional efficiency. **Chapter 4** deals with a recently developed tax incentive for research

[3]See Fehr and Schmidt (1999); Bolton and Ockenfels (2000) for theories on inequality aversion and Saez and Stantcheva (2016) for field-experimental evidence of social preferences towards redistribution and hence, demand for equality.

and development (R&D), namely the intellectual property box (IP-Box). Said to foster innovation by the implementing governments, critics accuse the IP-Box regimes of providing yet another profit shifting opportunity for multinational enterprises (MNEs). The study assesses the implications that the introduction of IP-Box regimes had on innovation and shifting behaviour, in order to judge on the efficiency and effectiveness of such a policy instrument.

To sum up, the thesis delves into aspects of both individual and corporate taxation at the national and international level, analyses tax schedules, tax incentives as well as disincentives and ultimately tries to provide guidelines for an efficient design of fiscal policy. Before the thesis turns to the in-depth analysis of modern day tax systems, it is important to emphasise that the problem of taxation and the question of how to tax are not new. To obtain a better understanding of taxation and the various policy instruments that governments have implemented in the past, a short overview of the history of taxation as well as philosophical considerations on the justness of taxation are provided in the following. This should enable the reader to put the results of this thesis into perspective and highlight that the topic of efficient and just taxation was as relevant 2000 years B.C. as it is 2000 years A.C..

The Roots of Taxation and Fiscal Policy[4]

If one wants to define the future, they must study the past

Confucius

The first records of taxation date back some 6000 years to the people of Lagash in the Sumer plain, located between Euphrates and Tigris in modern day Iraq. But the first extensively documented system of taxation was in place

[4] The narrative of this section draws heavily on Adams (1999)

in ancient Egypt. The pharaoh was the owner of all land and the land was leased to farmers, subject to the harvest tax. Alongside the taxation of agricultural production, also sales, businesses as well as trade and even people were taxed. In other words, the tax system of ancient Egypt was acquainted with an income tax, a VAT, corporate taxation as well as tariffs for international trade, all of which are tax policy instruments implemented to this day. To govern and monitor taxation, the ancient Egyptians had a vast armada of scribes[5] that worked for the pharaoh and ensured the payment of taxes. The scribes had great power in determining the taxes due and this inevitably led to corruption within the system, which in turn slowly led to the downfall of the civilisation. The downfall of the empire run by the pharaohs also documents the revolt of people against unfair taxation and the subsequent engagement in evasion and avoidance activities. The famous Rosetta Stone describes an amnesty offered as a means to end a civil war, which was sparked by an uprising against high taxation. In summary, ancient Egypt knew nearly all tricks of the trade when it came to levying taxes, monitoring payments and dealing with delinquents.

The Egyptians succeeded at times in constructing a very efficient tax system, especially before the scribes possessed too much power and corruption undermined the system. But for the peasants, taxation was hardly fair. The concept of justness in the matter of taxation was introduced in an unprecedented and in its thoroughness never again achieved form by the ancient Greeks, especially in the state of Athens. Although Confucius and his successor Mencius in China developed ideas of just taxation in the form of a 10% flat tax and the Well-Field tax system[6], as well as setting out principles that we would still consider as just

[5] The literal translation would be *writers*.

[6] The idea of the system is to divide land into nine equal squares. The eight squares on the outside are for the farmers and their families, but the central square is public land, cooperatively cultivated by the farmers that work on the surrounding squares. The produce of this ninth square is the tax paid by the farmers.

today[7], it was the ancient Greeks who were able to lay the foundations for any equity concerns in taxation that are still present nowadays.

The Athenians abolished direct taxation on themselves and advocated indirect taxation instead. Taxation fell on the use of public infrastructure in the form of tolls, on sales and imports. This also ensured that most taxes were paid by foreigners, who benefited from the infrastructure provided by the Athenians. Athens also controlled the tax system of the Delian League, a defense league of almost two hundred city-states, with each of its members paying a tribute to the League. Aristides, an Athenian general, was appointed to oversee the collection of the tribute. He determined the taxes due on a city-by-city basis and assessed everyone according to ability and worth, hence introducing the first form of vertical and horizontal tax equity. Moreover, Aristides succeeded in creating a tax system that was considered as just and fair by all relevant parties, the Athenians as well as the taxed city-states.

The concept of justness was also debated by the philosophers of the time. Plato writes the following in Book 1 of *The Republic*: In a dialog between Socrates and Thrasymachus, the latter explains that the just man will pay more and the unjust man less on the same amount of money if there were an income tax. Likewise, the just man will receive little and the unjust much if there is something to give. Socrates argues against Thrasymachus that it is always advantageous to be just (Bloom et al., 1991). Here, the justness of the taxpayer is described and the argument is brought forward that the mode of taxation doesn't matter, for as long as taxpayers have heterogeneous levels of justness, injustice, especially horizontal injustice, will be inherent to any tax system.

A second feat achieved by the Athenians was the collection and sharing of public revenue without bureaucracy. The tax system was progressive and there

[7]Mencius advised that taxes should be paid on income, not gross production, double taxation should be banished and tariffs should be abolished.

were tax exemptions for the poor, like in many countries today. But the main source of public revenue came through the liturgy, a voluntary contribution to a public good, which was enforced only by tradition and a great love for the city. Whenever the public was in need of something, the richest men were called upon to provide it, e.g. build a bridge, host athletic games or donate military equipment. The liturgy thus worked in a similar way as modern day taxation in that the rich provide something for the general public, but the fundamental difference is that the government played no role in the management of the liturgy. Nowadays, a significant portion of tax money is spent on governing it, on the revenue but mostly on the expenditure side. This causes inefficiencies that were not present with the liturgy.

For many decades, taxation in various forms was laid upon defeated nations by the victors of a war. It was often seen as compensation for the losses suffered during the war. But in the second century B.C., the Romans used taxation as a means to topple the commercially powerful Rhodians. Rhodes was in a key geographic location for any merchant that wanted to reach either Greece or Rome and each ship that stopped in the harbour was subject to a 2 percent harbour tax based on the value of their cargo. When Rhodes infuriated the Romans by not supporting them in the war against Macedonia, the Roman Senate established a free port on the Isle of Delos, thus creating the world's first tax haven. In one year, trade volumes passing through Rhodes declined by 85% and this led to the economical destruction of Rhodes.

In summary, the ancient Egyptian, Greek and Roman civilisations already faced the topics of taxation that this thesis intends to discuss in the subsequent chapters: How should an efficient tax system be designed? How can equity concerns be incorporated? And most predominantly, how does taxation alter behaviour? Ancient though these problems are, the current state of tax practices suggests that humanity has since not succeeded in finding a suitable answer to any of these questions. As an example, the smorgasbord of tax systems in place in Europe is described in the following, focussing especially on the different aspects of personal and corporate taxation as well as tax incentives.[8] This should lay a foundation for the in-depth analysis to come in the subsequent chapters. It can also show how the results derived in this thesis apply to a broader range of countries that have similar tax systems in place.

Overview of Taxation in Europe

If the ultimate tax system already existed, we would expect to see countries with similar tax systems. But even in a geographically close and economically integrated area like Europe, where most countries are part of the same union, vast differences between tax systems can be observed.

Table 1.1 summarises the different tax schemes in place in the EU. Whilst almost all countries have a single flat tax rate in place for corporate income, individual taxation follows more diverse patterns. Some of the former Soviet countries (e.g. Bulgaria, the Baltic countries or Hungary) have followed Russia in experimenting with a flat tax rate for individual income, some countries implement tax brackets with progressive marginal income tax rates (e.g. Belgium, the Netherlands and the United Kingdom) and there are also countries that have a (more or less) smoothly increasing marginal tax rate (e.g. Denmark, Germany and Sweden).

Next to the structure of the tax scheme, the tax rates also vary considerably between the EU member states. Those countries implementing a flat tax rate

[8] Note that the sole existence of differences in taxation across countries does not necessarily show that taxation is inefficient. In open economies however, like in the EU, tax systems should aline to minimise the possibilities that taxees have to avoid taxation.

Table 1.1: Tax Systems in the EU

Country	Corporate Tax	Individual Tax
Austria	25% flat rate	progressive rates
Belgium	33.99% flat rate; under certain conditions (SME) progressive rate	progressive rate, 5 brackets
Bulgaria	10% flat rate	10% flat rate
Croatia	20% flat rate	progressive, 4 brackets
Cyprus	12.5% flat rate	progressive, 5 brackets
Czech Republic	19% flat rate	15% flat rate
Denmark	23.5% flat rate	progressive
Estonia	20% flat rate	20% flat rate
Finland	20% flat rate	progressive, 6 brackets
France	33.33% flat rate; 15% for SME	progressive, 5 brackets
Germany	15% flat rate	progressive marginal rate increases smoothly with income
Greece	29% flat rate	progressive, 4 brackets
Hungary	10% below HUF 500 million of the positive tax base and 19% over HUF 500 million of the positive tax base	16% flat rate
Ireland	12.5% on trading income, 25% on non-trading income	progressive, 3 brackets
Italy	27.5% flat rate	progressive, 6 brackets
Latvia	15% flat rate; 9% for SME	23% flat rate
Lithuania	15% flat rate; 9% for SME	15% flat rate
Luxembourg	21% flat rate	progressive, 19 brackets
Malta	35% flat rate	progressive, 4 brackets
Netherlands	first bracket: 20%, second bracket 25% (above € 200,000)	progressive, 4 tax brackets
Poland	19% flat rate	progressive, 3 brackets
Portugal	21% flat rate; 17% for SME	progressive, 6 brackets
Romania	16% flat rate	16% flat rate
Slovakia	22% flat rate	progressive, 3 brackets
Slovenia	17% flat rate	progressive, 5 brackets
Spain	28% flat rate; 25% for SME	progressive, 6 brackets
Sweden	22% flat rate	progressive
United Kingdom	20% flat rate	progressive, 4 brackets

Notes: This table shows the anatomy of the different corporate and individual income tax systems in place across the EU-28. SME refers to special tax rates for small and medium enterprises. The data are taken from European Commission et al. (2015).

for personal income exhibit relatively low tax rates, Latvia's (23%) being the highest. In contrast, the Scandinavian countries all have very high top marginal tax rates surpassing 50%, which is due to the financing of extensive redistribution systems. So another problem related to the structural differences between the tax schemes of EU countries is the difference in redistribution systems.[9] This has to be kept in mind when discussing optimal tax policies.

For the corporate tax rate, Bulgaria, Hungary, Cyprus and Ireland are at the lower end of the scale with statutory tax rates of 10-12.5%. At the top, France (33.33%), Belgium (33.99%) and Malta (35%) stand out. As will be shown in more detail in Chapter 3, companies react sensitively to taxation with respect to their location decision. Therefore, the significant differences in corporate tax rates (25 percentage points) within the EU suggest that the corporate tax rate is, at least partly, utilised to attract MNEs.

Strategically setting corporate tax rates is one possibility to use the tax system as an incentive device. But across the EU, several other tax incentives exist, both in personal and corporate taxation. For personal income, deduction possibilities such as for children or training, are present in all EU member states. Those deductions could be used strategically by individuals to lower taxable income, especially in countries that exhibit tax brackets. Such a case is analysed thoroughly in Chapter 2.

Corporations also face several tax incentives, such as the deductability of costs from the tax base and incentives to foster R&D. R&D is a very risky, yet crucial element of success for a corporation, because it enables the growth of a firm. To make the risk economically bearable, many countries implement tax incentives for R&D, recognising that R&D has the characteristics of a public

[9] Tiebout (1956) argues that differences in redistributive systems reflect individual preferences and are therefore no sign of inefficient taxation. But given that redistribution is a public good, free riding becomes a major problem, especially with the Freedom of Movement Act in the EU.

Table 1.2: Popularity of R&D Tax Incentive Instruments

Countries	Tax Credits	Enhanced allowance	Accelerated depreciation	IP-Box
Austria	x			
Belgium	x		x	x
Bulgaria	x		x	
Croatia		x		
Cyprus		x		x
Czech Republic	x	x		
Denmark	x	x	x	
Estonia				
Finland		x	x	
France	x			x
Germany				
Greece		x		x
Hungary		x		x
Ireland	x			
Italy	x		x	
Latvia		x		
Lithuania		x	x	
Luxembourg				x
Malta	x			x
Netherlands	x	x		x
Poland	x	x		
Portugal	x			x
Romania		x	x	
Slovakia	x			
Slovenia		x	x	
Spain	x			x
Sweden	x			
United Kingdom	x	x	x	x

Notes: This table shows different tax incentives for R&D in the EU-28. IP-Box refers to intellectual property boxes. The data are taken from CPB et al. (2015).

good and without government intervention, an underprovision of R&D is a likely outcome. Innovations coming from R&D, unless strongly protected by patents, can be utilised by anyone without incurring additional costs and there will be non-rivalry in consumption. This creates further disincentives for R&D from a firm's perspective.

The responses of firms specifically to such tax incentives are analysed in Chapter 4, where effects of intellectual property boxes (IP-Boxes) are discussed. Table 1.2 provides an overview of the types of R&D tax incentives that are in place across the EU. Although there are several arguments in favour of incentivising R&D, Estonia and Germany implement no incentive instrument at all. On the other end of the scale, the United Kingdom utilises all four instruments to increase investments into R&D. There is a fundamental difference, however, between tax credits, enhanced allowances accelerated depreciation (the first three incentives in Table 1.2) and IP-Boxes (the fourth incentive in Table 1.2). Tax credits, enhanced allowances and accelerated depreciation all aim at reducing the costs of R&D, either the employment costs or the general development costs. Thus they subsidise the inputs of the R&D process and are therefore also available to unsuccessful projects. IP-Boxes on the other hand target the output of the R&D process and thus by definition only subsidise successful R&D. Whilst the first three instruments thus directly target an increase in innovative activity, this link remains in the dark for IP-Boxes and Chapter 4 is dedicated to identifying this link.

This section provided a very brief overview of some of the characteristics of the tax systems that are in place across the EU. It showed that no country taxes like the other, which indicates that the optimal tax system is yet to be discovered. Neither the tax rates themselves, nor the number or extent of tax incentives are the same in any two member states of the EU. Through the in-depth analysis of certain aspects of tax systems, this thesis aims at providing a better understanding of those tax systems and potentially provide guidelines that could lead to an alignment of tax systems across the EU.

Chapter 2

A Data-Driven Procedure to Determine the Bunching Window

2.1 Introduction

A central topic in public economics is the assessment of welfare losses caused by behavioural responses to income taxation. Following the seminal paper by Feldstein (1995), a large literature emerged where welfare losses are inferred from the elasticity of taxable income (ETI).[1] Notwithstanding the large variation in identification strategies and data used in these studies, a common finding is that the elasticities and thus the tax-induces welfare losses are modest. Recent studies hint at different explanations for these modest estimates, such as optimisation frictions (Bastani and Selin, 2014; Chetty et al., 2011), shifting of income over time (Le Maire and Schjerning, 2013) or shifting across tax bases (Harju and Matikka, 2016). More fundamentally, other papers claim that the

[1]See Saez et al. (2012) for a comprehensive overview.

structural parameter cannot be retrieved from these estimates, because the ETI depends on the institutional framework, such as the exact definition of taxable income (Slemrod, 1998; Saez et al., 2012; Doerrenberg et al., 2017).

A growing strand of the literature utilises the bunching method to obtain a non-parametric estimate of the ETI (Saez, 2010; Chetty et al., 2011).[2] This method exploits the clustering behaviour of individuals at kinks in a non-linear tax system[3] to identify the ETI by the number of individuals that adjust their income to stay below the threshold of a tax bracket. Using the bunching method is attractive as it builds on a sound theoretical foundation and is not susceptible to endogeneity biases, a problem suffered by previous ETI literature (Saez, 2010; Gruber and Saez, 2002; Weber, 2014).

The aim of our study is twofold: First, we tackle the issue of finding an optimal bunching window, which is necessary as individuals are not able to perfectly adjust their taxable income to the tax threshold. The large number of robustness checks in previous studies already hints at the uncertainty regarding the optimal choice of the bunching window and the appropriate counterfactual model. However, using the correct, potentially asymmetric bunching window is crucial for the unbiasedness of the ETI as it does not only enter the estimation of the excess mass of individuals around the threshold but also directly affects the estimation of the counterfactual density following Chetty et al. (2011) that is needed to derive the elasticity. We therefore propose a simple, data-driven procedure to determine the bunching window. This improves on the visual inspection that determines the bunching window in the previous literature. As a consequence, our method explicitly allows the bunching window to be asymmetric around the threshold and to be more flexible. It also enhances the reproducibility of studies implementing the bunching approach. Second, we

[2]For an overview of the recent advances in the bunching literature see Kleven (2016).
[3]Kinks appear at thresholds in a tax schedule, where marginal tax rates jump up.

estimate the compensated elasticity of taxable income with respect to the net-of-tax rate in the Netherlands using the refined bunching approach. We employ a unique longitudinal data set containing exact declared taxable income and tax deductions for a representative sample of the Dutch population (IPO data from 2003 to 2014). Information on taxable income and deductions is provided by the Dutch tax authority and, therefore, free of measurement errors – something that is vital to obtain reliable estimates with the bunching method. Since we observe the exact taxable income, we do not need to rely on imputation techniques. The data also contains covariates, such as gender and marital status as well as information on self-employment, which enable us to analyse various sub-samples. In addition, we are able to analyse the anatomy of responses using information on mortgage interest deductions.

Our main findings are as follows. First, Monte Carlo simulations show our refinement of the bunching method to be robust to various changes in key parameters, such as different binwidths, sample sizes, tax rate changes and degrees of optimisation frictions. This consistency of the approach is a valuable addition to the literature that has previously relied on visual inspection. Second, in our empirical application, we estimate an ETI with respect to the net-of-tax rate of 0.023 at the highest tax threshold, significant at the one-per-cent level. This result is in line with some of the bunching literature, such as in Chetty et al. (2011), who find an elasticity below 0.02 for their full sample of Denmark, but differs from, for example, Bastani and Selin (2014) who report an elasticity of only 0.004 for Swedish tax payers. Third, we find significantly higher compensated ETIs for women and self-employed individuals. However, contrary to most other studies, we are also able to identify a non-zero elasticity for individuals in paid employment. Fourth, by analysing the anatomy of response, we find that most employees reduce their taxable income by utilising mortgage interest deductions. Further exploration reveals that this effect is driven by joint filers that have the possibility to shift these deductions between them.

The paper proceeds with Section 2.2, which introduces the bunching methodology as well as our improvements. Subsequently, the institutional setting and the data are presented in Sections 2.3 and 2.4, respectively. Section 2.5 presents our estimation results. Section 2.6 provides the conclusions.

2.2 Methodology

2.2.1 Bunching

There are three potential ways in which individuals can respond to taxation. The first is a real response. As suggested by standard microeconomic theory, the distortion of prices and wages in the economy due to taxation induces individuals to adjust their working hours and effort as well as their educational or training decisions. The second response is legal tax avoidance, such as using deductions or moving income to other time periods to reduce the taxable income in the current period. The third type of response is illegal tax evasion. To test the prediction from microeconomic theory and to quantify the responses, we follow the literature and identify the compensated ETI in the spirit of Feldstein (1995), which is a summary statistic for all kinds of behavioural responses. This central parameter is defined as the percentage change in taxable income z due to an increase in the net-of-tax rate $(1-\tau)$ of one percent:

$$e(z) = \frac{dz}{z} \Big/ \frac{d(1-\tau)}{(1-\tau)}. \tag{2.1}$$

Theoretically, the introduction of a kink in the budget set of individuals induces bunching behaviour within a certain income range, provided that preferences are convex and smoothly distributed among the population. This will lead to a spike in the density exactly at the kink, but due to adjustment costs and optimisation frictions, a bunching window around the kink is observed more often in reality (Chetty et al., 2011). Comparing the income density with a counterfactual scenario without a kink, the excess mass of taxpayers can be used to determine the elasticity $e(z)$. A detailed derivation of the bunching estimator can be found in Saez (2010) and Chetty et al. (2011). The compensated ETI, identified locally at the threshold k, is then given by

$$e(k) = \frac{b}{k \cdot \log(\frac{1-\tau_1}{1-\tau_2})}, \tag{2.2}$$

where the net-of-tax rate changes by $\log(\frac{1-\tau_1}{1-\tau_2})$ per-cent.[4] The relative excess mass of taxpayers at the threshold k is given by b, which is the only parameter that needs to be estimated. To estimate b, Chetty et al. (2011) propose to determine the counterfactual density by running a local polynomial regression on binned data, while excluding data bins within the bunching window.

A major drawback of the bunching method is that it is sensitive to the choice of bunching window (Adam et al., 2015). A commonly used approach is that of selecting the window by visual inspection, which makes it vulnerable as it is selected at the researcher's discretion. Furthermore, recently published papers select the counterfactual by trial-and-error and the model seems to be chosen ad libitum. Neither visual inspection, nor this selection of the counterfactual model are optimal for efficiency, reliability and reproducibility of the results.

[4]It is identified if and only if the derivative of the counterfactual density function $h_0(z)$ with respect to z is continuous in z $\forall z$.

2.2.2 Extension

Motivated by the drawbacks of the usual implementation of the bunching approach, we extend the estimation procedure by relying on the data at hand rather than on visual inspection to determine the bunching window. Removing the researcher's discretion in this matter is preferable in its own right, but we also argue that our method produces more efficient estimates of the elasticity. The optimal situation would be for the bunching window to comprise all the individuals who would adjust their taxable income as a response to the tax change at the threshold and *only* those. The bunching window should not be too small, for fear of omitting some taxpayers that attempt to bunch at the kink, nor should it be too large, which would bias the results by also including non-bunchers. The existing literature implements symmetric bunching windows around the kink with varying sizes that are determined by graphical inspection. We propose the use of a possibly asymmetric bunching window with an endogenously determined size.[5] One argument in favour of an asymmetric bunching window is that risk-averse individuals are expected to be more likely to over-adjust their income to make sure they realise an income which is below the threshold. Furthermore, recent empirical evidence shows that individuals have difficulties in understanding the difference between average and marginal tax rates, leading to suboptimal behaviour (Ito, 2014). These cognitive and psychological components might lead to an asymmetric bunching window. A graphical intuition is given in Figure 2.1 for the simple case of a linear counterfactual model.[6] The figure plots taxable income (relative to the threshold)

[5]Using an optimal window renders robustness checks with different bunching windows obsolete. To show the gain in efficiency, Figure 2.B.3 of the Appendix depicts a comparison with two different model specifications.

[6]For illustrative purposes, we depict a linear counterfactual model. In our empirical application, the counterfactual is allowed to be a higher order polynomial model, determined by the minimum BIC value.

Figure 2.1: Data-driven Procedure to Determine the Bunching Window

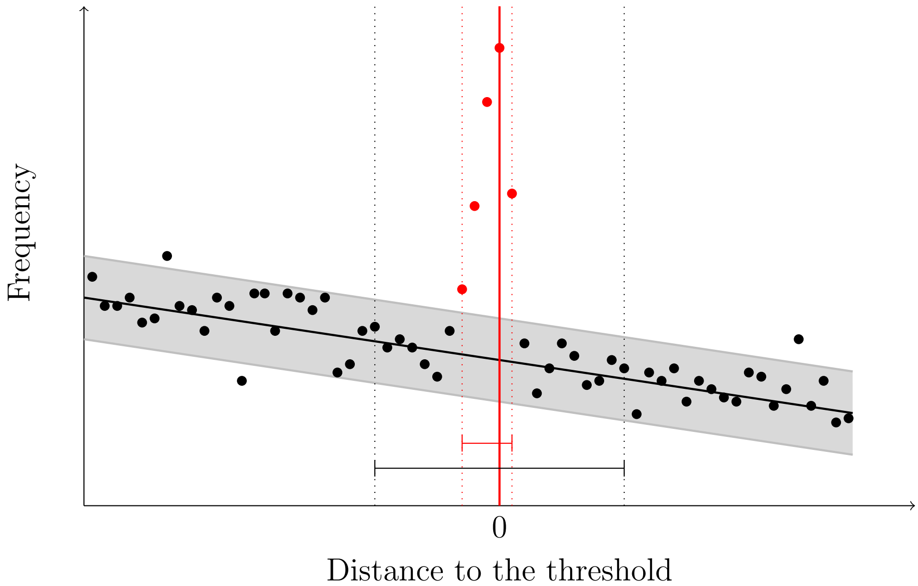

Notes: This figure shows the bin midpoints as well as the fitted values of a linear regression. The grey confidence band is calculated with the standard errors of the point prediction. Here, five subsequent bin midpoints around the threshold lie outside and therefore determine the relevant (asymmetric) bunching window.

against the number of individuals within certain income bins. All subsequent binpoints around the threshold that have a higher actual number of taxpayers than predicted (coloured in red) are then used to determine the bunching window. In order to determine its optimal width, we propose the following step-wise procedure:

1. Set an excluded region around the threshold.
2. Run a local regression through all data bins outside the excluded region and predict the frequencies.
3. Compute a confidence interval around the prediction.
4. Subsequent bin midpoints outside the confidence interval comprise the bunching window.

In general, the excluded region can be set arbitrarily; however, we propose to iterate through different combinations of upper and lower bounds of the excluded region. This hedges against concerns that the chosen excluded region could affect the determination of the bunching window.[7] The confidence band should be set using standard significance levels, with higher levels tending to lead to a smaller bunching window. In other words, the probability that we erroneously include non-bunchers decreases. Depending on the setting and the data, this will lead to more conservative estimates of the elasticity.

The bunching window is formally derived as follows: Let $x_- \in \{-X, (-X+1), \ldots, 0\}$ and $x_+ \in \{0, 1, \ldots, X\}$ be the respective lower and upper bound of the excluded region, where X represents the midpoint of an income bin. Furthermore, define $l(x_-, x_+)$ as the lower bound of the bunching window and $u(x_-, x_+)$ as the upper bound, given the excluded region from $[x_-, x_+]$. For every tuple (x_-, x_+), run a local regression of polynomial order q:

$$\tilde{N}_j^{BW} = \sum_{i=0}^{q} \beta_i X_j^i + \varepsilon_j \qquad \forall \;\; j \notin [x_-, x_+]. \tag{2.3}$$

Then predict the counterfactual values $\hat{N}_j^{BW}$ and the associated standard error of the forecast:

$$\hat{N}_j^{BW} = \sum_{i=0}^{q} \hat{\beta}_i X_j^i \qquad \forall \;\; j \tag{2.4}$$

As a next step, calculate the upper value of the confidence interval CI_j^+ for a given t-value using standard procedures. To determine whether there are more individuals than predicted in an income bin j, subtract the CI_j^+ from the

[7] Our results indicate virtually no sensitivity of the bunching window to the size of the excluded region.

observed number of taxpayers for each j:

$$E_j = N_j - CI_j^+. \tag{2.5}$$

A positive E_j means that the number of individuals in income bin j exceeds the predicted number of individuals, as estimated by the polynomial regression. Put differently, if all E_j are negative, no bunching is present in the respective sample. Otherwise, the lower bound of the bunching window is given by:

$$l(x_-, x_+) = j_l^* + 1, \qquad \text{where} \quad j_l^* = \max\{j \in \mathbb{Z}_- :\, E_j < 0\} \tag{2.6}$$

which is the smallest subsequent income bin j that still satisfies the condition $E_j > 0$. Similarly, the upper bound is given by:

$$u(x_-, x_+) = j_u^* - 1, \qquad \text{where} \quad j_u^* = \min\{j \in \mathbb{Z}_+ :\, E_j < 0\} \tag{2.7}$$

which is the largest subsequent income bin j that still satisfies the condition $E_j > 0$.

By following this procedure, depending on the number of iterations with respect to the choice of the excluded region, different values are obtained for the lower and upper bounds of the bunching window. Several possibilities arise for which values of $l(x_-, x_+)$ and $u(x_-, x_+)$ to use as the limits of the bunching window, but we advocate using the mode of all estimated values. This ensures that, in most cases, the exact bounds of the bunching window will be obtained.[8] To estimate the ETI from Equation (2.2), the excess mass b is the only parameter that needs to be estimated, as the other parameters are known policy

[8]Other possibilities would be to use the minimum, maximum or mean, although we do not find large variations among these choices.

parameters. b is estimated in the following way:

$$\hat{b} = \frac{\hat{B}}{\frac{\sum_l^u \hat{N}_j}{(u-l+1)}}, \tag{2.8}$$

with $\hat{B}$ being the number of individuals within the bunching window. $\hat{N}_j$ represents the counterfactual number of individuals within an income bin j that are determined by local polynomial regression of the form:

$$\hat{N}_j = \sum_{i=0}^{q} \beta_i \cdot X_j^i + \sum_{s=l}^{u} \gamma_s \cdot I[X_j = s] + \varepsilon_j. \tag{2.9}$$

Next to the data driven procedure for the determination of the bunching window, we rely on the Freedman-Diaconis rule to determine the optimal bin size in each estimation. It states that the optimal binwidth is given by: $2 \cdot IQR(x) \cdot n^{\frac{1}{3}}$ (Freedman and Diaconis, 1981).[9] Prior works have chosen a single binwidth for their analysis and subsequently altered the binwidth in robustness checks to show that the estimates are robust to the choice of the binwidth. The size of the optimal binwidth, however, should depend on the number of observations around the threshold and therefore be determined for each threshold separately. In addition, we use the BIC criterion to determine the optimal number of polynomials when running the local regression.

2.2.3 Evaluation

Our preferred method of validation for the endogenous procedure to determine the bunching window would be to replicate previous studies that detected the bunching window by eyeballing. In the taxation literature, however, most bunching studies rely on administrative datasets of personal taxable income,

[9]IQR stands for interquartile range.

which are not freely available to other researchers. For example, Danish micro data as used by Chetty et al. (2011) can only be accessed through a Danish partner institution and the restricted PSID files used by Saez (2010) significantly differ from the public use files and are hard to obtain outside the US. Therefore, we assess the validity of our endogenously determined bunching window by Monte Carlo simulations and evaluate the performance for two predictions: how well the approach can recover the true elasticities and how well it can identify the bunching individuals. Moreover, we test the robustness of our approach by varying the key parameters of the model. We especially examine the variations in binwidth, amount of frictions, sample size and size of the tax change at the threshold. In a second exercise, we also vary the threshold and compare our results to two standard approaches from the literature.

The baseline simulation consists of $N = 1,000,000$ observations drawn from a triangular distribution.[10] It has a threshold k at $z = 50,000$, a binwidth of 100, and a tax change of 10 percentage points.[11] We run estimations for three true elasticities: $e = 0.02$, $e = 0.1$ and $e = 0.5$. As the bunching literature tends to find small elasticities, the paper only reports the detailed results for $e = 0.02$, although the results for the larger elasticities are in line with $e = 0.02$.

A comparison of the income distribution exhibiting a kink with a counterfactual scenario without a kink can be used to determine the elasticity (see Section 2.1). Potential incomes z_0 are used to calculate pre- and post-reform taxable incomes z_1 and z_2 respectively, where $z_1 = z_2$ for all individuals who would be at or below the kink, as they would not be affected by the new tax system.[12] We

[10]Because we draw from a triangular distribution, we know that the counterfactual model is best approximated by a linear model

[11]More specifically, the change is from 42% to 52%, resembling the change at the top tax threshold in the Netherlands.

[12]The choices of z_1 and z_2 come from maximising a quasi-linear utility function. The approach is similar to the approach taken in the working paper version of Bastani and Selin (2014).

identify all individuals as bunchers that have their highest post-reform utility at income level k, provided they had $z_1 > k$. To model optimisation frictions, we introduce a random component in the income of the bunching individuals, described by $\varepsilon \sim N(0, 142.3)$ in our baseline specification.[13]

Table 2.1 shows the results of our Monte Carlo simulations. The columns present the difference between the true and simulated elasticity as well as the ratio of identified bunchers to actual bunchers, which resembles an estimation error. Each row represents a different specification. In the baseline setting, the estimated elasticities have a mean very close to the true elasticity of $e = 0.02$. At the same time, we are able to identify 98.92% of the bunchers using our data-driven procedure. To assess the robustness of our approach we change the size of key parameters.

Throughout the bunching literature, various binwidths are implemented. Many studies alter the binwidth in robustness checks and show limited sensitivity to changes in the binwidth (Chetty et al., 2011; Bastani and Selin, 2014). The results in Table 2.1 show no significant changes regarding the estimated elasticities. A greater binwidth naturally would improve the identification of the number of bunchers by up to almost 100 %, but the number of individuals wrongly assumed as bunching would also rise with an increased binwidth (bias-efficiency trade off).

Next to the binwidth, the variance that represents optimisation frictions could affect the performance of our data-driven procedure. Indeed, increasing the variance term in the randomised component has a severely negative effect on the performance of the bunching estimator. The bias increase to around -0.007 which is far off the true elasticity and are only able to identify 59.65% of the

[13] The variance component comes from the working paper version of Bastani and Selin (2014) and is adjusted for Euro values. It is a function of working hours and the average wage. To check for sensitivity, it is altered in a later specification.

Table 2.1: Monte-Carlo Simulations for $e = 0.02$

	Bias		Ratio	
	Mean	Std. Dev.	Mean	Std. Dev.
Baseline	-0.0001	0.0003	0.9892	0.0038
Binwidth				
200	0	0.0004	0.9905	0.1085
400	0.0003	0.0004	0.99998	0.00003
Variance (t=1.96)				
100	0.0014	0.0006	0.9876	0.0017
300	-0.0067	0.0005	0.5956	0.0069
Variance (t=1)				
100	0	0.0007	0.9917	0.0043
300	-0.0005	0.0011	0.9622	0.0189
Observations				
$550,000$	-0.0006	0.0011	0.9606	0.0233
$2,050,000$	-0.0002	0.0005	0.9867	0.0066
Tax Change				
$42\% - 48\%$	-0.0008	0.0013	0.9466	0.0231
$42\% - 60\%$	-0.0002	0.0004	0.9887	0.0039

Notes: This table shows the results from the Monte Carlo simulations running 600 repetitions. The baseline consists of binwidth 100, variance 142.3, observations 1,000,000 and tax change 42%-52%. All specifications use a t-value of 1.96, except the third, which uses a t-value of 1. Note that a bias of zero indicates that the bias is less than 1/10000.

bunching taxpayers. A potential driver behind this could be the choice of confidence interval. A high confidence interval should provide a narrow bunching window. But because the optimisation frictions are so high, we would expect a much wider range of the bunching window as well as a flatter area of excess mass around the kink point. Therefore, for the third specification, we use a t-value of 1 instead of 1.96. The results improve significantly, and our procedure is able to identify 96.22% of all bunching individuals when the variance term is 300. In light of this, researchers should take the anticipated amount of optimisation frictions into account when setting the t-value for the confidence interval. For example, a more complex or dynamic tax system should lead to more optimisation frictions.

Because of its non-parametric nature, the bunching estimator relies on a large sample size. We test the impact of different sample sizes on the efficiency of our estimation procedure. As expected, we find that an increased sample size increases efficiency, although the gains asymptotically decrease to zero. This is in line with earlier findings, but our approach is also able to estimate the true elasticity in smaller samples with little bias. In addition to the sample size, the size of the tax change might also influence the estimation of the ETI. As larger tax rate changes have more severe consequences for individuals, we should observe more precise bunching with greater tax rate differences, as the costs of adjusting taxable income are increasingly outweighed by the benefits (Chetty et al., 2011). The true elasticity can be identified more precisely by increasing the size of the difference between the two marginal tax rates, which confirms the results by Bastani and Selin (2014) and Chetty (2012) that larger jumps in the marginal tax rate are more informative of the true ETI.

All of the above variations have been made using a linear counterfactual and a threshold of $50,000$. In a second step, we model a more realistic, log-normal income distribution to ensure that the results do not hinge on the location of the threshold within the income distribution, which at $k = 50,000$ is in the descending part. We rerun the simulation for different thresholds, namely $k = 10,000$, $k = 20,000$ and $k = 40,000$, which are in the ascending, flat and descending part of the income distribution, respectively (see Figure 2.B.1 of the Appendix). The bias at the lower threshold of $10,000$ is larger than in the baseline (-0.0021), but outperforms a small as well as a large symmetric bunching window on the basis of the RMSE.[14] The results at the middle and upper threshold are very close to the baseline results and suggest that the endogenous procedure to determine the bunching window is robust to the location of the threshold. Again, the RMSE suggests that the endogenous detection of the bunching window delivers the best results.

2.3 Institutional Background

The Dutch tax system is almost fully individualised and tax liabilities mainly depend on individual worldwide income. However, there are a few exceptions, two of which are relevant for our analysis. The first exception is that of means-tested subsidies, such as on health tax, child care and rent, which are all based on taxable household income. The second is that personal tax-favoured expenditures are transferable between partners, thus reducing taxable income. This last possibility is attractive under a progressive tax schedule such as that of the Dutch tax on labour income.[15]

Since 2001, income from different sources is treated in three different “boxes”,

[14] Small is defined as going from three binpoints below the threshold to three binpoints above, whilst large covers seven binpoints below and above the threshold.

[15] From a labour supply perspective, a third exception is also relevant. A non-working spouse can transfer the lump-sum tax credit to his or her partner. The moment this spouse starts working, their income will be taxed starting at the marginal tax rate. This, however, is not the focus of our study.

each with its own taxable income concept and tax schedule. In Box 1, income from profits, employment and home ownership is taxed. This includes wages, pensions and social transfers. Box 2 consists of income from substantial shareholding such as dividends and capital gains. Any other income from savings and investments is taxed in Box 3. Income in Box 1 is taxed at progressive rates that jump up at certain thresholds and thus create kinks in the tax schedule, whereas income in Box 2 and Box 3 is subject to a flat tax, that, in 2014, was 25% and 30% respectively.[16] For our analysis, we exploit the kinks in the Box 1 tax schedule for identification. It is furthermore worth noting that income losses in one box cannot be used to counterbalance taxable income in one of the others.

Income in Box 1, minus personal deductions, is taxed at progressive tax rates. It is important to note that the tax rates in the first and second tax bracket also include a social security contribution of around 31% for old-age pensions and exceptional medical expenses. Figure 2.2 provides an overview of the Dutch tax schedule in 2014. The marginal tax rate is represented by the solid line. In 2014, there is an increase in the marginal tax rate of 8 percentage-points at the first threshold. The social security contribution in the third tax bracket is compensated by an equal rise in tax rate in the second tax bracket, implying that marginal tax rates in the second and third brackets are exactly the same. However, there is a large jump in the marginal tax rate, from 42% to 52%, in the last bracket.

While the upper two tax rates stayed constant over the whole sample period, there exists some variation over time for the lower two tax rates. Figure 2.B.2 graphically illustrates the development of each marginal tax rate over time, again

[16] We are aware of the possibility of moving income between the boxes, which could be especially pronounced for self-employed individuals. For information on the importance of shifting between tax bases see Harju and Matikka (2016). Because of data limitations, we are unable to extend our analysis in that way.

Figure 2.2: Tax Schedule of 2014, Box 1

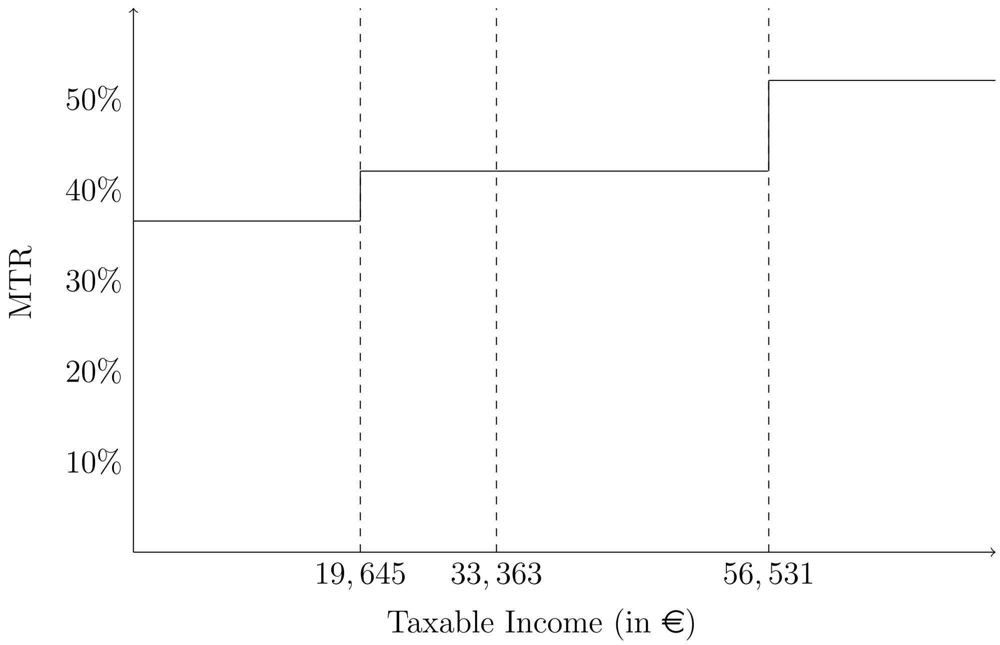

Notes: The figure shows marginal tax rates for the year 2014. At each threshold, denoted by the dashed lines, the marginal tax rate jumps up, except for the second threshold, where the tax rate and the social security contributions in the lower bracket equal the marginal tax rate in the higher bracket.

adding the social security contributions for the lower two tax brackets. Due to the stability of the two upper tax rates, this large jump of 10 percentage-points was existent in all years since 2001. Finally, is worth noting that for the considered time period, the income thresholds were adjusted upwards to account for inflation and to avoid the phenomenon of "cold progression" (see Table 2.B.1 in the Appendix).[17]

To assess the welfare losses due to taxation and better understand how individuals adjust their taxable income, it is essential to know the exact definition of taxable income. An overview of the computation of taxable income in the Netherlands is given in Table 2.2. One important channel of adjusting taxable income is legal tax avoidance by utilising deductions (Chetty et al., 2011). In our setting, these deduction possibilities include alimonies paid, charitable givings, health expenditures or mortgage interest deductions.[18] In the Netherlands, the mortgage interest deduction is quite high and common among house-owners. More importantly, all of these deductions can be shifted between partners.

Finally, important for any analysis looking on bunching is the exact tax payment procedure which has an influence on the technical possibilities to avoid or evade taxes. Three things are worth noting here. First, it should be emphasised that for people in paid employment, their employer withholds income tax from the income taxed under Box 1, which can be seen as a prepayment credited against the final tax amount payable at the end of the year. This "third-party reporting" is important for the interpretation of the results as it makes systematic tax evasion – one way of adjusting taxable income – more difficult (Kleven et al., 2011). Final income taxes are determined after the end of the fiscal year, when tax deductions and income from other sources are all taken into

[17]Given the specific values of the thresholds, which are never a multiple of one hundred, we are less concerned with round number bunching (Kleven and Waseem, 2013).

[18]These are (at least in parts) common in other countries like Great Britain or Germany.

Table 2.2: Definition Taxable Income Box 1

Gross Wage	
	– Pension fund and unemployment insurance contributions employee
	+ Health insurance contribution employer
= Taxable Labour Income	
	+ Income from housing
	+ Freelance earnings
	– Alimony/maintenance paid
	– Charity donations
	– Mortgage interest deductions
	– Health expenses deduction
	– Other personal deductables
= Taxable Income Box 1	

Notes: This table shows the computation of Box 1 taxable income. Gross wage includes pension benefits and received social transfers.

account.[19] Second, an important distinction is single filing or joint filing of tax returns. Even though the Dutch tax system is rather individualised, married couples can choose to file their returns jointly. In addition, cohabiting couples are also allowed to jointly file their tax return, provided they have lived together for more than six months. Third, taxes can be filed digitally (computer-assisted) or on paper. The share of digital filers has increased dramatically from about 30 percent in 2003 to almost 95 percent in 2015. Digital filing of tax returns is not only helpful when deducting certain personal expenditures, but also facilitates the optimal shifting of income. The exact threshold becomes more salient and enables people to locate at the threshold more precisely.

In sum, the Dutch tax system can induce bunching behaviour because of a combination of three things: 1) partners can move deductions between them and

[19]The tax thresholds in the Netherlands are known before the start of each fiscal year, as these are published together with the governmental budget which is presented each year, on the third Tuesday of September (Prinsjesdag).

this is most attractive in the highest income tax bracket; 2) mortgage interest deductions are quite large; 3) the digital filing of tax returns clearly reveals tax thresholds as well as the related benefits of shifting certain deductions. As is shown below, these specific features of the Dutch tax system result in sharp bunching.

2.4 Data

The data used in this study is the Income Panel Data (IPO) provided by Statistics Netherlands. This longitudinal data set covers the period from 2001 to 2014. It contains administrative data on all possible sources of income, on an individual level, as well as a very detailed account of possible deductions from the tax base. The panel is updated with new information on marital status and include other, randomly selected individuals in every period to account for people who are no longer observable. Most importantly for this study, Statistics Netherlands provides the information on relevant taxable income for Box 1 (see Table 2.2). The taxable income variable is obtained from the tax department, representing the exact taxable income per individual. This circumvents the problem of measurement error, which is vital for analyses that use the bunching method. As our income measure includes all tax deductions, we do not have to rely on tax simulators that are used in other studies (Gruber and Saez, 2002; Chetty et al., 2011) to determine the tax liabilities, thus mitigating bias that could stem from this exercise.

In addition to the information on taxable income and deductions, the dataset also includes demographic characteristics, which we exploit to study heterogeneity in the bunching behaviour of different socio-economic groups. We provide separate estimates for self-employed individuals, who theoretically would be more prone to bunching because of the lower costs and greater possibilities of adjusting their taxable income. Furthermore, we distinguish people according to gender and filing status. Our estimation sample is restricted as follows: We exclude students as well as all people receiving governmental benefit payments, as most of them receive similar amounts, thus creating an artificial mass point. Because the tax is different for individuals aged 65 and over, we also exclude them from our estimation, as well as those below the age of 18. We omit the years 2001 and 2002 to avoid the inclusion of any after effects of the 2001 major Dutch tax reform.[20] Furthermore, we only retain individuals with a positive reported taxable income. The pooled sample consists of $N = 1,219,572$ individuals, which is roughly 1% of the Dutch population per year. The sample is evenly balanced with respect to gender (55% male) and married individuals (65%). Furthermore, the sample contains 14% self-employed individuals, including CEO's, who would be in a position to decide on their own salary and are able to adjust it.

2.5 Results

2.5.1 Bunching Evidence

Figure 2.3 gives a first glance at the bunching behaviour. It displays the income distribution for the most recent year of our sample where the data is collapsed into income bins of 200 euros. The income thresholds of 2014 are indicated by vertical lines. Clear bunching behaviour can be seen at the first and third threshold. Note that there is no change in the marginal tax rate at the second threshold in 2014 and so there should be no incentive to adjust taxable income.

[20] The tax reform substantially changed the thresholds and marginal tax rates and introduced the system of income boxes.

Figure 2.3: Income Distribution in 2014

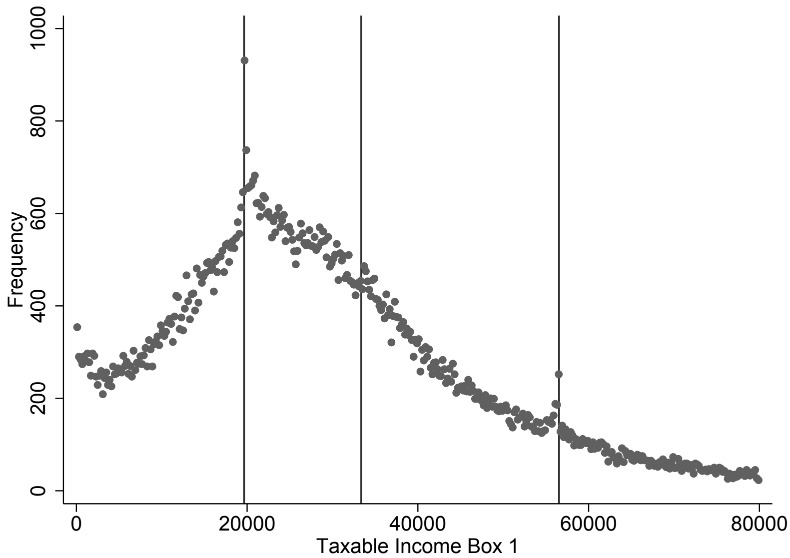

Notes: This figure shows the sample distribution of income below 80,000 € in the Netherlands for 2014. The data is collapsed into 200-euro bins. The vertical lines represent the first, second and third threshold of the Dutch tax system respectively.

We start out with the upper threshold, where the change in the marginal tax rate is largest with 10 percentage-points (23.81%). here, the incentive to bunch is most pronounced. Figure 2.4 reports the results for our pooled sample from 2003 to 2014 showing the number of observations per bin, relative to the threshold value. For the pooled years, our method to endogenously determine the bunching window provides an asymmetric bunching window ranging from -483 to +207 euros. We implement a 95% confidence interval for determining the bunching window throughout this study.[21] In addition, the BIC criterion suggests a 7th order polynomial counterfactual model for the upper threshold. In order to calculate an elasticity according to Equation (2.2), a weighted average threshold value is used ($54,163$ euros). The weights are constructed by the

[21]We also tested a smaller confidence level, i.e. a one-standard deviation increase, which corresponds to a 68% confidence interval. The results are slightly larger but less precisely estimated.

number of taxpayers within the bunching window in each year. Standard errors are calculated using bootstrapping techniques.

Figure 2.4: Bunching at the Third Threshold - Pooled Sample

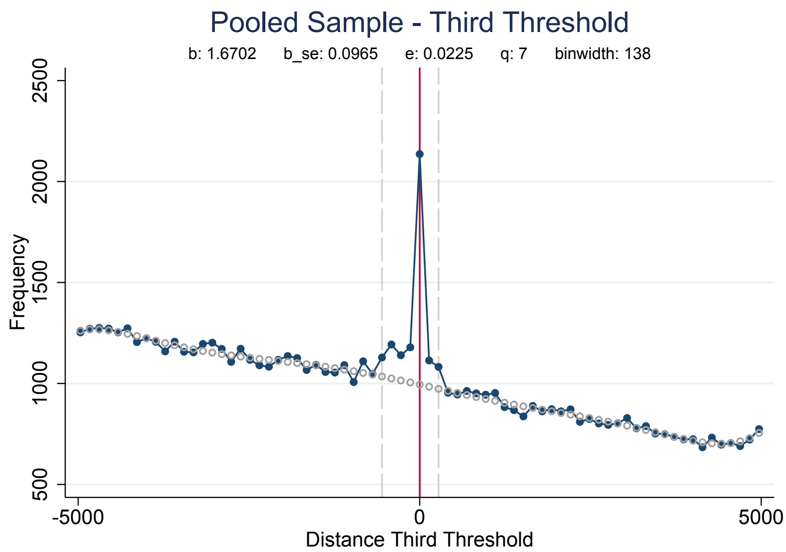

Notes: In this figure, bin counts are plotted relative to the third threshold for the pooled sample from 2003 to 2014. The bunching window is between -483 and +207 euros and the counterfactual model is a 7th order polynomial.

We observe sharp bunching at the third threshold and estimate an excess mass of $b = 1.67$, which corresponds to 1.67 times more individuals being at the threshold than would have been in the absence of any change in the marginal tax rate. The estimated excess mass translates into an ETI of 0.023, which is statistically significant at all usual significance levels. Quantitatively, a 10% decrease in the net-of-tax rate would induce a 0.23% reduction in taxable income. From an economic point of view, the tax response at this threshold is small, but is in line with the findings of the bunching literature for similar tax increases.

We compare our data-driven procedure to determine the bunching window with the estimates of using two different, symmetric bunching windows, to assess whether our procedure provides an improvement. Figure 2.B.3 reports the results using a large bunching window from −750 euro to +750 euro and a small bunching window from −350 euro to +350 euro. Compared to Figure 2.4, both the large as well a the small window delivers smaller ETIs combined with larger standard errors.

Finally, we estimate the excess mass for taxpayers and the ETI at the third threshold for all years separately. This eases any concerns about using a weighted average threshold in the pooled sample and hedges against possible bias stemming from observing individuals multiple times. The results are provided in the Appendix (Figures 2.B.5 to 2.B.16). One striking observation is that the bunching behaviour of individuals is increasing and becoming more precise over time. We ascribe this to learning effects, as taxpayers became more familiar with the tax system that fundamentally changed in 2001. For the year 2003 we still observe delayed effects from the major tax reform of 2001 and, therefore, the bunching behaviour is fuzzy and small. It then increases in the subsequent years until the elasticity reaches a level of around 0.025. Another explanation for this increase in the amount and precision of bunching could be the emergence of digital filing of tax returns, which made the threshold more salient to the general public.

A case could be made for bunching at the other thresholds of the Dutch tax system as well, albeit that the change in the net-of-tax rate is much smaller and therefore, we would expect less reaction. Figure 2.5 shows the results for the pooled sample for the first and second threshold, respectively. Surprisingly, we observe clear bunching behaviour of individuals at both thresholds.

At the first threshold, the income levels are quite low, which might suggest that individuals are more dependent on their income and should therefore show little real responses to a change in the marginal tax rate. However, the estimate for the ETI is about four times higher when compared to the third threshold

Figure 2.5: Bunching at the First and Second Threshold - Pooled Sample

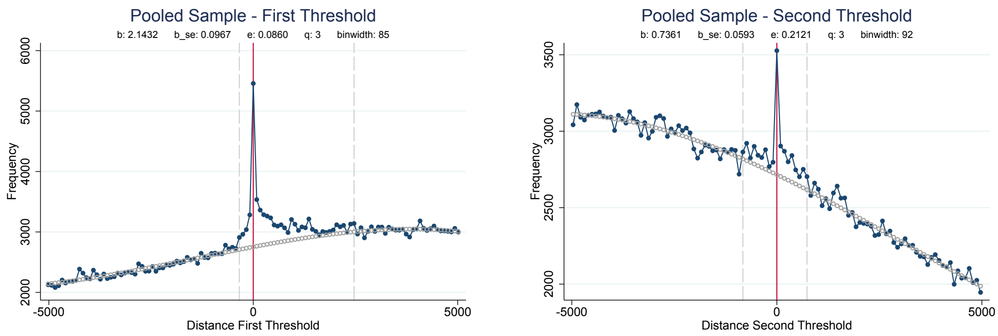

Notes: The figures show bunching at the first and second thresholds for the pooled sample from 2003 to 2014. The bunching window is between -297.5 and +2422.5 euros for the first threshold and between -782 and +690 euros for the second threshold. The counterfactual model is a 3rd order polynomial in both cases.

($e_{th1} = 0.086$ vs. $e_{th3} = 0.023$). Note that exact estimates for the ETI at both thresholds cannot be depicted, because we have changing tax differences over time, in addition to the changing threshold values. To calculate the ETI for these two thresholds, we used a weighted average tax change on top of the weighted average threshold. However, taking the average of the single-year ETI estimates as a sensitivity check delivers similar results: an elasticity of 0.085 at the first threshold. Contrary to our hypothesis, individuals with an income around the lower threshold seem to be more engaged in all kinds of tax-optimizing behaviour in order to relocate at the threshold. In addition, the bunching behaviour is less precise and there is slightly more mass to the right of the threshold suggesting that individuals are either less informed or less able to accurately adjust their income.

From an economic perspective, the second thresholds might be of special interest for two reasons. First, the total change in the net-of-tax rate (tax rate plus social security contributions) is comparably small, with 3.35 percentage points as a maximum in 2003. The gain of manipulating taxable income thus may not be larger than potential adjustment costs, which would lead to less bunching. Second, the jump in marginal tax rates vanished in some years due to the adjustments of the tax rate in tax bracket 2.[22] Especially in these years, there is no incentive to bunch at the respective threshold. Despite these small incentives, the right graph of Figure 2.5 clearly shows that there is bunching behaviour at the second threshold. For the pooled sample, the estimated ETI amounts to 0.212 and is much higher compared to the other thresholds. This can be partly attributed to the small tax changes (frequently below 1 percentage point) that are used to derive the ETI. As expected, the result is largely driven by early periods of the sample, where the jump in the marginal tax rate was still noticeable. Also, there is no bunching evidence in the years where the change in marginal tax rates is exactly zero. Figure 2.B.4 shows, for example, the results for the years 2008 and 2009, with 2009 being the first year where there was no jump at the second threshold. In 2008, we could still observe a small excess mass of 0.29, but in 2009, where the incentive to adjust taxable income vanished, we estimate a negative excess mass of merely -0.04, which is statistically insignificant.[23] For the year 2010, we are again able to identify a small excess mass of 0.18, which is significant at the 5-per-cent level. We believe that in accordance to the overestimation of small probabilities known from the behavioural economics literature (Kahneman and Tversky, 1979), individuals overestimate benefits from small changes in the net-of-tax rate and subsequently adjust their taxable income even though the economic gain is minimal.

Comparing the estimates from the first, second and third threshold, we conclude that the behavioural responses to taxation are heterogeneous and depend on the location of taxable income within the income distribution. This is an

[22]This phenomenon occurred in 2009, 2013 and 2014, respectively.

[23]Note that because the tax change is zero in 2009, we cannot compute a value for the ETI and therefore argue via the excess mass.

aspect that deserves a more thorough discussion in the future, although it is beyond the scope of this paper to shed further light on this matter.

2.5.2 Subgroup Analysis

To analyse if our results are driven by sub-groups, we split the sample according to gender, employment status and filing status. The bunching literature has shown that women tend to react more sensitive than men to changes in the net-of-tax rate. To investigate this in our setting, we split the sample according to gender and rerun the analyses for all thresholds. The results are shown in Figures 2.B.17, 2.B.18 and 2.B.19. At all thresholds, the ETI is larger for women than for men, although the difference at the first threshold is relatively small. Here, the number of women exceeds the number of men at the threshold, but at the second and especially the third threshold, the number of men is substantially higher than the number of women. This indicates that the ETI at the third threshold in the pooled sample is predominantly driven by males. Given that there are roughly three times as many men than women around the third threshold, the ETI here in the pooled sample can also be described by the weighted sum of the elasticities of men and women, i.e. $(3 \cdot e_{men} + e_{women})/4 = 0.023$.

We then split the sample by employment status. Self-employed individuals have better possibilities to adjust their taxable income and are therefore more prone to bunching. The results, shown in Figures 2.B.20, 2.B.21 and 2.B.22, confirm this hypothesis for all thresholds. At the upper threshold, the ETI for self-employed is 0.042, which is about twice the size of the estimate in the full sample. In contrast to findings in many other studies, we also find a significant elasticity for wage earners ($e = 0.017$); therefore the baseline result is not purely driven by self-employed individuals and a similar weighting exercise as before can be implemented here to obtain the ETI of the full sample.

One explanation for the significant bunching behaviour of employed individuals could be that of trade unions jointly setting wage levels for groups of individuals. Collective wage bargaining is very common in the Netherlands. Both agreements at national level and industry or company-wide agreements are made. First, at the national level, representatives of employers and employees advice on wage mutations. Then, negotiations take place at the industry-level where this advice is taken into account. Specific for the Dutch wage setting is that these agreements are extended by Law to non-unionized employees within a firm and to all employers in that industry by the Minister of Employment and Social Affairs. In addition to company-wide agreements, many large-scale companies have their own agreements and as such are not subject to the industry-wide agreements. Even though union membership is declining in the Netherlands as in other countries, the extension prevents the coverage rate from falling. The observed bunching for wage earners might also be an indication of collusion between employers and employees and of contracts being specifically designed to achieve a taxable income at the threshold (Chetty et al., 2011). In the Netherlands, employers and employees have the possibility to decide upon wage changes at the individual level. However, the degree of flexibility depends on the wage system. A majority of Dutch employees' wage payments are based on an industry-wide or company-wide wage schedule that resembles a staircase with fixed starting salaries and upper ceilings and fixed wage increments for each step in between (Deelen and Euwals, 2014). On top of this general increases, employees and employers can decide on performance-related pay, which are at the heart of the individual wage differences.

Finally, we split the sample by filing status. Although the Dutch tax system is rather individualised, cohabiting people (married or unmarried) can file a tax return jointly and are then known as fiscal partners. Fiscal partners have the

Figure 2.6: Subsamples by Filing Status - Threshold 1

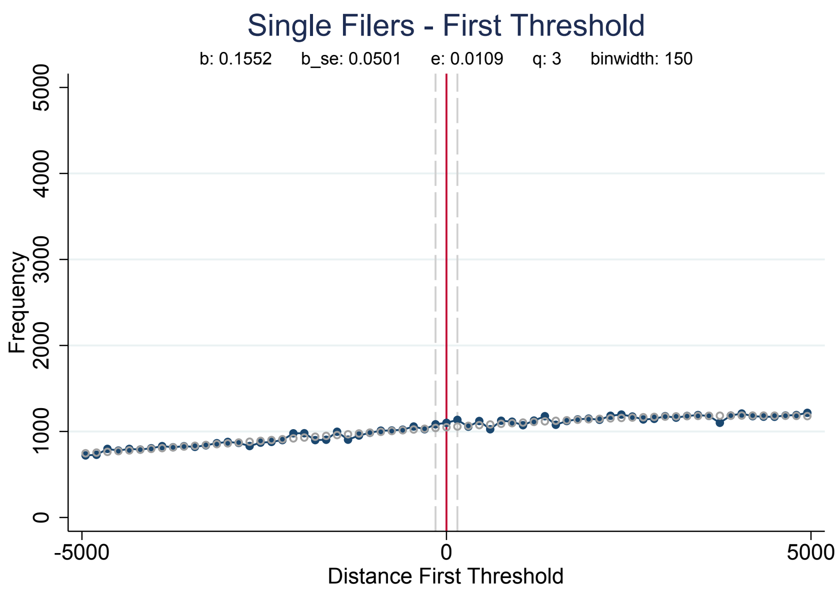

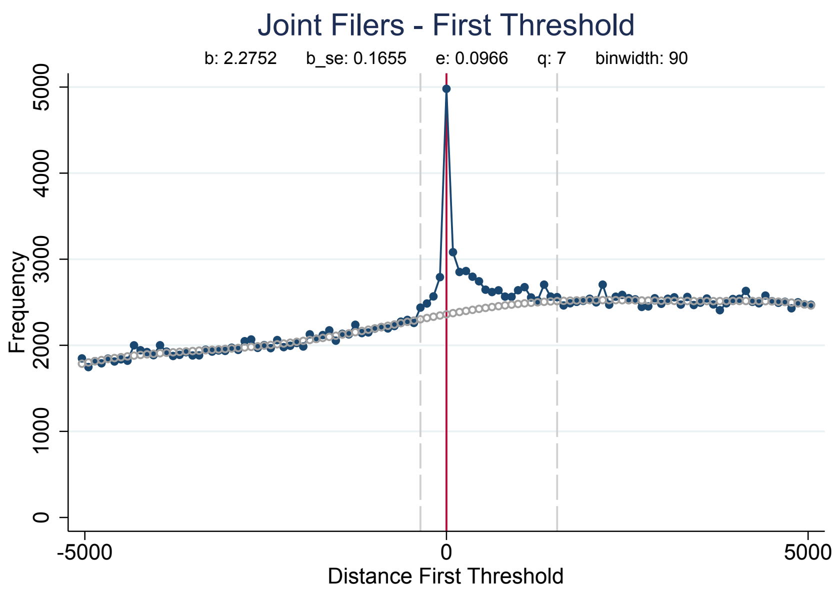

Notes: The figures show bunching behaviour by filing status for the first threshold. Single filers are defined as having no possibility to file a tax return with another individual. Joint filers are those individuals that have the possibility to file taxes together with a fiscal partner.

Figure 2.7: Subsamples by Filing Status - Threshold 2

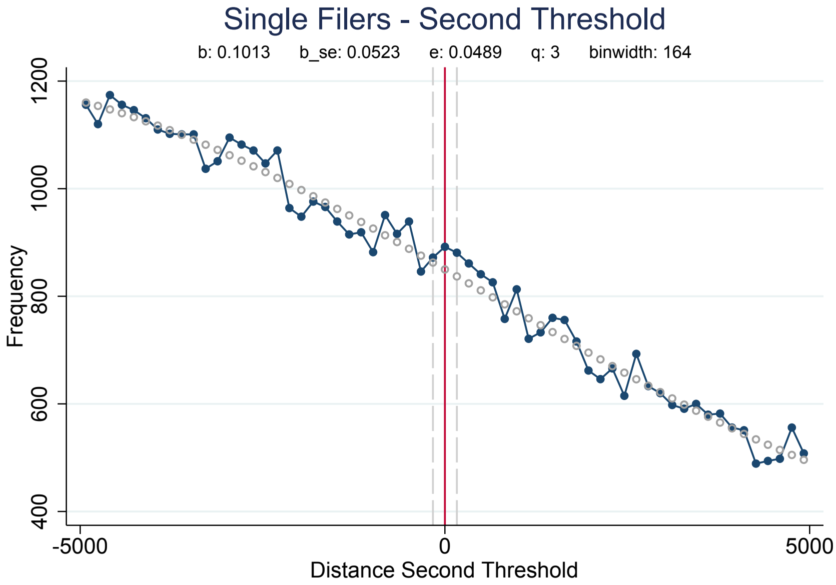

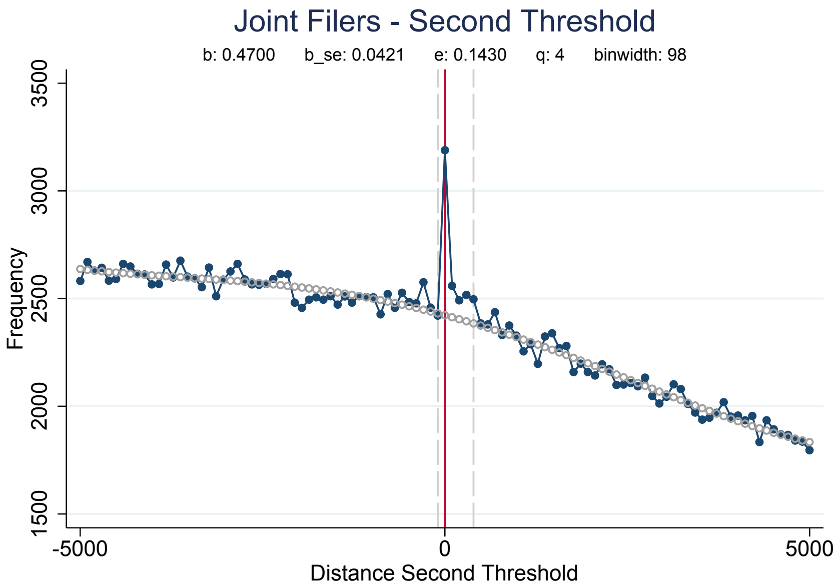

Notes: The figures show bunching behaviour by filing status for the second threshold. Single filers are defined as having no possibility to file a tax return with another individual. Joint filers are those individuals that have the possibility to file taxes together with a fiscal partner. Given the small number of single filers, a smaller scale had to be used.

Figure 2.8: Subsamples by Filing Status - Threshold 3

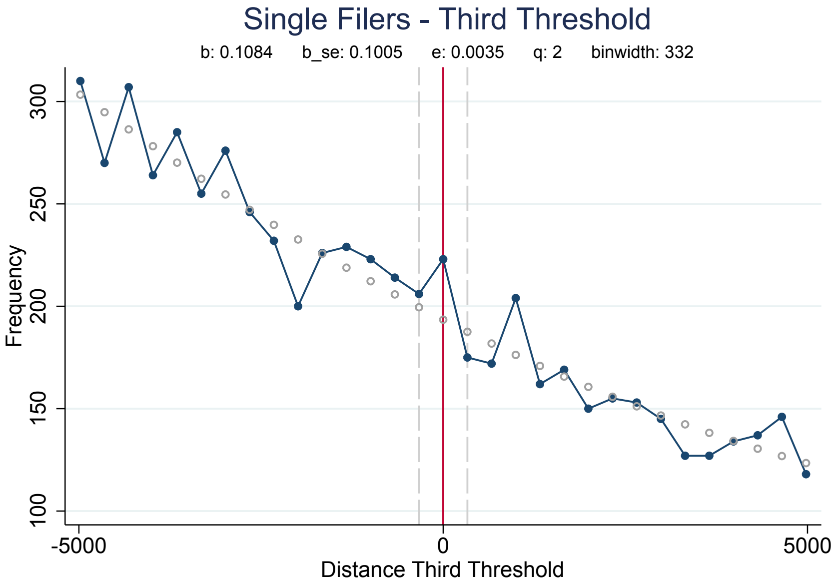

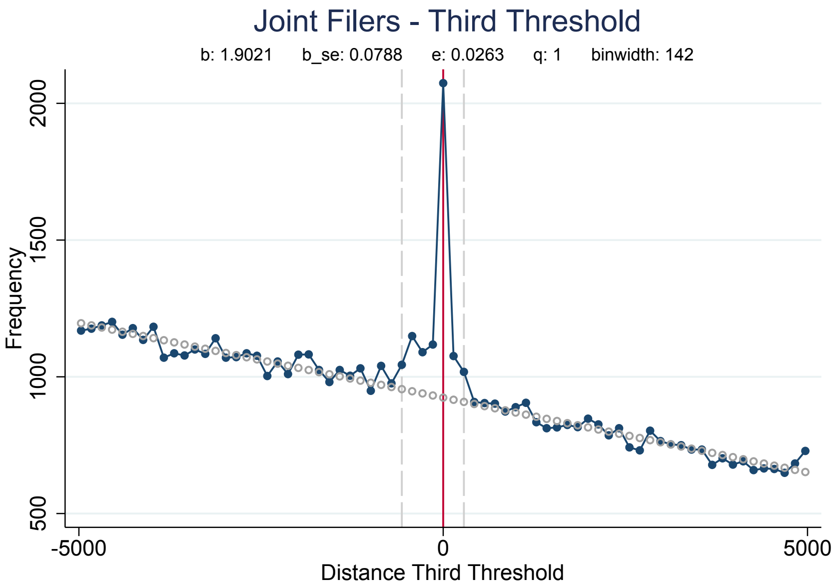

Notes: The figures show bunching behaviour by filing status for the third threshold. Single filers are defined as having no possibility to file a tax return with another individual. Joint filers are those individuals that have the possibility to file taxes together with a fiscal partner. Given the small number of single filers, a smaller scale had to be used.

possibility to shift deductions between them, an option that is not available to single filers. If individuals exploit the shifting of deductions as a main channel to adjust taxable income, we would expect to see no spike for single filers, whilst the estimates of the ETI for the joint filers should be close to the estimates from the pooled sample. The results shown in Figures 2.6, 2.7 and 2.8 support this hypothesis. We can see no bunching behaviour of single filers at any threshold, whilst for individuals that have the possibility to jointly file their tax returns, we clearly see bunching behaviour. The elasticity estimates are also similar to the estimates from the full sample, at the first threshold 0.097 compared to 0.086 and at the third threshold, the ETI is 0.026 compared to 0.023 in the full sample. At the second threshold, the deviation is more substantial (0.143 compared to 0.212), which could be driven by the relatively large number of single filers located around the second threshold. The results indicate that the shifting of deductions plays a crucial role in adjusting taxable income and this relationship is examined more thoroughly in the following subsection.

2.5.3 Anatomy of Response

The channels through which individuals bunch at the thresholds in a tax system are manifold. A recent study by Doerrenberg et al. (2017) shows the importance of tax deductions for welfare analyses with the ETI. As pointed out by Slemrod (1996), one way to reveal the channel that drives bunching is to look at the "anatomy of the behavioural response" (Saez et al., 2012). Due to the high presence of mortgage interest deductions in the Netherlands, it is interesting to examine this special kind of deduction, which can only be claimed for one, usually the main mortgage. In total, mortgage interest deductions are by far the biggest deduction claimed in the Netherlands with 9.9 billion € in 2015, as reported by the Ministry of Finance.[24] It can be claimed by single filers as well as by joint filers. We analyse the anatomy of response of wage earners for 2011 for which year we have additional information on the shifting behaviour. Interestingly, almost 88% of all wage-earners in the vicinity of the third threshold claim mortgage interest deductions. A much smaller fraction is partly self-employed and a few individuals claim other expenditures such as for health expenditures or charity donations. We are unable to identify the source of bunching for about 5% of all wage-earners within the bunching window. Their bunching response could be driven by a real response, such as a reduction in working hours.

As shown in Figures 2.6, 2.7 and 2.8, only joint filers engage in bunching behaviour and therefore, we examine the mortgage interest deductions claimed by joint filers. Because of the progressive tax system, shifting the full deduction to the highest earning partner will reduce tax liabilities the most. However, the actual incentive depends on the distance of taxable income to the threshold. In cases where the higher earning partner claims the deduction to reduce his taxable income but, in doing so, will cross the threshold, he has two options once his income level reaches that threshold: He can then either deduct the rest of the amount at a lower marginal tax rate, or he can shift the remaining deduction to his partner. If his partner's income is taxed at a lower marginal tax rate, then the fiscal partners have no incentive to shift any of the deductions. If, however, his partner now faces a higher marginal tax rate, the remaining part of the tax deduction should be shifted. This last case, where individuals around the thresholds do not utilise the full amount of mortgage interest deduction, will result in sharp, "negative" bunching at the third threshold.

The sharp shifting is visible in Figure 2.9. The graphs show the share of

[24] This compromises the mortgage interest deduction and taxed fictional income from housing.

Figure 2.9: Share Mortgage Interest Deductions

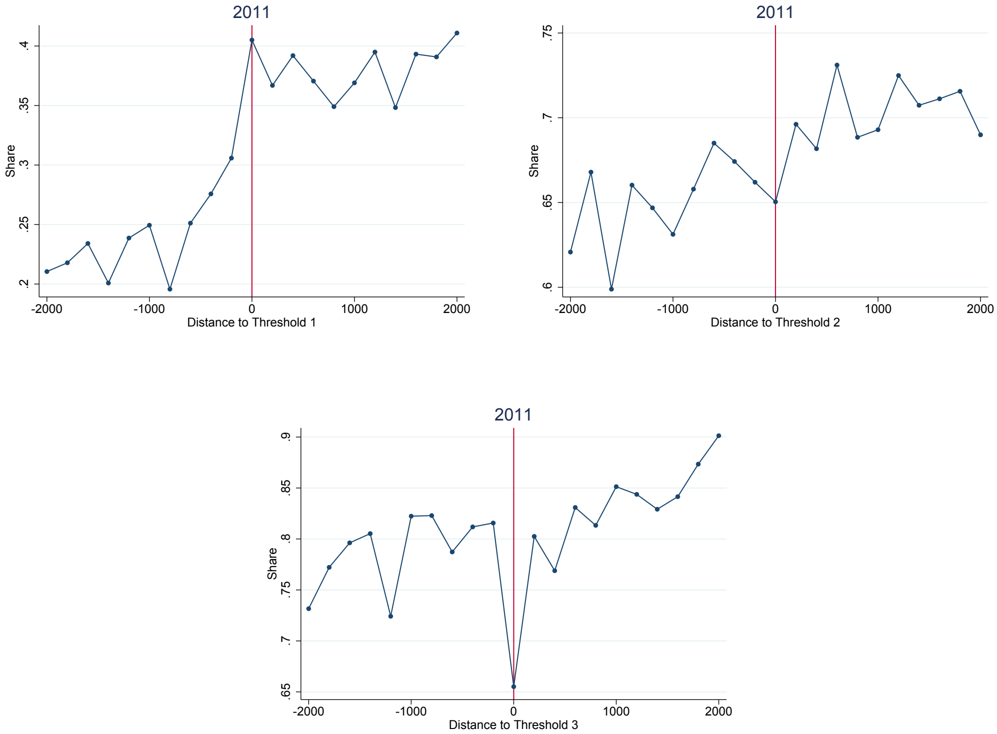

Notes: The top figures show the share of mortgage interest deduction claimed by one partner within couples around the first threshold (left) and the second threshold (right). The bottom figure shows the share of mortgage interest deductions within couples around the third threshold. The binsize is 200 euro.

total mortgage interest deduction claimed by one partner within a couple. For high-earners around the third threshold, the average share is around 80% of total mortgage interest deduction of the couple. The higher the income, the higher the share of the mortgage interest deduction claimed by the high-earning partner, which is in line with the common expectation that the tax advantage is higher for the highest earner. This can be seen by the jump in the share at the first threshold, where the share of mortgage interest deduction claimed by an individual below the threshold is between 20-25% and above between 35-40%. At the second threshold, where the tax change is minimal, no clear jump in the share is visible, but in line with the expected behaviour in the absence of a threshold, we see a gradual increase in the share. Following this line of argumentation, we would expect to see another jump at the third threshold, but we find a sharp dip in the share of mortgage interest deduction claimed by an individual at the upper threshold. This suggests that individuals strategically shift the mortgage interest deduction to their partners, as soon as they have located their taxable incomes at the threshold. Especially if both fiscal partners earn more than the third threshold, this splitting of the mortgage interest deduction reduces the overall tax burden of the fiscal partners. In summary, the graphs show that the mortgage interest deductions combined with the opportunity to file jointly are an important channel for reducing taxable income to reach thresholds of the tax system, especially at the third threshold.

A second possible channel is that of the real response, for example in hours worked. Due to the structure of our data, identification of these type of responses in hours could only be done indirectly, for example, via hourly wages. We do, unfortunately, not observe actual working hours. As bunchers come from above the threshold and hourly wage can be assumed to increase with taxable income[25], an individual that bunches should have a higher hourly wage than other individuals that obtain a similar taxable income. However, looking at data from 2006 to 2011, we cannot detect any significant difference between bunchers and non-bunchers left or right of the bunching window in terms of hourly wages, suggesting that real responses do not play a significant role in adjusting taxable income.

2.5.4 Relation to the literature

Our results relate to the literature in several ways, although cross-country comparisons of elasticities might be difficult due to different institutional features (Bastani and Selin, 2014). In line with other studies that implement the bunching approach, we find small but precise estimates of the compensated ETI with respect to the net-of-tax rate at the top tax threshold of 0.023. Chetty et al. (2011) find an elasticity at the upper threshold below 0.02 for their full sample on Denmark, while Bastani and Selin (2014) find close-to-zero elasticities on Sweden at the top tax threshold. Evidence on the United States, published by Saez (2010), indicates an elasticity of between 0.1 and 0.2, depending on the methodology, at the first threshold of the federal income tax schedule. He finds a smaller response for married individuals than for singles. This is in stark contrast to our findings indicating significant bunching by cohabiting individuals in the Netherlands, even at the first threshold of the tax system.

One structural difference that may explain this deviation between the United States and the Netherlands is the social acceptance and federal legitimation of part-time work. Employees in the Netherlands are arguably more free to choose their working hours than workers in other countries because of the existence

[25]This can be justified for example by the higher skill level that high earners have compared to low earners.

of the Dutch Working Hours (Adjustment) Act. They can file a request for amendment (increase or reduction) of their working hours that the employer cannot refuse. In the United States, only 19% of the working population was working part-time in 2013, whereas in the Netherlands this figure was almost twice as high, with 36%. The significantly larger proportion of women bunching can also be explained by this. In the United States, 26% of the female workforce worked part-time, whereas in the Netherlands, this was 58% and these women would likely earn an income close to the first threshold.[26] This could be a reason for the differences in the results. Alternative explanations are differences in other institutional features, such as the possibility to shift tax deductions between partners and the presence of digital filing of tax returns.

Earlier studies for the Netherlands find larger elasticities. Jongen and Stoel (2013) find an elasticity of around 0.1 for the short run and 0.2 for the medium run, with larger elasticities for women. The aforementioned study employs a panel approach and uses instrumental variable techniques to correct for endogenous taxes in line with Gruber and Saez (2002). In contrast to our study, they had to rely on a tax simulator to obtain marginal tax rates and determine taxable income. This can potentially cause measurement error, which could explain some of the deviation between the results. Another explanation would be that the bunching approach identifies a local elasticity as opposed to an average elasticity derived from the IV approach (Chetty, 2012).

A recent study by Bettendorf et al. (2016) for managing directors that own at least 5% of a corporation[27] finds elasticities between 0.06 and 0.11 for the upper threshold of the Dutch tax schedule, using bunching techniques. This is slightly larger than the elasticity of 0.04 that we identify for self-employed

[26]Shares are calculated from the OECD Statistics database, where the labour force is measured by national criteria.

[27]These so-called DGAs (*Directeur-Grootaandeelhouder*) face a special tax scheme. In our study, this sub-group belongs to that of the self-employed.

individuals, and could suggest that our results are partly driven by the DGA sub-group. Unfortunately, the limited number of DGAs in our sample prevents us from running the estimation separately for this group.

In previous bunching studies, a distinction is made between real response and income shifting. In a study on the self-employed in Denmark, Le Maire and Schjerning (2013) show that about 50% to 70% of the bunching in taxable income is due to income shifting over time. In a similar study for business owners in Finland, Harju and Matikka (2016) attribute two thirds of the ETI to income shifting between tax bases. However, we find that a large share of bunching is driven by tax deductions in combination with shifting them between partners. The presence of deduction possibilities confounds welfare analyses using the ETI (Doerrenberg et al., 2017). Our results confirm the significance of deduction possibilities for optimising taxable income. This finding also mirrors earlier findings on itemised deductions (Saez, 2010).

Furthermore, our results show little evidence of collusion between employers and employees. As final income taxes are based on taxable income and not on broad income on the payslip, it is harder for employers and employees to determine the exact taxable income. The same holds for the response in hours worked. Although, compared to employees in other countries, those in the Netherlands can more easily adjust the number of hours they work, adjusting the number of hours in such a way that the income stays below a certain taxation threshold is very difficult. This requires an extensive knowledge on those thresholds and the amount of all the deductions that turn labour income into taxable income in Box 1. Nevertheless, responses in hours worked could play a role when analysing self-employed, but due to a lack of data on the hours worked by self-employed, we are unable to test this hypothesis. Our results indicate that the shifting of deductions, particularly between partners, is the key channel of bunching behaviour in the Netherlands.

2.6 Concluding Remarks

We implemented a purely data-driven procedure to find an optimal, potentially asymmetric bunching window. Applying this extension to our data, we found elasticities that are quantitatively similar, yet more precisely estimated than elasticities estimated with the traditional approach. Our modification thus forms a valuable contribution to the literature, as it allows for a more precise calculation of the excess mass around the kink and the subsequent estimation of the ETI.

We have estimated the elasticity of taxable income with respect to the net-of-tax rate in the Netherlands. Using a unique tax data set on indiviudals from Statistics Netherlands, containing exact taxable incomes, we exploited bunching behaviour at kink points in the Dutch tax schedule to identify the ETI. We found an elasticity of 0.023 in the full sample at the upper threshold of the tax system, where the tax change is largest. Women respond more to taxation and our analyses suggest that their response is mainly driven by working in part-time employment. Self employed individuals also respond more, which is in line with better (legal and illegal) adjustment opportunities. Unlike Chetty et al. (2011) and Bastani and Selin (2014), we find a statistically significant ETI for wage earners in the Netherlands at some parts of the income distribution.

Explorations into the anatomy of responses by wage earners revealed that bunching is caused by shifting tax deductions, especially the mortgage interest deduction, between joint filers. The shifting is facilitated by the digital filing of tax returns, which makes the thresholds more salient to the general public. Our results thus corroborate earlier studies claiming that the ETI is not a structural parameter but rather depends on institutional settings.

2.A Derivation of the Bunching Estimator

For illustration, assume that each individual has the following utility function

$$u(c, z) \qquad \text{with} \quad u_c > 0, \quad u_z < 0,$$

where $c = (1 - \tau_1)z$ is after-tax income used for consumption, z is before-tax income and τ_1 gives the marginal tax rate in the pre-reform period. Now, the marginal tax rate changes from τ_1 to τ_2 with $\tau_1 < \tau_2$ at a specific earnings level k so that all individuals with $z > k$ face a higher marginal tax rate.[28] Individuals earning an income less than k are not affected and therefore do not react at all. Those with $z > k + dz$ will reduce their taxable income, but they will not exactly bunch at k in response to the reform (see Saez (2010) for a more detailed illustration here). For all taxpayers between k and an income level $k + dz$ though, it is no longer optimal to supply the given amount of labour and they re-optimise their income to k, i.e. they bunch at k. This situation is depicted in Figure 2.A.1.

In the pre-reform period with a constant marginal tax rate τ_1, individual before-tax incomes z are smoothly distributed according to a smooth density distribution $h_0(z)$. Bunching behaviour induces a spike in the post-reform income distribution $h_1(z)$ and the density above the interval $(k,k + dz]$ shifts to k. Because of optimisation errors and imperfect control over taxable income, we are more likely to observe a mass around the threshold, rather than a sharp spike. This is depicted in Figure 2.A.2.

To derive the compensated elasticity of taxable income, we can use that for small tax changes, the mass of taxpayers bunching is given by:

[28] To abstract from any income effects, the utility function is assumed to be quasi-linear. However, simulations by Bastani and Selin (2014) show no economically significant bias of income effects on compensated elasticities, which would make it superfluous to assume such a utility function.

Figure 2.A.1: Utility under a Kinked Tax Schedule

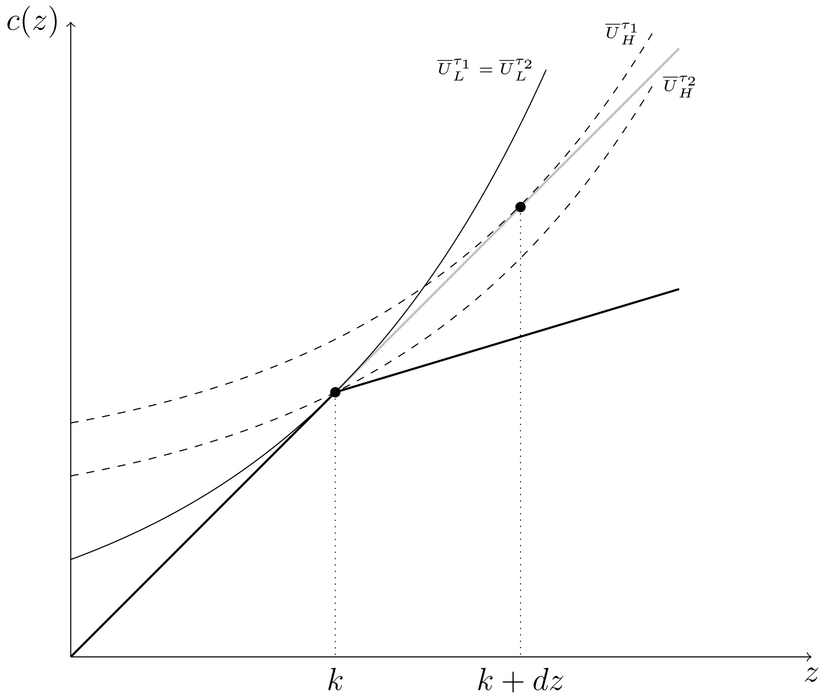

Source: Based on Saez (2010)

Figure 2.A.2: Imperfect Bunching

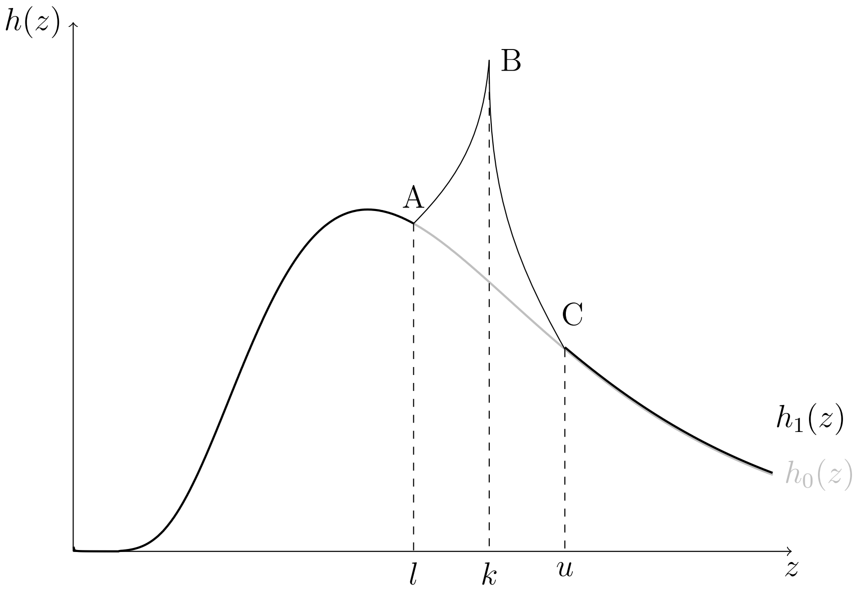

Source: Based on Weber (2014)

$$B = \int_{k}^{k+dz} h_0(z)\,\mathrm{d}z \approx h_0(z)\,\mathrm{d}z, \tag{2.10}$$

where the approximation follows from the mean value theorem for integration. Substituting $dz = \frac{B}{h_0(z)}$ from (2.10) into the elasticity formula given by

$$e(z) = \frac{\frac{dz}{z}}{\frac{d(1-\tau)}{(1-\tau)}} \tag{2.11}$$

at $z = k$ leads to

$$e(k) = \frac{B(dz)}{k \cdot h_0(z) \cdot \log(\frac{1-\tau_1}{1-\tau_2})}. \tag{2.12}$$

Thus, the elasticity of taxable income is non-parametrically identified by expression (2.12) if and only if the derivative $h_0(z)$ with respect to z is continuous in $z\ \forall\ z$, which means there should be no peak in the pre-reform distribution at the kink point. However, although the counterfactual pre-reform density is not observable in reality, in most applications it seems reasonable to hold on to that assumption if there are no obvious violations.[29] Note, that the expression in (2.12) shows that the elasticity parameter is proportional to the number of taxpayers who bunch at the kink point k.

While k and $\log(\frac{1-\tau_1}{1-\tau_2})$ from (2.12) are directly observable policy parameters, the relative excess mass of taxpayers defined by $b = \frac{B(dz)}{h_0(k)}$ has to be estimated. We follow the estimation approach developed in Chetty et al. (2011) and estimate the counterfactual density $h_0(k)$ – the density in absence of any kink – directly by local polynomial regression. However, for plausible reasons, the

[29]Besides, there exist some smoothness checks depending on the respective setting. Le Maire and Schjerning (2013) propose to exploit shifting of kinks over time to examine the smoothness of the density at the respective threshold in the after-shift period. Moreover, it seems plausible to assume smoothness if the distribution to the left and right of the bunching window does not show any jumps. Nevertheless, the assumption cannot be tested directly.

empirical post-reform distribution will not have a single spike at $z = k$ even if there is behavioural reaction in response to a tax change in the population. Therefore, the theoretical spike at k will become an empirical bunching window around k ("imperfect bunching"). Some of the reasons for imperfect bunching are uncertainty in income due to random income components (income volatility) or the inability to perfectly adjust labour supply in case of contracted hours constraints from the employer (Chetty et al., 2011).

Additionally, to account for changing thresholds over time when using taxable income from several years, we need to re-centre the income variable by calculating the difference between taxable income and the respective threshold before pooling the data. Individuals are then grouped into bins of length δ, with Z_j being the midpoint of the distance interval. Thus, Z_j is the absolute distance between income bin j and the threshold k where k differs over years (e.g., due to inflation adjustments). To estimate the counterfactual distribution, Chetty et al. (2011) propose a two-step procedure using a local polynomial regression while excluding all observations within the bunching window $[l; u]$:

$$\tilde{N}_j = \sum_{i=0}^{q} \beta_i \cdot Z^i + \sum_{i=l}^{u} \gamma_i \cdot I[Z_j = i] + \varepsilon_j \tag{2.13}$$

with N_j denoting the number of individuals in income bin j and q denoting the order of the polynomial. I is an indicator function equal to one if the bin point lies within the bunching window. ε_j denotes the error of the polynomial regression. An initial simple estimate of the number of bunching individuals is $\tilde{B} = \sum_l^u (N_j - \tilde{N}_j)$. However, this estimate would overestimate the excess mass as the integration constraint is not satisfied. The area under the counterfactual density would not sum to one as the observations within the bunching window (with an expected higher density) are left out from the regression. Therefore, the counterfactual density has to be shifted to the upper right for those observations that lie to the right of the bunching threshold. $\hat{B}$ is thus equally distributed on the total right tail of the income distribution following

$$\hat{N}_j(1 + I[j > u]\frac{\tilde{B}}{\sum_{u+1}^{\infty} N_j}) = \sum_{i=0}^{q} \beta_i \cdot Z^i + \sum_{i=l}^{u} \gamma_i \cdot I[Z_j = i] + \varepsilon_j. \tag{2.14}$$

To get the relative excess mass $\hat{b}$, the estimated bunching mass has to be related to the (average) height of the counterfactual density at the kink in the following way

$$\hat{b} = \frac{\hat{B}}{\frac{\sum_l^u \hat{N}_j}{u-l+1}}. \tag{2.15}$$

Along with the policy parameters k, expressed in terms of the binwidth (Bastani and Selin, 2014), and $\log(\frac{1-\tau_1}{1-\tau_2})$, the compensated elasticity of taxable income with respect to the net-of-tax-rate can be estimated by inserting all values into (2.12). With regard to the calculation of standard errors, we decided to use a parametric residual bootstrap for our analysis.

2.B Additional Graphs and Tables

Figure 2.B.1: Sample Income Distribution for Monte-Carlo Simulation

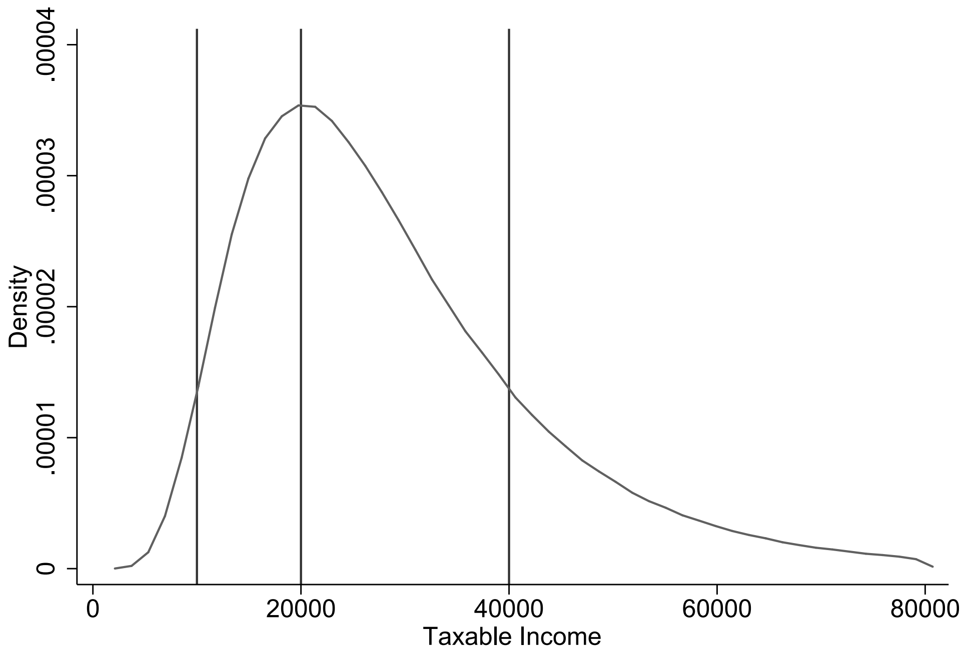

Notes: The figure shows the income distribution used in the Monte-Carlo simulations. The analysed thresholds at 10,000, 20,000 and 40,000 Euros are depicted by the vertical lines. The distribution is cut off at 80,000 Euros, which is roughly the 99th percentile of the simulated income distribution.

Table 2.B.1: Development of Thresholds

	First Threshold	Second Threshold	Third Threshold
2001	14,870	27,009	46,309
2002	15,331	27,847	47,746
2003	15,883	28,850	49,464
2004	16,265	29,543	50,652
2005	16,893	30,357	51,762
2006	17,046	30,631	52,228
2007	17,319	31,122	53,064
2008	17,579	31,589	53,860
2009	17,878	32,127	54,776
2010	18,218	32,738	54,367
2011	18,628	33,436	55,694
2012	18,945	33,863	56,491
2013	19,645	33,363	55,991
2014	19,645	33,363	56,531

Notes: The table reports the development of the thresholds of the Dutch tax system from 2001 to 2014.

Figure 2.B.2: Development of Marginal Tax Rates

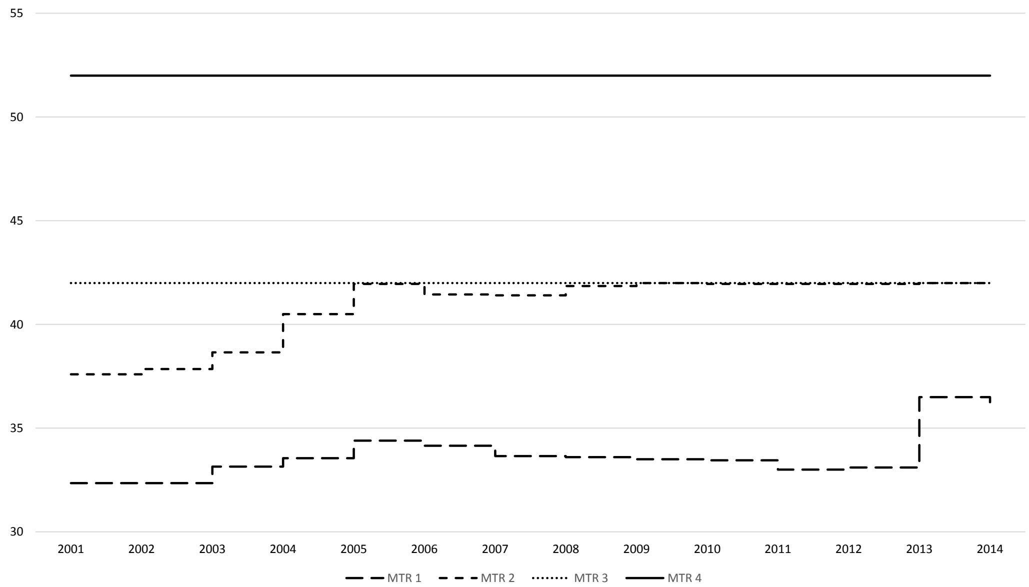

Notes: The figure depicts the development of the marginal tax rates in the Netherlands from 2001 to 2014.

Figure 2.B.3: Bunching at the Third Threshold - Large and Small Windows

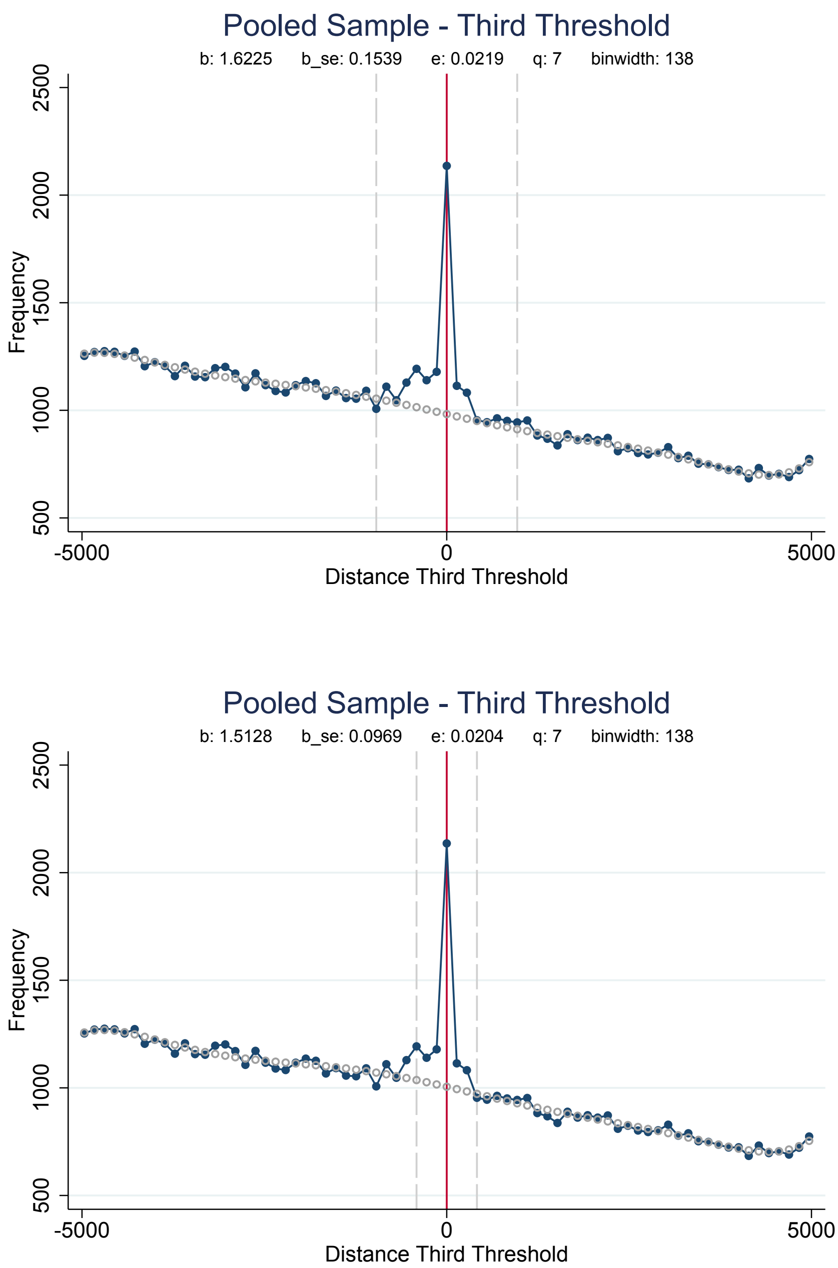

Notes: The figures show bunching at the third threshold using two alternative, symmetric bunching windows. The left graph shows the results using a large window and the right graph show the respective results for the small window.

Figure 2.B.4: Bunching at the Second Threshold - 2008 and 2009

Notes: The figures show (non-)bunching behaviour at the second threshold of the Dutch tax system for the years 2008 and 2009. Because the tax change was zero in 2009 and it is part of the denominator in the elasticity formula, no value for the ETI can be estimated in 2009.

Figure 2.B.5: Single Year Estimates 2003

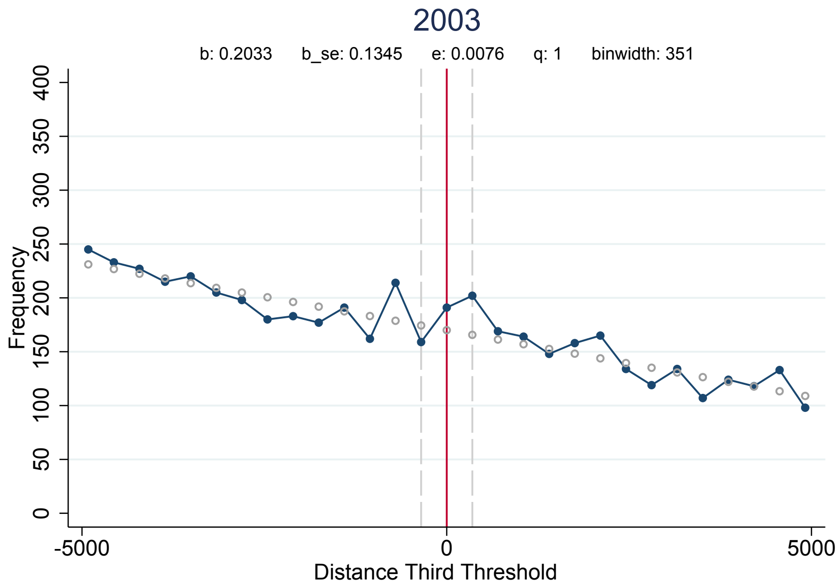

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2003.

Figure 2.B.6: Single Year Estimates 2004

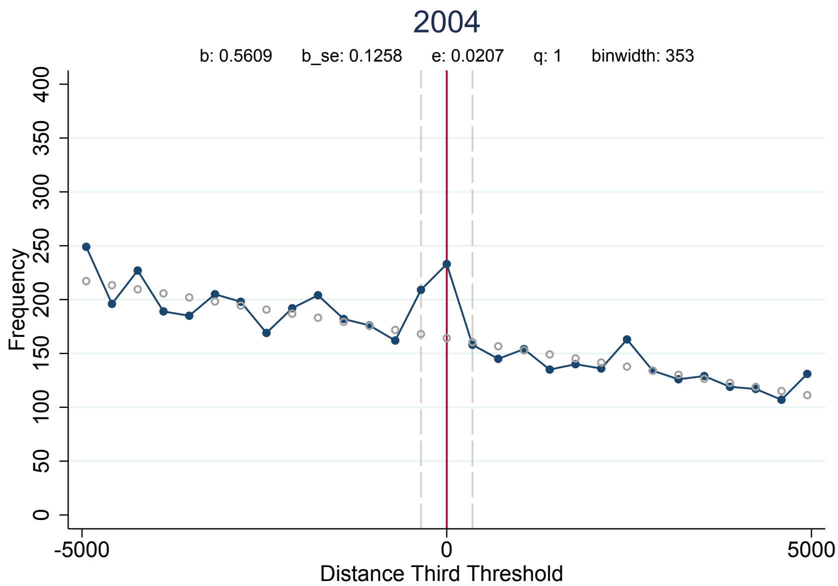

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2004.

Figure 2.B.7: Single Year Estimates 2005

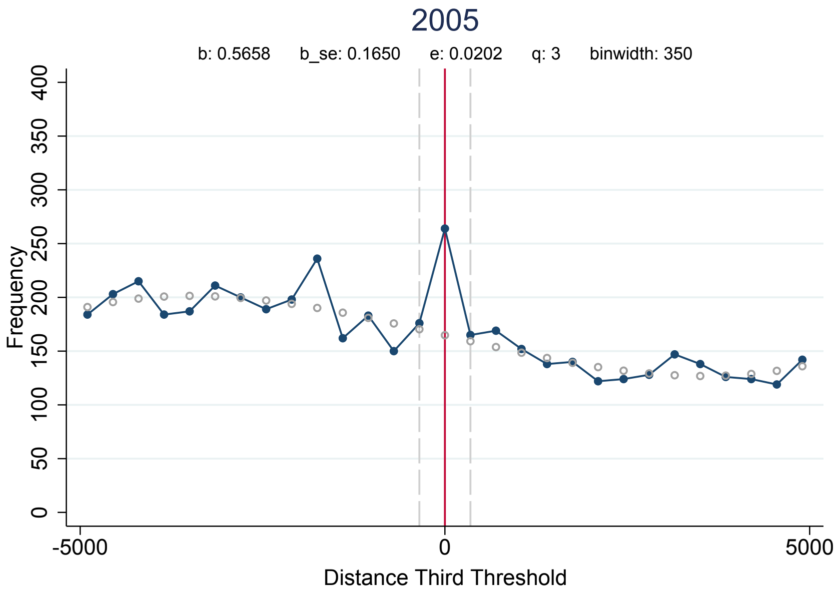

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2005.

Figure 2.B.8: Single Year Estimates 2006

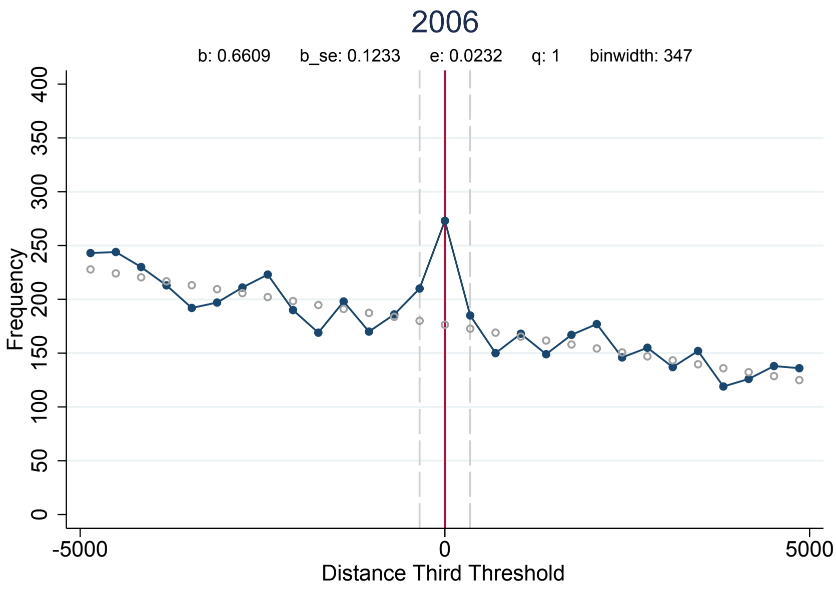

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2006.

Figure 2.B.9: Single Year Estimates 2007

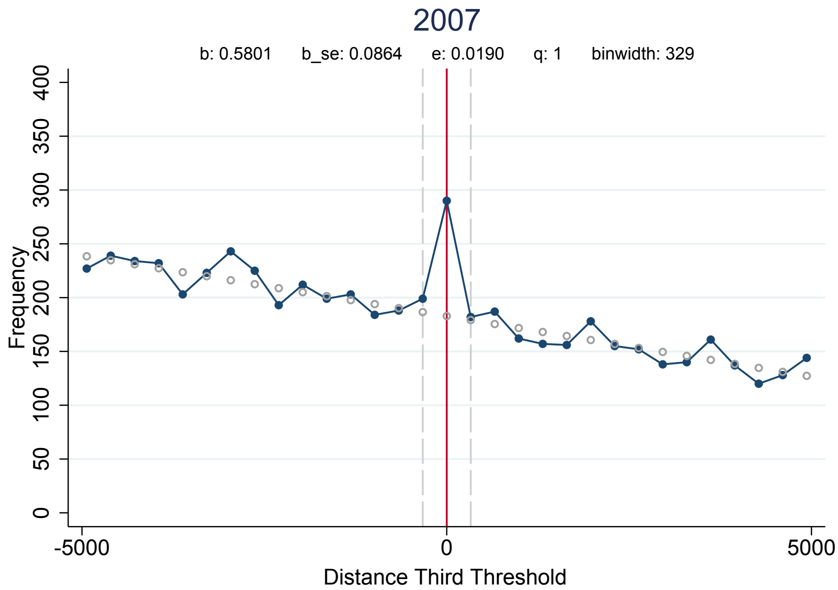

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2007.

Figure 2.B.10: Single Year Estimates 2008

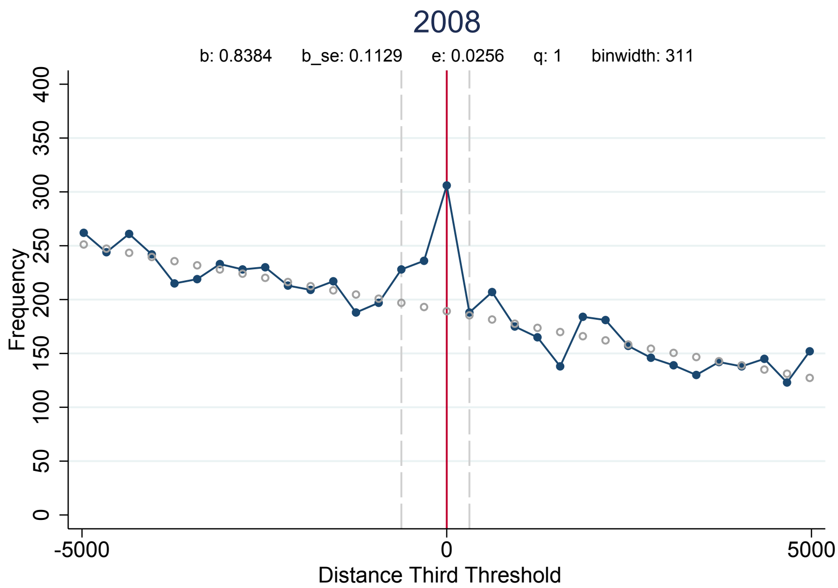

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2008.

Figure 2.B.11: Single Year Estimates 2009

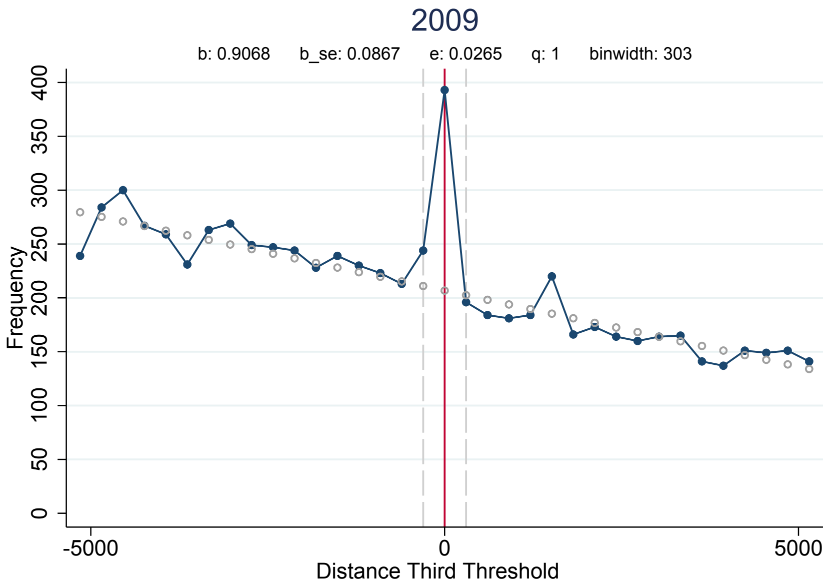

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2009.

Figure 2.B.12: Single Year Estimates 2010

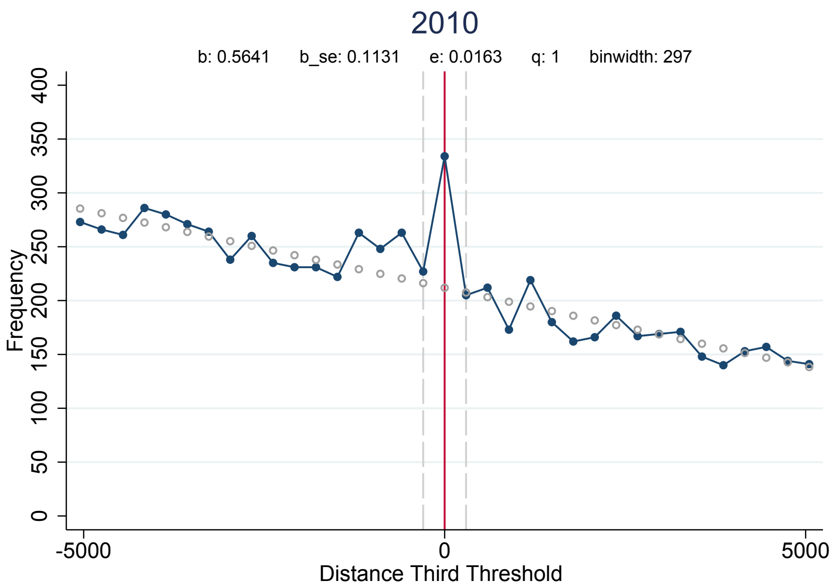

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2010.

Figure 2.B.13: Single Year Estimates 2011

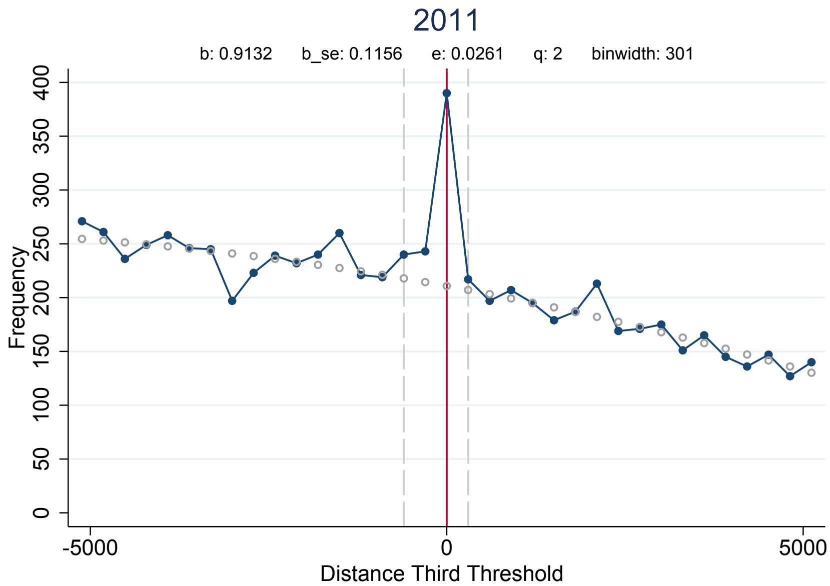

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2011.

Figure 2.B.14: Single Year Estimates 2012

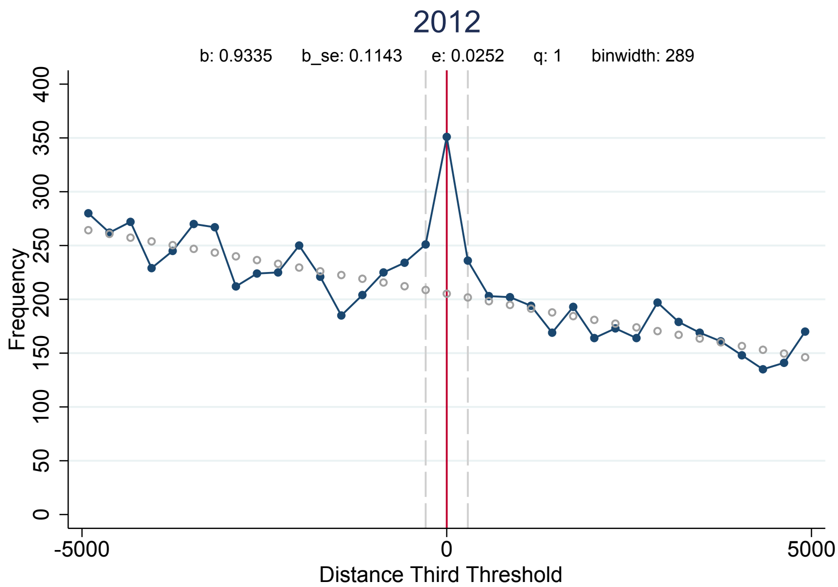

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2012.

Figure 2.B.15: Single Year Estimates 2013

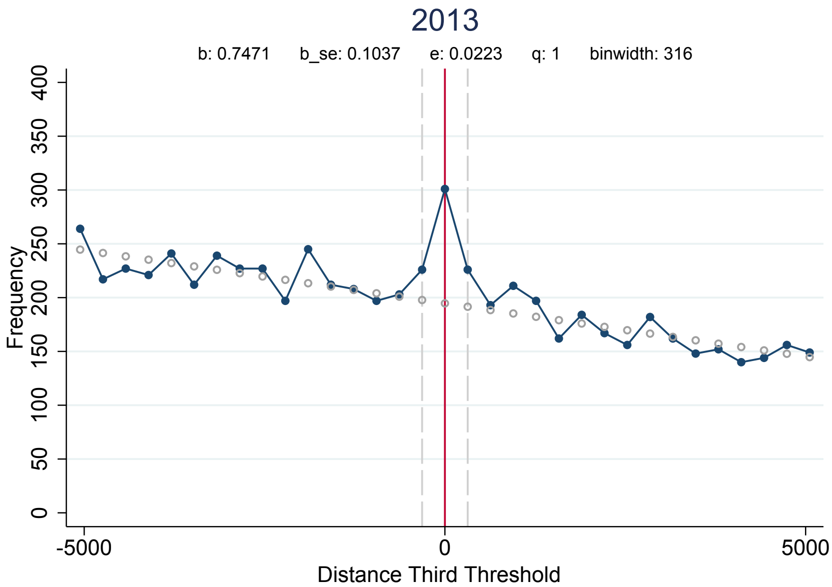

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2013.

Figure 2.B.16: Single Year Estimates 2014

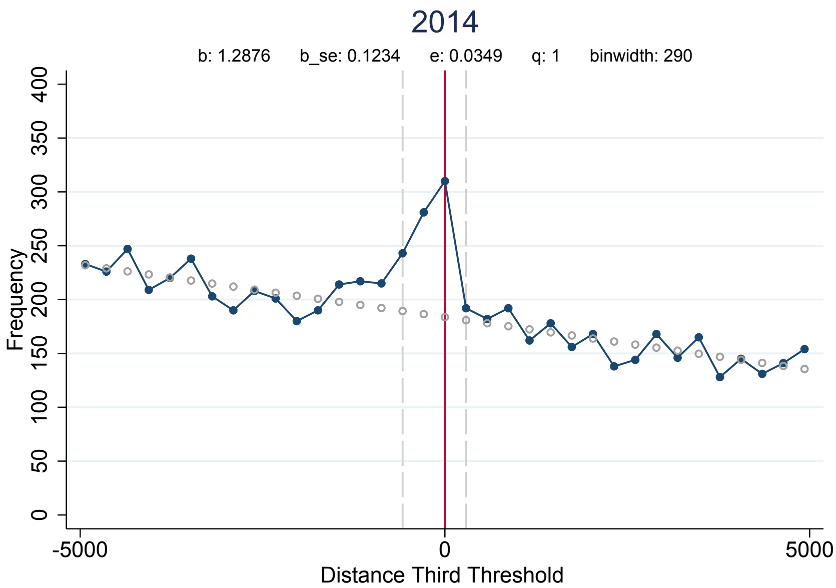

Notes: The figure shows the bunching behaviour at the third threshold of the Dutch tax system for the year 2014.

Figure 2.B.17: Subsamples by Gender - Threshold 1

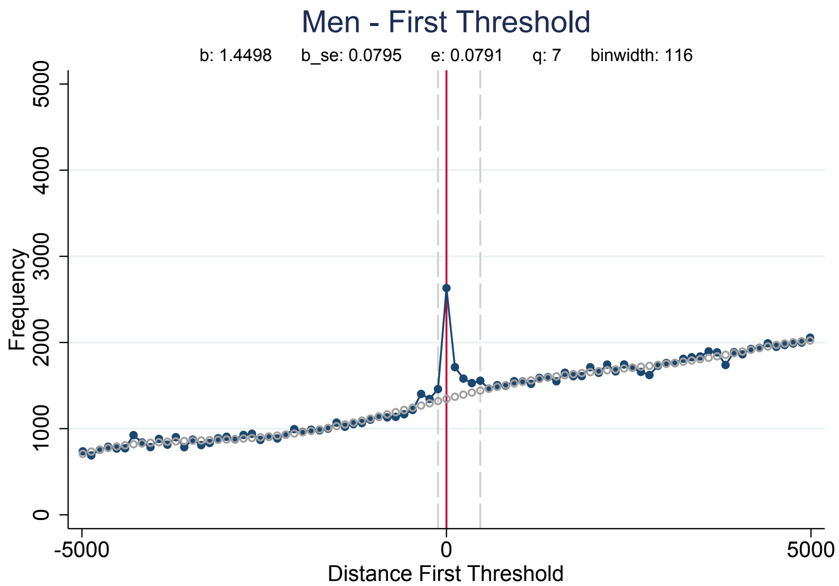

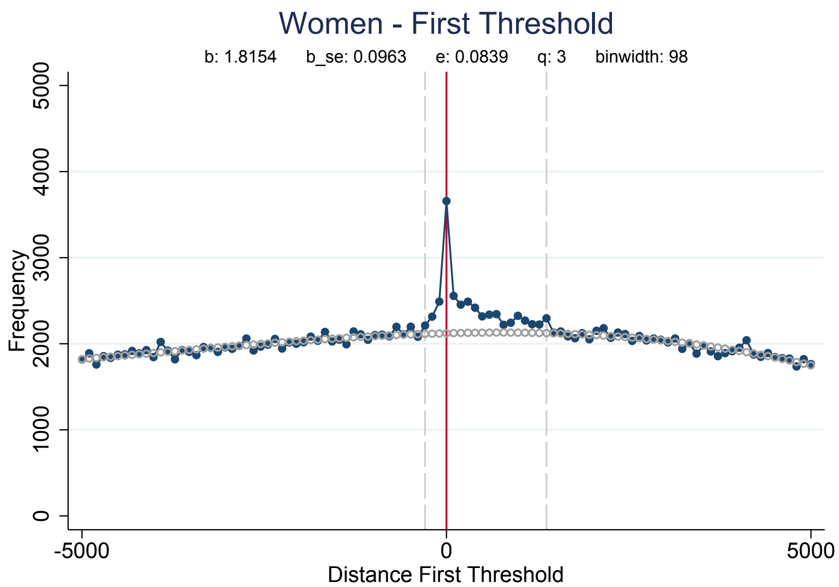

Notes: The figures show bunching behaviour by gender at the first threshold. The upper graph shows men, the lower graph shows women.

Figure 2.B.18: Subsamples by Gender - Threshold 2

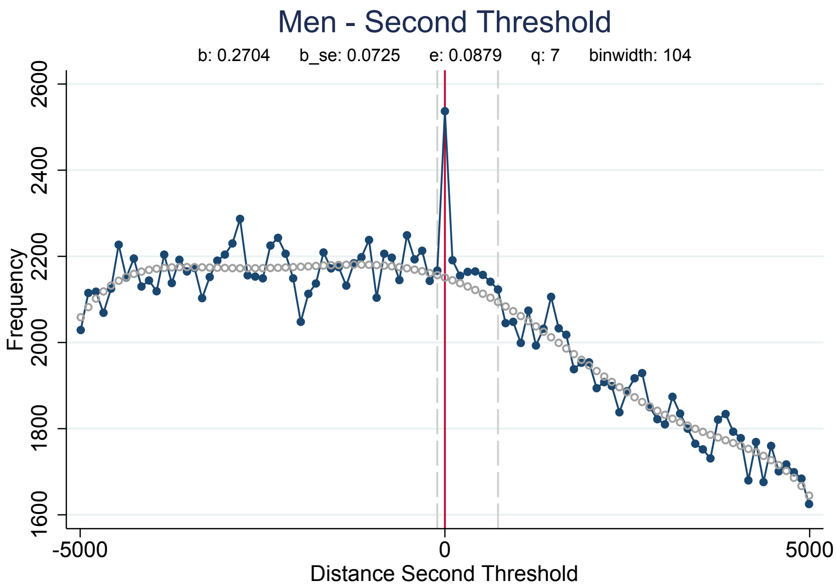

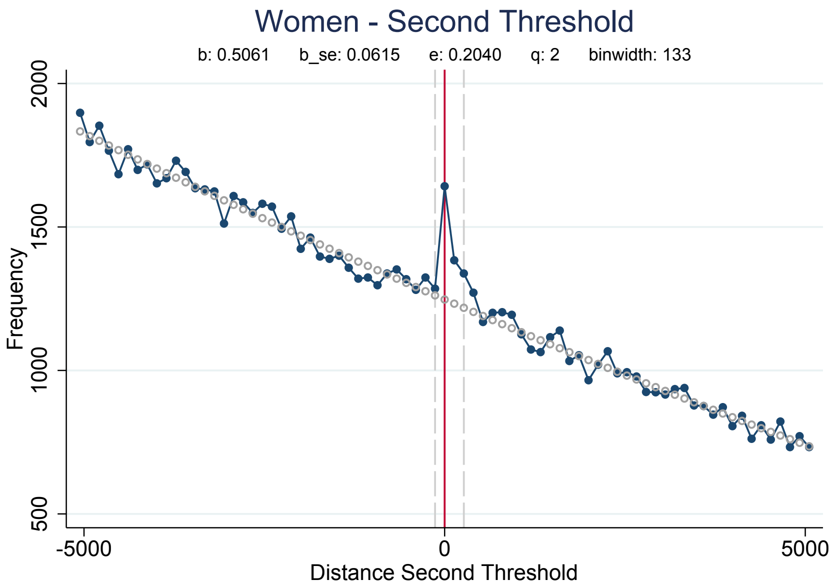

Notes: The figures show bunching behaviour by gender at the second threshold. The upper graph shows men, the lower graph shows women.

Figure 2.B.19: Subsamples by Gender - Threshold 3

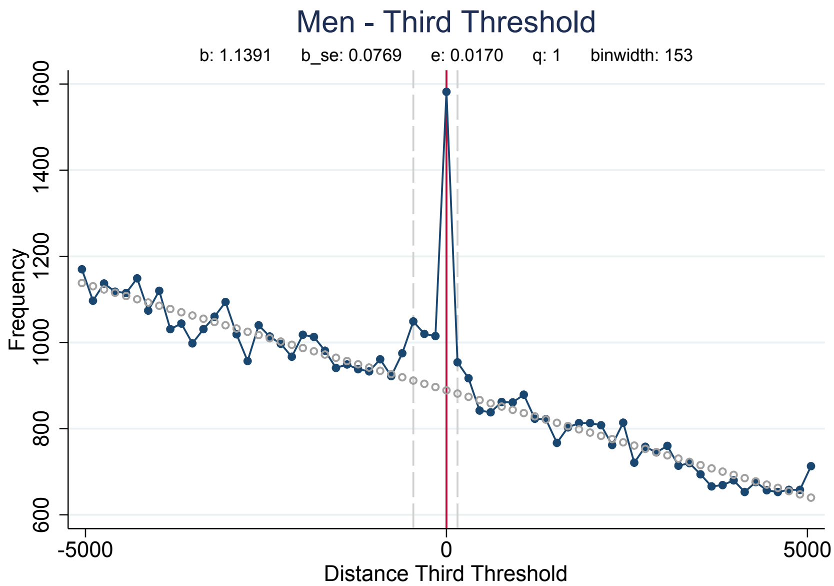

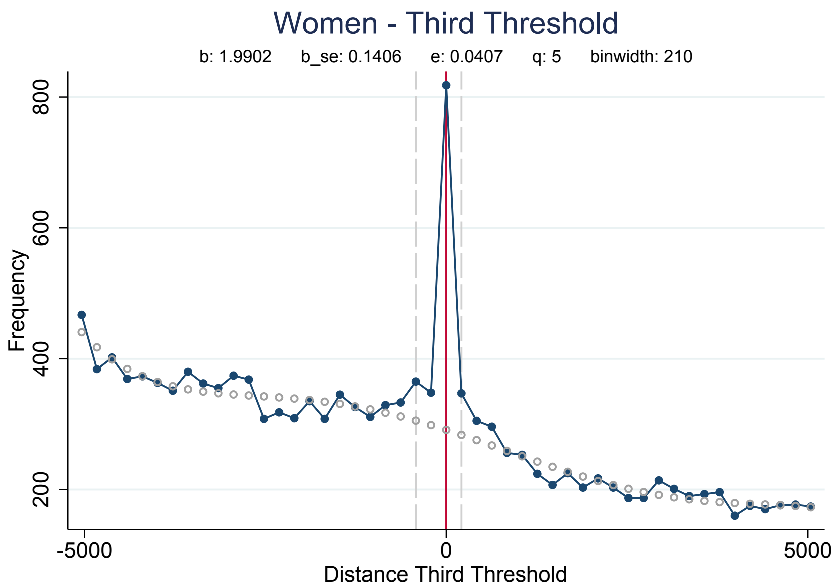

Notes: The figures show bunching behaviour by gender at the third threshold. The upper graph shows men, the lower graph shows women.

Figure 2.B.20: Subsamples by Employment Status - Threshold 1

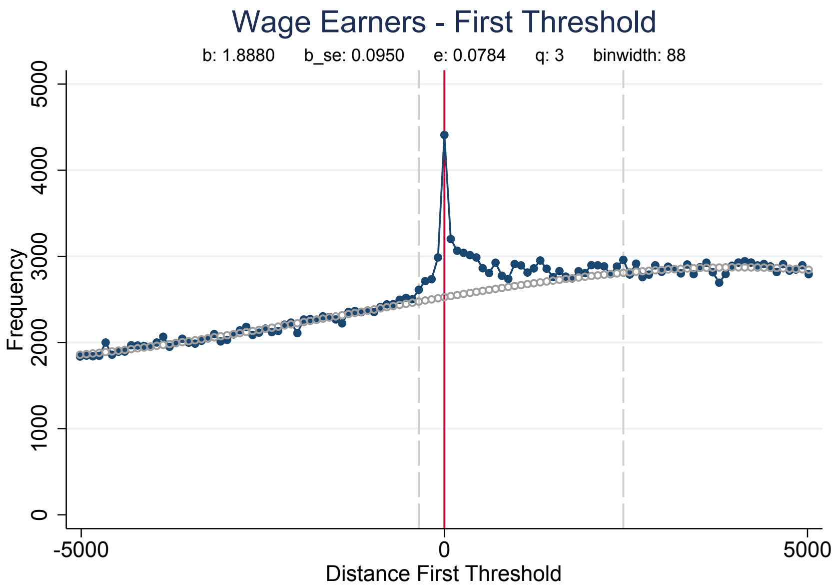

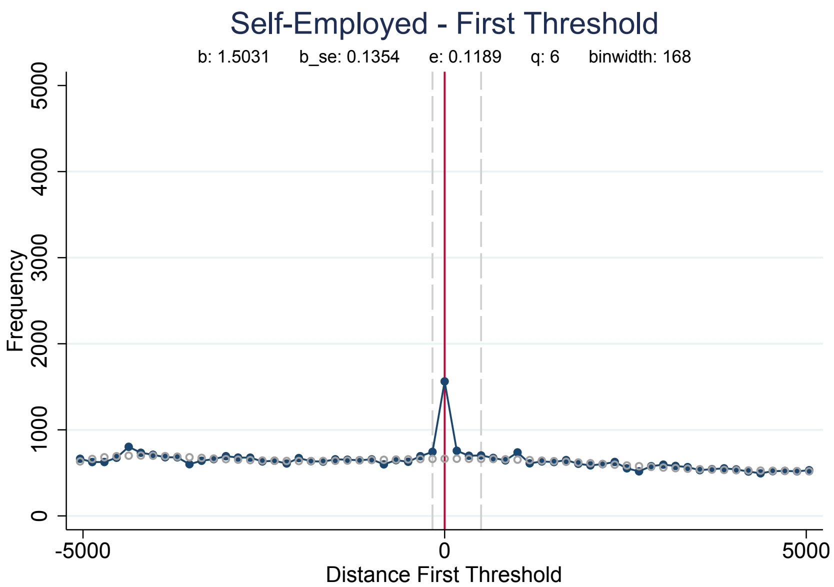

Notes: The figures show bunching behaviour by employment status at the first threshold. The upper graph shows wage earners, the lower graph shows self-employed individuals.

Figure 2.B.21: Subsamples by Employment Status - Threshold 2

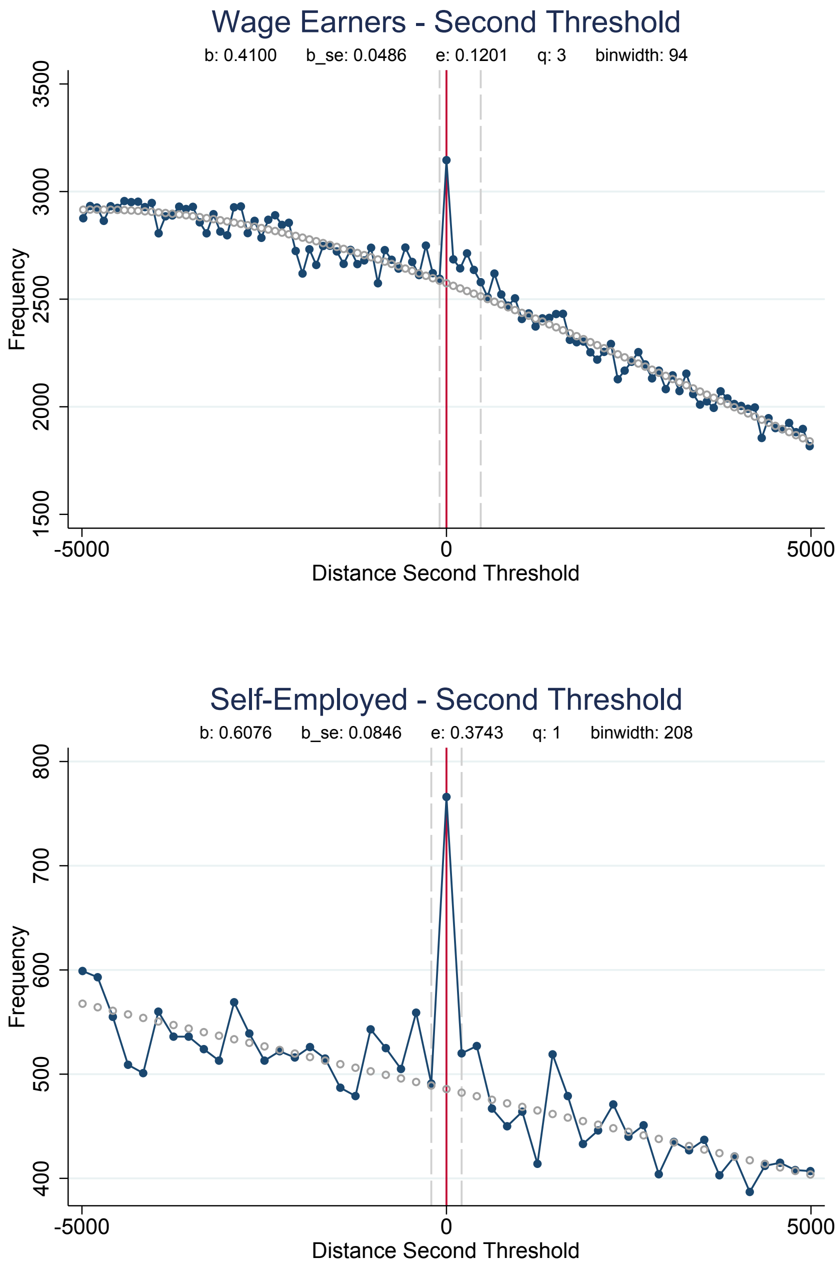

Notes: The figures show bunching behaviour by employment status at the second threshold. The upper graph shows wage earners, the lower graph shows self-employed individuals.

Figure 2.B.22: Subsamples by Employment Status - Threshold 3

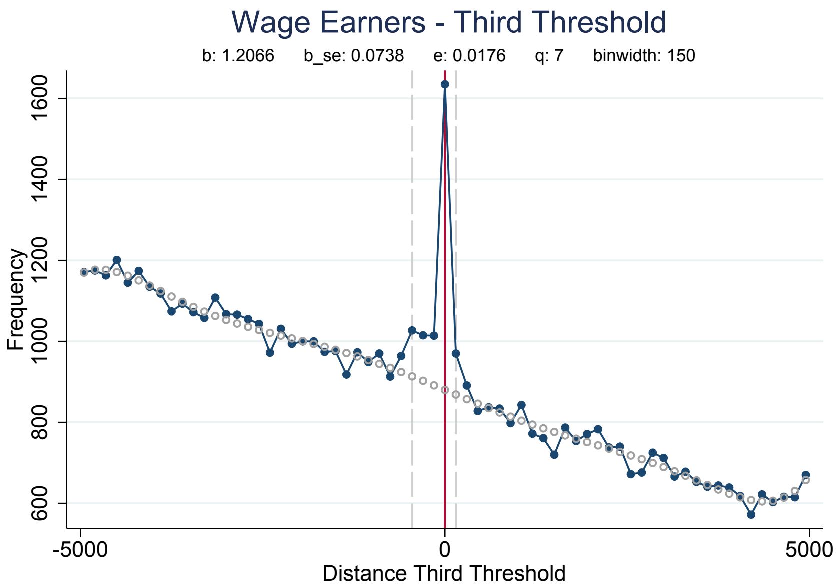

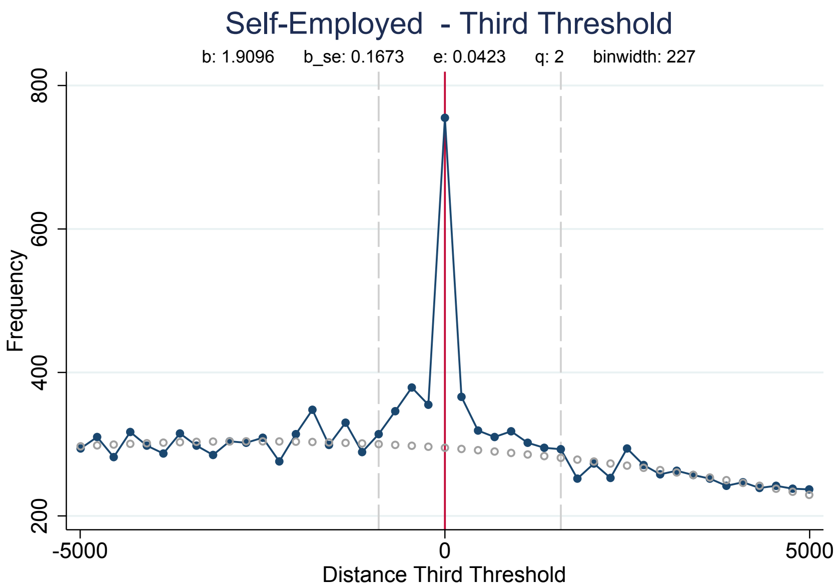

Notes: The figures show bunching behaviour by employment status at the third threshold. The upper graph shows wage earners, the lower graph shows self-employed individuals.

Chapter 2 showed how individuals react to changes in the marginal tax rate. The estimated elasticities of taxable income were low, indicating that individuals do not react excessively to increases in the MTR. This may be due to optimisation frictions, hours constraints or a lack of understanding of the tax system. Especially for financially constrained individuals, costly tax optimisation is not feasible and a combination of the aforementioned factors prevents large reactions to increases in the MTR. Firms on the other hand are assumed to be less constrained than individuals and more concerned with maximising shareholder value, thus also more likely to engage in tax optimisation. This behaviour intensifies with the size of the firm and is especially pronounced for MNEs that have the possibility to utilise tax differences across countries. As described in the introduction, the corporate tax rate of the countries within the EU varies between 10% and 35% as well as the tax code, which creates significant incentives to shift income between jurisdictions.

To show this, Chapter 3 investigates the introduction of transfer pricing regulations (TPR) by analysing data on the value and quantity of exports. TPR aim at stopping the misuse of transfer prices to minimise (maximise) taxable income in high-tax (low-tax) jurisdictions. Prior to the introduction of TPR, firms have an incentive to utilise the tax rate differentials that exist between two countries. Because a distorted world is observed prior to the introduction of TPR, I expect to find counterbalancing reactions to their introduction, i.e. the overvaluation (undervaluation) of exports to high-tax (low-tax) jurisdictions should be reduced. The pricing effect of TPR has been studied extensively in previous research, but evidence on quantitative reactions that hint at allocative inefficiencies is scarce. The chapter intends to tackle this problem and provide a rationale for the effects that the quantitative responses might have for the efficiency of the tax system.

Chapter 3

The Effect of Transfer Pricing Regulations on Intra-Industry Trade

3.1 Motivation

n October 2015, the OECD announced the Base Erosion and Profit Shifting (BEPS) action plan, which aims at taxing corporations at the location where production takes place and preventing MNEs to internationally shift profits. Multinational firms have several possibilities to shift profit to affiliates in low tax countries. One of the predominant tax planning strategies is the manipulation of intra-firm prices. Many countries have introduced transfer pricing regulations to keep corporate profits within their borders. Overall, these laws attempt to tie intra-firm prices to arm's-length equivalents. The empirical literature has shown that firms manipulate transfer prices for tax optimising purposes (Bartelsman and Beetsma, 2003; Huizinga and Laeven, 2008; Clausing, 2003; Cristea and Nguyen, 2013; Davies et al., 2015) and that they react sensitive to transfer pricing regulations (Lohse and Riedel, 2013). While the literature has provided evidence for the manipulation of intra-firm transfer prices exploiting the pricing mechanism, little is known about how trade volumes are affected.

The main objective of this paper is to analyse the quantity effects on trade flows that transfer pricing regulations might have. Under the premise that intra-firm prices did indeed deviate from their corresponding arm's-length prices, we assume that prices are adjusted towards the arm's-length price when transfer pricing regulations are introduced. Although this change in prices influences the profits and therefore the tax bases of firms in the respective countries, which on its own could lead to severe economic consequences, we argue that a change in trade volume could amplify this effect. In other words, only considering the pricing channel would lead to an underestimation of the effectiveness of transfer pricing regulations. In the extreme case, if an affiliate was only held for tax optimising purposes and intermediate goods were merely imported and reexported, reoptimisation by the multinational would lead to the abandonment of the affiliate in that country. An effect of transfer pricing regulations on traded quantities also hints at allocative inefficiencies that are induced by taxation.

For a long time, the empirical literature on anti-avoidance strategies by-and-large focusses on analysing one measure at a time. Broadly, three instruments can be defined: Transfer pricing regulations (TPR), thin capitalisation rules (TCR) and controlled foreign company (CFC) regulations. Among others, Lohse and Riedel (2013), Zinn et al. (2014) and Beer and Loeprick (2015) focus on TPR, Buettner et al. (2012), Merlo et al. (2015) and Mardan (2017) discuss TCR and Voget (2011) as well as Buettner and Wamser (2013) analyse CFC rules.

Only recently, the literature has started to analyse the interplay of different anti-avoidance measures. For instance, the BEPS Action Plan 3 states in its final report that CFC rules somewhat complement transfer pricing regulations and address the same income, but neither policy is able to fully substitute the other (OECD, 2015). Therefore, considering both TPR as well as CFC rules could lead to a different interpretation of the effectiveness of TPR. This is because part of the effect attributed to the introduction of TPR could stem from the effective CFC rules already in place. Schindler and Schjelderup (2013) provide a theoretical framework for the interdependency of debt shifting and transfer pricing strategies. Empirical evidence is limited to two studies that explicitly address the interconnectivity of different profit shifting activities. Based on firm-level data, Saunders-Scoot (2015) finds that the two most prominent profit shifting channels of debt shifting and transfer pricing are substitutes. Nicolay et al. (2016) provide evidence for Europe that substitution effects between both channels exist but that TCR are ineffective in limiting profit shifting behavior if strict TPR are missing. However, to the best of our knowledge, there is no study examining whether anti-profit-shifting activities affect traded volumes.

We base our estimation strategy on the workhorse model of international trade, the gravity model. First applied by Tinbergen in 1962, it shows that trade between two countries can be explained by their relative force of gravity, using GDP as the weight of the countries and distance between them as a negative correction. Baltagi et al. (2003) develop a panel data gravity model that we adopt here for the use of intermediate goods trade. Transfer pricing regulations enter the model as one form of trade costs. We estimate the model using UN COMTRADE data on bilateral exports of intermediate goods in the car industry for the period 1995 until 2012. The car industry is characterised by high specialisation and a low share of trade between unrelated third parties, thus providing ample opportunities to manipulate transfer prices. The approach taken in this paper allows us to analyse the effect at the level where the variation in regulations takes place. Because not all trade activities between countries can be attributed to profit shifting behaviour, it is important to keep in mind that we will not find results at the extensive margin of trade but rather at the intensive margin.

Our main findings are as follows: Under the assumption that transfer prices have been manipulated, the introduction of TPR reduce (increase) trade with countries that have a lower (higher) tax rate than the exporting country. The effect is driven by the size of the tax rate differential. We find evidence for the interplay of different anti-avoidance measures, especially TPR and CFC. In a back-of-the-envelope calculation utilising the value of trade, we can show that the pricing reaction in our data is in line with the literature and that the quantity effects found in our study amplify the reaction of firms to TPR.

The rest of the paper is structured as follows: We first describe our data and derive our hypotheses, before explaining the estimation strategy. Section 3.5 presents and discusses the results. The final section concludes.

3.2 Data

The data used in this study comes from the UN COMTRADE database as harmonised by the CEPII in their BACI database. It provides bilateral trade data at a disaggregated, 6-digit HS goods classification level. We observe all intermediate goods, which enter the production for motorised vehicles weighing less than 3.5t.[1] Considering intermediate goods trade has an appealing advantage over trade in final goods: It allows us to abstract from demand shocks that are less pronounced compared to final goods. We focus on the car industry, which relies on highly specialised intermediate goods and is characterised by frequent

[1] HS-codes ranging from 870600 to 870899.

Table 3.1: Transfer Pricing Regulations

Country	**TP Regulations**
Brazil	1997
Canada	1998
China	2008
France	1996
Germany	2003
India	2001
Italy	2010
Japan	1986
Mexico	1997
Russia	–
South Korea	1996
Spain	2006
United Kingdom	1999
United States	1994

trade between related parties. Therefore, the manipulation of transfer prices is relatively easier and we expect a significant reaction to transfer pricing regulations. The data were merged with information on corporate tax rates coming from Loretz (2008) as well as KPMG. Data on GDP and economic integration were taken from the World Bank and the EIA Database respectively.

We hand collected data on introduction of transfer pricing regulations from Deloitte (2015); Ernst&Young (2014); KPMG (2014) and PWC (2016). For some countries, the publications offer different years of enactment of transfer pricing regulations and we chose the most common. When several dates were offered in one publication, we chose the point in time where mandatory documentation requirements came into place. Table 3.1 shows the years in which the exporting countries in our sample introduced transfer pricing regulations.

We concentrate on data from 1995 until 2012, a period in which many countries introduced transfer pricing regulations, which gives us the variation for identification. We focus on the 14 most important exporting countries with the largest car industries as measured by production in 1995 and assess their trade with all other countries of the world. This gives us confidence that our model captures a great share of worldwide intermediate goods trade in the car industry.[2] The BACI data provides us with information on actual trade flows, but omits zero trade flows. For our analysis however, it is important to also account for trade flows that did not take place, as these could potentially be caused by the existence of transfer pricing regulations. Therefore we rectangularise our data set so that we have observations for each exporter-importer-year combination. All new observations are assigned a quantity and a value of zero, assuming that if we do not observe a trade flow, there was none. In an extension, we could explore the possibility of misreporting by the exporting country by analysing imports of the partner country, which we can use to verify our procedure. We are left with a total of 153 importing countries for which we have obtained all variables.

3.3 Hypotheses

To stress the quantity effect and why allocative distortions could exist, consider a representative MNE active in the three countries A, B and C. To produce the final good, only trade in intermediaries from A to B is required, but for tax optimising purposes, intermediate goods are exported from A to C and then from C to B. This situation can arise, whenever tax rate differences exist between countries. Figure 3.1 shows such a situation. In order to minimise the tax burden, the MNE would like profits to accrue in country C. This can be

[2]The 14 countries had a share in excess of 95% of total car production in 1995(`https://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/national_transportation_statistics/html/table_01_23.html_mfd`. Last accessed 18.06.2017).

Figure 3.1: Three Country Model with Tax Rate Differences

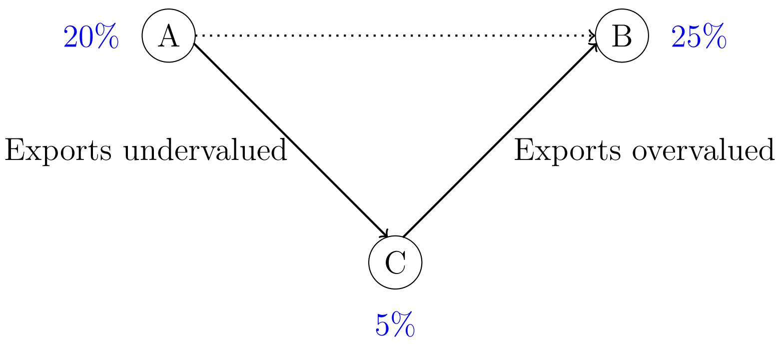

facilitated by undervaluing exports from A to C and overvaluing exports from C to B.

Now suppose country A introduces transfer pricing regulations. Because C has a lower tax rate, the exports from country A to country C are undervalued to minimise profits in A and maximise profits in C. The regulations will force companies in A to increase their prices, therefore again making the trade via C relatively more expensive. This is the situation that is analysed in the empirical application of this paper.[3] This simple illustration shows that no matter which country imposes transfer pricing regulations, there are incentives for firms to alter the volumes traded between countries, given sufficiently high trade costs.

A priori, we would expect MNEs to utilise tax rate differences and therefore export more to countries that have a lower tax rate than the home country. This would indicate that we should see a negative coefficient for the tax rate difference, when it is non-positive and a positive coefficient when the tax rate difference is positive. The introduction of transfer pricing regulations should

[3] Transfer pricing regulations in B will require the company in B to pay an arm's-length price for the intermediate good imported from C. This will increase profits in B and decrease profits in C, which also lowers the attractiveness of exporting via country C. Country C benefits from the manipulation of transfer prices and therefore has no incentive to introduce TPR.

bring the effects closer together and we expect this to be strongly driven by the size of the tax rate difference. This corresponds to a positive coefficient on the interaction of tax rate difference and transfer pricing regulations for non-positive tax rate differences and a negative effect for this interaction for importers with a lower tax rate than the exporting country.

3.4 Estimation Strategy

The gravity model of trade was first introduced by Tinbergen (1962). It applies Newton's law of gravity from physics to economics and explains trade between two countries as proportional to their economic size (GDP) and inversely proportional to the distance between the two countries. Although the gravity model in its simplest form is able to explain a substantial amount of trade already, several extensions to the base model exist, most notably the introduction of trade costs as a second deterring factor of trade and the inclusion of multilateral resistance terms (Anderson and Van Wincoop, 2003). In empirical applications, these are controlled for by adding country fixed effects to the estimation.

We analyse the relationship between transfer pricing regulations and trade flows using a gravity-type model of the following form:

$$X_{ijt} = \alpha l_{it}^{\lambda} m_{jt}^{\mu} t_{ijt}^{\beta}, \tag{3.1}$$

which follows Baltagi et al. (2014) and where X_{ijt} denotes the exports (quantity or value) from country i to country j at time t. Baltagi et al. (2014) define l_{it} and m_{jt} as exporter-time-specific and importer-time-specific factors respectively. t_{ijt} broadly reflects trade costs that can possibly vary across all dimensions and λ, μ andβ are measures for the partial elasticity of trade flows with respect to the respective parameters.

To estimate the model, we impose a logarithmic transformation on Equation (3.1) and follow Baltagi et al. (2014) in defining the following set-up for our estimation equation:

$$ln(X_{ijt}) = \alpha + \mathbf{t}'_{\mathbf{ijt}}\beta + \lambda l_{it} + \mu m_{jt} + u_{ij} + \delta_t + \varepsilon_{ijt}, \tag{3.2}$$

where u_{ij} are country-pair fixed effects, δ_t represent time fixed effects and ε_{ijt} is an error term. To avoid endogeneity problems, it is important that the equation consists of some components that are exporter-time-specific and importer-time-specific, i.e. l_{it} and m_{jt} from Equation (3.1) respectively. Baltagi et al. (2003) propose a generalised model that accounts for this endogeneity by adding exporter-time and importer-time fixed effects instead of the simple year fixed effects in Equation (3.2). A drawback from this approach, however, is that it does not allow to separately identify the effect of time-invariant variables. In our application, the variable of interest - transfer pricing regulations - varies at the exporter-year or importer-year level, that is it would drop from the estimation due to collinearity. We account for exporter-time-specific and importer-time-specific factors by including control variables that are either independent of the i dimension, representing importer-year-specific factors, or the j dimension, representing exporter-year-specific factors. Furthermore, trade costs are represented by the tax rates and a transfer pricing regulation dummy, as well as an interaction term between the two. Our estimation equation thus takes on the following form:

$$\begin{aligned} ln(X_{ijt}) = &\alpha + \beta_1 \ln \tau_{it} + \beta_2 \ln \tau_{jt} + \beta_3 TPR_{it} + \beta_4 \ln \tau_{it} \cdot TPR_{it} + \\ &EIA_{ijt} + \lambda \ln GDP_{it} + \mu \ln GDP_{jt} + u_{ij} + \delta_t + \varepsilon_{ijt}, \end{aligned} \tag{3.3}$$

where τ_{it} and τ_{jt} represent the statutory corporate tax rates in country i and j

at time t, respectively, TPR_{it} is a dummy indicating whether transfer pricing regulations were in place in the exporting country at time t and EIA_{ijt} controls for the strength of economic integration between i and j through a set of dummies. GDP_{it} and GDP_{jt} are GDP in the exporting and importing country at time t.[4]

A regression of Equation (3.3) using the full sample could be problematic, because the incentives to shift profits (and possibly quantities) differ depending on the sign of the tax rate difference between two trading partners. If a positive tax rate difference exists, firms in country i will have an incentive to shift profits to country j by manipulating the transfer price downwards, whilst in the case that a negative tax rate difference exists, the transfer prices should be manipulated upwards. Looking at both cases together could cancel out any effects of transfer pricing regulations, as we expect opposite reactions depending on the sign of the tax rate difference. We therefore split the sample into cases where $\tau_i > \tau_j$ (positive tax difference) and cases where $\tau_i < \tau_j$ (negative tax difference) and explicitly exclude the case where the tax rates are equal, as profit shifting opportunities only arise, when a tax rate difference can be utilised.

When analysing worldwide trade, the number of country pair fixed effects to be estimated increases rapidly. Also, log-linearisation of the model that is common in the literature could lead to biased estimates, for example through the mishandling of zero trade flows. Silva and Tenreyro (2006) propose a Poisson Pseudo Maximum Likelihood (PPML) estimator that incorporates the multilateral resistance terms and circumvents the problems arising from log-linearising the model. The results suggest some fundamental differences with parameters estimated using the traditional fixed effects method on log-linearised data. For

[4] We do not include distance or other time-invariant country pair characteristics that are familiar from earlier gravity models, because they are collinear to the country pair fixed effect u_ij.

example, the effect of GDP is not close to one but significantly lower and the effect of geographical distance as well as colonial ties are greatly exaggerated in the classical log-linearised model. We will therefore report results of estimating a traditional fixed effects log-linear gravity model as well as a PPML model and discuss potential differences between the models.

3.5 Results

3.5.1 Baseline Results

The baseline results of estimating Equation (3.3) are shown in Table 3.2. The observational unit is a bilateral trade flow from exporting country i to importing country j. In total, we observe $N = 26,419$ such pairs. In the first two columns, we regress the logarithm of quantity only on the logarithm of the tax rates in i and j as well as the control variables.[5] Column (1) shows the results for all country pairs, where the exporting country i has a higher tax rate than the importing country j. The positive and significant coefficient on the own tax rate indicates on the one hand that countries with higher tax rates tend to export more and on the other hand also shows that ceteris paribus i.e. for a given tax rate of the importing country, an increase in the exporting country's tax rate leads to more trade. From an economic point of view, this is a plausible result, because as we are looking at the sample where $\tau_i > \tau_j$, any increase in τ_i will ultimately increase the tax difference and therefore the incentive to shift profits. The reverse is true for the case where we have a negative tax difference ($\tau_i < \tau_j$) and this can be seen in column (2) with the positive coefficient for τ_j. The cross

[5]Throughout all estimations in this section, the coefficients for $\ln GDP_{i/j}$ are positive and significant, as suggested by economic theory. We cannot report estimates for other common control variables in the trade literature such as distance, common language or contiguity, because they are captured by the country-pair fixed effects.

Table 3.2: The Effect of Transfer Pricing Regulations on Quantity

	(1) $\tau_i > \tau_j$	(2) $\tau_i < \tau_j$	(3) $\tau_i > \tau_j$	(4) $\tau_i < \tau_j$	(5) $\tau_i > \tau_j$	(6) $\tau_i < \tau_j$
	Dependent Variable: Log of Exports					
$\ln \tau_i$	1.472***	-0.618	1.770***	-0.644	3.357***	-0.750
	(0.538)	(1.836)	(0.543)	(1.825)	(0.604)	(1.809)
$\ln \tau_j$	-0.365	3.215**	-0.380	3.190**	-0.289	4.461**
	(0.478)	(1.323)	(0.474)	(1.313)	(0.459)	(1.751)
TPR_i			0.154***	0.062	2.545***	0.871
			(0.045)	(0.085)	(0.309)	(0.604)
$\ln \tau_i \cdot TPR_i$					-7.984***	
					(1.009)	
$\ln \tau_j \cdot TPR_i$						-2.583
						(1.826)
N	19,300	7,119	19,300	7,119	19,300	7,119
R^2	0.227	0.123	0.228	0.124	0.237	0.124
$F - Test$	–	–	–	–	0.000	0.035

Notes: Fixed effects regressions of Equation (3.3). Cluster-Robust standard errors on the country-pair level are in parentheses. All estimations include the logarithm of GDP for both countries, a set of dummies controlling for economic integration and year fixed effects as control variables. Coefficients are omitted for brevity. F-Test shows the p-value for a test of joint significance of the tax rate and the interaction term. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

tax rates are both negative, which further enhances the effect of the tax rate difference, but both coefficients are statistically not distinguishable from zero, thus suggesting that trade flows depend more on the country with the higher tax rate. The results furthermore show that trade flows differ in a world where taxation exists from trade flows in a world without (distortive) taxation.

In columns (3) and (4), we introduce a dummy indicating whether TPR were in place in country i at time t. In case of a positive tax rate difference, transfer pricing regulations lead to a slight increase in trade. This seems puzzling at first, because we would have assumed that more quantity was traded than optimally required, prior to the introduction of TPR and therefore, a reduction of the traded quantities should be expected. However, the reduction should take place especially with countries that have a substantially lower tax rate than the exporting country and as trade with the very low tax rate importing countries declines, trade with the importing countries that are close to the exporting country in terms of the tax rate could increase. This suggests that the effect is possibly driven by the tax rate and therefore in columns (5) and (6), we interact the TPR dummy with the tax rate of the exporting country (5) and importing country (6) respectively.

The reaction is indeed driven by the tax rate, as indicated by the statistically significant negative coefficient on the interaction variable in column (5). Moreover, we find that the positive effect of the tax rate on exported quantities found in columns (1) and (3) was biased downward by country-pairs affected by TPR. An increases in τ_i of one percent leads to an increase in exported quantities of 3.36 percent when no TPR are in place. Likewise, when TPR are in place, an increase of one percent in τ_i leads to a decrease in traded quantities of 4.63 percent. When $\tau_i < \tau_j$ though, TPR in country i do not seem to affect traded quantities as shown in column (4). The interaction effect with the tax rate of the importing country shows a slight decrease in the quantities traded when TPR are in place: An increase in the profit shifting incentive of one percent is associated with an increase in traded quantities of 1.88 percent as opposed to 4.46 percent without TPR, but in contrast to the case where $\tau_i > \tau_j$, the effect remains positive. A test for joint significance reveals that τ_j and the interaction term are jointly significant at the five percent level.

To account for the bilateral trade flows over time, we have relied on clustering the data at the country-pair level. However, as the variation in our treatment dummy comes from the exporting country side only, we ideally would like to use cluster-robust standard errors at the exporting country level. Due to the small number of exporting countries in our sample, standard clustering techniques are not suited and we employ randomisation inference in the spirit of Fisher (1935) to calculate valid p-values. The idea behind this inference technique is based on standard permutation methods: by reassigning treatment randomly across clusters many times, one gets a self-made distribution under the null hypothesis of no effect.[6] This distribution is then used to calculate p-values. Note that in our study, it is the *sequence* of the treatment dummy that has to be resampled and not the treatment dummy alone to avoid unrealistic transfer pricing histories where countries randomly seem to switch transfer pricing regulations on and off.

Using the specifications from columns (5) and (6) of Table 3.2, we ran 10,000 replications to construct a valid distribution of the t-values. From the results, we cannot identify a robust effect for importing countries that exhibit a higher tax rate than the exporting country, with the p-value for the interaction between the tax rate and the transfer pricing regulations being 0.602. But for the opposite case, we find a highly statistically significant ($p = 0.005$) reduction of the traded quantities from the exporting countries to countries with a lower tax rate once transfer pricing regulations are in place.

3.5.2 Interplay of Anti-Avoidance Measures

Some authors have argued that there could be an interplay of different anti-avoidance measures. When analysing the effect of TPR on the firm level and using operating profits as dependent variable, CFC rules and TCR should not play a role, as they both affect financial profits (Lohse and Riedel, 2013). But there is no reason to believe that these rules could not have an effect on traded quantities and therefore, we investigate the relationship between traded quant-

[6] For a detailed description see Barrios et al. (2012); Ho and Imai (2006).

Table 3.3: Interaction with other Anti-Avoidance Instruments

	(1)	(2)	(3)	(4)	(5)	(6)
	$\tau_i > \tau_j$	$\tau_i < \tau_j$	$\tau_i > \tau_j$	$\tau_i < \tau_j$	$\tau_i > \tau_j$	$\tau_i < \tau_j$
	Dependent Variable: Log of Exports					
$\ln \tau_i$	3.548***	-0.751	3.856***	-1.584	3.635***	0.576
	(0.608)	(2.072)	(0.613)	(2.037)	(0.594)	(1.970)
$\ln \tau_j$	-0.296	3.121*	-0.275	3.437**	-0.290	3.653**
	(0.457)	(1.662)	(0.453)	(1.635)	(0.457)	(1.660)
TPR_i	2.573***	0.168	3.069***	0.846	2.643***	0.628
	(0.309)	(0.566)	(0.329)	(0.576)	(0.298)	(0.556)
$\ln \tau_i \cdot TPR_i$	-8.243***		-8.900***		-8.333***	
	(1.010)		(1.016)		(0.990)	
$\ln \tau_j \cdot TPR_i$		-0.826		-1.205		-1.653
		(1.737)		(1.712)		(1.712)
CFC_i	0.163**	0.723***	0.315***	1.184***	0.145**	0.620***
	(0.067)	(0.106)	(0.070)	(0.141)	(0.072)	(0.111)
TCR_i	-0.200***	0.005	-0.211***	-0.071	-0.150**	0.273**
	(0.048)	(0.101)	(0.048)	(0.101)	(0.065)	(0.134)
$TPR_i \cdot CFC_i$			-0.398***	-0.884***		
			(0.096)	(0.131)		
$TPR_i \cdot TCR_i$					-0.071	-0.346***
					(0.070)	(0.125)
N	19,300	7,119	19,300	7,119	19,300	7,119
R^2	0.240	0.142	0.243	0.152	0.240	0.144
$F-Test$	0.000	0.098	0.000	0.071	0.000	0.072

Notes: Fixed effects regressions of Equation (3.3). Cluster-Robust standard errors on the country-pair level are in parentheses. All estimations include the logarithm of GDP for both countries, a set of dummies controlling for economic integration and year fixed effects as control variables. Coefficients are omitted for brevity. F-Test shows the p-value for a test of joint significance of the tax rate and the interaction term. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

ities and anti-avoidance measures more thoroughly in Table 3.3. In the first two columns, we additionally control for CFC rules and TCR being in place in the exporting country. When $\tau_i > \tau_j$, we find that both measures have an impact on the exports from country i to country j, but the main result from Table 3.2 remains unchanged: a one percent increase in the tax rate now leads to a 4.69 percent decrease in traded quantities once TPR are in place opposed to 4.63 percent before. In the second column, however, CFC rules show a significant influence also on the main results, as the effect of τ_j is smaller in the absence of TPR and once TPR are introduced, a one percent increase in the τ_j leads to an increase of 2.3 percent (compared to 1.88 percent before). In other words, CFC rules, in cases where $\tau_i < \tau_j$, dampen the effects of TPR.

In order to further investigate the interplay of CFC and TPR, we include an interaction of both dummies in a next step. The interaction effect is negative, suggesting that if both TPR and CFC rules are in place, less quantity is traded in column (3), which suggests that CFC rules enhance the effectiveness of TPR. This is also evident from the absolute increase in the effect of the tax rate conditional on TPR, which now indicates a reduction of 5.04 percent of exports for a one percent increase in τ_i. Column (4) reports the results for cases where the exporting county has a lower tax rate than the importing country. CFC rules alone still exert a statistically significant positive effect on exports, but this effect is smaller if TPR are in place, as shown by the negative interaction term. The effect of τ_j and the TPR, CFC interaction term remain quantitatively similar to the results from column (2), which shows that there is limited sensitivity to the simultaneous presence of TPR and CFC rules if the importing country has a higher tax rate than the exporting country.

Finally, columns (5) and (6) investigate the interplay between TPR and TCR. Column (5) shows no additional effect of the interaction between TPR and TCR and the reaction to an increase of one percent in τ_i when facing TPR is -4.7 percent, which is again close to the baseline effect from Table 3.2. Although the interaction effect between TPR and TCR is statistically significant in column (6), the magnitude of the effect is small. Also, the baseline effect of the tax rate and the interaction with TPR remains similar in size (2 vs. 1.88 percent) compared to the results in Table 3.2. TCR therefore do not seem to alter the effects that TPR alone have on quantities, neither in the case of $\tau_i > \tau_j$, nor in the reverse case.

3.5.3 PPML Regressions

The estimations in Tables 3.2 and 3.3 were conducted on log-linearised data and as pointed out by Silva and Tenreyro (2006), due to Jensens inequality, estimates could be severely biased as additional heterogeneity is introduced into the data that not only affects the variance but also biases the point estimate. Therefore, we follow the recent advances in the literature and estimate the gravity model in its multiplicative form, utilising the PPML estimator proposed by Silva and Tenreyro (2006). The results from this exercise are presented in Table 3.4. The negative coefficient for τ_i in column (1), though insignificant, seems puzzling. But as shown in column (3), this result is driven by the negative effect that active TPR have on the tax rate. The result from the first two tables is confirmed that τ_i has a negative effect once TPR are in place, indicating that less trade commences when the tax rate difference increases. The effect is smaller in size, -2.89 percent compared to the -4.63 percent from column (5) of Table 3.2 for one percent increase in τ_i, but still statistically and economically significant. The difference in coefficient size is a common finding when comparing results from log-linearised models to models estimated by PPML and has been at the heart of the critique by Silva and Tenreyro (2006). Column (4) indicates no

Table 3.4: Effect of Transfer Pricing Regulations on Quantity - PPML Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	$\tau_i > \tau_j$	$\tau_i < \tau_j$	$\tau_i > \tau_j$	$\tau_i < \tau_j$	$\tau_i > \tau_j$	$\tau_i < \tau_j$
			Dependent Variable: Exports			
τ_i	-0.139	2.147	0.449	2.146	0.635	3.425***
	(0.420)	(1.328)	(0.461)	(1.305)	(0.475)	(1.265)
τ_j	-0.038	1.053	-0.002	1.058	0.015	0.399
	(0.494)	(0.947)	(0.477)	(1.494)	(0.477)	(1.450)
TPR_i	0.048	-0.077	1.198***	-0.073	1.215***	-0.567
	(0.045)	(0.106)	(0.288)	(0.547)	(0.291)	(0.519)
$\tau_i \cdot TPR_i$			-3.334***		-3.422***	
			(0.841)		(0.849)	
$\tau_j \cdot TPR_i$				-0.008		1.064
				(1.223)		(1.192)
CFC_i					0.119	0.619***
					(0.0833)	(0.108)
TCR_i					-0.075	0.171*
					(0.048)	(0.097)
N	18,827	6,670	18,827	6,670	18,827	6,670
$F-Test$	–	–	0.000	0.339	0.000	0.043

Notes: PPML regressions for the gravity model in multiplicative form. Cluster-Robust standard errors on the country-pair level are in parentheses. All estimations include the logarithm of GDP for both countries, a set of dummies controlling for economic integration and year fixed effects as control variables. Coefficients are omitted for brevity. F-Test shows the p-value for a test of joint significance of the tax rate and the interaction term. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

effect of TPR on quantities for cases where $\tau_i < \tau_j$. This somewhat contrasts the earlier findings but as becomes evident from column (6), CFC rules play a major role, which is again in line with the findings from Table 3.3. In summary, the PPML results confirm the qualitative results from the log-linearised model, namely the importance of TPR as an anti-avoidance measure when $\tau_i > \tau_j$ and the dominance of CFC rules when $\tau_i < \tau_j$. The coefficients are smaller in size, which confirms the findings of previous authors regarding differences between log-linearised models estimated via OLS and multiplicative models estimated via PPML.

3.5.4 Effect on Value

The COMTRADE data also include the value of an export. Value is defined as the product of price and quantity and therefore, the results should show a mixture of the price and the quantity reaction to the introduction of TPR. Given the way previous studies have identified the pricing reaction (Clausing, 2003; Lohse and Riedel, 2013), we would expect to see an decrease in the magnitude of the effect for value when compared to the quantity reaction alone. This is because the pricing and quantity reaction should have opposite signs: Following the introduction of TPR, prices should be reduced (increased) and quantities increased (reduced) when exports were overvalued (undervalued), Table 3.5 shows a replication of the quantity regressions from Tables 3.2 and 3.3. The negative influence of TPR is also visible in column (3) where $\tau_i > \tau_j$, but no value reaction can be inferred for the opposite case. Here, column (6) indicates again that CFC rules are the predominant anti-avoidance measure to influence exports, also when measured by value.

Given the differences found between OLS estimates and PPML estimates, we proceed to re-estimate the PPML regressions from Table 3.4 for the value

Table 3.5: The Effect of Transfer Pricing Regulations on Value

	(1) $\tau_i > \tau_j$	(2) $\tau_i < \tau_j$	(3) $\tau_i > \tau_j$	(4) $\tau_i < \tau_j$	(5) $\tau_i > \tau_j$	(6) $\tau_i < \tau_j$
	Dependent Variable: Log of Value of Exports					
$\ln \tau_i$	2.683***	2.008	4.402***	1.903	4.593***	2.055
	(0.577)	(1.969)	(0.657)	(1.947)	(0.664)	(2.287)
$\ln \tau_j$	-1.580***	2.626*	-1.526***	3.515**	-1.534***	2.111
	(0.561)	(1.384)	(0.545)	(1.773)	(0.543)	(1.680)
TPR_i			2.182***	0.650	2.200***	-0.084
			(0.332)	(0.631)	(0.329)	(0.595)
$\ln \tau_i \cdot TPR_i$			-6.584***		-6.759***	
			(1.082)		(1.070)	
$\ln \tau_j \cdot TPR_i$				-1.859		-0.008
				(1.940)		(1.846)
CFC_i					0.194***	0.736***
					(0.073)	(0.116)
TCR_i					-0.129**	0.025
					(0.057)	(0.124)
N	19,300	7,119	19,300	7,119	19,300	7,119
R^2	0.309	0.145	0.316	0.145	0.317	0.159
$F-Test$	–	–	0.000	0.131	0.000	0.285

Notes: Fixed effects regressions of Equation (3.3) using the logarithm of value instead of quantity as the dependent variable. Cluster-Robust standard errors on the country-pair level are in parentheses. All estimations include the logarithm of GDP for both countries, a set of dummies controlling for economic integration and year fixed effects as control variables. Coefficients are omitted for brevity. F-Test shows the p-value for a test of joint significance of the tax rate and the interaction term. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.6: Effect of Transfer Pricing Regulations on Value - PPML Estimates

	(1) $\tau_i > \tau_j$	(2) $\tau_i < \tau_j$	(3) $\tau_i > \tau_j$	(4) $\tau_i < \tau_j$	(5) $\tau_i > \tau_j$	(6) $\tau_i < \tau_j$
	Dependent Variable: Value of Exports					
τ_i	-0.404	0.327	-0.138	0.278	-0.0431	1.791
	(0.336)	(1.167)	(0.372)	(1.218)	(0.378)	(1.314)
τ_j	-0.480	0.144	-0.452	-1.066	-0.443	-1.522
	(0.474)	(0.793)	(0.467)	(1.055)	(0.468)	(0.970)
TPR_i	0.0269	-0.183**	0.505**	-0.924**	0.538**	-1.206***
	(0.0408)	(0.0849)	(0.242)	(0.361)	(0.243)	(0.349)
$\tau_i \cdot TPR_i$			-1.370**		-1.462**	
			(0.689)		(0.690)	
$\tau_j \cdot TPR_i$				1.863**		2.535***
				(0.838)		(0.828)
CFC_i					0.0769	0.537***
					(0.0701)	(0.107)
TCR_i					-0.00534	0.264**
					(0.0416)	(0.107)
N	18,827	6,670	18,827	6,670	18,827	6,670
$F-Test$	–	–	0.052	0.063	0.057	0.004

Notes: PPML regressions for the gravity model in multiplicative form. Cluster-Robust standard errors on the country-pair level are in parentheses. All estimations include the logarithm of GDP for both countries, a set of dummies controlling for economic integration and year fixed effects as control variables. Coefficients are omitted for brevity. F-Test shows the p-value for a test of joint significance of the tax rate and the interaction term. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.7: Price Reaction to TPR

	PPML		OLS	
	$\tau_i > \tau_j$	$\tau_i < \tau_j$	$\tau_i > \tau_j$	$\tau_i < \tau_j$
Quantity Reaction	-2.787	1.463	-4.695	2.295
Value Reaction	-1.505	1.013	-2.166	2.103
p_{crit}	0.54	0.69	0.46	0.92
Price Reaction	> 0	< 0	> 0	< 0

Notes: The table shows estimated signs of the price reaction to the introduction of TPR. The price reaction is conditional on the price being larger than p_{crit}.

of exports. The results in Table 3.6 differ from their OLS counterparts, most notably in the joint significance of the tax rate and TPR in cases where $\tau_i < \tau_j$. Columns (4) and (6) both show significant positive effects on the value of exports that are attributable to the introduction of TPR.

Having estimated the effects of TPR on the value of exports, we would like to validate that the pricing effect, which is implicitly in the value effect, has the expected sign. Thus, we would like the pricing effect to be positive (negative) when $\tau_i > \tau_j$ ($\tau_i < \tau_j$). The magnitude of the value effects is smaller in absolute terms than the quantity effects in both settings. Table 3.7 shows the estimated effects that a tax rate increase has, given that TPR are in place. The coefficients for the PPML columns are taken from columns (5) and (6) of Tables 3.4 (Quantity) and 3.6 (Value) respectively, where we also control for other anti-avoidance measures. The OLS results are taken from columns (1) and (2) of Table 3.3 (Quantity) and columns (5) and (6) of Table 3.5.[7] The difference between the quantity and the value effect is driven by the pricing

[7]We focus on the discussion of the PPML results, although the results for the OLS results are similar and the arguments brought forward apply idem dito.

reaction to taxation. Because

$$v(\tau) = p(\tau) \cdot q(\tau), \tag{3.4}$$

we can decompose the effect of an increase in τ[8] on $v(\tau)$ by the total differential to obtain:

$$dv(\tau) = \frac{\partial v}{\partial p}\frac{\partial p}{\partial \tau} + \frac{\partial v}{\partial q}\frac{\partial q}{\partial \tau}. \tag{3.5}$$

The partial derivatives of the value with respect to price and quantity are just $q(\tau)$ and $p(\tau)$ respectively and both strictly non-negative. Implicitly, we abstract from cross elasticities of prices and quantities, which should play a smaller role in intermediate goods trade than in trade in final goods. Additionally, we can justify this by the large share of intra-firm trade in intermediate goods trade, which amassed to around 45% for the US in 2013.[9] The reaction of the quantity to an increase in τ, measured by $\frac{\partial q}{\partial \tau}$ in Equation (3.5), is known from the quantity regressions and displayed in the row *Quantity* in Table 3.7. Likewise, the total reaction of the value is also known and displayed in the row $Value$ of the table. Taking all information together, we can show that the pricing effect in the case of $\tau_i > \tau_j$ is given by

$$\frac{\partial p}{\partial \tau_i} = \frac{2.787p - 1.505}{q} \tag{3.6}$$

and for the opposite case by

$$\frac{\partial p}{\partial \tau_j} = \frac{1.013 - 1.463p}{q}. \tag{3.7}$$

As soon as p surpasses its critical value of $p_{crit} = 0.54$ in Equation (3.6)

[8] Subscripts i and j have, without loss of generality, been dropped for brevity. When $\tau_i > \tau_j$, the exposition refers to τ_i and in the opposite case to τ_j.

[9] Own calculation on the basis of data from BEA (intra-firm trade in manufacturing) and WITS (total trade in intermediate goods).

or $p_{crit} = 0.69$ in Equation (3.7), we can unambiguously derive the sign of the pricing reaction. In the former case, $\frac{\partial p}{\partial \tau_i}$ is positive, indicating that prices are corrected upwards when τ_i increases, following the introduction of TPR. This is in line with the ex ante manipulation of prices and the undervaluation of exports in cases where the tax rate of the partner country is lower than the tax rate of the exporting country. In the latter case, $\frac{\partial p}{\partial \tau_j}$ is negative, which hints at an overvaluation of exports prior to the introduction of TPR. Back-of-the-envelope calculations of the average prices in our sample (we simply assume $p = v/q$) indicate that only 0.39% of all prices are below the critical value when $\tau_i > \tau_j$ and 3.9% in the opposite case. In other words, for nearly all observed trade flows, we find that TPR must have the expected effect on pricing behaviour, namely price increases when exports are undervalued and price decreases when exports are overvalued.

3.5.5 Discussion

The results presented here indicate a substantial response in traded quantities following the introduction of TPR. They show the presence of allocative inefficiencies through exploitation of tax rate differentials by MNEs. Our results are consistent with the hypothesis that TPR help to partly correct these misallocations, as we find a reduction in traded quantities with countries that have a lower tax rate than the exporting country. This suggests that part of the trade in intermediate goods between two countries was purely driven by tax considerations. Following the introduction of TPR, this channel, became unprofitable and we observe relatively more trade with countries that exhibit a higher tax rate than the exporting country.

Our findings confirm the results from earlier studies on the misuse of transfer prices on the firm level: Clausing (2003) reports significant distortions of prices and Davies et al. (2015) find low internal prices for low tax trading partners, especially very low tax trading partners such as tax havens. We add to this by showing that quantities were distorted as well, amplifying the effect of TPR. Whilst the manipulation of transfer prices is de facto a manipulation of book-keeping, manipulation of quantities can have real economic consequences such as labour market responses and any regulatory changes should therefore carefully consider the reaction in quantities.

In line with Lohse and Riedel (2013) and Beer and Loeprick (2015), we can identify that TPR have a dampening effect on profit shifting behaviour, which is also visible in traded quantities. Thus, we are able to show that following the introduction of TPR, quantities exported to lower tax countries are reduced. This could in turn lead to negative real responses in the respective low tax countries, at least from a global welfare perspective.

We analysed the interplay of different anti-avoidance measures and are able to find evidence for concerns raised by some authors that the isolated consideration of an anti-avoidance measure could lead to biased estimates as the joint effect is neglected. Especially for TPR and CFC rules, we can identify some substitutional effects, but in line with the findings from OECD (2015), we do not find TPR and CFC rules to be perfect substitutes. Where TPR are the dominant force in controlling profit shifting behaviour when $\tau_i > \tau_j$, CFC rules have the strongest effect in when $\tau_i < \tau_j$. TCR play no role when analysing quantities but because they only affect the financing structure of firms, which does not require any trade to take place, this could be expected from the data we look at.

Next to analysing the effect of TPR on quantities in a log-linearised model, we also showed results from estimations of a multiplicative model via PPML. As advocated by Silva and Tenreyro (2006), OLS estimates of the log-linearised model could be severely biased and although we found significant quantitative differences in the estimated coefficients arising from this, the qualitative results remain the same. This is in line with several previous studies from the international trade literature that compared OLS and PPML results, such as Baltagi et al. (2014), Gómez-Herrera (2013) and Silva and Tenreyro (2011).

Utilising the value of exports, we were able to identify the sign of the pricing reaction. In cases where there was an incentive to undervalue exports in order to minimise profits accrued in the exporting country, we find a positive pricing reaction following the introduction of TPR. Likewise, when an incentive existed to overvalue exports, because the exporting nation was the country with the lowest tax rate, TPR seem to correct prices downward. Both results resemble the findings of earlier studies, for example Lohse and Riedel (2013), Zinn et al. (2014) or Cristea and Nguyen (2013).

3.6 Conclusion

We analyse the effect of transfer pricing regulations on international trade flows in intermediate goods. We exploit bilateral trade data for the automobile industry from the BACI database for the years 1995 to 2012, as well as information on the introduction of transfer pricing regulations from Deloitte, Ernst&Young, KPMG and PWC. We find evidence that is in line with the ex ante manipulation of transfer prices for tax optimising reasons. This reduces trade quantities for importing countries with higher tax rates than the exporting country and significantly increases trade quantities with countries that exhibit lower tax rates. The effects are strongly driven by the tax rate difference, which is as expected given that the tax rate difference represents the incentive to manipulate transfer prices for profit shifting purposes.

Furthermore, our results indicate that different anti-avoidance measures interact and especially TPR and CFC rules can be seen as partial substitutes. This implies that countries that already sustain either TPR or CFC rules should consider the adverse effects that the introduction of the second anti-avoidance measure could have on the existing one. Utilising the value of exports, we were also able to identify the commonly acknowledged signs of the pricing reaction to the introduction of TPR. TPR act as a corrective both to under- as well as overvalued exports.

A potential shortcoming of our study is the focus on one industry. The automotive industry is characterised by highly specialised products that are seldomly traded with unrelated third parties, thus providing ample opportunities for the manipulation of transfer prices. We would thus expect the effectiveness of TPR to decrease or even vanish when looking at less specialised or more open sectors. Especially when looking at intermediate goods trade as a whole, the positive effects of TPR on the reduction of transfer mispricing in cases where the opportunities are manifold could be confounded by the insignificance of TPR for other sectors of the economy. Given the significant economic burden on companies and the tax administration, social desirability of TPR depends on the extent of sectors present in an economy that have the opportunity to excessively manipulate transfer prices. Our study adds to the discussion that the allocative distortions through the quantity reactions need to be considered on top of the pricing reactions.

Chapter 3 documented a significant reaction of firms to the introduction of TPR in line with prior manipulation of transfer prices. The chapter provided evidence both for manipulated prices and distortions in traded quantities that were reduced following the introduction of TPR. Furthermore, the point was stressed that the effects of TPR are multi-dimensional and that - just as in the case of individual taxation and the equity-efficiency trade-off - a trade-off between the price and the quantity of an exported good exists, which needs careful balancing in order to keep the tax base broad and the taxable income high.

We showed that tax policy can be utilised to correct misallocations in the market via **disincentives** such as TPR, leading to a more efficient allocation of goods. But the question remains whether **tax incentives** also lead to the desired effect of increased efficiency. By analysing the effect of a tax incentive for R&D, namely the intellectual property box (IP-Box), Chapter 4 is dedicated to this question. IP-Boxes are by and large a European phenomenon and thus well suited in the context of this thesis that has analysed tax policies in place in the EU. They are designed to encourage R&D by setting lower tax rates on income derived from patents or other intellectual property. The coverage and the extent of the tax reductions vary between the countries implementing IP-Box regimes. Especially the development condition, which declares that only newly developed patents are eligible for the reduced tax rate under the IP-Box regime, is only present in three countries, thus raising questions over the effectiveness of the IP-Box as a R&D incentive.

Chapter 4

Opening Pandora's Box - Do Intellectual Property Boxes Foster Innovation?

4.1 Introduction

The concept of an intellectual property box (IP-Box) is to grant tax exemptions or reductions for intellectual property such as patents. Countries implementing a patent box regime argue that lower tax rates on income from the use of patents foster innovation and increase employment in the R&D sector.[1] Critics accuse theses regimes of creating incentives for multinational enterprises (MNE) to strategically shift their intangible assets into the countries implementing an IP-Box regime to minimise their effective tax burden (Evers et al., 2015; Griffith et al., 2014).

A government that incentivises patent application can do this for two reasons:

[1]See, for example, the Dutch tax authorities `http://www.belastingdienst.nl/wps/wcm/connect/bldcontentnl/belastingdienst/zakelijk/winst/vennootschapsbelasting/innovatiebox` (only available in Dutch). Last accessed 18.06.2017.

1. Stimulate R&D within the country and 2. Attracting foreign R&D, i.e. the research is conducted in a foreign country and the patent is applied for tax saving purposes in the country offering incentives. Both could be viable goals of a government, as both effects should increase tax revenue. However, especially in recent times, public opinion has fallen out with the attraction of foreign R&D as many MNE extensively use such tax incentives to minimise their effective tax burden.[2] Therefore, governments implementing IP-Box regimes have a strong incentive to highlight the real response (more domestic R&D).

I exploit the potential difference between the location of the inventor and the location of the applicant of a patent to disentangle the innovation response from the shifting responses. If an inventor resides in the same country as the applicant, the patent is considered home grown. An increase in the number of home grown patents following the introduction of an IP-Box regime would be evidence in favour of the real response channel. If an inventor of a patent resides in a different country from the applicant, however, the patent is considered as shifted. An increase in the number of shifted patents following the IP-Box introduction will be in favour of the shifting channel. Note, though, that both channels could potentially coexist and from the perspective of a budget maximising government, both effects can be desirable.

The paper relates to several strands of the literature. First, tax incentives play a crucial role in fostering R&D (Bloom et al., 2002; Lokshin and Mohnen, 2013), especially for smaller firms (Bronzini and Iachini, 2014). Furthermore, R&D tax incentives play a role in the location of R&D (Grubert, 2003; Hines Jr, 1994). The tax incentives analysed in many previous studies are ex-ante incentives, i.e. incentives that are given *before* the innovation process has finished. These include direct incentives, such as grants, as well as indirect incentives,

[2] Cases of minimising tax burdens by reallocation of intellectual property have been widespread in the media for the likes of Apple, Starbucks, IKEA and other global MNEs.

such as wage subsidies for workers in R&D or reduced wage tax rates for R&D employees.[3] IP-Boxes are an ex-post incentive, as only successful R&D (in the form of a patent) is rewarded. In this sense, whilst the classical R&D incentives aim at encouraging any type of R&D, IP-Box regimes should only affect successful R&D, which are considered to be a signal of quality (Czarnitzki et al., 2014).

A second strand of the literature has addressed IP-Box regimes, with a comprehensive overview of IP-Box regimes and their theoretical effects on the effective average tax rate given in Evers et al. (2015) and a more detailed analysis of different IP-Box characteristics in Alstadsater et al. (2015), though only for three industries. Both studies find that there is no direct impact of IP-Box regimes on innovation activities and shifting of patents is the predominant reaction. In a recent study, Schwab and Todtenhaupt (2016) show that IP-Box regimes can have at least partial positive spill-over effects on R&D investments across countries and therefore could globally lead to more R&D.[4] Such a positive effect of tax differentials on (immobile) R&D investment is also documented in Beer (2015).

Third, there exist some prior work on the fact that MNEs react to tax incentives with reallocation of (in-)tangible assets (Böhm et al., 2015; Dischinger and Riedel, 2011; Griffith et al., 2014), for an overview of recent empirical evidence on tax-motivated profit shifting see Dharmapala (2014).

The literature has thus provided arguments in favour of the positive effect of IP-Box regimes on R&D investments as well as evidence that IP-Boxes could facilitate profit shifting. The question remains, which channel potentially dominates and is at the heart of this study. I add to the literature by analysing both

[3]For an overview of R&D tax incentives that are currently in place in Europe see CPB et al. (2015).

[4]Note that I do not refer to knowledge spill-overs but to spill-overs through the tax incentive and the expected increase in net return.

channels simultaneously on the country level and providing, based on the main results, a rationale for the crowding out of home grown patents following the introduction of an IP-Box regime. The analysis provides evidence in favour of the shifting channel but finds negative effects of the IP-Box regime on home grown patents, thus not only rejecting the claim that IP-Boxes foster innovation in the country they are implemented in but actually suggesting that local innovative activities are reduced. A positive effect on the total number of patents cannot be identified, further suggesting a replacement of home developed patents by foreign developed patents.

The remainder of the paper is structured as follows: Section 4.2 describes the institutional setting, especially the different characteristics of IP-Box regimes. The methodology as well as the data are presented in section 4.3. Section 4.4 discusses the results of the baseline regressions and of several robustness checks. Section 4.5 concludes.

4.2 Institutional Setting

An intellectual property box is designed to allow income derived from patents or other intellectual property to be taxed more beneficially. The scope of other intellectual property that is also governed by the IP-Box regime differs by country, as does the design and extent of the tax benefit. All IP-Boxes include income from the use of patents (royalties), but some also include trademarks, designs or copyrights. The tax benefit can be a reduced tax rate on income that falls into an IP-box but can be granted as wage tax deductions for workers in R&D as well or as a reduction in the tax base. An overview of the different designs of IP-Box regimes in European OECD countries is given in Table 4.1. The main feature of an IP-box is its effective tax rate, which is substantially lower than

Table 4.1: IP-Box Regimes in European OECD Countries until 2014

	Year	Tax rate	Qualifying Income	Acquired IP	Existing IP	Larger Scope
Belgium	2007	0.068	Royalties, Sales income, Notional royalties	No	No	No
France	2000	0.1676	Royalties, Capital gains	Yes	Yes	No
Hungary	2003	0.095	Royalties, Capital gains	Yes	Yes	Yes
Ireland	2008$^{a)}$	-	-	-	-	-
Luxembourg	2008	0.0441	Royalties, Capital gains, Sales income, Notional royalties	Yes	Yes	Yes
Netherlands	2007	0.05	Royalties, Capital gains, Sales income, Notional royalties	No	No	Yes
Portugal	2014	0.15	Royalties, Capital gains	No	No	Yes
Spain	2008	0.12	Royalties, Capital gains	No	Yes	Yes
United Kingdom	2013	0.10	Royalties, Capital gains, Sales income, Notional royalties	Yes	Yes	No

Notes: Year indicates the year of enactment. Tax rate refers to the tax rate under the IP-Box regime. a) Ireland abolished the IP-Box system (and reintroduced it in 2016). Qualifying income indicates all possible qualifying income. Larger scope refers to more kinds of IP protection than patents and supplementary protection certificates. Data taken from (Evers et al., 2015), Alstadsater et al. (2015) and European Tax Handbooks.

the headline corporate tax rates that apply in the respective countries. Belgium for example has a difference of 27.2 percentage points between the corporate tax rate and the IP-Box rate, whilst Hungary only offers a 9.5 percentage points reduction. Therefore, next to analysing the extensive margin effects of introducing an IP-Box, the tax benefit it offers (i.e. the intensive margin) can also deliver important insights.

IP-Boxes are designed to foster innovation and therefore should not allow existing IP to qualify for the preferential tax rate. This is only the case in Belgium, the Netherlands and Portugal. So in order to qualify for the IP-Box tax rate, innovation must necessarily take place, at least from a global perspective. The treatment of existing IP has also been one of the main criticisms of IP-Box regimes and has resulted in the emergence of the Modified Nexus Approach (Action 5 of the OECD's BEPS Action Plan). It suggests minimum levels of economic activity that should be carried out in the country where the applicant resides, if the applicant wants to benefit from the IP-Box tax rate.[5] Inside the EU however, any level of research that is to be undertaken in the country of residence of the applicant is allowed to be undertaken in any member state, to be in accordance with the Freedom of Movement Act.

The development conditions in place link the purely ex-post incentives of the IP-Box to the original, ex-ante, research. Ex-post incentives for research and innovation, such as general protection from a patent or the beneficial tax rate under an IP-Box regime, only apply to successful research, i.e. if a patentable idea arose from the R&D activity. Therefore, it is unclear how they stimulate research in general. Including a development condition in the setup of an IP-Box regime circumvents this problem as the beneficial tax rate is only granted if some innovative activity took place. So for a government that aims at increasing R&D, a development condition of some sort in the IP-Box design is necessary in order to be credible. If the aim of the IP-Box regime is to attract foreign IP, though, a development condition is not necessary and can even harm this goal. In light of this, France, Hungary, Luxembourg and the UK are expected to be particularly attractive locations for foreign developed patents.

4.3 Data and Methodology

To uncover the effect of IP-boxes on innovation activities, measured by the number of patents, I exploit the (possible) difference between inventor and applicant of a patent. Inventors are the actual developers of the patent, i.e. the group

[5] Although the BEPS Action Plan consists of suggestions, several countries have altered or will change their IP-Box regimes to incorporate the modified nexus approach. These include Belgium (2016), Hungary (2016), the Netherlands (2017), Portugal (2016), Spain (2016) and the UK (2016). Luxembourg has abolished its IP-Box regime in 2016 following the BEPS Action Plan.

Figure 4.1: Relationship between Inventor and Applicant

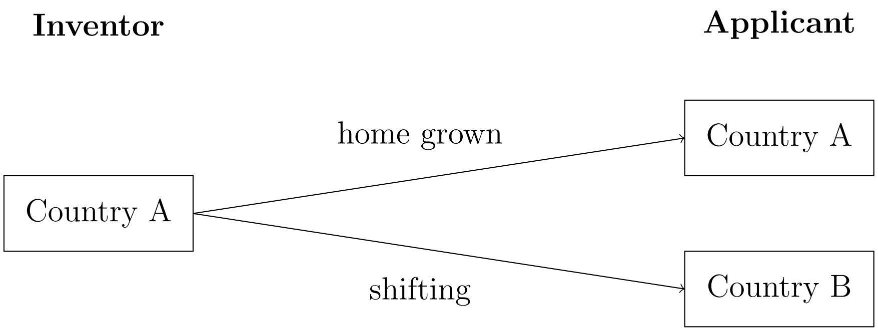

that creates the innovation. Applicants can either be legal or natural persons that apply for the patent to be granted. It is possible to develop a patent in one country and have the application to the patent office in a different country. Especially for tax saving purposes, it can be beneficial to locate patents in jurisdictions that offer preferential tax rates, for example in the form of an IP-Box. Patents that are shifted can be identified as those patents, where the inventor is located in another country as the applicant.[6] On the contrary, if inventor and applicant are located in the same country, this patent is considered to be home grown, i.e. not shifted. This relationship is shown in Figure 4.1. As innovation boxes are described by policy makers to foster innovation, one would assume to detect a rise in home grown patents, with the effect on shifted patents being a priori ambiguous.

Data on patents come from the EPO's PATSTAT database (version Autumn 2014)[7], which covers worldwide patent applications. Specifically, data on the applicant and the inventor are used to identify whether a patent is home grown or shifted. I use the years 1990 to 2009 as sample period and focus on

[6]It is important to mention that this is the country of residence of the inventor or applicant and thus the country, where tax benefits can accrue. This comparison gives no indication on the location of the patent office, where the patent was applied for, which is irrelevant for this study.

[7]Access was kindly provided by Datenlabor Hohenheim.

OECD countries, because they are a relatively homogeneous group in terms of industrial development and innovation. Furthermore, I focus only on patents that are applied for in countries that are member states of the European Patent Organisation.[8] All member states grant patents in accordance with the European Patent Convention and therefore, similar monetary and legal costs apply. It is important to stress that the scope of the patent, i.e. its coverage, can be set similarly at all patent offices. Thus, there is no substantive difference between the rights a firm can obtain when applying for a patent in any of the member states. This makes the patents more comparable and hedges against concerns that firms strategically relocate their patents for other than tax related motives.

Several difficulties arise in constructing the data set from the raw PATSTAT data, which stem from one feature of a patent: more often than not, patents are developed by several inventors, applied for by more than one applicant at a number of different authorities. Therefore, each case needs to be treated separately, in order to avoid counting some patent applications twice whilst leaving out others. Per se, the different authorities a patent is applied at do not affect the outcome of this paper, as the country of residency of the applicant is the determinant of whether tax benefits are granted or not. But the patent office dimension still needs to be considered in order to avoid wrong aggregation of the data. In total, eight possible cases can arise, of which four are easy to treat and four are harder to disentangle. For ease of comprehension, I shall introduce the following notation to separate the cases: (N(inventor),N(applicant),N(authority)). An overview of the cases is shown in Table 4.2.

The first and simplest case is a patent that has one inventor, one applicant

[8] A list of all 38 member states and their date of entry into the organisation can be found in the appendix.

Table 4.2: Combinations of Inventors, Applicants and Authorities

	Inventor	Applicant	Authority
Case 1	1	1	1
Case 2	1	1	many
Case 3	1	many	1
Case 4	1	many	many
Case 5	many	1	1
Case 6	many	1	many
Case 7	many	many	1
Case 8	many	many	many

Notes: The table shows the different combinations of inventors, applicants and authorities considered. Many refers to more than one in each case.

and one authority, thus making it case (1,1,1). To obtain the number of patents per country code of the inventor, country code of the applicant and year, one can simply count the number of observations that satisfy these criteria. The second case considered is (1,1,m). Here, the applicant has applied at several authorities but naturally can only claim the tax benefit for one patent. Therefore, for each unique combination of inventor and applicant, the year the patent was first applied for is kept, as this allowed for the tax benefit to be claimed.[9] For case (1,m,1), we have several applicants for one patent. Potentially, these could all claim tax benefits for the patent, hence they need to be treated as separate observations and I proceed like in case (1,1,1). In a combination of the cases (1,1,m) and (1,m,1), for case (1,m,m) I keep all the applicants, but if an applicant

[9] As an example, a patent could be applied for at the German patent office in 1995 and at the US patent office in 1996 to extend its coverage. Relevant for this analysis is then the first date that the patent was applied for, hence the application to the German patent office in 1995.

has several authorities, I again keep only the oldest observation.

Up until now, I have only considered the cases where one inventor is present. These were straightforward to handle but as soon as more than one inventor is present, the situation becomes more complex. The underlying question is, how a patent can be classified as home grown or shifted if inventors potentially reside in different countries? Three scenarios are possible: Classify a patent as a home grown patent as soon as one of the inventors resides in the same country as the applicant ("Domestic"), classify a patent as shifted patent as soon as one of the inventors resided in a different country from the applicant ("Foreign") and construct weights from the number of home and foreign inventors and treat the patent as being partly home grown and partly shifted ("Weighted"). Because each of the three scenarios can potentially affect the conclusions, the analysis is carried out for each scenario separately. Under the different peculiarities of scenarios Domestic, Foreign and Weighted, cases (m,1,1), (m,1,m), (m,m,1) and (m,m,m) are dealt with in the same manner as their counterparts with one inventor, i.e. (1,1,1), (1,1,m), (1,m,1) and (1,m,m) respectively.

Next to all patent related data from PATSTAT, data on corporate income tax rates from the OECD as well as data on further control variables from the World Development Indicators of the World Bank are exploited to **estimate the effect of IP-Boxes** on the creation and location of intangible assets in the form of patents. Specifically, I run two different regression models. The estimation equation of the first model is given by:

$$\ln Patents_{lit} = \alpha + \beta_1 ipbox_{it} + \mathbf{X_{it}}'\gamma + \theta_i + \delta_t + \varepsilon_{it}, \tag{4.1}$$

and for the second model by:

$$\ln Patents_{lit} = \alpha + \beta_2 taxbenefit_{it} + \mathbf{X_{it}}'\gamma + \theta_i + \delta_t + \varepsilon_{it}. \tag{4.2}$$

$\ln Patents_{lit}$ is the logarithm of one of the l categories, with l being either the total number of patents, the number of patents classified as home developed or the number of patents classified as shifted. $ipbox_{it}$ is a dummy specifying whether country i had an IP-Box regime in place in year t. $taxbenefit_{it}$ refers to the actual reduction in the tax rate for intellectual property and is defined as corporate tax rate less IP-Box tax rate. Naturally, this will be zero for countries that do not offer an IP-Box, as their tax rate on intellectual property is the same as on all other corporate activities. $\mathbf{X_{it}}$ is a vector of country level controls that include the logarithm of GDP per capita and the unemployment rate. They capture the time-varying heterogeneity between the countries that could influence at least the location decision for a patent. Because the countries in the sample differ in certain characteristics that are fixed over time, a set of country fixed effects θ_i is included. To account for a general upward trend in patent development, δ_t is included as a set of time fixed effects. ε_{it} is an error term.

4.4 Results

This section shows the results from estimating Equations (4.1) and (4.2). Next to the baseline results, I show several robustness checks.

4.4.1 Baseline Results

In the baseline regressions, the fixed effects estimator is used to estimate the model in order to account for the heterogeneity between the different OECD countries. Especially Germany, Japan and the United States have substantially higher numbers of patents than the other countries, which can be captured by the country fixed effect. Table 4.3 shows the results of estimating Equations (4.1)

Table 4.3: Baseline Results for Case Domestic

	(1)	(2)	(3)	(4)	(5)	(6)
	Total		Home		Foreign	
IP-Box	-0.100		-0.286**		1.056*	
	(0.173)		(0.130)		(0.540)	
Tax Benefit		-1.104*		-1.623**		3.201
		(0.552)		(0.609)		(2.053)
Corporate Tax Rate	-0.044	-0.003	-0.317	-0.321	0.906	1.147
	(1.715)	(1.737)	(1.625)	(1.635)	(1.410)	(1.446)
GDP per Capita	1.877***	1.868***	1.992***	1.966***	0.934**	1.042**
	(0.350)	(0.346)	(0.333)	(0.329)	(0.402)	(0.507)
Unemployment	0.055**	0.055**	0.023	0.023	0.081**	0.083**
	(0.022)	(0.022)	(0.020)	(0.020)	(0.033)	(0.034)
N	569	569	569	569	569	569
Clusters	34	34	34	34	34	34
R-squared	0.619	0.620	0.702	0.703	0.208	0.167

Notes: Total refers to the total number of patents filed, Home considers only patents developed in the same country as the patent applicant resides in and Foreign indicates that the patents were developed abroad. All specifications include a full set of time- and country-fixed effects. Standard errors, clustered at the country level, are reported in parenthesis. *,** and *** indicate significance at the 10-, 5- and 1-percent level respectively.

and (4.2) for case Domestic. Because in this case a patent is only considered to be foreign when all inventors of the patent are from outside the country of the applicant, this should provide a precise estimation for the effect of IP-Boxes on shifted patents. The effect on home-grown patents can be viewed as a lower bound, because the count of home-grown patents includes patents that were partly developed outside of the country of the applicant. The standard errors across all specifications are clustered at the country level to allow for correlation of the error terms within a given country over the observational period.

In column (1) the effect of an IP-Box as a whole is not statistically different from zero, implying that the introduction of an IP-Box does not significantly affect the number of new patents developed and hence not the global innovation activity. Column (2) shows a marginally significant negative effect of the height of the tax benefit under the IP-Box regulation on the number of patents developed. It states that a one percentage point increase in the tax benefit under the IP-Box regime decreases the number of patents by 1.1 percent. Note that this result should be viewed with caution as the number of changes in the tax rate on income from patents has varied only sporadically over the sample period.

More strikingly, the results in column (3) show that the introduction of an IP-Box regime has a statistically and economically significant negative effect on the number of home-grown patents. All things equal, the introduction thus reduces the number of home-grown patents by 24.87 percent. So instead of increasing innovative activities in the country, the introduction of an IP-Box regime has substantially harmed R&D in that country. This result is supported by column (4) when looking at the tax rate on intellectual property: a one percentage point increase reduce home-grown patents by 1.62 %. Columns (5) and (6) show the estimation results for foreign patents. The number of foreign developed patents increased nearly fourfold following the introduction of an IP-

Box regime. Taken together with the reduction in home-grown patents, the IP-Box regimes have managed to attract foreign patents that in return crowded out patents developed at home, leading to no significant increase in the total number of patents. The beneficial tax rate again seems to be the most important driver and for every percentage point increase in the tax benefit, 3.2 percent more patents are attracted from abroad, although this effect falls just short of being significant at the 10%-level. These results give clear evidence of a domination of the shifting channel and even suggest that through crowding out of home development, innovation is hampered rather than encouraged by an IP-Box regime.

Table 4.4 shows the results for case Foreign, where patents are classified as foreign as soon as one of the inventors comes from a different country as the applicant, as well as case Weighted, where a weighted average is used. The general results from Table 4.3 hold for case Foreign as well, although the positive effect for the foreign developed patents disappears. When using a weighted average of foreign and home inventors of a patent, a similar picture to Table 4.3 prevails, which is due to the relatively higher number of home developed patents over foreign developed patents. The total number of patents remains insignificantly affected by the introduction of an IP-Box regime, whilst home developed patents seem to be crowded out by patents developed in a foreign location. The remainder of the paper will report the results for case Weighted only, which is the preferred specification as it takes the relative weights of domestic and foreign inventors into account.[10]

The reduction of home grown patents might seem puzzling at first. To understand the mechanism behind this result, it is necessary to understand that only MNEs can be responsible for this effect. Applicants that reside only in the

[10] All results for the other two cases are available from the author upon request.

Table 4.4: Baseline Results for Case Foreign and Case Weighted

Case Foreign

	(1)	(2)	(3)	(4)	(5)	(6)
	Total		Home		Foreign	
IP-Box	-0.091		-0.304**		0.686	
	(0.170)		(0.136)		(0.420)	
Tax Benefit		-1.034*		-1.727**		1.847
		(0.549)		(0.642)		(1.546)
Corporate Tax Rate	-0.130	-0.090	-0.379	-0.384	0.654	0.826
	(1.710)	(1.732)	(1.689)	(1.699)	(1.259)	(1.290)
GDP per Capita	1.883***	1.875***	2.114***	2.086***	1.173***	1.239***
	(0.356)	(0.353)	(0.354)	(0.349)	(0.356)	(0.412)
Unemployment	0.057**	0.056**	0.023	0.023	0.083***	0.083***
	(0.022)	(0.022)	(0.020)	(0.020)	(0.028)	(0.028)
N	569	569	569	569	569	569
Clusters	34	34	34	34	34	34
R-squared	0.619	0.621	0.685	0.685	0.327	0.304

Case Weighted

	(1)	(2)	(3)	(4)	(5)	(6)
	Total		Home		Foreign	
IP-Box	-0.095		-0.304**		1.056*	
	(0.177)		(0.136)		(0.540)	
Tax Benefit		-1.097*		-1.727**		3.201
		(0.568)		(0.642)		(2.053)
Corporate Tax Rate	-0.081	-0.039	-0.379	-0.384	0.906	1.147
	(1.781)	(1.804)	(1.689)	(1.699)	(1.410)	(1.446)
GDP per Capita	1.967***	1.959***	2.114***	2.086***	0.934**	1.042**
	(0.368)	(0.364)	(0.354)	(0.349)	(0.402)	(0.507)
Unemployment	0.058**	0.058**	0.023	0.023	0.081**	0.083**
	(0.023)	(0.023)	(0.020)	(0.020)	(0.033)	(0.034)
N	569	569	569	569	569	569
Clusters	34	34	34	34	34	34
R-squared	0.600	0.602	0.685	0.685	0.208	0.167

Notes: Total refers to the total number of patents filed, Home considers only patents developed in the same country as the patent applicant resides in and Foreign indicates that the patents were developed abroad. All specifications include a full set of time- and country-fixed effects. Standard errors, clustered at the country level, are reported in parenthesis. *,** and *** indicate significance at the 10-, 5- and 1-percent level respectively.

country that introduces the IP-Box can only benefit from the lower tax rate and have no incentive to reduce R&D. A possible rationale for the crowding out of home developed patents (by MNEs) can be seen in the following example.

Consider a MNE active in two countries A and B, with A having a lower tax rate than B ($\tau_A < \tau_B$). Because of the tax differential, the MNE prefers to hold the patent in country A (Karkinsky and Riedel, 2012). Shifting of intangible assets is assumed to be free of cost. To develop a patent that reaps an expected net payoff $E[\pi_i(1-\tau_A)]$, the MNE has to incur costs c_i in country i, with $i = A, B$. c_i is an increasing function of the quality of the researchers, so that better quality patents come at the cost of more expensive research. Furthermore, due to tax exemptions, the costs are not affected by the tax rate.[11] Assume that a patent developed in country A has a lower expected net payoff, as well as a lower cost, such that:

$$E[\pi_A(1-\tau_A)] - c_A > E[\pi_B(1-\tau_A)] - c_B \tag{4.3}$$

in the initial situation. Therefore, the MNE will develop the patent in country A, even though $E[\pi_B] > E[\pi_A]$. Now country A implements an IP-Box regime with a tax rate $\tau_A^{IP} < \tau_A$. This will increase expected net payoffs as $\frac{\partial E[\pi_i(1-\tau_A)]}{\partial(1-\tau_A)} > 0$ for $i = A, B$. For sufficiently low τ_A^{IP}, keeping π_i and c_i constant for $i = A, B$, it holds that

$$E[\pi_A(1-\tau_A^{IP})] - c_A < E[\pi_B(1-\tau_A^{IP})] - c_B \tag{4.4}$$

and the decision of the initial situation will be reversed. It is now optimal for the MNE to develop the patent in country B and shift it to country A. This is consistent with the observation that R&D in country A is reduced and the number of shifted patents increases.

[11] A derivation with the tax rate affecting the cost of developing a patent is relegated to the appendix.

For simplicity and ease of notation, the line of argumentation followed before considers only costless shifting. It is possible to extend the analysis to incorporate shifting costs, which will require an even stronger tax incentive to make shifting profitable. Especially if the IP-Box regime contains a development condition[12] or does not allow acquired IP to be taxed under the IP-Box tax rate, the costs of developing a patent in country B with the intention to shift it to country A increases. The higher the fraction of R&D that needs to be carried out locally, the lower is the probability that the MNE will benefit from shifting the patent from country B to country A, instead of developing it in country A straightaway. Because only three countries in the sample have a development condition in place, the assumption of (near) costless shifting seems plausible.

4.4.2 Sectoral Heterogeneity

Patents and the expected benefit thereof can potentially differ, depending on the sector they were developed for. This is due to the heterogeneous technological sophistication of the respective sectors. To analyse this heterogeneity between sectors, Tables 4.5 and 4.6 show the estimation results from rerunning the analysis per sector of the IPC classification scheme. In total, eight sectors are defined and as shown in the 4th column of both tables, the number of patents applied for in the sample period varies considerably from 297,315 in Sector D: Textiles and Paper to 2,037,537 in Sector B: Performing Operations and Transporting.

Column (1) shows the estimated coefficient for the IP-Box dummy (Table 4.5) and the tax benefit (Table 4.6) in the regression of total number of patents. Compared to the baseline results for the full sample, considerable heterogeneity

[12]This requires that (part of) the research for the patent has to be carried out in the country that grants the preferential tax rate under an IP-Box regime, which has also been put forward by the OECD as the modified nexus approach.

Table 4.5: Per Sector Analysis (Case Weighted)

	(1) Total	(2) Home	(3) Foreign	No. of Patents
Coefficient:	*IP-Box Dummy*			
A: Human Necessities	-0.165	-0.124	0.537**	801,291
	(0.179)	(0.192)	(0.245)	
B: Performing Operations; Transporting	0.174	-0.025	0.376	2,037,537
	(0.315)	(0.369)	(0.245)	
C: Chemistry; Metallurgy	-0.164	-0.177	0.502	1,321,558
	(0.170)	(0.193)	(0.356)	
D: Textiles; Paper	0.387**	0.005	0.962**	297,315
	(0.181)	(0.221)	(0.440)	
E: Fixed Constructions	0.272	0.188	0.781***	538,571
	(0.330)	(0.296)	(0.284)	
F: Mechanical Engineering; Lighting; Heating	0.227	0.089	0.575***	1,193,086
	(0.248)	(0.236)	(0.167)	
G: Physics	-0.028	-0.164	0.670**	1,595,945
	(0.292)	(0.310)	(0.319)	
H: Electricity	0.030	-0.058	0.597**	1,432,381
	(0.211)	(0.213)	(0.279)	

Notes: The table shows the value of the coefficient on the IP-Box dummy from the baseline regression (Equation (4.1)), carried out for a sub sample of each sector separately. Sectors are defined as per top-level IPC classification. No. of patents gives the number of patents in the sample in the respective sector. Total refers to the total number of patents filed, Home considers only patents developed in the same country as the patent applicant resides in and Foreign indicates that the patents were developed abroad. All specifications include a full set of time- and country-fixed effects. Standard errors, clustered at the country level, are reported in parenthesis. *,** and *** indicate significance at the 10-, 5- and 1-percent level respectively.

Table 4.6: Per Sector Analysis (Case Weighted)

	(1) Total	(2) Home	(3) Foreign	N
Coefficient:	*Tax Benefit*			
A: Human Necessities	-1.362*	-0.721	2.087	801,291
	(0.715)	(0.950)	(1.248)	
B: Performing Operations; Transporting	-0.229	-1.241	1.339	2,037,537
	(1.240)	(1.678)	(1.503)	
C: Chemistry; Metallurgy	-1.288*	-1.049	1.761	1,321,558
	(0.697)	(0.891)	(1.659)	
D: Textiles; Paper	2.002**	-0.088	5.137**	297,315
	(0.821)	(1.108)	(2.001)	
E: Fixed Constructions	0.149	0.181	3.005*	538,571
	(1.048)	(0.980)	(1.573)	
F: Mechanical Engineering; Lighting; Heating	0.376	-0.168	2.646***	1,193,086
	(0.843)	(0.952)	(0.828)	
G: Physics	-1.289	-1.816	2.486	1,595,945
	(1.235)	(1.385)	(1.718)	
H: Electricity	-0.190	-0.399	2.579*	1,432,381
	(0.869)	(0.946)	(1.372)	

Notes: The table shows the value of the coefficient on tax benefit from the baseline regression (Equation (4.1)), carried out for a sub sample of each sector separately. Sectors are defined as per top-level IPC classification. N gives the number of patents in the sample in the respective sector. Total refers to the total number of patents filed, Home considers only patents developed in the same country as the patent applicant resides in and Foreign indicates that the patents were developed abroad. All specifications include a full set of time- and country-fixed effects. Standard errors, clustered at the country level, are reported in parenthesis. *,** and *** indicate significance at the 10-, 5- and 1-percent level respectively.

is observable, especially in the regressions for the tax benefit. Whilst the sectors Human Necessities and Chemistry; Metallurgy show a similar negative and statistically significant effect of the tax benefit on the total number of patents, the coefficient turns insignificant for the other sectors and even positive and significant for the Textiles; Paper sector.

The effect on home developed patents remains negative (except for Fixed Constructions) in Table 4.6, but loses its significance compared to the aggregate estimation. The pattern for foreign developed patents remains the same in the sectoral analysis and especially sectors Fixed Constructions and Mechanical Engineering; Lighting; Heating react strongly to the introduction of the IP-Box regime (as measured by the dummy) as well as to changes in the tax benefit.

Over all sectors, I find a positive coefficient for the *Foreign* specification and a statistically insignificant effect on the total number of patents developed. The Textiles; Paper sector seems to be an exception here, but the results could also be driven by the considerably smaller sample size as compared to all other sectors. In summary, the results from this section do not indicate that the baseline results are driven by one or a few sectors.

4.4.3 Robustness Checks

The number of clusters in the sample is 34, corresponding to the 34 OECD countries. This is close to the critical number of 30 clusters, for which Cameron and Miller (2015) propose the use of bootstrapping techniques to obtain correct standard errors. Implementing such percentile-t bootstrapping procedures does not change the significance of the results, leading to the conclusion that the cluster robust standard errors of the baseline model are valid.[13] In the following, two alternative models are introduced to confirm the robustness of the main

[13]Results are available from the author upon request.

results.

Count Data Model

The data at hand are discrete count data, so taking the logarithm and treating the variable as continuous could lead to biased results. To hedge against this concern, the analysis is rerun using a negative binomial count data model. The negative binomial model is preferred to a poisson model because it allows the data to be overdispersed, whilst the poisson model assumes equality of mean and variance (equidispersion). More specifically, a NegBin II model, as proposed by Cameron and Trivedi (1986), is estimated using the counts of the total number of patents, the number of home developed patents and the number of foreign developed patents.[14]

Table 4.7 shows the results of estimating the negative binomial model for case Weighted, using either the IP-Box dummy or the tax benefit as explanatory variable, alongside the same control variables as before. The results support the pattern found in the baseline regressions: The effect on the total number of patents is negative but not statistically significant, the introduction of an IP-Box regime reduces the number of home developed patents and increases the number of foreign developed patents. Overall, the results are very similar to the baseline specification, which indicates that the log-linearisation of the data introduces no bias in the full sample.

Given the sectoral heterogeneity in the data, as displayed in Tables 4.5 and 4.6, Table 4.8 shows the negative binomial regressions per sector for the IP-Box dummy. The pattern observed in Table 4.5 is also present and even intensified, leading to the conclusion that the treatment of the count data as continuous might have been wrong in the per sector analysis. The reason for this is the

[14]Results from a poisson model with robust standard errors to account for overdispersion are available from the author upon request.

Table 4.7: Count Data Model - Full sample (Case Weighted)

	(1)	(2)	(3)	(4)	(5)	(6)
	Total		Home		Foreign	
IP-Box	-0.101		-0.317**		1.108*	
	(0.192)		(0.162)		(0.574)	
Tax benefit		-1.095*		-1.604**		3.457
		(0.596)		(0.739)		(2.288)
Corporate Tax Rate	-0.210	-0.171	-0.329	-0.286	0.334	0.498
	(1.701)	(1.705)	(1.706)	(1.702)	(1.175)	(1.250)
GDP per Capita	1.898***	1.887***	2.050***	2.018***	0.878**	0.978**
	(0.437)	(0.429)	(0.430)	(0.425)	(0.360)	(0.475)
Unemployment	0.069***	0.069***	0.039**	0.039**	0.075**	0.076**
	(0.020)	(0.020)	(0.020)	(0.020)	(0.030)	(0.031)
N	569	569	569	569	569	569

Notes: Total refers to the total number of patents filed, Home considers only patents developed in the same country as the patent applicant resides in and Foreign indicates that the patents were developed abroad. All specifications include a full set of time- and country-fixed effects. Standard errors, clustered at the country level, are reported in parenthesis. *,** and *** indicate significance at the 10-, 5- and 1-percent level respectively.

smaller number of counts in the per sector analysis and the increase in the number of zeros that are not modeled correctly in the log-linearised baseline model. In most sectors, the number of foreign developed patents increased after the introduction of an IP-Box regime. The effects on the number of home developed patents as well as the total number of patents are mostly insignificant and in case of the total number of patents ambiguous between the sectors.

Binomial Regression Model

In the previous estimations, the effects were measured for home and foreign developed patents separately and it was difficult to control for a change in the number of patents induced by the IP-Box, because both home and foreign developed patents constitute a fraction of total patents. But this relationship can be exploited in a binomial regression framework, where it is possible to account

Table 4.8: Per Sector Analysis (Case Weighted)

	(1) Total	(2) Home	(3) Foreign	N
Coefficient:	*IP-Box Dummy*			
A: Human Necessities	-0.114	-0.051	0.707**	801,291
	(0.211)	(0.219)	(0.306)	
B: Performing Operations; Transporting	0.180	0.122	0.575**	2,037,537
	(0.314)	(0.379)	(0.268)	
C: Chemistry; Metallurgy	-0.158*	-0.162	0.671	1,321,558
	(0.199)	(0.228)	(0.516)	
D: Textiles; Paper	0.338**	0.091	1.212***	297,315
	(0.149)	(0.234)	(0.347)	
E: Fixed Constructions	0.245	0.235	0.935***	538,571
	(0.292)	(0.273)	(0.275)	
F: Mechanical Engineering; Lighting; Heating	0.181	0.082	0.677***	1,193,086
	(0.221)	(0.241)	(0.167)	
G: Physics	0.054	-0.013	0.913**	1,595,945
	(0.291)	(0.344)	(0.375)	
H: Electricity	0.086	-0.007	0.795**	1,432,381
	(0.213)	(0.238)	(0.368)	

Notes: The table shows the value of the coefficient on the IP-Box dummy for the negative binomial model, carried out for a sub sample of each sector separately. Sectors are defined as per top-level IPC classification. N gives the number of patents in the sample in the respective sector. Total refers to the total number of patents filed, Home considers only patents developed in the same country as the patent applicant resides in and Foreign indicates that the patents were developed abroad. All specifications include a full set of time- and country-fixed effects. Standard errors, clustered at the country level, are reported in parenthesis. *,** and *** indicate significance at the 10-, 5- and 1-percent level respectively.

for changes in the number of patents by conditioning on the total number of patents. In essence, the outcome variable is seen as the sum of n successful Bernoulli trials and the coefficients will indicate how the chances of success are influenced by the explanatory variables (Wooldridge, 2010).

In the setting at hand, the total number of Bernoulli trials is the total number of patents applied for in year t in country i. For ease of interpretation, success is defined as the number of times that the patent is developed abroad.[15] Therefore, the probability that a given patent is developed in a foreign country is given by $p = (Foreign/Total)$ and $1 - p$ is the probability that a patent, conditional on the total number of patents, is developed at home. The coefficients presented below thus show the effect of the explanatory variables on the probability that a patent is developed in a foreign country. The results are given in odds ratios, so any coefficient greater than 1 will indicate that the variable positively influences the chances of success (i.e. of a foreign developed patent). Likewise, a ratio smaller than 1 indicates that the chances of success are hampered by that variable and the probability of a home developed patent rises.

Table 4.9 shows the estimation results from the binomial regression model. Column (1) displays the coefficient estimates for the IP-Box dummy in the various samples. In the full sample, the chances that a newly developed patent is from abroad are increased by 73% in countries that have an IP-Box regime in place. By the complementarity of the probabilities, the chances that a new patent is developed in the home country are reduced by 73%. In column (2), the odds ratios for a unit change in the tax benefit are presented. Because the tax benefit takes on a value between 0 and 1, the odds ratio is hard to

[15] Defining the number of home developed patents as success will lead to the exact same conclusions, as it is the complementary probability. The results will be in odds ratios and the interpretation of an odds ratio < 1 is more cumbersome than the interpretation of a ratio > 1. Given the baseline results, a pattern where the odds ratio is greater than 1 is expected for the foreign applications.

Table 4.9: Binomial Regression Model (Case Weighted)

	(1) IP-Box	(2) Taxbenefit	(3) Taxbenefit*100	N
Full Sample	1.730**	9.296*	1.023*	569
A: Human Necessities	1.938**	13.705**	1.027**	568
B: Performing Operations; Transporting	1.442	3.947	1.014	568
C: Chemistry; Metallurgy	1.677	6.048	1.018	568
D: Textiles; Paper	1.831	27.374*	1.034*	499
E: Fixed Constructions	2.335**	26.862**	1.033**	544
F: Mechanical Engineering; Lighting; Heating	1.497	4.803	1.016	552
G: Physics	1.818*	9.592	1.023	557
H: Electricity	1.516	4.867	1.016	542

Notes: The coefficients represent odds ratios for the different samples. Taxbenefit*100 is the tax rate differential multiplied with 100, which allows the coefficient to show the effect of a one percentage point change, rather than the 100 percentage point change that is shown in column (2). *,** and *** indicate significance at the 10-, 5- and 1-percent level respectively.

interpret. Therefore, in column (3), the tax benefit variable in the regression is multiplied by 100, to allow the interpretation in terms of a one percentage point change in the tax benefit. For the full sample, a one percentage point increase in the tax benefit leads to a 2.3% higher probability that a new patent is developed abroad. Across the different sectors, considerable heterogeneity is again observable, but not a single odds ratio is smaller than one, indicating that the probability of a home developed patent is never positively affected by the IP-Box regime. Either, the IP-Box regime significantly enhances the chances of a foreign developed patent or there is no significant effect of the IP-Box regime on the probabilities that a patent is developed at home or in a foreign country, conditional on the total number of patents.

4.5 Conclusion

Governments implementing IP-Box regimes highlight that the IP-Boxes are designed as an ex-post incentive to foster R&D. This study has shown that, at least within the implementing country, IP-Boxes have failed to spur on innovation. More specifically, evidence for a crowding out of R&D from the country that implemented an IP-Box regime was found. A rationale for this can be seen in the ability of a MNE to differentiate between the location of innovation and the location of application of a patent. Schwab and Todtenhaupt (2016) and Beer (2015) report positive spillover effects of IP-Boxes and R&D tax incentives for other countries and so globally, it is unclear whether the effect of IP-Boxes on innovation is truly negative. This should be the subject of future research.

A second channel that was identified in the literature is supported by the results of this study. IP-Boxes that grant a substantially lower tax rate on intangible assets like patents and other intellectual property, attract foreign developed patents. This channel of attracting foreign R&D is lucrative for governments as it brings in new tax revenues, which previously accrued somewhere else. Therefore, from an economic point of view, introducing an IP-Box regime can be a viable measure for a government seeking to increase its budget.

As a policy recommendation, IP-Boxes should not be presented as a tool to foster innovation. They seem to have a negative effect on innovation in the home country and the effect on a global level remains unclear. Some improvements to the design IP-Boxes can be made: Including a development condition is a crucial asset of an IP-Box regime, as this at least requires some form of innovation to take place and could help to improve the image of IP-Boxes as a R&D incentive. The BEPS Action Plan and especially the modified nexus approach have kick-started a range of modifications to existing IP-Box regimes, all aimed at ensuring a minimum level of economic activity in the country where the applicant resides.

This could ultimately improve the capabilities of IP-Boxes to foster innovation. In the current form, though, IP-Boxes are just another profit shifting device that quite rightly came under scrutiny of the European Commission in recent times.

4.A Loss in Home Developed Patents with Cost

If the cost of developing a patent is not tax deductable, the profit functions in countries A and B are given by

$$P_A = (E[\pi_A] - c_A)(1 - \tau_A)$$

and

$$P_B = E[\pi_B](1 - \tau_A) - c_B(1 - \tau_B).$$

An increase in the net-of-tax rate (corresponding to a decrease in the tax rate itself) will yield an increase in profits from a patent developed in country A of $E[\pi_A] - c_A$ and an increase of profits from a patent developed in country B of $E[\pi_B]$. For any positive cost of developing a patent ($c_A > 0$), the gain in country B is strictly greater than the gain in country A, given that we have assumed $E[\pi_B] > E[\pi_A]$. Therefore, the results prevailing when regarding tax exemption of costs provide an upper bound for the net-of-tax rate, in the sense that any higher net-of-tax rate will always induce the MNE to develop the patent in country B. The lower the tax deductibility of expenses, the higher the critical τ_A^{IP} will be.

4.B Member States of the European Patent Organisation

Country	Date of Entry	Country	Date of Entry
Belgium	7 October 1977	Turkey	1 November 2000
Germany	7 October 1977	Bulgaria	1 July 2002
France	7 October 1977	Czech Republic	1 July 2002
Luxembourg	7 October 1977	Estonia	1 July 2002
The Netherlands	7 October 1977	Slovakia	1 July 2002
Switzerland	7 October 1977	Slovenia	1 December 2002
United Kingdom	7 October 1977	Hungary	1 January 2003
Sweden	1 May 1978	Romania	1 March 2003
Italy	1 December 1978	Poland	1 March 2004
Austria	1 May 1979	Iceland	1 November 2004
Liechtenstein	1 April 1980	Lithuania	1 December 2004
Greece	1 October 1986	Latvia	1 July 2005
Spain	1 October 1986	Malta	1 March 2007
Denmark	1 January 1990	Croatia	1 January 2008
Monaco	1 December 1991	Norway	1 January 2008
Portugal	1 January 1992	FYR Macedonia	1 January 2009
Ireland	1 August 1992	San Marino	1 July 2009
Finland	1 March 1996	Albania	1 May 2010
Cyprus	1 April 1998	Serbia	1 October 2010

Chapter 5

Conclusion

Summary

The thesis at hand set out to comprehensively analyse a broad variety of tax policy instruments. The intention was a multidimensional analysis, spanning both individual and corporate taxation, tax schedules, incentives and disincentives, whilst focussing on the EU as the playing field. Chapter 2 was dedicated to the individual tax schedule in the Netherlands, the third chapter showed how anti-avoidance measures like transfer pricing regulations (TPR) as a disincentive device affect firm behaviour and the fourth chapter was dedicated to the analysis of a tax rate reduction in the form of an intellectual property box (IP-Box), thus covering corporate taxation and incentives.

From the analysis in the second chapter, it became evident that individuals react to jumps in the marginal tax rate. The bunching behaviour of individuals at those kink points in the budget set was exploited to identify the elasticity of taxable income (ETI) with respect to the net-of-tax rate. The chapter first provided an extension to the classical bunching approach introduced by Saez (2010) and extended by Chetty et al. (2011). Because individuals face optimisation frictions, perfect bunching at the kink as predicted by theory was not observable. Rather, a window around the kink, known as the bunching window, was used in the analysis. Where prior research had relied on visual inspection to determine the size of the bunching window, the chapter proposed a data-driven procedure instead, which was shown to be robust to variations in various parameters and took away the researchers discretion in that matter. Thus, the chapter provided a methodological contribution to a comparably young, but growing field of research.

The results from the bunching analysis of the Dutch tax schedule revealed that individuals react to changes in the marginal tax rate, albeit very inelastically. Subsample analyses, however, showed substantial heterogeneity between employed and self-employed individuals as well as between men and women. Self-employed and women exhibit a far larger ETI than their respective comparison groups. Having identified bunching behaviour of Dutch individuals, the study proceeded with analysing the question of how taxable income was adjusted. Of the four potential channels[1] to adjust taxable income, the utilisation of deduction possibilities and in particular the shifting of mortgage interest rate deductions between fiscal partners was identified as main instrument. By contrast, an analysis exploiting the hours worked in 2011 did not reveal any significant reactions, thus suggesting that the adjustment of labour supply plays no significant role for reducing taxable income, at least in paid employment. Income shifting over time and tax evasion remain potential channels that were not subject to analysis due to data constraints.

After empirically showing the reaction of individuals to changes in the marginal tax rate, the thesis set its sights on corporate taxation and the efficiency of anti-avoidance measures. The empirical literature has shown that MNEs util-

[1] adjustment of labour supply, utilisation of deductions, income shifting and tax evasion

ise transfer prices to shift profits into (out of) low-tax (high-tax) jurisdictions. Several countries have implemented regulations that limit the manipulation of transfer prices, most often demanding that they should be set at an arm's-length, i.e. be similar to prices agreed with unrelated third parties. Evidence was given in prior literature that MNEs react sensitively to the introduction of TPR in reducing (increasing) their prices when they were overvalued (undervalued) before the implementation of regulations. Surprisingly, a reaction in quantities, i.e. shifts in production and trade flows, had not been analysed in the literature before. Chapter 3 aimed at providing a first analysis of the reaction of quantities to TPR. The results indicated a substantial quantity reaction and also a pricing reaction, which was shown to be in line with the literature. This suggests that before the introduction of TPR, firms shift more exports to low tax countries and less exports to high tax countries for tax optimising purposes. Following the introduction of TPR, especially the reduction in quantities traded with low tax countries was identified.

Another aspect pointed out by the analyses in Chapter 3 was the importance of the interplay between different anti-avoidance measures. Especially controlled foreign company (CFC) rules and TPR were identified as partial complements, although none of the instruments can fully replace the other. Thin capitalisation rules (TCR), on their own or in combination with other instruments, play a lesser role in the shifting of quantities. This was to be expected as they focus on certain minimum requirements of capitalisation within a firm, thus purely affecting financing of a firm and not its production decision. In summary, TPR seemed to ease allocative inefficiencies that arose from the manipulation of transfer prices by MNEs on top of the changes in the tax base due to pricing adjustments.

Whilst Chapter 3 focussed on the reaction of firms to a disincenitve, Chapter 4 asked the question of how firms react to incentives of a tax system. The introduction of this thesis already highlighted the importance and widespread implementation of R&D incentives across Europe. Whilst most incentives are ex ante tax incentives, i.e. incentives that act during the innovation process and before the innovative product was developed, IP-Boxes, analysed in Chapter 4, are an ex post tax incentive, thus only benefiting successful R&D. The analysis attempted to explore whether IP-Boxes are a local innovation enhancing device, as propagated by the countries implementing IP-Boxes, or merely facilitate profit shifting for MNEs by offering a substantially lower tax rate on income from intellectual property.

The results clearly showed that the shifting channel dominates the home innovation channel. Some evidence was found that home developed patents were crowded out by foreign developed and subsequently shifted patents. The total number of patents did not seem to react to the introduction of IP-Boxes, thus even questioning the global innovation enhancing effect of IP-Boxes. Given the nature of the data, it was not possible to investigate the different designs of IP-Boxes more thoroughly, although the implementation of a development condition should be part of every IP-Box regime. This would ensure that, at least from a global or even European perspective, innovation must take place somewhere.

Implications and Extensions

Individuals as well as corporations react to any form of taxation. Except for some very special cases[2], these behavioural responses lead to inefficiencies. In order to be able to minimise these inefficiencies, it is important to study the extent of the behavioural responses, which was at the heart of this thesis. The

[2]Perfectly inelastic demand/supply or non-distortionary taxes such as lump sum taxation.

results indicate that corporations react more elastically to changes in the tax schedule, which is in line with the common finding that firms are more rational than individuals, less prone to optimisation errors and have more shifting possibilities. This suggests that distortions are larger when firms form the tax base and therefore, from an efficiency point of view, taxation of firms should be lower than taxation of individuals. This rule is implemented in nearly all European tax schedules[3], where the top marginal tax rate for individuals is as high or higher than the statutory corporate tax rate.

Chapter 2 analysed the Dutch tax system that is characterised by four tax brackets. The respective elasticities found locally around the upper threshold are modest at most. Further explorations into the anatomy of response showed that this is due to the shifting of deduction possibilities between partners and a real response in terms of hours worked could not be inferred from the data for the population in paid employment. This allows me to conclude that a tax system with tax brackets is not likely to be less efficient than a smoothly increasing progressive tax schedule like in Germany or a flat tax regime that is in place in several Eastern European countries. The story would be different, however, if the tax schedule exhibited notches instead of kinks. As Kleven and Waseem (2013) showed for Pakistan, notches induce a more severe behavioural response and therefore, such tax systems can be seen as less efficient.

Another aspect of a tax system analysed in this study is the scope of anti-avoidance measures for corporate taxation. Especially when the tax base is mobile and spread out over several countries, firms can utilise the differences in tax rates across the locations of their subsidiary firms to minimise the effective tax burden. A predominant way of shifting profits between entities was the manipulation of transfer prices. Several countries have implemented measures

[3] Czech Republic and Hungary being the notable exceptions.

to counter the excessive manipulation of transfer prices, mostly by tying them to an arm's-length price. Although the results of Chapter 3 indicate that TPR were effective in reducing the misreporting of prices as well as the inefficient allocation of quantities, the arm's-length price remains hard to determine in certain specialised industries that are characterised by high shares of intra-firm trade, such as the automobile industry. As some studies have shown, firms also alter their arm's-length prices to hide the true extent of their profit shifting activities (Cristea and Nguyen, 2013). Future research, especially on the design of tax legislation, should therefore concentrate on finding alternative and more efficient measures to determine a comparable price for the transfer prices set by firms.

Tax schedules also offer a wide range of incentives, either by tax cuts or targeted subsidies. These are, from an efficiency point of view, desirable when the market outcome is unsatisfactory. For example R&D can be classified as a public good and given the characteristics of public goods, underprovision through the free riding mechanism is likely. Because R&D is substantial to economic development, a government wants to increase spending on R&D by using the tax scheme to set incentives for investment in R&D by firms. One of these incentives is the IP-Box, which allows income from patents and sometimes other intellectual property to be taxed at a considerably lower rate than other income. The results of Chapter 4 clearly indicate that the main effect of IP-Boxes is the attraction of foreign developed patents, in most cases at the expense of patents developed within the country that offers the IP-Box. If the goal of IP-Boxes is to foster innovation, locally or globally, it is necessary that a development condition becomes mandatory for the design of an IP-Box regime. In the sample under consideration, only Belgium, Portugal and the Netherlands had such a condition in their definition of the IP-Box and simultaneously did not grant the lower tax rate for existing or acquired intellectual property. A solution to the tax competition problem within the EU that is aggravated by the IP-Box regimes could be an EU-wide IP-Box, where the tax revenue could flow into the EU budget. This will have positive effects once it is established that IP-Boxes increase innovation at least somewhere in the EU. On the other hand, it would also mean that national taxation rights would need to be transferred to the EU, which has not been feasible in the past.

Final Words

Taxation is at the heart of modern welfare states. The quest for an efficient tax system remains an important topic in public finance and tax systems have changed little since the early days of civilisation. The Ancient Egyptians as well as subsequent civilisations through to the Middle Ages relied mostly on flat taxes such as the tithe and flat tax systems have survived to this day, especially in Eastern Europe. A problem with flat taxes is that they are proportionate rather than progressive, so when equity concerns matter, a progressive tax system with increasing marginal tax rates, as implemented in most other European countries, should be preferable. The Ancient Romans literally set up a tax haven to cripple the economic prosperity of Rhodes. Tax havens still exist around the world today and operate on the same principles as 2,200 years ago: They offer a substantially lower or even zero tax rate on corporate income to attract foreign firms. Recently, countries have started to combat tax havens by introducing transfer pricing regulations, as well as controlled foreign company or thin capitalisation rules. The thesis at hand showed that these rules are effective in reducing excessive profit shifting. Last but not least, the thesis provided evidence that tax incentives, especially for risky investments such as investments in R&D, where an underprovision in absence of the tax incentive is likely, are desirable. By analysing the effects of intellectual property boxes, however, serious questions arose over the effectiveness of ex post tax incentives and the focus of public policy should be on input related R&D incentives such as direct subsidies or wage subsidies for employees in R&D, if the government desires to increase innovation.

Next to the efficiency of tax systems that has been at the heart of this exploration, equity concerns play a crucial role in finding an optimal tax system. Through the equity-efficiency trade-off, any redistributional considerations induce inefficiencies into the tax system. The goal of any optimal public policy should therefore be to fix the desired level of redistribution in a first stage and subsequently, conditional on the desired scope of redistribution, find the taxation instrument that comes with the minimum additional inefficiency. In other words, to find an optimal tax system, policymakers should accept finding a second best solution in terms of efficiency at the benefit of a desired level of equity.

Bibliography

Adam, S., Browne, J., Phillips, D., and Roantree, B. (2015). Adjustment costs and labour supply: evidence from bunching at tax thresholds in the UK. *mimeo*.

Adams, C. (1999). *For good and evil: The impact of taxes on the course of civilization*. Madison Books, 2 edition.

Alstadsater, A., Barrios, S., Nicodème, G., Skonieczna, A., and Vezzani, A. (2015). Patent boxes design, patents, location and local R&D. Technical report, Oxford University Centre for Business Taxation.

Anderson, J. E. and Van Wincoop, E. (2003). Gravity with gravitas: a solution to the border puzzle. *The American Economic Review*, 93(1):170–192.

Baltagi, B. H., Egger, P., and Pfaffermayr, M. (2003). A generalized design for bilateral trade flow models. *Economics Letters*, 80(3):391–397.

Baltagi, B. H., Egger, P., and Pfaffermayr, M. (2014). Panel data gravity models of international trade. *CESifo Working Paper Series No. 4616*.

Barrios, T., Diamond, R., Imbens, G. W., and Kolesár, M. (2012). Clustering, spatial correlations, and randomization inference. *Journal of the American Statistical Association*, 107(498):578–591.

Bartelsman, E. J. and Beetsma, R. M. (2003). Why pay more? Corporate tax avoidance through transfer pricing in OECD countries. *Journal of Public Economics*, 87(9):2225–2252.

Bastani, S. and Selin, H. (2014). Bunching and non-bunching at kink points of the Swedish tax schedule. *Journal of Public Economics*, 109:36–49.

Beer, S. (2015). Productive and Unproductive Reactions of MNEs to International Tax Differentials. *mimeo*.

Beer, S. and Loeprick, J. (2015). Profit shifting: drivers of transfer (mis) pricing and the potential of countermeasures. *International Tax and Public Finance*, 22(3):426–451.

Bettendorf, L., Lejour, A., and van 't Riet, M. (2016). Tax bunching by owners of small corporations. CPB Discussion Paper 326, CPB Netherlands Bureau for Economic Policy Analysis.

Böhm, T., Karkinsky, T., and Riedel, N. (2015). Corporate Taxes and Strategic Patent Location within Multinational Firms. *mimeo*.

Blaufus, K., Bob, J., Hundsdoerfer, J., Sielaff, C., Kiesewetter, D., and Weimann, J. (2015). Perception of income tax rates: evidence from Germany. *European Journal of Law and Economics*, 40(3):457–478.

Bloom, A. D. et al. (1991). *The Republic of Plato*. Basic Books.

Bloom, N., Griffith, R., and Van Reenen, J. (2002). Do R&D tax credits work? Evidence from a panel of countries 1979–1997. *Journal of Public Economics*, 85(1):1–31.

Bolton, G. E. and Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American economic review*, pages 166–193.

Bronzini, R. and Iachini, E. (2014). Are incentives for R&D effective? Evidence from a regression discontinuity approach. *American Economic Journal: Economic Policy*, 6(4):100–134.

Buettner, T., Overesch, M., Schreiber, U., and Wamser, G. (2012). The impact of thin-capitalization rules on the capital structure of multinational firms. *Journal of Public Economics*, 96(11):930–938.

Buettner, T. and Wamser, G. (2013). Internal debt and multinational profit shifting: Empirical evidence from firm-level panel data. *National Tax Journal*, 66(1):63–96.

Cameron, A. C. and Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372.

Cameron, A. C. and Trivedi, P. K. (1986). Econometric models based on count data. comparisons and applications of some estimators and tests. *Journal of applied econometrics*, 1(1):29–53.

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., Olsen, T., and Pistaferri, L. (2011). Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *The Quarterly Journal of Economics*, 126(2):749–804.

Clausing, K. A. (2003). Tax-motivated transfer pricing and US intrafirm trade prices. *Journal of Public Economics*, 87(9):2207–2223.

CPB et al. (2015). A Study on R&D Tax Incentives-Final report. Technical report, Directorate General Taxation and Customs Union, European Commission.

Cristea, A. D. and Nguyen, D. X. (2013). Transfer Pricing by Multinational Firms: New Evidence from Foreign Firm Ownerships. *mimeo*.

Czarnitzki, D., Hall, B. H., and Hottenrott, H. (2014). Patents as quality signals? The implications for financing constraints on R&D. Technical report.

Davies, R. B., Martin, J., Parenti, M., and Toubal, F. (2015). Knocking on Tax Haven's Door: Multinational Firms and Transfer Pricing. *CBT Working Papers WP 15/01*.

Deelen, A. and Euwals, R. (2014). Do wages continue increasing at older ages? Evidence on the Wage Cushion in the Netherlands. *CPB Discussion Paper*, 282.

Deloitte (2015). Global Transfer Pricing Country Guide. Technical report.

Dharmapala, D. (2014). What Do We Know about Base Erosion and Profit Shifting? A Review of the Empirical Literature. *Fiscal Studies*, 35(4):421–448.

Diamond, P. A. and Mirrlees, J. A. (1971a). Optimal taxation and public production I: Production efficiency. *The American Economic Review*, 61(1):8–27.

Diamond, P. A. and Mirrlees, J. A. (1971b). Optimal taxation and public production II: Tax rules. *The American Economic Review*, 61(3):261–278.

Dischinger, M. and Riedel, N. (2011). Corporate taxes and the location of intangible assets within multinational firms. *Journal of Public Economics*, 95(7):691–707.

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

Ernst&Young (2014). Worldwide Transfer Pricing Reference Guide. Technical report.

European Commission et al. (2015). Taxation trends in the European Union. Technical report.

Evers, L., Miller, H., and Spengel, C. (2015). Intellectual property box regimes: effective tax rates and tax policy considerations. *International Tax and Public Finance*, 22(3):502–530.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868.

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.

Fisher, R. A. (1935). The design of experiments. *Oliver and Boyd, Edinburgh.*

Freedman, D. and Diaconis, P. (1981). On the histogram as a density estimator: L2 theory. *Probability Theory and Related Fields*, 57(4):453–476.

Gómez-Herrera, E. (2013). Comparing alternative methods to estimate gravity models of bilateral trade. *Empirical Economics*, 44(3):1087–1111.

Griffith, R., Miller, H., and O'Connell, M. (2014). Ownership of intellectual property and corporate taxation. *Journal of Public Economics*, 112:12–23.

Gruber, J. and Saez, E. (2002). The elasticity of taxable income: Evidence and implications. *Journal of Public Economics*, 84(1):1–32.

Grubert, H. (2003). Intangible income, intercompany transactions, income shifting, and the choice of location. *National Tax Journal*, pages 221–242.

Harju, J. and Matikka, T. (2016). The Elasticity of Taxable Income and Income-shifting: What is" real" and what is Not? *International Tax and Public Finance*, 23:640–669.

Hines Jr, J. R. (1994). No place like home: Tax incentives and the location of R&D by American multinationals. In *Tax Policy and the Economy, Volume 8*, pages 65–104. MIT Press.

Ho, D. E. and Imai, K. (2006). Randomization inference with natural experiments: An analysis of ballot effects in the 2003 California recall election. *Journal of the American Statistical Association*, 101(475):888–900.

Huizinga, H. and Laeven, L. (2008). International profit shifting within multinationals: A multi-country perspective. *Journal of Public Economics*, 92(5):1164–1182.

Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *The American Economic Review*, 104(2):537–563.

Jongen, E. L. W. and Stoel, M. (2013). Estimating the Elasticity of Taxable Labour Income in the Netherlands. *CPB Background Document.*

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the econometric society*, pages 263–291.

Karkinsky, T. and Riedel, N. (2012). Corporate taxation and the choice of patent location within multinational firms. *Journal of International Economics*, 88(1):176–185.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, 8:435–464.

Kleven, H. J., Knudsen, M. B., Kreiner, C. T., Pedersen, S., and Saez, E. (2011). Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark. *Econometrica*, 79(3):651–692.

Kleven, H. J. and 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.

KPMG (2014). Global Transfer Pricing Review. Technical report.

Le Maire, D. and Schjerning, B. (2013). Tax bunching, income shifting and self-employment. *Journal of Public Economics*, 107:1–18.

Lohse, T. and Riedel, N. (2013). Do transfer pricing laws limit international income shifting? Evidence from European multinationals. Technical report, Oxford University Centre for Business Taxation.

Lokshin, B. and Mohnen, P. (2013). Do R&D tax incentives lead to higher wages for R&D workers? Evidence from The Netherlands. *Research Policy*, 42(3):823–830.

Loretz, S. (2008). Corporate taxation in the OECD in a wider context. *Oxford Review of Economic Policy*, 24(4):639–660.

Mardan, M. (2017). Why countries differ in thin capitalization rules: The role of financial development. *European Economic Review*, 91:1–14.

Mason, R. and Calvin, L. D. (1984). Public confidence and admitted tax evasion. *National Tax Journal*, pages 489–496.

Merlo, V., Riedel, N., Wamser, G., et al. (2015). The Impact of Thin Capitalization Rules on the Location of Multinational Firms' Foreign Affiliates. Technical report, CESifo Group Munich.

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

Nicolay, K., Nusser, H., and Pfeiffer, O. (2016). On the interdependency of profit-shifting channels and the effectiveness of anti-avoidance legislation. *mimeo*.

OECD (2015). Designing Effective Controlled Foreign Company Rules, Action 3 - 2015 Final Report. Technical report, OECD/G20 Base Erosion and Profit Shifting Project.

PWC (2016). International Transfer Pricing Guide. Technical report.

Ramsey, F. P. (1927). A Contribution to the Theory of Taxation. *The Economic Journal*, 37(145):47–61.

Saez, E. (2010). Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2(3):180–212.

Saez, E., Slemrod, J., and 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. and Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *The American Economic Review*, 106(1):24–45.

Schindler, D. and Schjelderup, G. (2013). Transfer pricing and debt shifting in multinationals. *mimeo.*

Schwab, T. and Todtenhaupt, M. (2016). Spillover from the haven: Cross-border externalities of patent box regimes within multinational firms. Technical report, ZEW-Zentrum für Europäische Wirtschaftsforschung/Center for European Economic Research WP No. 16-073.

Silva, J. S. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and statistics*, 88(4):641–658.

Silva, J. S. and Tenreyro, S. (2011). Further simulation evidence on the performance of the poisson pseudo-maximum likelihood estimator. *Economics Letters*, 112(2):220–222.

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

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

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy*, 64(5):416–424.

Tinbergen, J. (1962). *Shaping the world economy; suggestions for an international economic policy.* Twentieth Century Fund, New York.

Voget, J. (2011). Relocation of headquarters and international taxation. *Journal of Public Economics*, 95(9):1067–1081.

Weber, C. E. (2014). Does the Earned Income Tax Credit Reduce Saving by Low-Income Households? *University of Oregon Working Paper.*

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.

Zinn, T., Riedel, N., and Spengel, C. (2014). The Increasing Importance of Transfer Pricing Regulations: A Worldwide Overview. *Intertax*, 42(6):352–404.

Declaration

I herewith declare that I have completed this doctoral thesis alone and without any further help other than indicated in the declaration of co-authorship. I have used the sources listed in the bibliography and only those sources in writing my dissertation. All information taken from these sources has been referenced accordingly.