Essays on International Tax and Firm Behavior

by

Lysle T. Boller

Department of Economics
Duke University

Date: _____

Approved:

_____
Allan Collard-Wexler, Advisor

_____
Juan Carlos Suárez Serrato

_____
Rosanne Altshuler

_____
James Roberts

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University

2023

ABSTRACT

Essays on International Tax and Firm Behavior

by

Lysle T. Boller

Department of Economics
Duke University

Date: _____

Approved:

_____
Allan Collard-Wexler, Advisor

_____
Juan Carlos Suárez Serrato

_____
Rosanne Altshuler

_____
James Roberts

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University

2023

Abstract

This dissertation examines unique tax incentives faced by multinational corporations (MNCs). Although nearly all businesses engage in profit-maximizing behavior, many do not have the ability to set up numerous subsidiaries in multiple jurisdictions. MNCs, on the other hand, often operate in a global context. The largest MNCs have hundreds of related subsidiaries that are distributed worldwide. Firms can utilize this geographic dispersion to take advantage of heterogeneity in tax laws across the jurisdictions in which they operate to minimize their tax burden. In the following chapters, I examine (1) common mechanisms used by some of the largest MNCs that facilitate the movement of income from high-tax jurisdictions to low-tax jurisdictions, (2) the extent to which researchers can trust different sources of data to measure the foreign activity of MNCs, and (3) how regulatory uncertainty concerning tax regimes can distort various aspects of firm behavior and financial reporting.

The second chapter studies the adoption and use of specific forms of tax planning by US multinational corporations (MNCs). Along with my coauthors, I use IRS data to identify hybrid tax planning affiliates (HTPs) that allow MNCs to avoid corporate income taxes by targeting mismatches between US and Irish, Dutch, and Luxembourg tax law. By 2016, we find that more than 35% of the foreign profits of US MNCs were linked to these HTPs. Difference-in-difference models reveal that, after adoption, corporations intensify behavior commonly linked to profit shifting, significantly increasing related-party loans, book values of foreign intangible assets, and profits held abroad relative to control MNCs. These changes coincide with stark reductions in foreign effective tax rates. MNCs that adopt HTPs also experience large increases in foreign tangible assets and in domestic R&D, payroll, and investment relative to other types of multinationals. To separate selection and treatment effects, we develop and estimate a model that rationalizes the selection of MNCs into HTPs and the changes that we observe in their reported economic activity.

The third chapter examines whether data that is commonly used to study MNCs can be trusted by researchers. Recent literature has noted the potential for significant measurement error, unique to MNCs, in tax and accounting data. I examine a key source of measurement error that can affect administrative tax data — aggregation error. I link data from tax filings and public disclosures to quantify the extent to which commonly-used aggregation techniques may result in double counting of foreign earnings. A comparison of book and tax data reveals that aggregation error has been increasing over time. The matched sample also reveals large inconsistencies in the reporting of corporate income tax across firms' books and tax filings, particularly in extractive and financial industries. I introduce a simple correction for aggregation error and examine the extent to which it harmonizes measurement of foreign income and tax rates across firms' books and tax filings. Applying this correction yields a 30% reduction in the magnitude of foreign earnings as measured in tax data in 2016 and significantly reduces book-tax differences. Furthermore, this correction breaks the systematic relationship between book-tax differences and the size of multinationals' foreign affiliate networks. Unadjusted book-tax differences are increasing over time. After applying the correction, this is no longer true. Finally, I replicate estimates from prior literature that do not correct for aggregation error and discuss robustness.

The fourth chapter examines how MNCs reacted to a regulatory loophole that emerged after the implementation of transfer pricing regulations passed in 1995. This loophole allowed US-based MNCs to shift labor costs from low-tax foreign affiliates to US parent companies by utilizing stock-based compensation. A 2005 tax court ruling created regulatory uncertainty around enforcement of the loophole, increasing the market value of firms that were in a position to take advantage of it. This period of regulatory uncertainty extended for over a decade. During this period, exposed firms increased their overall usage of stock-based compensation and its relative intensity as a share of their total labor costs. Exposed firms also increased their overall R&D activity and claimed domestic tax credits for larger portions of their overall R&D expense.

Contents

List of Tables

List of Figures

Acknowledgements

Thank you to my commitee members, Juan Carlos Suárez Serrato, Rosanne Altshuler, Allan Collard-Wexler, and Jimmy Roberts for their support.

To my coauthors, Juan Carlos and Rosanne. The start of the research phase of my doctorate program coincided with the onset of the pandemic, and the many Zoom meetings I spent with them kept me going. I consider myself incredibly lucky to have worked with them.

To Fiona Scott Morton, my first coauthor, for showing me the ropes.

To Mom and Dad, who taught me to distinguish between what is important and what is trivial, and to my siblings, Jack and Katie, who taught me resiliency.

To Emily, for everything.

1

Introduction

This research was partially conducted as part of a Joint Statistical Research Program with the Internal Revenue Service (IRS). Any views expressed are those of the authors and not those of the IRS. The statistics reported in this paper have been reviewed and cleared for disclosure by the IRS.

The study of taxation is at the root of the field of economics. At its heart, the subfield is born out of the failure of the second welfare theorem. Governments exist only because they are able to raise revenue, and feasible tax instruments distort the behavior of economic agents. This naturally leads to the question of optimal taxation, which in turn depends on the degree to which tax instruments generate unproductive activity. This dissertation does not try to determine which tax instruments are more distortionary than others. Instead, it examines some of the distortions that are generated by one such instrument — the corporate income tax.

This instrument is closely tied to the international context for a simple reason — the success of a corporate income tax requires coordination. This is implied by the seminal work on optimal taxation and production efficiency of Diamond and Mirrlees (1971*a*) and Diamond and Mirrlees (1971*b*), who showed that even with a failure of the second welfare theorem, optimal tax regimes should equate optimal rates of transformation across sectors. Countries, however, do not generally coordinate their tax systems. As a result, if capital is mobile, cross-jurisdictional differences in the tax treatment of corporations will generate deadweight loss. Even if governments try to set optimal corporate taxes, the interaction of their independent strategic behavior will only exacerbate public inefficiency (Kanbur and Keen, 1993; Zodrow and Mieszkowski, 1986; Keen and Konrad, 2013).

There are at least two ways to think about the mobility of capital. The first is across firms — investors are more likely to allocate capital to businesses where they can achieve higher post-tax rates of return. These rates are higher where the corporate income tax is lower, *ceteris paribus*. Firms in low-tax jurisdictions will therefore tend to grow more than firms in other jurisdictions even if they are not more productive. The second is within firms. Firms that have the ability to shift their internal capital across multiple jurisdictions will face incentives to allocate it unproductively. In the subnational context, this can occur when firms move or reallocate capital between states (Fajgelbaum, Morales, Suárez Serrato and Zidar, 2018; Suárez Serrato and Zidar, 2016). In the global context, within-firm capital mobility can be thought of as occuring within multinational corporations (MNCs). MNCs are defined by their ability to operate across multiple countries. Large MNCs can have hundreds of related subsidiaries distributed across the globe, and these agents, unlike domestic firms, are the ones best suited to exploit the main weakness of the corporate income tax.

MNCs, unlike domestic firms, can use their ability to move capital across different jurisdictions to minimize their tax burden. This advantage has only grown stronger since the early 1990s in tandem with global economic integration and a permissive regulatory environment (Altshuler and Grubert, 2006*a*,*b*; Mutti and Grubert, 2009). As a result, MNCs receive a natural competitive advantage that is unrelated to their productivity, distorting capital allocations. Worse, these MNCs may take a hit to their productivity if tax optimization is tied to real economic activity. In Chapter 2, a coauthored work, we study how multinationals use specific forms of tax planning. We use IRS data to identify hybrid tax planning affiliates (HTPs). These affiliates are "hybrids" because they are viewed as passthrough entities by one government and opaque entities by another.[1] We focus on entities that have structures in Ireland, the Netherlands, and Luxembourg because they have been widely covered in the press as common domiciles for HTPs, and the mechanisms

[1] I use passthrough in the sense that the activities of these entities are considered to be consolidated with their owners, and opaque to mean the opposite.

by which MNCs can shift profits through them are well known. There is a lack of rigorous information, however, on their prevalence and significance relative to the population of multinationals. By 2016, we find that more than 35% of the foreign profits of US MNCs were linked to these HTPs. We provide formal comparisons of MNCs that adopt these structures with non-adopting MNCs by estimating modern staggered difference-in-difference models. These estimates reveal that, after adoption, MNCs intensify behavior commonly linked to profit shifting, significantly increasing related-party loans, book values of foreign intangible assets, and profits held abroad relative to control MNCs. These changes coincide with stark reductions in foreign effective tax rates. Evidence suggests that these structures may also generate real economic responses by firms (as opposed to just reflecting the movement of accounts) — MNCs that adopt HTPs experience large increases in foreign tangible assets and in domestic R&D, payroll, and investment relative to other types of multinationals. While these estimates are useful comparisons that provide new information about MNC tax planning, we acknowledge that they depend on strong identifying assumptions for a causal interpretation. To separate selection and treatment effects, we develop and estimate a model that rationalizes the selection of MNCs into HTPs and the changes that we observe in their reported economic activity.

The third chapter examines whether data that is commonly used to study MNCs has fundamental flaws, and if so, whether it can be salvaged. Recent literature has noted the potential for significant measurement error, unique to MNCs, in tax and accounting data. I examine a key source of measurement error that can affect administrative tax data — aggregation error. I construct a proxy for this error by generating a link between IRS data, which provides information about US parents and the activity of their foreign affiliates, and data from public disclosures on the consolidated multinational. By comparing the disaggregated data to the consolidated figures, I can quantify the extent to which commonly-used aggregation techniques may result in double counting of foreign earnings in tax data. This book-tax comparison reveals that aggregation error has been increasing over time, likely

because the mechanism by which this error occurs suggests that it grows as MNCs create more complicated affiliate networks. The linked sample also reveals large inconsistencies in the reporting of corporate income tax across firms' books and tax filings. These inconsistencies are particularly noticeable in extractive and financial industries. Extractive industries often operate on concessions from foreign governments that are structured by contracts that share revenue, which can be misreported as a corporate income tax. To address aggregation error that affects foreign income, I introduce a simple correction proposed by Blouin and Robinson (2020), and examine the extent to which it harmonizes measurement of foreign income and tax rates across firms' books and tax filings. Applying this correction yields a 30% reduction in the magnitude of foreign earnings as measured in tax data in 2016 and significantly reduces book-tax differences. Furthermore, this correction breaks the systematic relationship between book-tax differences and the size of multinationals' foreign affiliate networks. Unadjusted book-tax differences are increasing over time. After applying the correction, this is no longer true. Both of these exercises indicate that the correction appears to significantly reduce measurement error. Finally, I replicate estimates from prior literature that do not correct for aggregation error and discuss robustness.

The fourth chapter examines how MNCs reacted to a regulatory loophole that emerged after the implementation of transfer pricing regulations passed in 1995. This loophole allowed US-based MNCs to shift labor costs from low-tax foreign affiliates to US parent companies by utilizing stock-based compensation. Unlike Chapter 2, which focuses mainly on how MNCs shift profits between foreign (non-US) countries, this chapter indicates that US multinationals also use related structures to shift income out of the United States, eroding its corporate tax base. I use IRS tax data to identify firms that were in a position to take advantage of this loophole. I then show that a 2005 tax court ruling created regulatory uncertainty around enforcement of the loophole and increased the market value of firms that were in a position to take advantage of it. This period of regulatory uncertainty extended for over a decade. During this period, exposed firms increased their overall usage of stock-

based compensation and its relative intensity as a share of their total labor costs. Exposed firms also increased their overall R&D activity and claimed domestic tax credits for larger portions of their overall R&D expense, suggesting another avenue for tax base erosion.

2

Tax Planning and Multinational Behavior

2.1 Introduction

The last quarter century has seen a remarkable increase in the complexity of tax planning by multinational corporations (MNCs). Document leaks and special government reports have revealed the existence of tax planning strategies that are designed to avoid corporate income taxes in multiple jurisdictions by exploiting mismatches in tax laws. While media attention following these revelations and dissatisfaction with the current system have motivated important international tax policy changes and multilateral projects, little is known about the importance of these tax planning strategies.[1] How prevalent are these strategies among MNCs? Do they facilitate profit shifting and lower foreign effective tax rates (ETRs)? Does tax planning influence the real economic activity of MNCs?

This paper uses IRS tax data to answer these questions by analyzing the adoption and use of a set of complex tax planning strategies that target mismatches in Irish, Dutch, and Luxembourgish tax law. We first use data from multiple IRS tax forms to reconstruct the ownership networks of the foreign affiliates of US MNCs. We then identify when a US MNC creates an ownership structure that can exploit mismatches in tax laws across these countries. Data from the tax returns of US MNCs is crucial to identifying the adoption of these tax planning structures and how MNCs use them to shift profits across countries. This paper is the first to systematically uncover these tax planning structures.

We use these administrative data to document the growth and prevalence of these tax planning strategies. Although these strategies were extremely rare in the early 1990s, they

[1] These changes and projects include important aspects of the recent Tax Cuts and Jobs Act of 2017, the European Anti-Tax Avoidance Directive, and the OECD's Base Erosion and Profit Shifting project.

were gradually adopted by MNCs following the 1997 regulations known as "Check the Box" (CTB), which facilitated this form of tax planning. By 2016, 16% of US MNCs in our sample had adopted at least one of the structures we identify, and these companies were responsible for more than 60% of foreign profits. Data on the foreign operations of adopting MNCs show that more than 50% of their foreign profits flow through one of these structures and that more than 35% of the foreign profits of MNCs in our sample are linked to these structures. MNCs that use these strategies are also responsible for significant shares of domestic economic activity, including 25% of domestic corporate payroll and 15% of domestic capital investment.

Relative to MNCs that never adopt these structures, those that do engage more intensely in the kinds of financial transactions that could be used to shift profits to low tax countries: they increase their foreign holdings of intangible capital, they have larger loan balances between related foreign affiliates, and they collect more royalty income and accumulate more cash abroad. Most strikingly, as shown in Figure 2.1, over the period during which these structures are gradually put in place, the foreign ETRs of adopting MNCs experience a dramatic decline. By 2010, the foreign ETR of adopters was roughly half that of other MNCs.

We use difference-in-differences event-study regressions to conduct more formal comparisons. These models measure changes in firm outcomes surrounding the adoption of these structures, relative to non-adopting firms. These models confirm that the marked increase in behaviors related to profit shifting coincides with the adoption of these structures. We then estimate that foreign ETRs see a relative decline of 18.4% ($p < 0.01$) six years after adopting one of these strategies. Concurrent with declining foreign ETRs, we estimate that adopting US MNCs have larger increases in foreign investment and accumulate 18% ($p < 0.01$) more depreciable capital in foreign affiliates than MNCs that do not adopt one of these structures. We also estimate that the adoption of these tax planning structures is followed by a 17% ($p < 0.01$) increase in domestic payroll and a 40% ($p < 0.01$) increase in

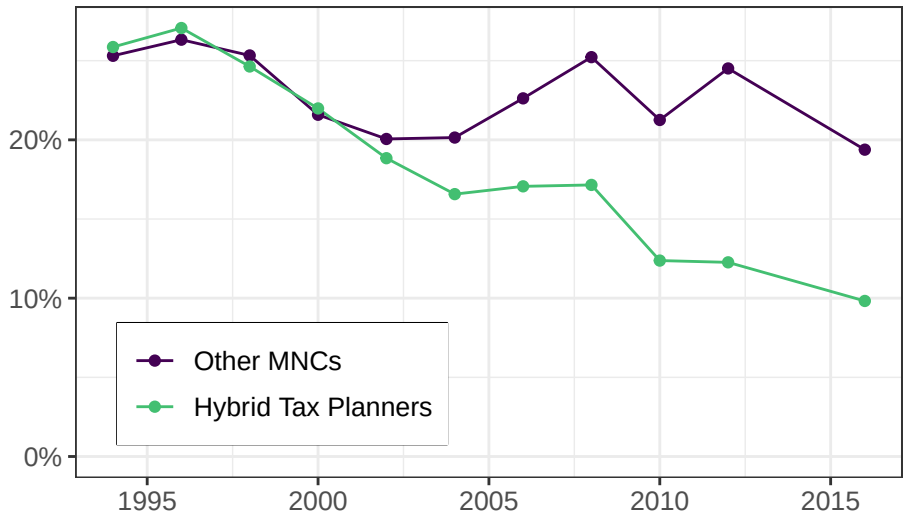

Figure 2.1: Comparison of Foreign Effective Tax Rates

Notes: This figure compares the aggregate annual foreign effective tax rate (ETR) for two groups of US multinationals (MNCs). The light green line shows the ETR for a group of MNCs that eventually adopt at least one of the hybrid tax structures described in Section 2.2. The dark blue line shows the ETR for MNCs that did not adopt any of these structures during the sample window. The combined sample includes most large US C corporations as described in Section 2.3.

research and development (R&D) tax credits.

We integrate these results by developing and estimating a model of MNC behavior under tax planning. In the model, firms deliver services in each destination country based on a shared, group-level productivity. MNCs can improve their global productivity by engaging in R&D. MNCs can also pay a fixed cost to engage in tax planning, which allows their foreign profits to be taxed at a lower rate. MNCs decide whether to adopt a tax planning strategy based on their profitability, their potential for R&D growth, and other determinants of tax aggressiveness. The model integrates our reduced-form results by allowing for both the selection of MNCs into tax planning and the reductions in foreign tax costs to impact their domestic and foreign activities. Our model estimates suggest that domestic outcomes are more likely to be driven by selection, while a significant share of foreign investment is the result of the lower taxes generated by tax planning. These results indicate that tax planning has small effects on the US economy and that limiting tax planning may reduce

foreign investment with little impact on tax revenue.

We develop our results in four steps. First, we reconstruct the foreign ownership structures of US MNCs using information from three key tax forms contained in the IRS data files. Parent-level data from Form 1120 and related tax forms provide information on domestic activity, including assets, payroll, domestic investment, and R&D tax credits. Parents also file an information form (Form 5471) for each of their controlled foreign corporations (CFC), which includes data on foreign assets, taxes, and earnings and profits (E&P), as well as related transactions between CFCs. With the advent of CTB, US MNCs could elect to "disregard" their foreign affiliates. These foreign disregarded entities (FDEs) are hybrid structures that are corporations in the host country but pass-through entities from the US perspective, making them transparent to the US Treasury. The IRS collects information about these entities in a separate information return, Form 8858. A novel aspect of this paper is the integration of FDE data with their CFC owners.

Second, with this information in hand, we flag MNCs that adopt tax planning structures that have previously only been revealed in leaks and special government reports. These planning structures use entities that facilitate tax planning by creating mismatches in their tax treatment between the United States and foreign countries. The first structure we identify, known as the Double Irish, uses FDEs to exploit aspects of Irish and US tax law and shift profits out of high-tax foreign countries into tax havens. The second and third structures use a strategy known as a Reverse Hybrid Mismatch to accumulate untaxed foreign profits in a foreign affiliate. We identify usage of Reverse Hybrid Mismatches for foreign affiliates located in the Netherlands and in Luxembourg. These three hybrid tax planning structures are well known and have been targeted by European countries through legal investigations.[2] They have also motivated policy agendas such as Action 2 of the OECD Base Erosion and Profit Shifting (BEPS) project, which aims to close down hybrid

[2]Companies such as Google (Drucker, 2010) and Apple (Duhigg and Kocieniewski, 2012) have been reported to have used a Double Irish structure, while companies such as Nike (Guardian, 2019), Amazon (Guardian, 2018), and Starbucks (Kleinbard, 2013) have been reported to have used Reverse Hybrid Mismatch arrangements through the Netherlands or Luxembourg.

mismatch arrangements.

Using these unique indicators of tax planning, we measure the growth, prevalence, and importance of these hybrid tax planning structures relative to the aggregate economic activity of US MNCs. As mentioned above, the three structures we identify are connected to large shares of aggregate foreign activity by US MNCs by 2016, the end of our sample period. Although there have been numerous investigations and case studies that reveal how these arrangements work, this paper is the first attempt to systematically measure their importance for foreign and domestic economic activity.

In our third step, we compare the foreign and domestic operations of US MNCs that ever adopted one of these strategies to those that do not. In the early 1990s, these two groups of firms have nearly identical patterns of ETRs and domestic and foreign activity. After the 1997 CTB regulations are put in place, foreign affiliates of MNCs that engage in hybrid tax planning (HTP) experience significantly larger declines in ETRs, as well as markedly larger increases in intangible assets, royalty income, and cash held abroad—all of which are consistent with profit shifting and tax deferral.

To ensure that these results are tied to the adoption of HTPs, we estimate staggered difference-in-differences models around the first year that a given MNC adopts an HTP. Estimates from these models provide convincing evidence that these structures were put in place for tax avoidance purposes.[3] Specifically, we show that MNCs engage in more financial transactions that can be used to shift profits and reduce foreign ETRs following the adoption of an HTP. We also find that adopting MNCs experience larger increases in foreign capital, domestic payroll, and R&D than non-adopting MNCs in the years surrounding the adoption of an HTP. One potential concern is that these complex tax structures are adopted by specific types of firms that were also subject to other macroeconomic shocks over this time period. We address this concern by showing that we obtain similar results when we include industry-by-year fixed effects, when we flexibly control for firm size bins interacted with year

[3]Throughout the paper we use tax avoidance to refer to legal strategies used to minimize tax obligations.

fixed effects, and when we flexibly control for differences in intangible assets across firms. We also obtain similar results when we use a two-way fixed effects specification and when we use the estimator of Sun and Abraham (2021) that is robust to potential contamination bias resulting from staggered adoption.

Estimates of our staggered difference-in-differences models have a causal interpretation under the assumption that the outcomes of MNCs that adopt an HTP would have otherwise trended similarly to those that did not adopt such a structure. Event study results generally show that important outcomes for adopting MNCs, such as growth in payroll, foreign assets, R&D, and foreign ETRs follow similar patterns to those of non-adopting MNCs prior to HTP adoption. These results suggest that MNCs do not select into HTPs based on prior trends in economic outcomes. However, it is possible that MNCs select into HTPs based on the gains from adopting an HTP, such that adopting MNCs would have different post-period trends absent an HTP.

The fourth and final step of the paper develops and estimates a model of tax planning and multinational behavior that incorporates selection on gains into tax planning. In the model, firms trade off heterogeneous costs and benefits of setting up tax planning structures against the potential benefits from lower foreign ETRs. The model relates the difference-in-differences estimates to model parameters and a selection term. The model parameters measure how foreign ETRs impact foreign investment and how R&D investment impacts productivity. The selection term captures differences in the underlying productivity of the firms that adopt hybrid tax planning structures. By separating these forces, the model quantifies the degree to which investment by hybrid tax planning firms is inframarginal, in the sense that some of their increased investment would have occurred in the absence of tax planning. Through the lens of our model, the changes in domestic economic activity are more likely to be driven by the selection of firms with potential to increase their R&D investment. This result would imply that there are smaller benefits to the US economy from foreign tax planning of US MNCs than the difference-in-differences estimates would suggest.

At the same time, the model also shows that increases in foreign capital are more likely to be driven by tax differentials, such that efforts to limit tax planning may significantly limit investment abroad.

Overall, our results provide the first systematic documentation of the prevalence of hybrid tax planning structures. By 2016, these structures accounted for about a third of foreign profits of US MNCs. These MNCs also play an important role in the domestic economy. We estimate that adoption of these strategies is followed by large increases in financial transactions that can be used to shift profits to low tax countries and by declines in foreign ETRs. We also estimate significant relative changes in domestic and foreign economic activity following the adoption of an HTP.

This paper contributes to studies that quantify the importance of profit shifting. In a seminal contribution, Hines and Rice (1994) describe and measure the importance of tax havens to the operations of US multinationals. Clausing (2016) uses the sensitivity of reported profits to tax rate differentials to estimate the magnitude of profit shifting of US MNCs. Tørsløv, Wier and Zucman (2018) use macroeconomic data and differences in the profitability of different affiliates to estimate the magnitude of profits shifted to tax havens. Bilicka (2019) uses tax data from the UK to argue that the large differences in the profitability between domestic UK firms and the affiliates of foreign MNCs (in the UK) are driven by profit shifting. In a recent survey, Dyreng and Hanlon (2021) highlight the cross-sectional variation in tax avoidance and conclude that a large portion of the variation in tax avoidance remains unexplained. Our focus on HTPs contributes to our understanding of the importance of specific tax planning strategies.

While several papers document the existence of profit shifting, the magnitude of this problem has been hard to pin down. Using tax data from US firms, Dowd, Landefeld and Moore (2017) argue that the sensitivity of reported profits to tax differentials can be non-linear and that accounting for non-linearities increases estimates of profits shifted to low tax countries. In contrast, Blouin and Robinson (2020) argue that prior estimates using

tax and survey data can be plagued by double counting of profits and that accounting for direct investment income between affiliates can significantly lower estimates of profit shifting. Following their suggestions, we subtract dividend income from related foreign corporations when computing aggregate foreign earnings for US MNCs. Rather than provide estimates of profit shifting, we document the prevalence of widely used tax planning structures, show that MNCs use them in transactions that are likely related to profit shifting, and find that close to one third of the foreign profits of US MNCs in our sample flow through these structures by 2016.

We also contribute to our understanding of CTB regulations by directly examining the adoption and consequences of complex tax structures facilitated by the policy. Using tax return data from US MNCs, Altshuler and Grubert (2006*b*) find reductions in foreign effective tax rates after the enactment of CTB in 1997 that are consistent with the use of the tax planning strategies we examine. Mutti and Grubert (2009) use multiple data sources to show that, after the implementation of CTB, MNCs increased profit shares in low-tax jurisdictions and conducted transfers of intangible assets abroad. Blouin and Krull (2014) show that MNCs had more tax haven affiliates and longer ownership chains after the enactment of CTB.[4] Faulkender, Hankins and Petersen (2019) study the 440 percent increase in cash held abroad by US MNCs between 1998 and 2008 and argue that this rise is driven by tax factors, including CTB. While prior research assumed that MNCs disregarded foreign affiliates following the CTB regulations, ours is the first paper to use tax information to confirm when an affiliate is disregarded and to systematically identify MNCs that use a specific set of tax planning structures. In a contemporaneous paper, Samarakoon (2022) uses tax data to identify firms that use a Double Irish structure and examines how the closure of this structure impacts repatriation of deferred earnings by MNCs.[5] Our results show that

[4]Grubert (2012) finds that foreign ETRs of US MNCs declined by 5 percentage points between 1996 and 2004 using tax return data. By examining CFCs that indicate that they have a disregarded entity he is able to determine that 1 to 2 percentage points of this decline is due to CTB.

[5]Hardeck and Wittenstein (2018) use data from the Luxembourg Leaks to identify firms with hybrid tax structures and find that hybrid tax structures reduce MNC tax rates, as measured by financial statements data.

the bulk of the decrease in foreign ETRs and increase in cash held abroad over the sample period was driven by MNCs that adopted a particular set of tax planning structures.

Finally, our paper contributes to the literature on how profit shifting impacts real behavior. Grubert and Slemrod (1998) study profit shifting opportunities through Puerto Rico and argue that US MNCs changed their investment decisions in response to these opportunities. Suárez Serrato (2018) studies the reduction in profit shifting opportunities through Puerto Rico and shows that US MNCs decrease their domestic investment in response.[6] Albertus (2019) uses Bureau of Economic Analysis data to compare US MNCs with different average foreign tax rates prior to the implementation of CTB. He finds that MNCs with higher initial tax rates experience a larger decline in average rates and increase their R&D intensity after 1996. By using tax data to identify specific tax planning structures and to demonstrate how they are used for profit shifting, we reveal substantial heterogeneity regarding which firms benefitted from CTB and shed light on the mechanisms through which CTB lowered foreign ETRs and affected real economic activity.

The remainder of the paper is organized as follows. Section 2.2 describes the three hybrid tax planning structures we examine and discusses how CTB facilitated their adoption by making it easy to create foreign disregarded entities. Section 2.3 provides an overview of the data. Section 2.4 discusses how the creation of foreign disregarded entities and hybrid tax planning structures has grown over time. Section 2.5 estimates the effects of adopting tax planning strategies on foreign tax rates and economic activity of US MNCs. Section 2.6 provides a model of international MNC investment in which firms select into tax planning by trading off heterogeneous setup costs that vary based on levels of IP. Section 2.7 concludes.

[6] de Mooij and Liu (2018) study the impact of transfer pricing regulations and show that these policies can reduce investment. Bilicka, Qi and Xing (2019) shows that a world-wide debt cap that limited interest stripping as a form of profit shifting also impact the investment decisions of UK MNCs.

2.2 Hybrid Tax Planning Structures and Check The Box

This section describes the three hybrid tax planning structures that we study. We first describe the "Check the Box" regulations that facilitated their adoption and then describe the structures in detail.

2.2.1 Check The Box Regulations

During the period we study, the United States imposed corporate tax on the worldwide income of US corporations with a credit for foreign taxes paid to avoid double taxation. The credit was limited to what US tax would have been on the foreign income. Taxes were not due on active foreign business income until it was repatriated to the parent corporation. This deferral feature of the US tax code made it attractive to hold income generated abroad in tax havens.

To prevent profit shifting, deferral was not extended to certain types of "tainted income" under what is generally referred to as Controlled Foreign Corporations rules. These rules are contained in Subpart F of the tax code and foreign income that is subject to current US tax is referred to as "Subpart F" income. Tainted income includes passive portfolio income and the payment of interest, dividends, and royalties from one CFC to a CFC in another jurisdiction.

In 1996, Treasury promulgated regulations effective on January 1, 1997 that made it easier for US corporations to change the entity classification (e.g., pass-through or corporate) of domestic and foreign affiliates. This policy change became known as Check the Box, referring to the ease with which US corporations could change entity classifications. CTB was originally intended to simplify tax filing for domestic firms. However, it also facilitated certain types of international tax planning strategies that exploit mismatches in tax laws across countries. These strategies make use of foreign affiliates referred to as "hybrid entities" that are treated differently for tax purposes at home and in host countries.

Below we discuss how tax planning structures that use hybrid entities allow US companies

to avoid US tax levied on intercompany payments such as dividends, interest, and royalties, and how these structures exploit mismatches in tax laws across countries to lower foreign tax bills.

2.2.2 Hybrid Tax Planning Strategies

To understand how CTB facilitates tax planning it is necessary to explain the concept of a hybrid entity. A hybrid entity is a business operation that is incorporated from the foreign country point of view and a pass-through (unincorporated branch of another corporation) from the US point of view (or vice-versa, in which case it is referred to as a "reverse" hybrid). Since 1997, an MNC can simply check a box on a tax form to create a pass-through entity. Importantly, the check the box regulations allowed the use of pass-through entities as "foreign disregarded entities". If a pass-through entity like a branch or partnership is disregarded, the transactions with their entity parent and with other disregarded entities owned by the same parent are disregarded. As a result, they become invisible to the U.S. Treasury as they are all part of one consolidated corporation.[7]

The simplest hybrid tax planning structure allows MNCs to take large deductions for interest in high-tax jurisdictions through the use of tax haven finance affiliates. Consider the following planning structure to finance a subsidiary in a high-tax country. Instead of funding the high-tax subsidiary directly, the parent injects equity into a tax haven affiliate which lends to the high-tax subsidiary. The high-tax subsidiary then pays interest to the tax haven affiliate. This profit shifting strategy is commonly known as "interest stripping." While the interest is deductible abroad against taxable income, it is subject to immediate US tax under the CFC rules.

Prior to 1997, CFC rules made the use of a tax haven financing affiliate unattractive for

[7]Prior to 1997, although it was possible for MNCs to create foreign disregarded entities, there were strict rules regarding what types of entities could be declared as such. In particular, such entities had to demonstrate that they possessed at least three of four characteristics associated with partnerships. In practice, we observe very few of these entities prior to the implementation of CTB in 1997.

tax purposes. Since 1997, the parent can check the box on the high-tax affiliate making it a hybrid FDE. From the US point of view, the high-tax affiliate is an unincorporated branch of the tax haven affiliate and the interest payment is invisible to the US Treasury, which regards the combined tax haven-high tax operation as one consolidated corporation. The interest payment therefore escapes Subpart F taxes and the company can defer US income tax in the tax haven.

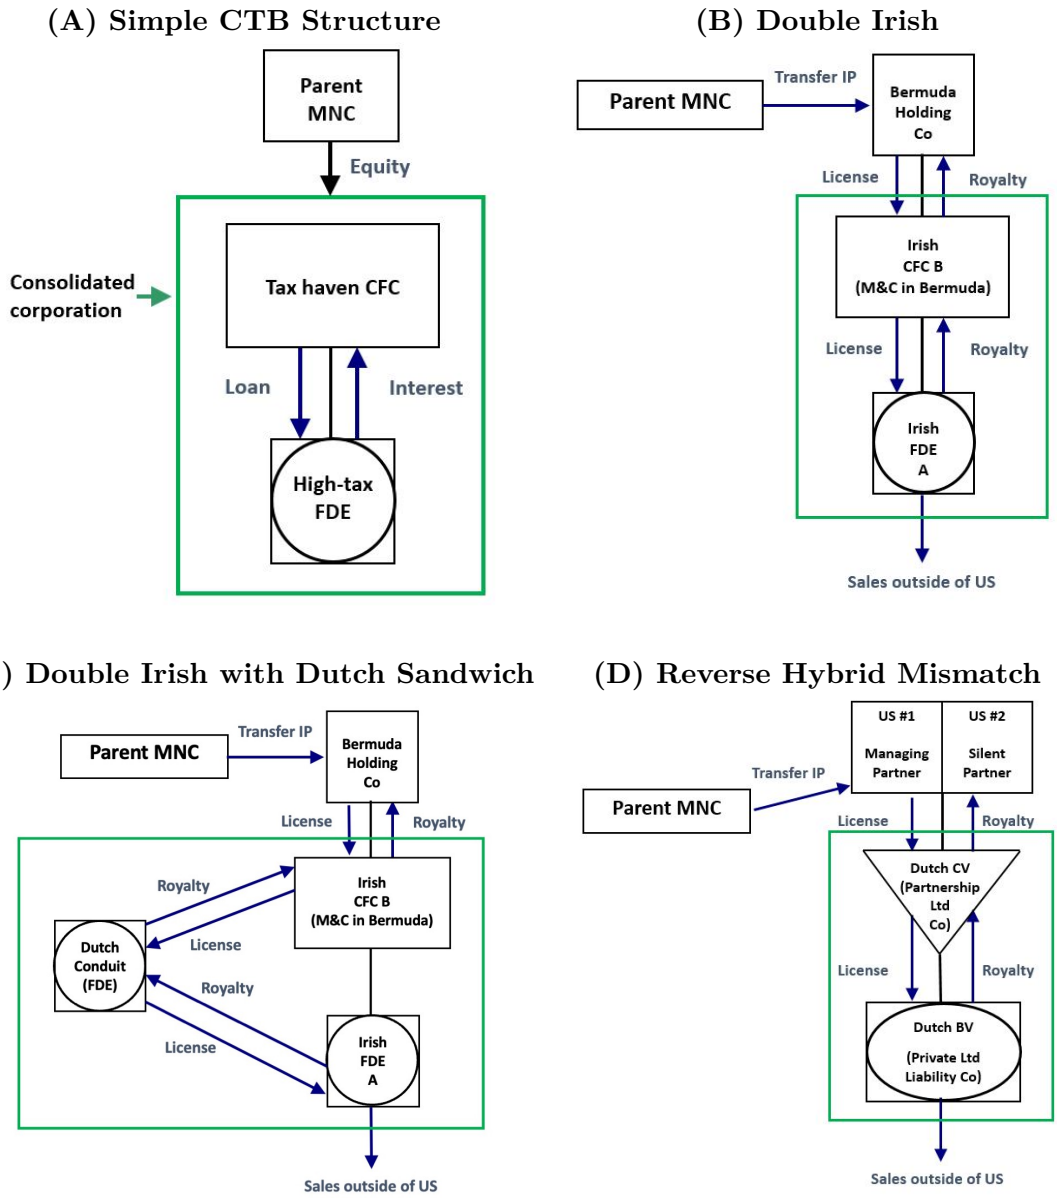

Figure 2.2: Diagrams of Hybrid Tax Planning Structures

Notes: Panel A of Figure 2.2 depicts a hypothetical financing CTB structure; Panel B describes a Double Irish structure; Panel C adds a Dutch sandwich; and Panel D describes a reverse hybrid mismatch structure, otherwise known as a CV-BV. In each of these diagrams, the green rectangles depict the combined structures as perceived by the IRS, corporations are denoted by squares, and disregarded entities are denoted by circles. Hybrid entities—which are corporation in the local country but are disregarded for US purposes—are denoted as squares with circles inside. In Panel D, the CV is shown as a triangle to denote that it is a reverse hybrid: is it a partnership for Dutch purposes but a corporation for US purposes.

Panel A in Figure 2.2 depicts this simple hybrid tax planning structure using a tax haven

affiliate. Squares in Figure 2.2 denote CFCs and circles within squares denote hybrids that are disregarded entities (FDEs). The green box around the two entities (the tax haven CFC and the high-tax FDE) indicates that the structure is consolidated from the US point of view. This simple structure allows the parent to capitalize a foreign affiliate through a tax haven while making intercompany payments invisible, thereby avoiding any current US tax on interest.

While this structure avoids Subpart F taxes and defers US income taxes, the MNC would still be subject to corporate income taxes in the tax haven (if they exist) and potentially to foreign withholding taxes on the interest payments between affiliates.[8] Moreover, to combat interest stripping, many countries have adopted "thin-capitalization" rules that limit the tax deductibility of interest payments, reducing the attractiveness of this option.

CTB and Cost Sharing Agreements

Another form of income shifting is available to MNCs with intellectual property (IP). This method uses cost sharing agreements (CSAs) to develop IP that can be licensed abroad. These agreements are particularly tax advantageous when combined with CTB.

Under a cost sharing agreement, the tax haven affiliate makes a "buy-in payment" that funds a part of the parent's R&D project. This gives the affiliate the right to license resulting IP to other foreign subsidiaries in exchange for royalty payments. Royalty payments are not subject to current tax under Subpart F if the parent checks the box on the affiliate making the payment.[9] The key is that with CTB, any payments for the use of the IP abroad are contained within one consolidated company from the view of the US Treasury.

It is important to note that determining the right arm's length payment for the buy-in is usually quite difficult. Typically the IP is not fully developed at the time the buy-in

[8] After the adoption of "look-through" rules passed as part of the Tax Increase Prevention and Reconciliation Act of 2005, MNCs could avoid Subpart F taxation on distributions of interest, rents and royalties across CFCs without relying on FDEs.

[9] Profits could be further accumulated in a tax haven if MNCs overprice the royalty. The absence of comparable transactions makes it hard for tax authorities to value intellectual property and correctly price royalty payments.

payment is made so there is risk regarding future profits. While the US has rules under which buy-in payments would be adjusted if the profits associated with the IP are too high relative to payments, it is possible for MNCs to underprice the IP. This allows US MNCs to shift income from IP to low-tax affiliates. MNCs are then able to use hybrid tax planning structures to minimize tax on the foreign profits generated from their IP.

As with interest stripping, royalty payments may still be subject to corporate income taxes in a tax haven and to withholding taxes that are meant to prevent profit shifting between countries. We now describe complex tax planning strategies that aim to reduce exposure to income and withholding taxes across multiple jurisdictions.

CTB, Ireland, and Intellectual Property

The first tax planning strategy we study is known as the Double Irish and involves setting up a network of affiliates in Ireland and tax havens. Consider a parent MNC that develops IP in the US that it wants to sell around the world. The parent can transfer the IP to a tax haven CFC using a cost sharing arrangement. The tax haven CFC could then license the IP to an Irish CFC which we show as CFC A in Panel B of Figure 2.2. The Irish CFC pays royalties to the tax haven CFC from the sales revenue it receives from selling the IP abroad.

The cost sharing structure described above creates three tax problems for the parent. First, the parent will be subject to Subpart F taxes (current US tax) on the royalties paid from Ireland CFC A to the haven. Second, taxes will be due in Ireland on any profits that remain in CFC A after royalties are paid to the haven affiliate. Finally, the parent will owe Irish withholding taxes on the royalty transfers to the haven.

These tax problems can be solved in two steps. In step one, the parent creates a second CFC, "CFC B" in Panel B of Figure 2.2, which is legally incorporated in Ireland but managed and controlled in Bermuda. Under Irish tax law, CFC B is a Bermuda company and therefore not subject to Irish tax. The parent can also check the box on CFC A to

avoid current US tax on the royalties. CFC A becomes a FDE and is fiscally transparent to the US Treasury. This first step eliminates Irish tax on any profits remaining in Ireland after royalties are paid, as well as current US taxes on the royalties. However, the transfer to the CFC B still generates Irish withholding taxes.

The second step eliminates these withholding taxes by establishing a Dutch conduit—a "Dutch Sandwich"—between the Irish affiliates. Panel C of Figure 2.2 shows this augmented structure. With the Dutch conduit in place, the parent owes no withholding tax on payments between the conduit and the Irish affiliate (FDE A), as no withholding taxes are due between European Union companies. Further, no withholding taxes will accrue on the royalties between the Dutch conduit and the Irish CFC B since no withholding tax is imposed on these transfers under Dutch law. To avoid Subpart F taxes on the royalty payments, the parent also checks the box on the Dutch conduit, making it an FDE. This "Double Irish with a Dutch Sandwich" hybrid tax planning structure solves all three tax problems we identified above.

Reverse Hybrid Mismatch

The second tax planning structure we examine uses a structure known as a Reverse Hybrid Mismatch. While these strategies can be employed using affiliates in different countries, we first describe a common structure using Dutch companies. To set up this structure, a US MNC creates two US-based affiliates to act as managing/silent partners in a Dutch closed limited partnership called a CV (*commanditaire vennootschap* in Dutch). The partnership is a reverse hybrid entity—it is treated as a pass-through company by the Netherlands and as a corporation by the US. The Dutch CV owns a private limited liability company, called a BV (*besloten vennootschap* in Dutch), which acts as a holding company. The BV owns foreign (non-US) subsidiaries (e.g., in Europe). The BV also holds the license for the US IP, sells it to foreign companies, and pays royalties to the CV. Panel D of Figure 2.2 depicts this structure.

The reverse hybrid mismatch allows the MNC to avoid taxes on foreign income by solving three tax problems. First, taxes may be due in the Netherlands. By Dutch tax law, the CV is a pass-through entity and, as a result, corporate taxes are not paid in the Netherlands. The CV is a reverse hybrid since the CV is a corporation from the US perspective. The BV is a hybrid and the US MNC disregards it by "checking the box." Second, payments from BV to CV can generate Subpart F taxes. Since the US sees a consolidated operation, no Subpart F tax will be due on the royalties. Finally, payments from the BV to the CV could trigger Dutch withholding tax. However, a 2005 decree by the Dutch Finance Ministry exempted US-based CV-BVs from withholding taxes. With this tax planning structure in place, profits from US-developed IP sold abroad are not subject to corporate tax in the Netherlands and enjoy indefinite deferral from US tax (under pre-TCJA law). The CV-BV structure is effectively a "sink" for foreign profits.

A similar structure can be set up through other countries. In particular, a combination of two types of Luxembourg companies, known as SCS and SARL yields a reverse hybrid mismatch structure.[10]

While the description of these structures emphasizes the potential to minimize tax obligations, it is important to note that tax planning is also costly. MNCs have to pay for accounting and legal advice and engage in transactions to form the structures. In addition, company executives differ in their perceived cost of adopting tax aggressive positions. Our model incorporates these fixed costs of tax planning into the decision to adopt an HTP. To the extent that MNCs incur these costs to avoid paying taxes, tax planning is distortionary from an economic perspective.

2.3 Data and Sample Construction

[10]SCS and SARL are short for *société en commandite simple* and *société à responsabilité limitée*, respectively. In this case, the SCS is the reverse hybrid company and the SARL is disregarded from the US perspective.

2.3.1 IRS Business Tax Data

We rely primarily on several IRS datasets for our analysis. These administrative datasets provide parent and CFC-level information disclosed in tax returns that allows us to measure the domestic and foreign activity of a large sample of US corporations, both private and public.

The first dataset, commonly referred to as the Statistics of Income (SOI) corporate sample, is an annual stratified sample of US corporations that SOI uses to produce publicly available aggregated business income statistics.[11] The SOI corporate sample contains information from unaudited tax returns for approximately 100,000 US corporations annually, and has been widely used in the business tax literature to study the behavior of domestic firms (e.g., as in Yagan, 2015; Zwick and Mahon, 2017). Our data focus on C corporations that were sampled between 1992 and 2016. The data primarily contain information from Form 1120, the US Corporate Income Tax Return, as well as some information from related forms. In our analysis, we also use information from Form 6765, which is used to claim the R&D tax credit, as well as Form 4562, which is used to calculate tax deductions for depreciation on capital assets.

The second dataset, which reports information related to foreign affiliates of US corporations, is used by SOI to publish aggregate statistics for international business taxes (IRS, 2022*a*). This dataset provides information from filings of Form 5471, an information return that provides financial information and activity of CFCs, as well as Form 8858, which provides similar information for FDEs, the entity type created by CTB. Unlike the SOI corporate sample, which is provided annually, CFC data is collected only in even years. FDE data is collected for four of the years in our sample period (2006, 2008, 2012, and 2016). Prior to 2004, SOI only collected information related to CFCs for large MNCs with more than $500 million in assets. Starting in 2004, the sampling procedure became much broader, and covers a superset of the SOI corporate sample.

[11] Statistics are available at IRS (2022*b*), and the sampling procedure is described in IRS (2011).

Table 2.1 summarizes the tax forms described above along with selected outcomes that we use in our analysis.

Table 2.1: IRS Forms and Selected Outcomes

SOI Corporate Sample

Form	**Description**	**Selected Outcomes**
1120	Corporate Income Tax Return	Domestic Assets Domestic Wages
6765	R&D Tax Credit	Domestic R&D Expenses Domestic R&D Wages R&D Tax Credit
4562	Depreciation and Amortization	Capital Investment

SOI International Business Tax Sample

Form	**Description**	**Selected Outcomes**
5471	CFC Information Return	Foreign E&P Foreign Taxes Foreign Assets Country of Incorporation CFC to CFC Transactions US Parent to CFC Transactions
8858	FDE Information Return	Foreign E&P Foreign Assets Country of Incorporation Date Disregarded Passthrough Owners Tax Owner

2.3.2 Sample Construction

Table 2.2 shows the size of several different samples of MNCs, along with their foreign affiliates (CFCs and FDEs). We consider a firm to be a MNC if it files Form 5471 for at least one CFC. Column (1) reports the number of MNCs that have coverage in our data from the SOI international business tax studies. Column (2) shows the number of MNCs

from Column (1) that we observe in the SOI corporate sample.[12] As mentioned above, there was a sampling change in the international business tax study starting in 2004 that resulted in a large increase in the sample, especially for smaller MNCs. To stabilize the firm sampling distribution between earlier and later years, we remove MNCs that did not have at least one CFC with $50 million in foreign assets as well as those with fewer than $500 million in domestic assets. Column (3) shows sample sizes after applying this filter. Column (4) shows sample sizes for a balanced panel of MNCs that appear in the data in 1996, the first period prior to the implementation of CTB, and that remain in the sample until at least 2002. To study firm-level outcomes, our analysis primarily uses the samples shown in Columns (3) and (4).

Table 2.2: Sample Sizes

	(1) Int'l. Business Sample	(2) SOI Corp. Sample	(3) Stable Sample	(4) Balanced
MNE Count	21411	17491	3638	460
CFC Count	247686	230874	40836	22648
FDE Count	53094	51579	46039	23278

Notes: This table provides the size of four different samples of US MNCs, along with their related controlled foreign corporations (CFCs) and related foreign disregarded entities (FDEs). Column (1) provides sample sizes using all MNCs in the SOI International Business Tax Sample. Column (2) provides sample sizes after removing MNCs that are not included in our extract of the IRS Statistics of Income (SOI) corporate sample. Column (3) applies a size filter that removes smaller MNCs from the sample so that the sample composition is similar in earlier and later years. Column (4) provides samples sizes for MNCs included in Column (3) that are also observed prior to the implementation of CTB in 1997 and that remain in the sample until at least 2002.

2.3.3 Measuring Foreign Effective Tax Rates

One possible concern when measuring the income of foreign affiliates of corporations relates to the measurement of foreign earnings. Blouin and Robinson (2020) suggest that aggregated IRS statistics may inadvertently double count foreign earnings. This is due to

[12] Our extract of the SOI corporate sample only includes C corporations, whereas our extract of the SOI international business sample includes other corporate forms. Thus, this restriction primarily excludes corporations that are not C corporations.

accounting quirks of MNCs. Consider a hypothetical US firm with two CFCs (A and B). Suppose that CFC A is a holding company that holds a 100% stake in CFC B and has no economic purpose other than to collect dividends from its subsidiaries, and further suppose that CFC B discloses E&P of $100 million, which is then issued as a dividend to CFC A. CFC B will then also report E&P of $100 million. A simple aggregation of the firm's foreign profits will result in an estimate of $200 million in foreign E&P even though the true figure is $100 million. If firms with hybrid structures tend to issue more dividends between their foreign affiliates, then this exercise would overestimate their E&P and in turn underestimate the ETR. However, each CFC must also file an attachment to Form 5471 that discloses transactions between the focal CFC and related CFCs, including any dividends that the CFCs may transfer to each other.

To ensure that there is no double-counting of foreign profits, we subtract these related dividends from our calculations. Specifically, we compute the firm-level foreign ETR as

$$ETR = \frac{\text{Foreign Taxes}}{\text{Foreign Taxes} + \text{Foreign Earnings and Profits}},$$

where total foreign tax payments are taken from Form 5471, Schedule E.[13] To calculate pre-tax earnings and profits, we obtain pre-tax E&P for each affiliated CFC (using Schedules H and E from Form 5471). Following the suggestion of Blouin and Robinson (2020), we remove dividends received from related CFCs from E&P (using Schedule M from Form 5471).

2.3.4 Measuring Changes in International Corporate Structures

Both the SOI corporate sample and the CFC sample have been used to study domestic and international business taxation. Relatively little work, however, has utilized the wealth of

[13]There does not appear to be any such double-counting concern related to the payment of foreign taxes. We exclude unprofitable firm-years from this calculation.

information contained in the FDE sample.[14] Although data from Form 8858 is collected less frequently than other samples, it allows us to observe two important features of US MNC structures. First, it reveals the date when an entity was first disregarded by a MNC, which allows us to measure adoption of CTB among MNCs. Second, it allows us to observe the tax ownership structure of each CFC along with its FDEs. These ownership structures reveal important cross-national linkages within MNCs and, most importantly, allow us to identify CFCs that have particular structures associated with the tax planning strategies described in Section 2.2.

Detecting the Double Irish

As described previously, the Double Irish involves two Irish entities—a top-level entity that is incorporated in Ireland, but managed and controlled in another low-tax foreign country, and a lower-level Irish entity that merchandises the IP and pays a royalty. Typically, the lower-level entity is "checked" and is classified as an FDE for US tax purposes. Alternatively, the MNC may "check" both types of entities which are then classified as FDEs under the "tax ownership" of a separate CFC. As a result, we flag two types of CFCs that could be used in a Double Irish arrangement. First, we flag any CFC that is incorporated in Ireland and that checks the box on an Irish FDE. Second, we flag any CFC that checks the box on two separate Irish FDEs. Note that this classification method flags "simple" Double Irish arrangements that involve a direct link between two Irish entities, but also more complex arrangements, such as the Double Irish with a Dutch Sandwich, that might involve intermediary affiliates through which profits are routed.

Detecting Reverse Hybrids

The other type of structures we consider are reverse hybrid arrangements common in the Netherlands and Luxembourg. This arrangement, also described in Section 2.2, involves a

[14] A notable exception is a recent working paper, Samarakoon (2022), that examines the impact of the closure of the Double Irish tax structure in Ireland.

top-level entity that is classified as a partnership and a bottom-level entity that is classified as an private limited company (PLC) in the associated country of incorporation. SOI data typically provide the acronym that is associated with the management form of foreign affiliates on Form 5471 and 8858. In the Netherlands, for example, partnerships are associated with the acronym CV and the equivalent form of a PLC is associated with the acronym BV. The equivalent acronyms in Luxembourg are SCS (for a partnership) and SARL (for a PLC). To classify potential reverse hybrid structures, we flag any CFC classified as a CV (incorporated in the Netherlands) or SCS (incorporated in Luxembourg) that check the box on an FDE classified as a BV (Netherlands) or SARL (Luxembourg). We also flag any CFC that checks the box on a CV-BV or SCS-SARL pair of FDEs.

For both the Double Irish and Reverse Hybrid Mismatch arrangements, we use the first date that all flagged FDEs were disregarded to measure the year that an MNC first adopted a particular structure.

2.4 Adoption and Prevalence of CTB and Hybrid Tax Planning Structures

In this section, we start by describing how US MNCs used CTB after its implementation in 1997. We show that after its implementation, usage of tax-transparent FDEs quickly became widespread among US MNCs. A large share of these entities are connected to well known tax havens.

Next, we focus on the tax planning structures described in Section 2.2. We show that MNCs gradually adopted these structures in the decade after the implementation of CTB. By the 2010s, MNCs that adopted at least one of these structures were generating a majority of E&P among firms in our sample and for these MNCs, a majority of foreign profits were connected to these structures. We show that these MNCs also represent a large share of domestic activity, with 20% of domestic payroll and 15% of domestic capital assets by 2016.

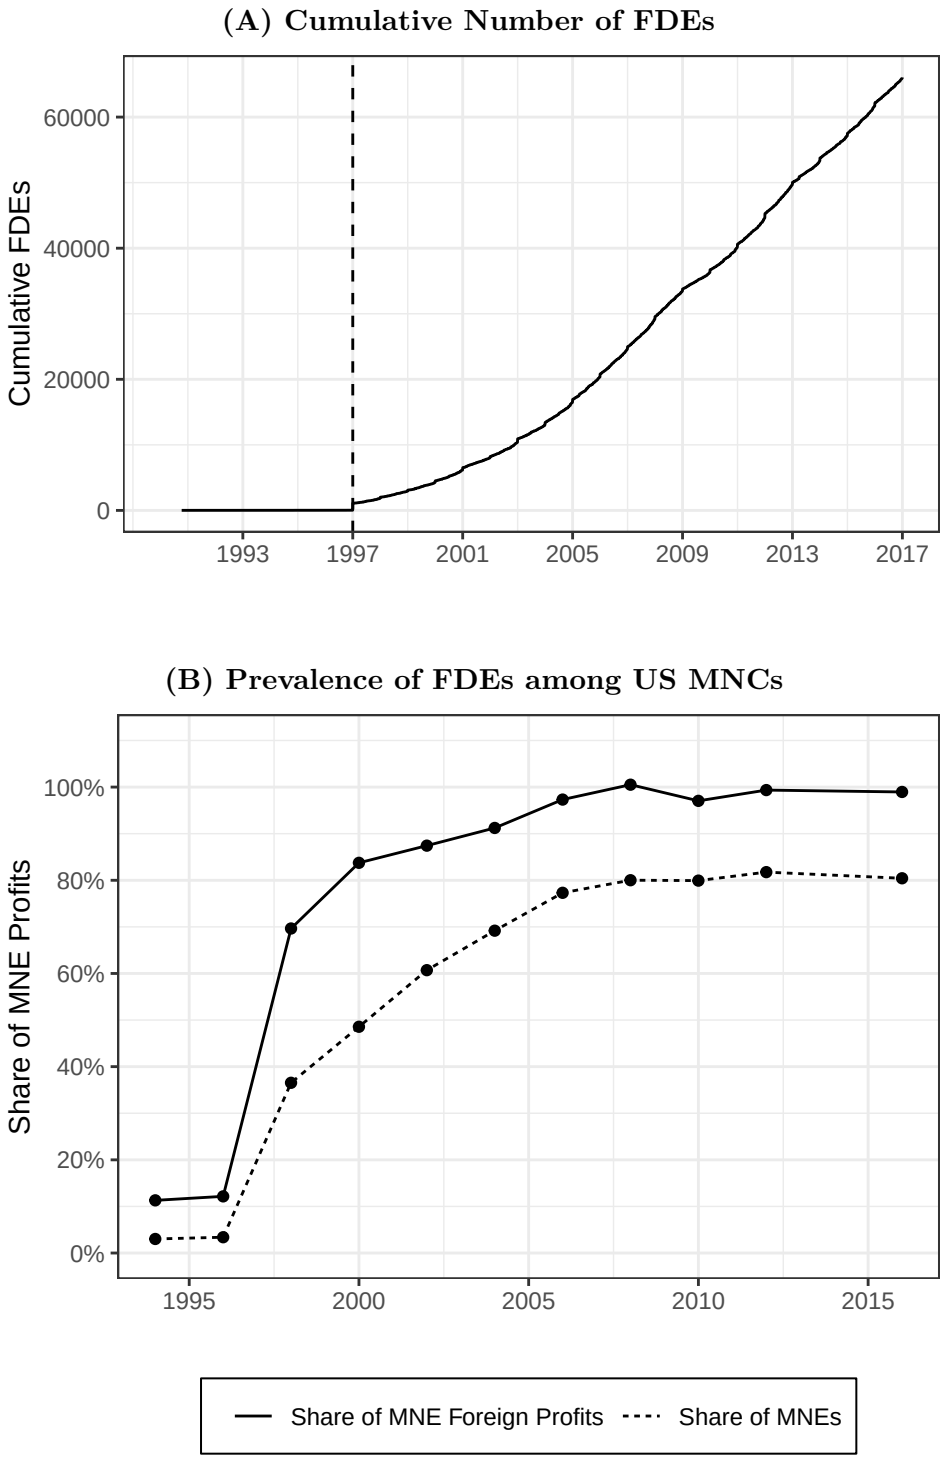

Figure 2.3: Adoption of Foreign Disregarded Entities

Notes: Panels A of Figure 2.3 plots the cumulative number of foreign disregarded entities (FDEs). There were fewer than 50 FDEs prior to 1997 and this number grew rapidly following the adoption of CTB regulations. Panel B plots the fraction of US MNCs with an FDE as well as the share of foreign profits that accrue to US MNCs with FDEs. By 2008, close to 80% of US MNCs have a FDE and these MNCs account for close to 100% of foreign profits.

Table 2.3: Disregarded Entities by Country of Incorporation

Country Name	Adj. E&P (billions)	Num. FDEs
Netherlands	69	12236
Ireland	67	4834
United Kingdom	66	26979
Cayman Islands	51	5276
Bermuda	46	2124
Singapore	38	3458
Switzerland	33	2134
Canada	28	7427
Australia	26	6780
China	22	6558

Notes: This table shows the largest ten countries by total foreign earnings generated by foreign disregarded entities (FDEs). Column (2) provides aggregate E&P, adjusted to remove related dividends, generated by FDEs in the the country listed in column (1). This includes E&P for all years that we observe Form 8858 filings (2006, 2008, 2012, and 2016). Column (3) shows the number of unique entities across all years of this sample.

2.4.1 Adoption of Check the Box

Panel A of Figure 2.3 shows the cumulative number of FDEs created between 1992 and 2016. Prior to 1997, usage of these transparent entities was relatively rare—the IRS used a resource-intensive system that required firms to show that their affiliates possessed a set of characteristics that were more consistent with either a partnership (transparent) or corporation (non-transparent) classification. Starting in 1997, the Treasury relaxed these restrictions, as described in Section 2.2. As a result, usage of FDEs became widespread over the next two decades, with over 60,000 foreign affiliates classified as FDEs by the end of 2016. Panel B shows that by 2008, about 80% of MNCs used CTB to declare at least one FDE, and that these MNCs generated nearly all of foreign E&P.

Table 2.3 shows the largest ten countries according to total earnings generated by FDEs. The Netherlands and Ireland are the largest domiciles for these types of foreign affiliates. FDEs also generate large amounts of earnings in well known tax havens, such as the Cayman Islands and Bermuda.

2.4.2 Adoption of Hybrid Tax Planning Structures

Panel A of Figure 2.4 shows the evolution of the share of MNCs in our sample that adopted particular hybrid tax planning structures. After the implementation of CTB, there was steady adoption of these structures, with more than 16% of MNCs adopting at least one by 2016. Panel B demonstrates that by 2008, these MNCs generated a majority of foreign E&P. Panel C shows that the CFCs linked to HTPs generated the majority of profits within MNCs that use them by 2016. Panel D shows the share of foreign E&P that is linked to HTP structures relative to all the foreign E&P of the MNCs in our sample. This figure shows that, by 2016, more than 35% of all foreign E&P of US MNCs was routed through an HTP.

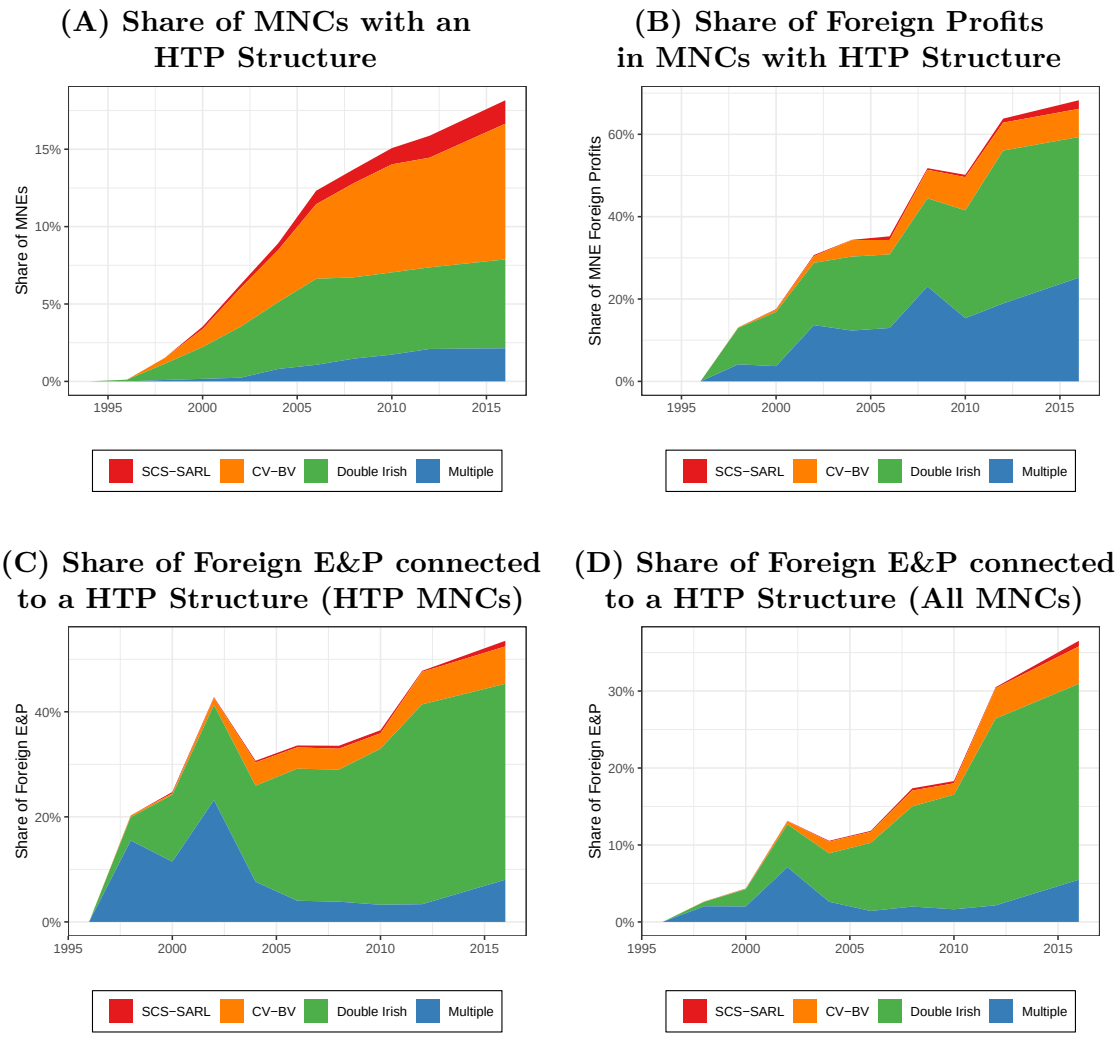

Figure 2.4: Adoption of Hybrid Tax Planning (HTP) Structures

Notes: These figures show that a growing share of US multinationals (MNCs) have adopted hybrid tax structures over time (Panel A) and that these MNCs are responsible for a large share of the overall foreign E&P of US MNCs (Panel B). Panel C shows that a large share of foreign E&P within adopting MNCs is connected to HTPs. Panel D shows the share of foreign E&P that is connected to an HTP relative to all the MNCs in our sample. In each panel, the blue area comprises MNCs that have adopted more than one structure.

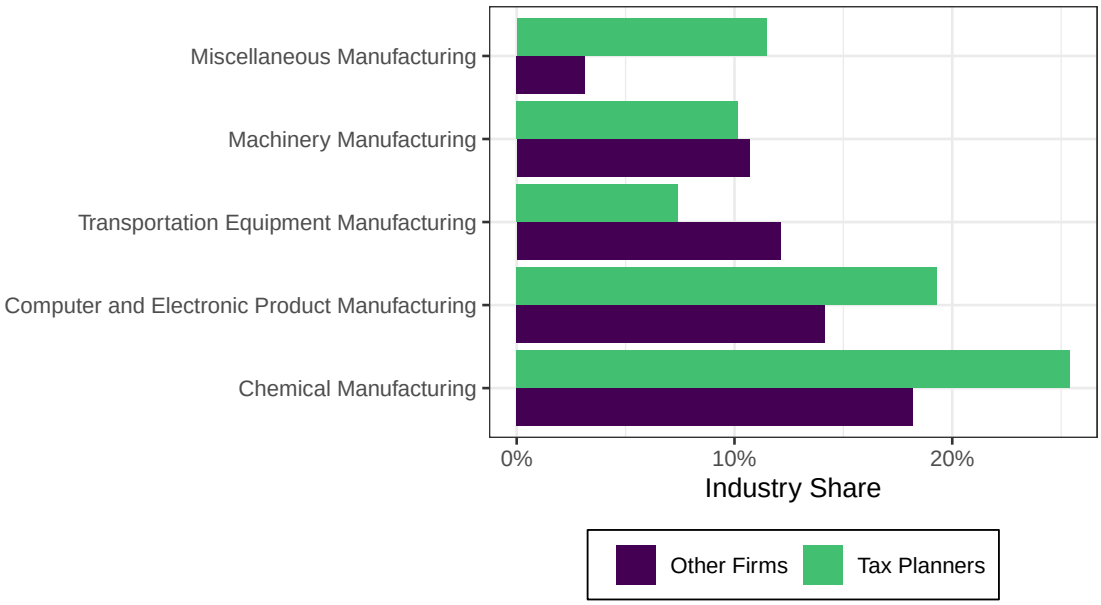

Figure 2.5: Comparison of Industry Shares by MNC Group

Notes: This figure compares manufacturing subindustry shares across two groups of US multinationals (MNCs) for the five largest manufacturing subindustries. The light green bars show the subindustry share for a group of MNCs that eventually adopt at least one of the hybrid tax structures described in Section 2.2. The dark blue bars shows the subindustry share for MNCs that did not adopt any of these structures during the sample window.

The bulk of our sample is concentrated in manufacturing. Figure 2.5 compares manufacturing subindustry shares among MNCs with tax planning structures to MNCs without these structures. MNCs with tax planning structures have noticeably larger industry shares within Chemical Manufacturing and Computer and Electronic Product Manufacturing. These subindustries include pharmaceutical and tech firms, IP-intensive businesses that are able to use intangible assets to relocate foreign earnings.

Figure 2.6 shows that MNCs with tax planning structures comprise a large share of domestic economic activity. As a share of all C corporations in the SOI corporate sample, which includes domestic corporations as well as MNCs, MNCs with tax planning structures paid more than 20% of domestic wages and account for about 15% of domestic investment by 2010.

The results in this section demonstrate that several hybrid tax planning structures became

widely adopted by US MNCs after the implementation of CTB, with large shares of foreign profits flowing through these structures in the decade after adoption.

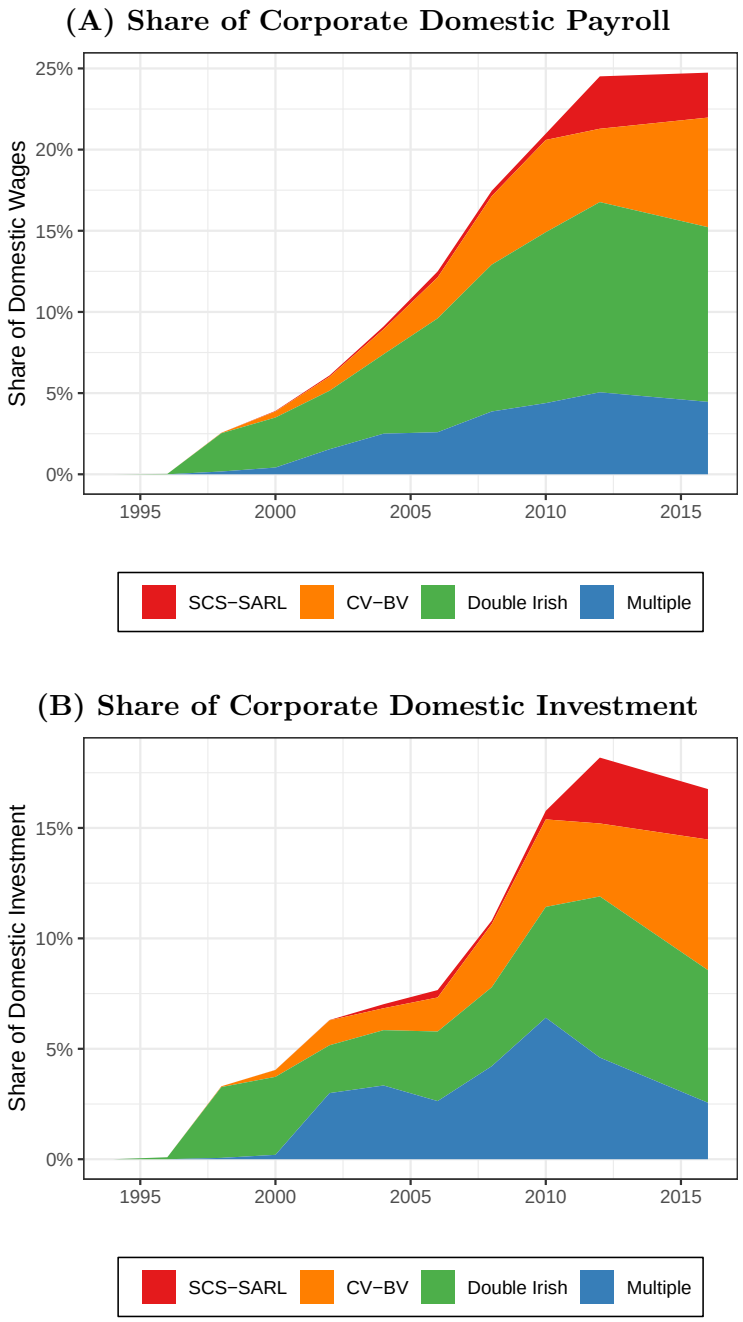

Figure 2.6: Hybrid Tax Planning Structures and Domestic Economic Activity

Notes: These figures show the share of domestic wages (Panel A) and domestic capital investment (Panel B) paid by US multinationals (MNCs) that adopt one of the hybrid tax structures described in Section 2.2. This share is computed as a fraction of all domestic wages and capital investment among C corporations in the IRS Statistics of Income corporate sample. See Section 2.3 and Table 2.2, Column 2 for a description of the sample.

2.5 Hybrid Tax Planning Structures and Multinational Activity

In this section, we examine changes in domestic and foreign economic activity of adopting MNCs. We start with a set of descriptive facts that compare select outcomes for MNCs that utilized HTPs, relative to other MNCs that do not rely on these arrangements. We then estimate staggered difference-in-differences models that control for firm characteristics to compare outcomes of MNCs that do and do not adopt hybrid tax structures post adoption.

2.5.1 Comparison of Aggregate Trends

While the structures we study had been suspected of being used for profit shifting, lack of tax data had prevented prior researchers from confirming their role in profit shifting. We thus start by examining whether the hybrid tax planning firms engage in the kinds of transactions that are associated with profit shifting. Figure 2.7 compares hybrid tax planners to other multinationals along a number of these dimensions. In Panel A we first document that hybrid tax planning firms see a much larger growth in the balance of loans between related CFCs than other MNCs. These transactions may be related to interest stripping strategies, as discussed in Panel A of Figure 2.2.

Panel B of Figure 2.7 shows that hybrid tax planning firms experience a much faster rise in the value of their intangible assets held abroad. This is consistent with using hybrid tax planning structures to exploit intangibles abroad. Panel C of Figure 2.7 shows that hybrid tax planning firms also had larger increases in compensation for services from CFCs to parents—which include cost sharing payments. Starting in 2008, we also observe whether a parent had any cost sharing agreements with one of its CFCs; Panel D shows that hybrid tax planning firms are also more likely to engage in these agreements.

Figure 2.1 compares the average foreign ETRs of hybrid tax planners with other US MNCs in our sample. At the beginning of the sample period, both types of MNCs paid taxes on foreign E&P at similar rates. Starting in 2002, however, there is a striking divergence in the evolution of each group's ETR. By 2016, hybrid tax planners faced a foreign ETR that was

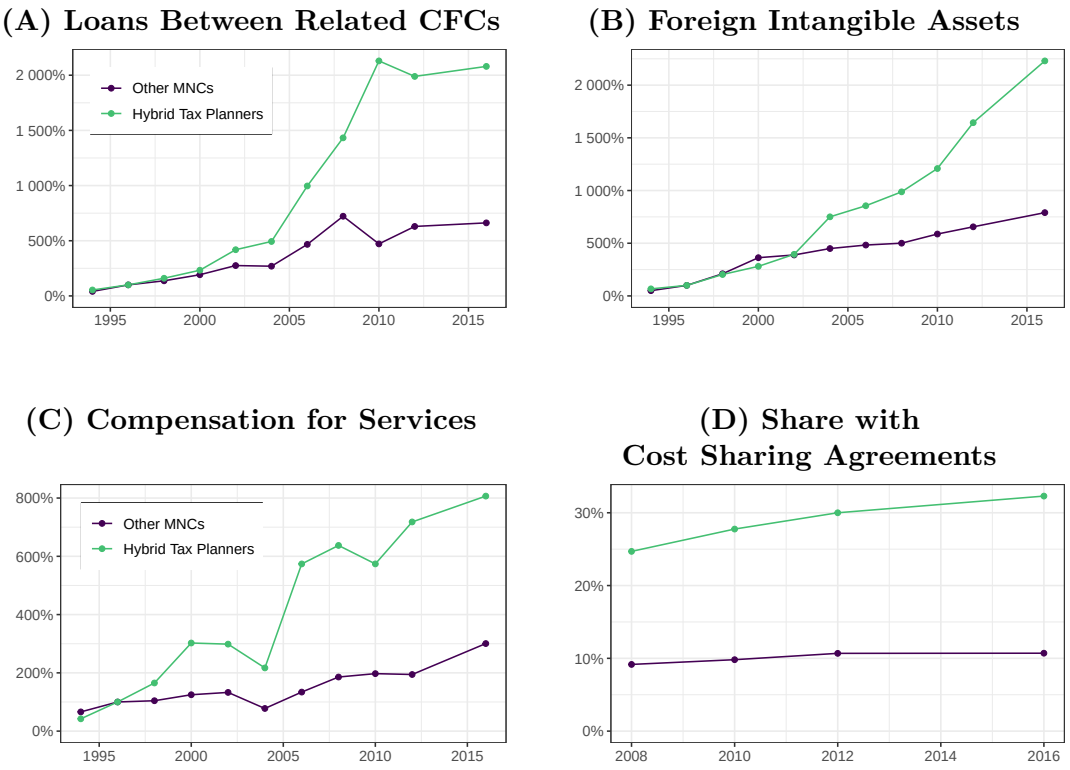

Figure 2.7: Mechanisms for Profit Shifting

Notes: These figures show the evolution of aggregate loans between related CFCs (Panel A), foreign intangible assets (Panel B), payments from CFCs to US parent companies for technical services (Panel C), and the share of MNCs in each group with active cost sharing agreements with a CFC. For comparability, aggregate values for both groups are normalized to 100% as of 1996 for Panels A through C. The light green line shows values for a group of MNCs that eventually adopt at least one of the hybrid tax structures described in Section 2.2. The dark blue line shows aggregate values for MNCs that did not adopt any of these structures during the sample window.

about half of the ETR incurred by other MNCs. Figure 2.7 provides important context for the decline in foreign ETRs. While one may suppose that declining statutory rates around the world may be responsible for this decline, Figure 2.7 shows that it was the use of hybrid tax planning structures that allowed these MNCs to obtain lower ETRs abroad. Indeed, as we see in Figure 2.1, non-HTP firms experienced a very small decline in foreign ETRs during our sample period.

Having shown that hybrid tax planning structures are likely used for profit shifting, we now examine whether hybrid tax planning firms also deferred more income abroad. Panel

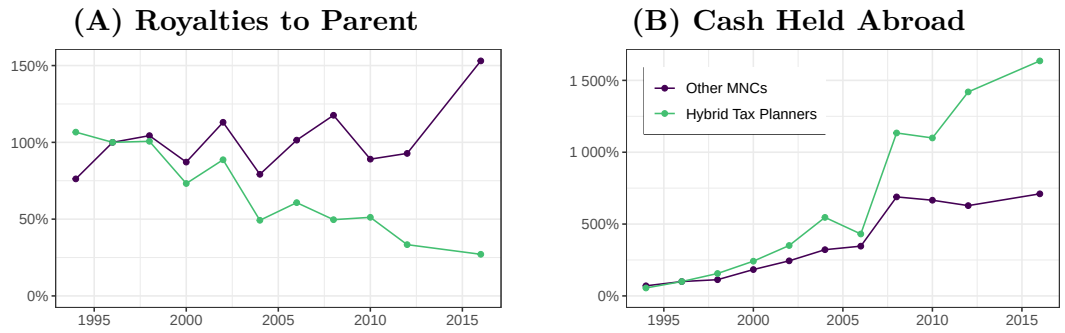

Figure 2.8: Evidence of Deferral

Notes: These figures show the evolution of royalty payments from CFCs to domestic parent entities as a share of foreign E&P (Panel A) and aggregate foreign cash balances (Panel B) for two groups of US multinationals (MNCs). For comparability, aggregate values for both groups are normalized to 100% as of 1996 in Panel B. The light green line shows values for a group of MNCs that eventually adopt at least one of the hybrid tax structures described in Section 2.2. The dark blue line shows aggregate values for MNCs that did not adopt any of these structures during the sample window.

A of Figure 2.8 shows that, relative to foreign E&P, hybrid tax planning firms saw faster declines in royalty payments from CFCs to parents. This result is consistent with a transition away from undeferred royalty income. Consistent with this interpretation, Panel B also shows that hybrid tax planning firms saw large increases in cash held abroad compared to MNCs that did not adopt HTPs. While the fact that MNCs accumulated cash abroad during the last two decades is well known, this figure shows that the bulk of the increase in cash held abroad occurred in the just over 300 firms that engaged in hybrid tax planning agreements through Ireland, Netherlands, and Luxembourg.

The results in Figures 2.7 and 2.8 provide *prima facie* evidence of the specific mechanisms through which hybrid tax planning strategies operate. The ability to shift profits to lower tax countries allowed MNCs to avoid taxes abroad, while locating intangibles abroad allowed US MNCs to defer paying US income taxes. We now study whether hybrid tax planning firms experience a different evolution in their real operations.

Figure 2.9 shows that, while hybrid tax plannings firms and other MNCs had similar patterns of economic activity prior to 1997, their economic activities diverged over the same time period that hybrid tax plannings strategies were adopted. Panel A shows that

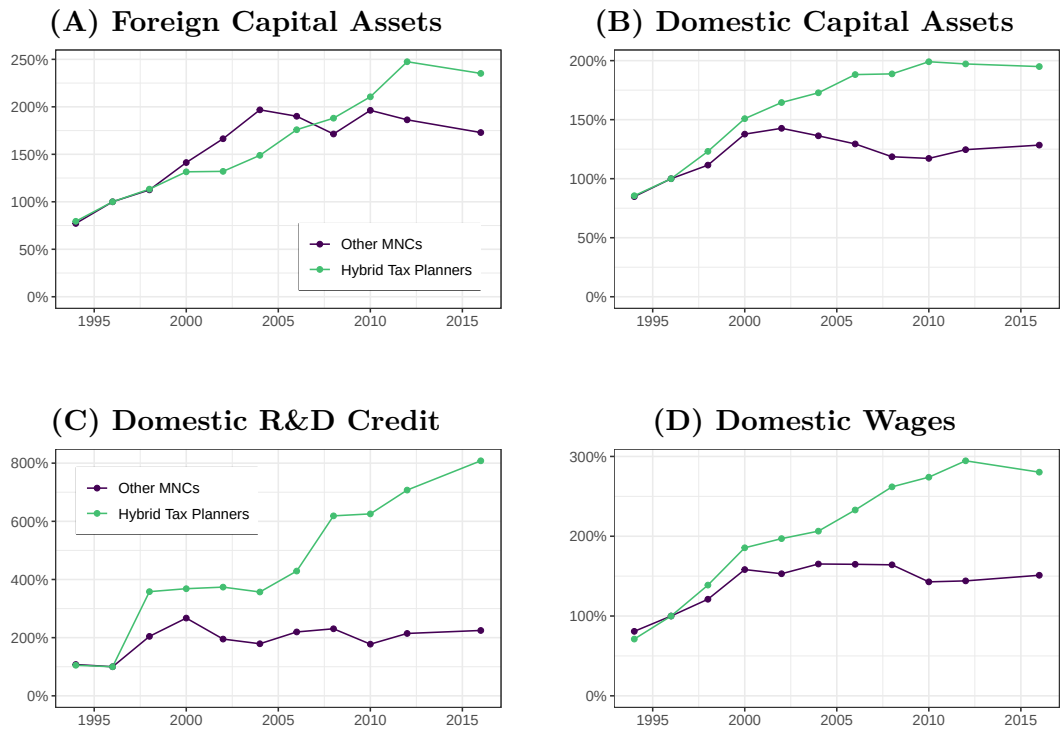

Figure 2.9: Hybrid Tax Planning and Real Economic Activity

Notes: These figures show the evolution of aggregate foreign capital assets (Panel A), domestic capital assets (Panel B), claimed domestic R&D credits (Panel C), and domestic wages (Panel D) for two groups of US multinationals (MNCs). For comparability, aggregate values for both groups are normalized to 100% as of 1996. The light green line shows aggregate values for a group of MNCs that eventually adopt at least one of the hybrid tax structures described in Section 2.2. The dark blue line shows aggregate values for MNCs that did not adopt any of these structures during the sample window.

hybrid tax planners had larger increases in foreign capital assets; Panel B shows that these firms also accumulated more domestic capital assets; Panel C shows larger increases in R&D conducted at home for hybrid tax planners; and Panel D shows larger increases in domestic wages. Across all of these measures, the declines in foreign ETRs were accompanied by increases in economic activity, both domestically and abroad.

2.5.2 Estimating Staggered Difference-in-Differences Models

The analysis in the preceding section demonstrates that, in aggregate, the foreign ETR among MNCs that adopted hybrid tax planning structures was nearly identical to the foreign ETR of other US MNCs prior to 2004. Furthermore, the aggregate trajectory of several foreign and domestic outcomes for both groups of MNCs were similar until the mid-2000s, at which point there appears to be substantial divergence. The timing of this divergence coincides with the period when foreign tax planning structures became widely adopted.

In this section, we examine whether this aggregate difference is present at the firm level and estimate relative changes in firm outcomes that adopt an HTP relative to those that do not. We estimate staggered difference-in-differences (DD) models that allow us to determine whether the timing of adoption of a hybrid tax structure aligns with changes in foreign and domestic outcomes.

Event Study Specification

Our main specification relies on a standard staggered differences-in-differences model,

$$Y_{i,t} = \alpha_i + \lambda_t + \sum_{\ell} \mu_\ell \mathbf{1}\left\{t - E_i = \ell\right\} + v_{i,t}, \tag{2.1}$$

where $Y_{i,t}$ is the selected outcome for MNC i in year t, α_i and λ_t are MNC and year fixed effects, E_i is the time when MNC i initially receives treatment, and ℓ is an indicator for

the relative number of periods after MNC i adopts a foreign tax planning structure. We focus on a 12-year relative period around the adoption of a foreign tax planning structure, i.e. we let ℓ vary between -6 and 6. In all specifications, we cluster standard errors at the MNC level.

We interpret results of Equation 2.1 as measuring dynamic changes in firm-level outcomes of adopting MNCs relative to non-adopting MNCs. Relative to the results in the prior section, these estimates help tie changes in firm outcomes to the timing of adoption. This approach also addresses the concern that firm outcomes are driven by concomitant shocks to firms with characteristics that are related to tax planning (e.g. larger firms, more IP-intensive firms, or firms in different industries). To do so, we estimate alternative specifications that interact year fixed effects with a set of pre-treatment covariates for MNCs to allow for time-varying heterogeneity across industries, across foreign and domestic firm sales bins, and across foreign and domestic bins for intangible assets.[15] Recent literature has pointed out the potential for contamination bias in the specification in Equation 2.1. As a robustness check, we also estimate effects using the methodology proposed by Sun and Abraham (2021).

While these firm-level comparisons are informative of the role of HTPs in driving the aggregate changes described in the previous section, a key question is whether the results of Equation 2.1 can be interpreted as causal effects of HTPs. The usual assumptions for a causal interpretation include parallel trends and no anticipatory behavior. The descriptive evidence provided in the previous section shows that the evolution of outcomes for adopting and non-adopting MNCs was strikingly similar prior to the bulk of adoptions in the mid-2000s. Additionally, pre-trend coefficients in the event study plots provided below are generally insignificant. Regarding anticipation, recall that our data only provides observations in even years, therefore MNCs would have to adjust behavior two years in advance for this form of bias to be present. Because the parallel trends assumption is inherently

[15]To be precise, this implies an augmented version of Equation 2.1 where λ_t is replaced by $\sum_{g \in G} \lambda_{gt}$, where G is a set of groups for which we include group-by-year fixed effects.

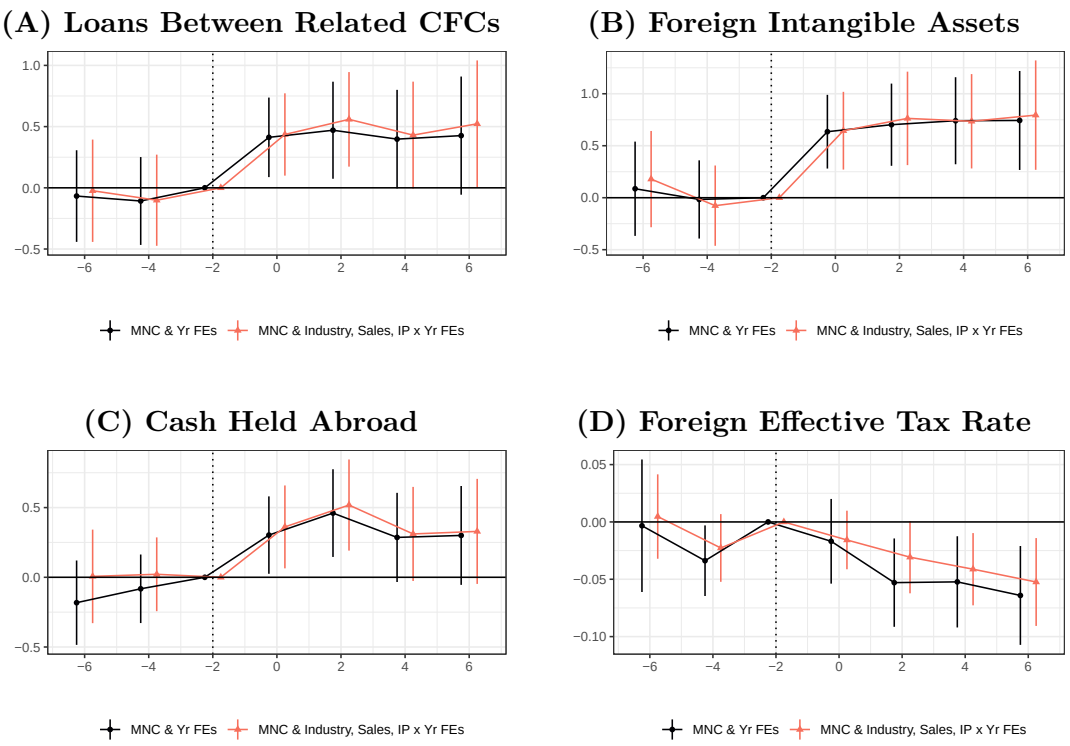

Figure 2.10: Event Studies: Profit Shifting Mechanisms and Foreign ETRs

Notes: This figure provides estimates of μ_ℓ from Equation 2.1 for the corresponding foreign outcome listed in each panel. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-industry fixed effects and year-by-group fixed effects, where groups include domestic and foreign sales quartiles and domestic and foreign intangible asset quartiles, and where quartiles are computed using pre-treatment values.

untestable, it is important to consider that firms may select into HTPs because they have more to gain from tax planning. Section 2.6 develops a model of selection on gains that interprets estimates from Equation 2.1 as arising from both selection into tax planning and from behavioral changes related to tax planning.

Estimates of Changes in Financial and Tax Outcomes

Figure 2.10 plots the results of estimating Equation 2.1 on a set of financial and tax outcomes. We report two specifications for each outcome. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-industry fixed effects and

Table 2.4: Profit Shifting Mechanisms and Foreign ETRs

	1	2	3	4
Panel A. Loans Between Related CFCs				
ATT	0.288*	0.346*	0.260+	0.334*
	(0.134)	(0.138)	(0.142)	(0.146)
Num. Units	1755	1755	1755	1755
Num. Treated	285	285	285	285
Panel B. Foreign Intangible Assets				
ATT	0.476**	0.492**	0.485**	0.495**
	(0.150)	(0.158)	(0.158)	(0.166)
Num. Units	2290	2290	2290	2290
Num. Treated	308	308	308	308
Panel C. Cash Held Abroad				
ATT	0.295**	0.255*	0.287*	0.263*
	(0.111)	(0.112)	(0.116)	(0.116)
Num. Units	3226	3226	3226	3226
Num. Treated	321	321	321	321
Panel D. Foreign ETR				
ATT	−0.043**	−0.029*	−0.042**	−0.030*
	(0.014)	(0.012)	(0.014)	(0.012)
Num. Units	3302	3302	3302	3302
Num. Treated	321	321	321	321
Industry by Year FEs	No	Yes	No	Yes
Size by Year FEs	No	No	Yes	Yes

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table provides ATT estimates for a staggered two-way fixed effects difference-in-differences model provided in equation 2.1 for the corresponding foreign outcome listed in each panel. Specification 1 does not include additional controls. Specification 2 includes year-by-industry fixed effects. Specification 3 includes year-by-group fixed effects, where groups include domestic and foreign sales quartiles and domestic and foreign intangible asset quartiles, and where quartiles are computed using pre-treatment values. The last specification includes all fixed effects used in specifications 2 and 3.

Table 2.5: Profit Shifting Mechanisms and Foreign ETRs (Sun & Abraham)

	1	2	3	4
Panel A. Loans Between Related CFCs				
ATT	0.448**	0.472***	0.421**	0.473**
	(0.147)	(0.143)	(0.156)	(0.153)
Num. Units	1755	1755	1755	1755
Num. Treated	285	285	285	285
Panel B. Foreign Intangible Assets				
ATT	0.602***	0.601***	0.599***	0.596***
	(0.148)	(0.151)	(0.155)	(0.159)
Num. Units	2290	2290	2290	2290
Num. Treated	308	308	308	308
Panel C. Cash Held Abroad				
ATT	0.402**	0.371**	0.386**	0.372**
	(0.126)	(0.131)	(0.130)	(0.134)
Num. Units	3226	3226	3226	3226
Num. Treated	321	321	321	321
Panel D. Foreign ETR				
ATT	−0.045***	−0.033**	−0.040***	−0.032**
	(0.011)	(0.011)	(0.012)	(0.011)
Num. Units	3302	3302	3302	3302
Num. Treated	321	321	321	321
Industry by Year FEs	No	Yes	No	Yes
Size by Year FEs	No	No	Yes	Yes

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table provides ATT estimates for a staggered difference-in-differences model using the estimator proposed by Sun and Abraham (2021) for the corresponding foreign outcome listed in each panel. Specification 1 does not include additional controls. Specification 2 includes year-by-industry fixed effects. Specification 3 includes year-by-group fixed effects, where groups include domestic and foreign sales quartiles and domestic and foreign intangible asset quartiles, and where quartiles are computed using pre-treatment values. The last specification includes all fixed effects used in specifications 2 and 3.

year-by-group fixed effects, where groups include domestic and foreign sales quartiles and domestic and foreign intangible asset quartiles, and where quartiles are computed using pre-treatment values.

Panel A of Figure 2.10 shows estimates for the log of the balance of loans between CFCs; Panel B shows results for the log of intangibles held abroad; and Panel C shows estimates for the log of cash held abroad. Across these three outcomes, we observe that these markers of profit shifting and deferral are similar between hybrid tax planning firms and other MNCs prior to the adoption of a hybrid structure. For all three outcomes, we observe significant relative increases for hybrid tax planning firms following the adoption of a hybrid tax planning strategy. Consistent with these mechanisms and the result of Figure 2.1, Panel D shows that the foreign ETRs of hybrid tax planning firms saw a significant relative decline following the adoption of a hybrid tax planning structure. For the case of foreign ETRs, Panel D of Figure 2.10 demonstrates that there was a gradual phase in of the ETR reduction. Six years after adoption, MNCs experience a reduction in their foreign ETR of between 5 and 7 percentage points. For all of these outcomes, we find that inclusion of size-bin-by-year fixed effects and industry-by-year fixed effects does not impact the estimated effects of adopting a hybrid tax planning strategy.

Table 2.4 displays coefficients from a difference-in-differences specification that summarizes the average changes in these outcomes. We estimate that loans between CFCs increased by 29 percent; foreign intangible assets increased by 48 percent; and cash held abroad increased by 30 percent. We also estimate declines in foreign ETRs of between 3 and 4.3 percentage points. These estimates are remarkably stable across our specifications, suggesting that they are not driven by comparisons across firms in different industries, in different domestic and foreign size categories, or that are more or less dependent on intellectual property. Table 2.5 reports estimates using the estimator proposed by Sun and Abraham (2021); we obtain similar estimates using this method.

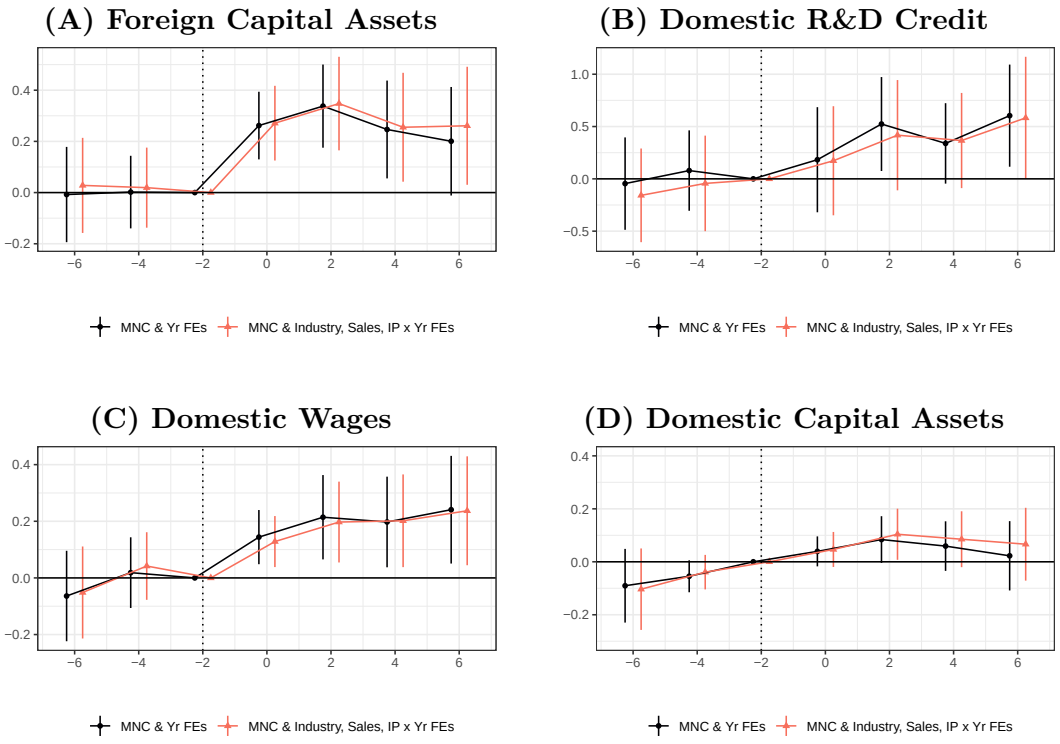

Figure 2.11: Event Studies: Hybrid Tax Planning and Real Economic Activity

Notes: This figure provides estimates of μ_ℓ from Equation 2.1 for the corresponding foreign outcome listed in each panel. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-industry fixed effects and year-by-group fixed effects, where groups include domestic and foreign sales quartiles and domestic and foreign intangible asset quartiles, and where quartiles are computed using pre-treatment values.

Estimates of Changes in Real Economic Activity

To measure changes in domestic activity, we estimate Equation 2.1 on a set of real economic outcomes, including foreign assets, wages, capital investment, and R&D. Foreign assets are reported on Form 5471, wages are reported on Form 1120, capital investment is measured as the sum of reported assets placed into service on Form 4562, and we proxy for R&D with the claimed R&D tax credit reported on Form 6765. Figure 2.11 reports the results of the event study analyses for these outcomes. As with the financial and tax outcomes, we do not estimate significant pre-trends in any of the real outcomes. Following the adoption of a hybrid tax planning structure, we observe significant relative increases in foreign and domestic capital assets, domestic R&D activity, and domestic wages.

Table 2.6 summarizes these estimates using a difference-in-differences specification. Panel A reports an average increase of 18 percent in foreign capital assets. Panel B shows that domestic R&D increases by 40 percent. Panel C reports a 17 percent increase in domestic wages, while Panel D reports a 9 percent increase in domestic capital assets. For all these outcomes, we find that including potential drivers of selection into a hybrid tax planning structure does not impact any of the estimates. Table 2.7 further shows that we obtain similar estimates when using the estimator of Sun and Abraham (2021).

To gauge the magnitude of these changes in real outcomes, it is useful to compare them to the estimated decrease in foreign ETRs. Specifically, we compute tax rate elasticities as the ratios of the percentages in each of these outcomes and the percentage change in foreign ETRs. Since Figure 2.7 showed that foreign ETRs declined more gradually, we compare the estimates on real outcome with the decline in foreign ETRs after six years of adoption. Relative to the average foreign ETR of close to 27 percent at the beginning of our sample period (Figure 2.1), we estimate that foreign ETRs declined by 18.4 percent. Compared to the impacts on real outcomes, we then estimate a foreign ETR elasticity of foreign capital of 0.98, a domestic payroll elasticity of 0.92, a domestic R&D elasticity of 2.2, and a domestic

capital elasticity of 0.5.[16]

2.6 A Model of International Tax Planning

This section develops a model in which MNCs provide services in different countries, engage in hybrid tax planning, and develop intellectual property through R&D. The model highlights two key aspects of tax planning. First, tax planning can impact production and investment abroad by lowering the cost of capital. In the model, R&D investment is nonrival within the MNC and impacts production everywhere by increasing productivity across the MNC. By lowering the cost of capital abroad, tax planning also incentivizes the MNC to increase R&D. Separately, the model allows for the possibility that MNCs select into a hybrid tax planning strategy based on unobserved shocks to their IP. By modeling the selection behavior of MNCs, we can decompose the reduced-form difference-in-differences estimates from the previous section into a "treatment effect" and a "selection effect."

2.6.1 Model Setup

The model has three stages, which we describe in descending order. In stage 3, the MNC takes taxes and group-level productivity ϕ as given and produces output—which is sold locally in each country—using local capital K. In stage 2, the MNC decides how much to invest in R&D taking taxes as given. In stage 1, the multinational chooses whether to set up a hybrid tax planning structure, taking into account idiosyncratic costs of tax planning, as discussed in Section 2.2.[17]

[16]These calculations use specification 4, where foreign ETRs declined by 5 percentage points. In specification 1, where foreign ETRs declined by 6.7 percentage points, we instead estimate a 37 percent decrease in foreign ETRs. This estimate corresponds to a foreign asset elasticity of 0.49, a domestic wage elasticity of 0.46, a domestic R&D elasticity of 1.09, and a domestic capital elasticity of 0.25.

[17]This model abstracts away from other behavior, such as the manipulation of transfer pricing in individual transactions (e.g., as in Hines and Rice, 1994). Such behavior can be analyzed by adding a fourth stage where MNCs optimally misreport the location of their profits.

Table 2.6: Estimates on Real Economic Activity

	1	2	3	4
Panel A. Foreign Capital Assets				
ATT	0.181**	0.177*	0.181*	0.189*
	(0.070)	(0.071)	(0.075)	(0.076)
Num. Units	2860	2860	2860	2860
Num. Treated	320	320	320	320
Panel B. Domestic R&D Credit				
ATT	0.395**	0.384**	0.388**	0.393**
	(0.143)	(0.147)	(0.143)	(0.151)
Num. Units	1744	1744	1744	1744
Num. Treated	262	262	262	262
Panel C. Domestic Wages				
ATT	0.169**	0.165**	0.161**	0.161**
	(0.059)	(0.061)	(0.060)	(0.061)
Num. Units	2823	2823	2823	2823
Num. Treated	311	311	311	311
Panel D. Domestic Capital Assets				
ATT	0.069	0.097*	0.062	0.092+
	(0.045)	(0.046)	(0.047)	(0.048)
Num. Units	2835	2835	2835	2835
Num. Treated	312	312	312	312
Industry by Year FEs	No	Yes	No	Yes
Size by Year FEs	No	No	Yes	Yes

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table provides ATT estimates for a staggered two-way fixed effects difference-in-differences model provided in equation 2.1 for the corresponding foreign outcome listed in each panel. Specification 1 does not include additional controls. Specification 2 includes year-by-industry fixed effects. Specification 3 includes year-by-group fixed effects, where groups include domestic and foreign sales quartiles and domestic and foreign intangible asset quartiles, and where quartiles are computed using pre-treatment values. The last specification includes all fixed effects used in specifications 2 and 3.

Table 2.7: Estimates on Real Economic Activity (Sun & Abraham)

	1	2	3	4
Panel A. Foreign Capital Assets				
ATT	0.263***	0.266***	0.248**	0.263***
	(0.073)	(0.075)	(0.077)	(0.079)
Num. Units	2860	2860	2860	2860
Num. Treated	320	320	320	320
Panel B. Domestic R&D Credit				
ATT	0.371*	0.400*	0.373*	0.392*
	(0.145)	(0.164)	(0.152)	(0.171)
Num. Units	1744	1744	1744	1744
Num. Treated	262	262	262	262
Panel C. Domestic Wages				
ATT	0.215***	0.203**	0.212***	0.202**
	(0.061)	(0.062)	(0.062)	(0.061)
Num. Units	2823	2823	2823	2823
Num. Treated	311	311	311	311
Panel D. Domestic Capital Assets				
ATT	0.088*	0.119**	0.083+	0.117*
	(0.042)	(0.044)	(0.044)	(0.046)
Num. Units	2835	2835	2835	2835
Num. Treated	312	312	312	312
Industry by Year FEs	No	Yes	No	Yes
Size by Year FEs	No	No	Yes	Yes

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table provides ATT estimates for a staggered difference-in-differences model using the estimator proposed by Sun and Abraham (2021) for the corresponding foreign outcome listed in each panel. Specification 1 does not include additional controls. Specification 2 includes year-by-industry fixed effects. Specification 3 includes year-by-group fixed effects, where groups include domestic and foreign sales quartiles and domestic and foreign intangible asset quartiles, and where quartiles are computed using pre-treatment values. The last specification includes all fixed effects used in specifications 2 and 3.

Stage 3: Production

For simplicity, we consider two countries: d for domestic and f for foreign. The multinational sells quantities q_c of its product in each country c. The prices for its product in each country is $p_c = I_c(q_c)^{\frac{1}{\varepsilon}}$, where I_c is a demand shifter (e.g., income) and $\varepsilon < -1$ is the elasticity of product demand. The MNC produces output with technology: $q_c = \phi K_c$, where K_c has a global price ρ, which we assume is not tax deductible.[18]

We assume the MNC faces a worldwide tax on corporate income from its home country d. If the MNC sets up a hybrid tax planning structure, the MNC still pays a tax rate of t_d on domestic profits but pays a lower tax rate of t_h by shifting foreign profits to a tax haven, which it defers indefinitely. In this way, the hybrid structure can be seen as a "do it yourself territorial regime" combined with a profit shifting strategy to avoid foreign country taxes.

The MNC solves the production problem:

$$\max_{K_c, K_f} \; (1-t_d) I_d (\phi K_d)^{1+\frac{1}{\varepsilon}} - \rho K_d + (1-t_h) I_f (\phi K_f)^{1+\frac{1}{\varepsilon}} - \rho K_f.$$

Note that assuming that $t_d = t_h$ yields the special case where the MNC does not set up a hybrid tax planning structure and instead faces worldwide income regime in which the domestic country has a higher tax rate $t_d > t_f$ and a credit for foreign taxes.

The optimal capital use in each country is given by: $K_c = (\phi)^{-(1+\varepsilon)} \left[\frac{1+\varepsilon}{\varepsilon} \frac{I_c}{\rho/(1-t_c)} \right]^{-\varepsilon}$. Because capital costs are not tax deductible, higher tax rates increase the effective cost of capital $\frac{\rho}{(1-t_c)}$. The hybrid tax planning structure does not impact domestic capital through the tax rate, but it can increase foreign capital by lowering the tax rate faced by foreign production.

As is the case in models of monopolistic competition, profits are given by a net markup,

[18] For simplicity, we assume the cost of capital is not tax deductible. In practice, the cost of capital may be partially tax deductible either because of a lack of expensing or from limits to the deductibilty of interest expense.

$\frac{-1}{1+\varepsilon}$, times costs, ρK_c, which yields country profits $\pi_c = \frac{-\rho}{1+\varepsilon}K_c$. Global profits are then:

$$\frac{-\rho}{1+\varepsilon}(K_d + K_f) = \frac{-\rho}{1+\varepsilon}K_d\left[1 + \left(\frac{I_d}{I_f}\right)^{\varepsilon}\left(\frac{1-t_d}{1-t_h}\right)^{\varepsilon}\right] = \pi_0\phi^{-(1+\varepsilon)}\left[1 + \left(\frac{I_d}{I_f}\right)^{\varepsilon}\left(\frac{1-t_d}{1-t_h}\right)^{\varepsilon}\right],$$

where $\pi_0 = \frac{-\rho}{1+\varepsilon}\left[\frac{1+\varepsilon}{\varepsilon}\frac{I_d}{\rho/(1-t_d)}\right]^{-\varepsilon}$.

The hybrid tax planning strategy increases the value of producing in the foreign country since $t_h < t_d$ and $\varepsilon < -1$ imply that $\left(\frac{1-t_d}{1-t_h}\right)^{\varepsilon} > 1$. Note that this term includes two forces. First, since the effective cost of capital is lower, the MNC invests more abroad, which increases profits. Second, the MNC keeps a larger fraction of its profits since it shifts profits to tax havens.

Stage 2: R&D Investment

The MNC decides how much to invest in R&D, which impacts ϕ. The investment decision takes into account the fact that its intellectual property is non-rival within the firm and will increase profits in both locations. For an MNC i, let $\phi_i = \zeta_i \times D_i^{\gamma}$, where $\gamma < \frac{1}{-(1+\varepsilon)}$, ζ_i is the MNCs existing stock of IP, and D_i is R&D. The MNC solves the R&D investment problem:

$$\max_{D_i} \ \pi_0\zeta_i^{-(1+\varepsilon)}D_i^{-\gamma(1+\varepsilon)}\left[1 + \left(\frac{I_d}{I_f}\right)^{\varepsilon}\left(\frac{1-t_d}{1-t_h}\right)^{\varepsilon}\right] - D_i - \frac{b_i}{2}D_i^2.$$

The term b_i is a firm-level adjustment cost of R&D investment. This term captures the fact that some firms may have a greater potential to engage in R&D projects. Since tax planning is complementary with R&D, firms with a higher potential to engage in R&D—i.e., those with lower values of b_i—may be more likely to set up HTPs.

To understand how differences in productivity ζ_i and tax rates impact R&D investment, consider the simpler case where $b_i = 0$. In this case, the optimal choice of $R\&D$ is given by:

$$D_i^* = \left\{-\gamma(1+\varepsilon)\pi_0\zeta_i^{-(1+\varepsilon)}\left[1 + \left(\frac{I_d}{I_f}\right)^{\varepsilon}\left(\frac{1-t_d}{1-t_h}\right)^{\varepsilon}\right]\right\}^{\frac{1}{1+\gamma(1+\varepsilon)}}.$$

It follows from this expression that the optimal choice of R&D is increasing in firm productivity ζ_i and in the relative reduction in foreign tax rates $\left(\frac{1-t_h}{1-t_d}\right)$. That is, MNCs conduct more R&D if they have a larger stock of IP and if setting up a hybrid tax structure will increase its returns from R&D investment abroad.[19] For this reason, MNCs with higher values of ζ_i have more to gain from establishing HTPs.

Let $\Pi(t_d, t_h, \zeta_i, b_i)$ denote the value function from conducting the optimal level of R&D using a hybrid tax planning structure. Similarly, let $\Pi(t_d, t_d, \zeta_i, b_i)$ denote the value of remaining in the worldwide tax system.

Stage 1: Tax Planning

Setting up a hybrid tax planning structure has a fixed cost c_i. This cost includes legal and accounting services related to tax planning, the cost of transactions to establish HTPs, as well as differences in executive attitudes toward aggressive tax planning structures.

The MNC decides to set up a hybrid tax planning structure when the profits it gains from doing so exceed the costs of setting-up the structure, i.e., if

$$\Pi(t_d, t_h, \zeta_i, b_i) - c_i > \Pi(t_d, t_d, \zeta_i, b_i).$$

This inequality captures the four forces driving selection into HTPs. First, the likelihood that a firms adopts an HTP is increasing in the tax wedge $t_d - t_h$. Second, increased productivity ζ_i leads to an increase in the productivity of adopting an HTP. Third, for firms with the same productivity ζ_i, those with better prospects for conducting R&D—i.e., lower b_i—will be more likely to adopt an HTP. Finally, firms with lower cost c_i will also be more likely to adopt an HTP. As we discuss below, these different sources of selection have implications for the interpretation of the difference-in-differences estimates as well as for the concentration of profits and R&D among adopting and non-adopting firms.

[19] Taking logs and differentiating shows that $\frac{\partial \ln D_i^*}{\partial \ln \zeta_i} = \frac{-(1+\varepsilon)}{1+\gamma(1+\varepsilon)} > 0$ and $\frac{\partial \ln D_i^*}{\partial \ln\left(\frac{1-t_h}{1-t_d}\right)} = \frac{-\varepsilon}{1+\gamma(1+\varepsilon)} > 0$. These elasticities are attenuated when $b_i > 0$.

2.6.2 Difference-in-Differences Estimates Through the Lens of the Model

Let β^D denote the difference-in-differences estimate of the relative change in log-R&D. From the discussion above, β^D combines a behavioral increase in R&D from exposure to lower tax rates, as well as differences in R&D that are attributed to selection based on productivity ζ_i and R&D potential b_i. From the equation describing the evolution of TFP, ϕ_i, we can write average relative changes in log-ϕ_i as $\Delta \ln \phi_i = \Delta \ln \zeta_i + \gamma \beta^D$, where the first term captures selection into tax planning based on ζ_i and the second captures how differences in R&D relate to differences in TFP ϕ_i.

Recall that optimal capital in a given country is given by: $K_c = (\phi_i)^{-(1+\varepsilon)} \left[\frac{1+\varepsilon}{\varepsilon} \frac{I_c}{\rho/(1-t_c)} \right]^{-\varepsilon}$. This expression and the preceding ones allow us to write the difference-in-differences estimates on domestic and foreign capital as follows:

$$\text{Domestic Capital:} \quad \beta^{K_d} = -(1+\varepsilon) \times (\Delta \ln \zeta_i + \gamma \beta^D) \tag{2.2}$$

$$\text{Foreign Capital:} \quad \beta^{K_f} = \beta^{K_d} - \varepsilon \ln \left(\frac{1-t_d}{1-t_h} \right). \tag{2.3}$$

Intuitively, we observe relative increases in domestic capital either because of differences in R&D or because of selection on ζ_i. The expression for foreign capital shares this TPF effect but additionally allows for capital investment to be influenced by tax differentials.

Letting the estimate on foreign ETR be given by: $\beta^{ETR} = \ln \left(\frac{1-t_d}{1-t_h} \right)$, we can then identify the parameter ε using the domestic and foreign capital effects as follows: $-\varepsilon = \frac{\beta^{K_f} - \beta^{K_d}}{\beta^{ETR}}$.[20] We can also recover the parameter γ by solving the expression for β^{K_d} above as follows:

$$\gamma = \left(-\frac{\beta^{K_d}}{(1+\varepsilon)} - \Delta \ln \zeta_i \right) / \beta^D.$$

This expression depends on the term $\Delta \ln \zeta_i$, which we will estimate using the selection model in the tax planning stage.

[20] The main assumption behind this result is that domestic and foreign productivity R&D are equally (or at least proportionally) affected by the underlying productivity of the MNCs.

Armed with an estimate of $\Delta \ln \zeta_i$, Equations 2.2 and 2.3 can be used to decompose the difference-in-differences estimates into behavioral responses to taxation and selection effects. Even without such an estimate, we can use Equations 2.2 and 2.3 to provide bounds on the selection and behavioral responses of HTP adoption. Under the assumption that $\Delta \ln \zeta_i = 0$, we would attribute all of the changes in domestic and foreign capital to behavioral responses to lower tax rates. In this case, domestic capital increases as a result of the productivity effects of the additional (tax-driven) R&D. Foreign capital additionally increases due a tax-investment effect. In contrast, consider now the largest possible role for selection. In this case, the entirety of the R&D and domestic capital changes would be due to selection.[21] Additionally, the behavioral response to lower taxes would then explain $1 - \frac{\beta^{K_d}}{\beta^{K_f}}$ of the changes in foreign capital. While these bounds may be informative, one of the benefits of the model is to obtain a more precise decomposition of the forces behind the difference-in-differences estimates.

2.6.3 Model Estimation

We now discuss how we parameterize and estimate the model. We assume that the drivers of selection (ζ_i, b_i, c_i) have a joint log-normal distribution $\mathcal{LN}(\mu, \Sigma)$. The vector $\theta = (\varepsilon, \gamma, \mu, \Sigma)$ includes parameters governing firm production and the distribution of drivers of selection.

We estimate θ using the method of simulated moments. For a given guess of θ, we simulate 30,000 firms and solve for their optimal capital use in each country, their optimal R&D, and their implied profits at home and abroad. Using these simulated firm choices, we compute moments of the firm size and profit distributions, the concentration of R&D in tax planning firms, and the fraction of tax planning firms. We also compute difference-in-differences estimates using the simulated data. Our estimate of θ is the value that minimizes the difference between these simulated moments and their empirical counterparts.

The discussion in the previous section shows how the difference-in-differences estimates

[21] Since $\beta^D > 0$ and $\beta^{K_d} > 0$ it follows that $\gamma = 0$. We therefore have that $\Delta \ln \zeta_i = -\frac{\beta^{K_d}}{(1+\varepsilon)}$.

Table 2.8: Estimated Model Parameters

	(1) Elasticity of Product Demand	(2) Productivity Elasticity of R&D	(3) Dispersion in	(4) Dispersion in	(5) Dispersion in	(6) Mean Cost
			ζ_i	b_i	c_i	
	ε^{PD}	γ	σ	σ_b	σ_c	$\bar{c}$
Estimate	-1.65***	0.70***	0.93	4.28	7.05**	8.19
SE	(0.46)	(0.27)	(0.95)	(2.88)	(3.42)	(31.12)

Notes: This table provides estimates of selected model parameters based on the simulated method of moments.

can help identify ε and γ. We now provide intuition for how the additional data moments can help pin down the importance of the different drivers of selection. If only variation in c_i drives selection, we should expect to see that firm size is unrelated to HTP adoption. This would be the case if in the data we observed that the share of profits of HTP firms equaled the share HTP firms and their share of R&D. In this case, we would also find that $\Delta \ln \zeta_i = 0$. Alternatively, if HTP firms have a much larger share of profits, we would expect to find that $\Delta \ln \zeta_i > 0$. Similarly, if the share of R&D of HTP firms is larger than the share of HTP firms, this would be indicative of selection on R&D potential b_i. While the data help identify the parameters of the distribution of (ζ_i, b_i, c_i), we are not able to idenfy all of the components of μ and Σ. We focus on estimating the mean and dispersion of c_i, the mean and dispersion of ζ_i, and the mean of b_i (we set $\mu_b = 0$).[22]

Table 2.8 reports the results of this estimation. We estimate a value of $\varepsilon = -1.65$. This value of the elasticity of product demand implies large markups, which may be reasonable in IP-heavy industries such as technology and pharmaceutical manufacturing. We estimate a value of $\gamma = 0.7$, which we interpret as the long-run effect of R&D on productivity. Of the drivers of selection, we estimate the largest value for σ_c. This result implies that idiosyncratic differences in the costs of adopting an HTP play a significant role in explaining selection. Nonetheless, as we show below, the model also attributes a large fraction of the changes in real economic activity to selection based on b_i and ζ_i.

[22]We set all correlations to zero, except for $\rho_{\zeta,c}$ which we initially calibrate to 0.15.

Figure 2.12 shows how the estimated model matches the data. The estimation targets the share of firms, profits, tax planning firms, and R&D at four ranges of the firm size distribution. Since tax planning firm are more likely to be larger, Panel A plots the model and data shares for firms in the Top 10% of the sales distribution. This figure shows that the model parameters do a good job at matching the importance of firms in the top of the distribution. Panel B plots aggregate moments and difference-in-differences estimates. While the model matches the overall share of planners and the estimates on domestic and foreign capital, the model has a hard time fitting the large fraction of foreign profits in HTP firms as well the large relative changes in R&D.

(A) Shares of Economic Activity of Hybrid Tax Planners Among Firms in the Top 10% of Sales

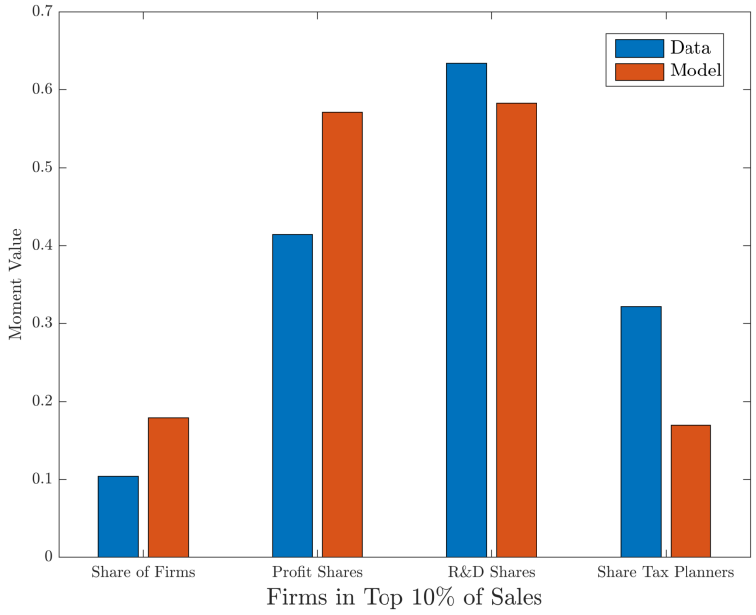

(B) Shares of Economic Activity of Hybrid Tax Planners Among All Firms and Difference-in-Differences Estimates

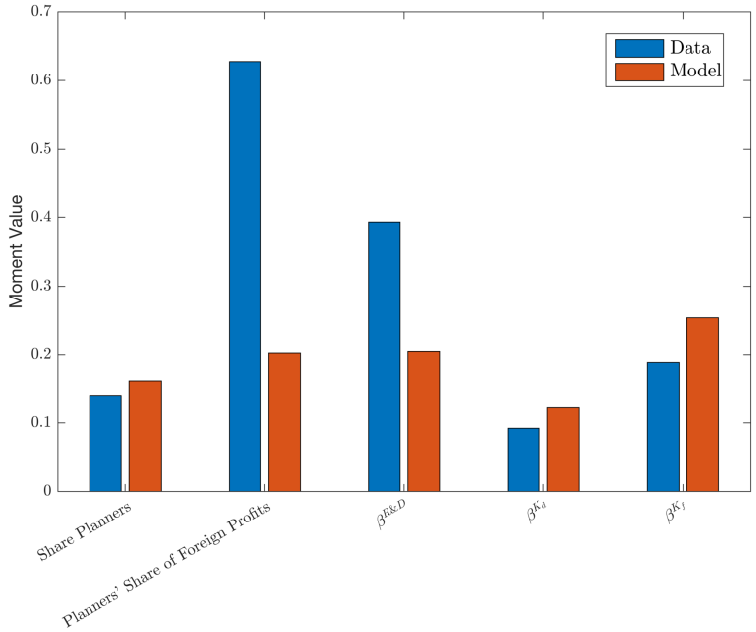

Figure 2.12: Comparison of Data and Model Moments

Notes: These figures compare simulated moments with data moments. Panel A compares share of economic activity of HTPs among firms in the top 10% of the sales distribution. Panel B shows aggregate shares and difference-in-differences estimates.

Table 2.9: Model Decomposition of Difference-in-Differences Estimates

	(1) DiD	(2) Model	(3) Tax Effect	(4) Selection	(5) Share due to ΔETR
R&D: $\beta^{R\&D}$	0.39	0.20	0.04	0.17	0.17
Domestic Capital: β^{K_d}	0.09	0.12	0.05	0.07	0.42
Foreign Capital: β^{K_f}	0.19	0.25	0.18	0.07	0.72

Notes: This table decomposes the difference-in-differences estimates into tax and selection effects. Column (1) lists the estimates from Section 2.5.2. Column (2) lists the model-implied values of these estimates. Column (3) lists the relative change in each outcome that can be attributed to the behavioral response to the tax change. Column (4) lists the selection effect and column (5) lists the ratio of columns (3) and (2).

2.6.4 Decomposition and Counterfactual Effects

As discussed in Section 2.6.2, we can use the model estimates to decompose the estimates of the difference-in-differences models into selection and behavioral responses of the tax effect of HTPs. Table 2.9 performs this decomposition. The first two columns list the data and model-based estimates of the difference-in-differences models. As discussed above, the model does a good job of matching the foreign and domestic capital estimates but undershoots the relative changes in R&D. Column (3) lists the contribution of the behavioral response to the tax effect to column (2) and column (4) lists the contribution of the selection term. Finally, column (5) lists the fraction of the effect in column (2) that can be attributed to the tax effect. In the case of R&D, the model implies that upward of 80% of the relative change is driven by selection. In the case of domestic capital, 40% of the relative change is driven by the effect of the tax advantage from HTPs. However, the model shows that about 70% of the change in foreign capital is driven by the response to lower taxes abroad.

This decomposition has meaningful implications for tax policy design. For the US, this result says that a significant fraction of the relative increase in the domestic activity of tax-planning MNCs is not driven by low tax rates. That is, the benefits to the US from tax planning are muted by selection.

To study the implications of our results for foreign governments, we use our estimated

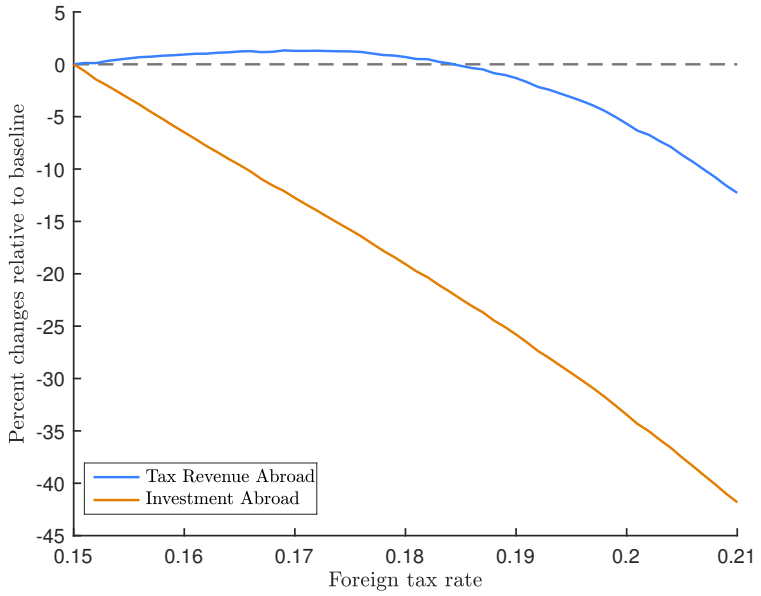

Figure 2.13: Counterfactual Analysis: Increasing Tax Haven Tax Rate t_h

Notes: This figure plots simulated percentage changes in tax revenues an investment in the foreign country as the tax rate in the tax haven is increased from 15% to 21%.

model to simulate the impact of reducing the tax advantage from HTPs from a 5pp. point differential to zero. In the model, we assume that the foreign tax rate under a HTP increases from 15% to 21%. Figure 2.13 plots the percentage changes in investment abroad and foreign tax revenues as this tax rate increases. The yellow line shows that investment abroad decreases significantly with the tax increase. This result is consistent with the decomposition in Table 2.9 showing that the tax advantage of HTPs explains a large share of the relative changes in foreign capital. The blue line shows that tax revenues initially increase with the tax rate. However, this increase is short-lived and tax revenue declines for foreign tax rates greater than 18.5%. Figure 2.13 paints a bleak picture for foreign governments wishing to increase the fiscal obligations of US MNCs, as it shows that limits to tax planning may significantly reduce investment and may lead to small increases or even declines in tax revenue.

2.7 Conclusion

Complex tax planning strategies have been a focus of media attention and have played an important role in motivating international tax reforms. Despite this focus, policy makers, practitioners and academics lack a comprehensive understanding of the prevalence of these strategies and their role in explaining some key facts surrounding international taxation.

Using an unique integration of tax data covering the domestic and foreign operations of US MNCs, we fill this important gap by reconstructing the ownership networks of foreign affiliates of US MNCs. This allows us to identify the adoption of three important tax planning strategies: the Double Irish and Reverse Hybrid Mismatch arrangements through the Netherlands and Luxembourg. We show that these structures account for a significant fraction of the foreign profits of US MNCs. We also link the use of these structures to financial transactions that companies could use to shift profits to low tax (and tax haven) countries. While only 16% of US MNCs adopt these structures, those that do obtain a significant tax advantage abroad over other MNCs. Remarkably, these few companies are responsible for the bulk of the increase in cash held abroad over the period we study.

Our analyses using the tax data reveal that the adoption of hybrid tax planning structures is accompanied by significant changes in real economic activity. We find that firms that adopted the structures we study also had significant increase in foreign tangible and intangible assets. While our differences-in-differences estimates measure large change in domestic payroll and R&D, our model attributes most of these changes to differential selection into tax planning. This result suggests that the gains to the domestic economy are smaller than the difference-in-differences estimates would imply. While our model accounts for selection, it also finds that foreign investment is sensitive to changes in tax planning opportunities. This result implies that policies that increase tax costs abroad may lead to small gains in tax revenues for foreign governments and significant declines in tangible investment.

The fact that few specific tax planning structures play an outsize role in explaining aggregate trends in MNC behavior that is associated with income shifting is important

for the design of the international tax system. Rather than focusing efforts on integrating policies over 137 countries, our results show that monitoring tax motivated structures set up in three specific countries may go a long way in curbing profit shifting.

3

Measurement of Multinational Activity in IRS Administrative Data

3.1 Introduction

Tax competition and corporate tax evasion have become central areas of focus for researchers and policymakers over the past several decades as the global economy has become increasingly integrated. Yet it is suprisingly difficult to reliably nail down even basic descriptive information regarding key elements of multinational corporate activity. For example, estimates of US tax revenue lost to profit shifting can differ by an order of magnitude. Clausing (2016) provides an upper bound estimated at $111 billion whereas Blouin and Robinson (2020, henceforth B&R) estimate a figure of only $10 billion.

A key reason for these discrepancies relates to what is a growing understanding of the potential for significant measurement error in various commonly used datasets of multinational activity (the first mention that I am aware of is Altshuler and Grubert, 2006*a*). B&R distinguish between two types of measurement error, **aggregation error** and **misattribution error**.

In the context of measurement of corporate earnings, aggregation error can occur when the same earnings appear on the books of multiple affiliates within a conglomerate. A simple approach to aggregation that sums the earnings of all affiliates within a conglomerate will then double count these earnings, biasing estimates upward.

Researchers also often attempt to determine the location where profits are earned. Earnings might be reported, however, not at the level of the affiliate or branch where a dollar of income was transacted but instead at the level of a holding company further up the ownership chain. If the holding company is domiciled in a different jurisdiction than the

lower-level affiliate, then researchers may incorrectly attribute the earned dollar of income to the wrong country. In the extreme, this type of error can distort international economic statistics. A well known example is that of Ireland, a popular domicile for holding companies, where GDP statistics are thought to be highly distorted relative to the true level of economic activity within the country.

As multinational corporations have become increasingly complex, with potentially hundreds of affiliates distributed globally, the potential for both types of measurement error have increased. This time-varying error can undercut the reliability of analyses that rely on temporal variation to measure changes in corporate activity. While B&R provide several recommendations that researchers can use to guard against measurement error when using accounting data from the BEA, they provide only cursory discussion of the potential for measurement error in IRS business tax data based on publicly available aggregate statistics and suggest that it may not be possible to fully correct for such measurement error in tax data.

This chapter takes a closer look at the potential for measurement error when aggregating tax filings by foreign affiliates of US multinationals. By comparing aggregated income from tax filings with consolidated foreign income data as reported in public filings, I show that adjustments to correct for aggregation error in tax data appear to be quite promising.

Why do we need yet another source of data to measure US multinational activity if BEA data (properly adjusted) give reliable measurements? The answer is that there are a number of advantages that tax data have over their BEA counterpart. First, BEA surveys typically focus only on large US multinationals. Tax data, on the other hand, are comprehensive, and samples compiled by the IRS contain information about smaller firms that are active abroad. Second, as I show in Section 3.3, measurement of corporate income tax appears to be more reliable in IRS data when compared with data from corporations' public filings. This may be particularly true in certain industries, such as oil and gas, where the concept of what qualifies as income tax can be fuzzy.

This chapter proceeds as follows. Section 3.2 provides an overview of several different commonly used data sources that are used to measure the activity of US multinationals, as well as a discussion of the relative advantages and disadvantages of tax data compared with other sources of accounting data for US multinationals. Section 3.3 describes how I construct a matched sample of multinationals that appear in both tax data and public filings compiled by Compustat. I use this sample to examine differences in how foreign earnings and income taxes are reported, to measure the extent of aggregation error that may be generated by naive aggregation of foreign affiliates' tax filings, and to assess the performance of a simple correction that aims to remove aggregation error from tax data. In Section 3.4, I reproduce estimates from earlier studies that attempt to measure tax semi-elasticities, i.e. how the allocation of foreign earnings by US multinationals correlates with average and statutory foreign tax rates. I also provide corrected estimates of average foreign effective tax rates by jurisdiction. Section Section 3.5 concludes.

3.2 Overview of Data Sources for Measuring Tax Outcomes of US Multinationals

Below, I provide an overview of several different datasets that researchers have used to study US multinational activity. Much of the literature discusses measurement issues with a focus on data compiled from surveys run by the Bureau of Economic Analysis (BEA). I review the main issues encountered in this data. I then discuss the extent to which these measurement issues may translate to the sources used in this study, Compustat and SOI.

3.2.1 Compustat

Compustat is a database that, among other things, compiles information about firms based on their public filings. While Compustat is a useful source of information for a large set of firms, it is not comprehensive—private firms are not covered by the database. Furthermore,

information about MNCs' foreign operations is relatively limited in Compustat. It does, however, provide some information about the aggregate consolidated foreign income and foreign tax expense for a set of MNCs.

3.2.2 BEA Multinational Data

The BEA provides two different sources of data on US multinationals that are frequently used in the literature. The first is typically referred to as "Direct Investment Income". The second source, which provides two different foreign income measures, is published under "Activities of US Multinational Enterprises."

Direct Investment Income

The BEA conducts quarterly surveys that collect information on investment positions and transactions with directly-owned foreign affiliates of US multinationals. Research that examines foreign earnings by US multinationals have used the "Direct Investment Income" (DII) measure contained in this data (e.g., Zucman, 2014, 2015). One advantage of DII is that it does not appear to suffer from issues related to double-counting of equity income that are present in other data sources, as discussed below. On the other hand, on a country-by-country basis, DII may not reflect the true geographic distribution of where earnings are first realized. This is due to the fact that income generated by indirectly-held foreign affiliates will only be observed at the level of a directly-owned affiliate. In addition, DII does not provide measures of foreign corporate income tax. As discussed in B&R, although researchers may attempt to combine DII with other data sources that provide information about tax payments, this can result in underestimation of average effective rates for havens to the extent that taxes on equity income do not match the jurisdiction in which they appear in DII data.

Activities of US Multinational Enterprises

The BEA also conducts annual surveys that collect detailed information on the activity of foreign affiliates directly or indirectly owned by US multinationals. This data is compiled into annual reports titled "Activities of US Multinationals" (AMNE). Much of the international tax literature focuses on statistics provided on the Income Statement (Tables starting with "D"). AMNE data provide disaggregated information that separates income from indirectly and directly owned foreign affiliates, thereby avoiding the misclassification error present in the DII data. As discussed in B&R, however, simple aggregation of foreign net income may double count so-called "equity income" that is included on the income statement of affiliates that are parents of other affiliates. B&R provide a simple adjustment to prevent such double-counting by subtracting related dividend income, which is also reported by the BEA.

Another advantage of AMNE data is that, unlike DII data, measures of corporate income tax are also provided. This allows for more reliable estimates of taxes paid and average tax rates by jurisdiction.

3.2.3 Administrative Tax Data

The Statistics of Income (SOI) division of the IRS produce two sets of corporate data that have been used by researchers to study US multinationals. These datasets provide information disclosed in tax returns for a large sample of US corporations. They provide information not only about the activity US parent companies, but also their foreign affiliates. The sample includes public and private corporations as well as information about relatively small firms that are not captured in BEA surveys.

SOI Corporate Sample

The first dataset, the SOI corporate sample, is an annual stratified sample of US corporations that SOI uses to produce publicly available aggregated business income statistics.[1] The SOI corporate sample provides information from unaudited tax returns for approximately 100,000 US corporations annually and has been widely used in the business tax literature to study domestic firms (Yagan, 2015; Zwick and Mahon, 2017). The sample used in this chapter contains C corporations that were sampled between 1992 and 2016. This analysis uses Schedule M-3, an attachment to Form 1120 that is mandatory for corporations with total assets of at least $10 million.[2]

SOI International Business Tax Statistics

The second dataset, which provides information related to foreign affiliates of US corporations, is used by SOI to publish aggregate statistics for international business taxes IRS (2022*a*). This dataset provides information from filings of Form 5471, an information return that provides financial information and activity of opaque foreign affiliates, as well as Form 8858, which provides similar information for passthrough foreign affiliates.[3] Unlike the SOI corporate sample, which is provided annually, information about opaque foreign affiliates are only collected in even years (with the exception of 2014). Information about passthrough foreign affiliates is collected for four of the years in our sample period (2006, 2008, 2012, and 2016). Prior to 2004, SOI only collected information for foreign affilaites of large multinationals with greater than $500 million in assets. Starting in 2004, the sampling procedure became much broader, covering a superset of the SOI corporate sample.

Like the BEA's AMNE data, SOI international tax data include information about foreign affiliates whether or not they are directly or indirectly owned. One key difference from

[1] Statistics are available at IRS (2022*b*), and the sampling procedure is described in IRS (2011).
[2] See `https://irs.gov/pub/irs-pdf/f1120sm3.pdf`
[3] Passthrough entities are often discussed in reference to a policy change known as "Check the Box." See Chapter 2.2 for a detailed discussion of this policy.

AMNE data is that, as alluded to above, foreign affiliates may be classified for tax purposes as opaque entities, which the IRS calls "controlled foreign corporations" (CFCs), or passthrough entities, which it calls "foreign disregarded entities" (FDEs). US parents can control their foreign affiliates directly or indirectly through another CFC or FDE. CFCs that are structured as parents of FDEs report consolidated financial information on Form 5471. This method of reporting can cause misclassification error similar to that observed in the BEA's DII series.

To see how misclassification error can occur in the tax data, consider a US multinational with two foreign affiliates — affiliate A is a holding company located in a tax haven, e.g. Bermuda, and affiliate B is an affiliate that sells merchandise to customers, e.g. in Germany. If the US parent classifies affiliate B as a passthrough (FDE) and affiliate A as an opaque entity (CFC), then the information contained on affiliate A's Form 5471 will provide foreign income and tax data that include its German operations. As noted above, since information covering passthrough entities is available only for a handful of years, in general it is not possible to "undo" this misclassification error. Unlike DII data, however, the SOI also reports information about foreign tax payments.

The fact that the SOI contains misclassification error is not necessarily a "disadvantage" of the tax data. This instead depends on the researcher's goal. If researchers are attempting to accurately measure effective tax rates imposed by countries on businesses that operate in their jurisdiction, SOI data will generally be unreliable when compared with BEA data. If researchers are instead interested in the effective tax rate faced by combined foreign structures, i.e. in the consolidated operation of Bermudan-German affiliates, then the tax data may be preferable as the BEA does not provide information about whether foreign affiliates are classified as passthroughs or opaque entities for tax purposes. Thus, while misclassification error may be present in SOI data, this may not be problematic as long as researchers have the correct interpretation of the data.

SOI data may also suffer from aggregation error in a manner that is similar to the BEA's

AMNE dataset. This may occur if there is indirect ownership of one opaque entity by another. In this case, foreign earnings that are distributed as dividends to the parent affiliate may appear as foreign earnings on each affiliate's Form 5471. However, the SOI also provides information that can be used to correct for this type of error. Schedule M of Form 5471 reports transactions between related foreign affiliates, including dividends. B&R suggest removing these dividends to accurately aggregate foreign income. Based on a comparison with BEA data, however, they question whether this adjustment is sufficient to fully correct for aggregation error.

Aggregation error potentially poses a much more serious problem for researchers. If SOI data overstates foreign earnings, estimates from research that examines foreign earnings in tax data may not be reliable. Worse, to the extent this error is growing over time, which one might suspect if multinationals' ownership structures are also becoming more complex, standard economic approaches to control for such error (e.g. unit fixed effects) might only exacerbate the problem (Bound, Brown, Duncan and Rodgers, 1994; Bound and Krueger, 1991).

3.3 Assessing Aggregation Error in Foreign Earnings and Taxes

In this section, I use data from Compustat and SOI to measure the extent to which SOI data may suffer from aggregation error, and examine whether the correction proposed by B&R appears to sufficiently correct for this error. First, I describe how I create a matched sample of US multinationals that appear both in SOI and Compustat. Compustat reports foreign income and foreign tax figures on a consolidated basis for a set of multinationals.[4] The consolidated nature of Compustat's reporting implies that it should not be contaminated by aggregation error. This provides a baseline comparison that can be used to measure aggregation error for a common set of firms, avoiding the possibility that aggregate differences may simply be due to differences in sample composition as in the analysis contained

[4]The Compustat sample typically includes large, publicly-traded firms.

in B&R.[5]

3.3.1 Sample Construction

Figure 3.1 shows attrition associated with the sample construction process. The top line in the left panel shows the annual sample size for US multinationals contained in the SOI sample, which ranges between 7,597 and 9,157. US multinationals are defined as US companies that file a corporate tax return and that own at least one opaque foreign affiliate that files Form 5471. SOI includes a firm identifier (EIN) that can be used to link the corresponding firm from Compustat. The size of this linked sample, shown in the middle dashed line in the left panel of Figure 3.1, ranges between 1845 and 2334. The considerable attrition rate relative to the full SOI sample is not surprising given that Compustat does not include information about privately-held firms or smaller multinationals.

To ensure that the match is of high quality, I use information from Form 1120, Schedule M-3, which must be filed by large US multinationals to reconcile book and tax financial information. Firms must report book earnings on their M-3 that in theory should be the same as those disclosed in their public filings. These earnings should be directly comparable to the Net Income figure reported in Compustat, therefore I drop firms from the matched sample that report M-3 Net Income with a different sign from Compustat or firms that report Net Income that is not within 1% of the Compustat Net Income figure. The size of the final linked sample, which ranges from 1680 to 2023, is shown on the bottom line of the left panel of Figure 3.1. The figure demonstrates that only a small share of firms are dropped when using Schedule M-3 to match on net income, which is encouraging and indicative that the comparison is of high quality.

The right panel measures attrition by examining aggregate pretax income for each step of the sample construction. A much smaller proportion of foreign earnings are dropped relative to the attrition in sample size, confirming that Compustat does not provide coverage of

[5]This is not a criticism of that analysis. B&R only have access to highly aggregated data made public by SOI, and therefore the concern about differences in sample composition is unavoidable.

relatively small US multinationals.

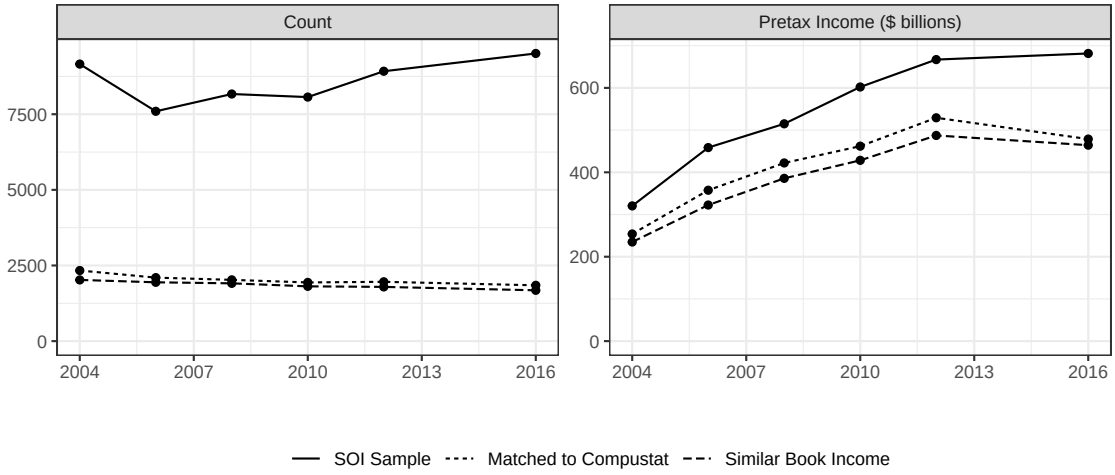

Figure 3.1: SOI MNCs vs Matched Sample

Notes: This figure demonstrates the attrition associated with the sample construction process. The "SOI Sample" contains all US MNCs from the SOI data sample. The "Matched to Compustat" sample contains SOI Sample MNCs matched to Compustat via unique firm identifier (EIN). The final linked sample, "Similar Book Income," was created by dropping Compustat-matched MNCs whose M-3 Net Income, as reported on IRS Form 1120, Schedule M-3, has a different sign than or is not within 1% of the firm's Compustat Net Income. The left panel displays sample sizes in terms of number of MNCs, while the right panel shows sample sizes in terms of pretax income. These panels showcase Compustat's lack of coverage for small and/or private MNCs, but underscore that the match quality of MNCs that are able to be linked is high.

Figure 3.2 provides an overview of attrition starting with the full Compustat sample. An additional filter is applied to remove Compustat observations that report missing pretax foreign income. As shown in the left panel, a much smaller portion of Compustat multinationals are dropped in the match, which is unsurprising given that SOI is the more comprehensive database. In 2016, for example, the matched simple includes 68.8% of Compustat multinationals. In the right panel, attrition is presented in terms of aggregate foreign earnings as reported in Compustat, with the matched sample including 84% of aggregate foreign earnings reported in Compustat in 2016.

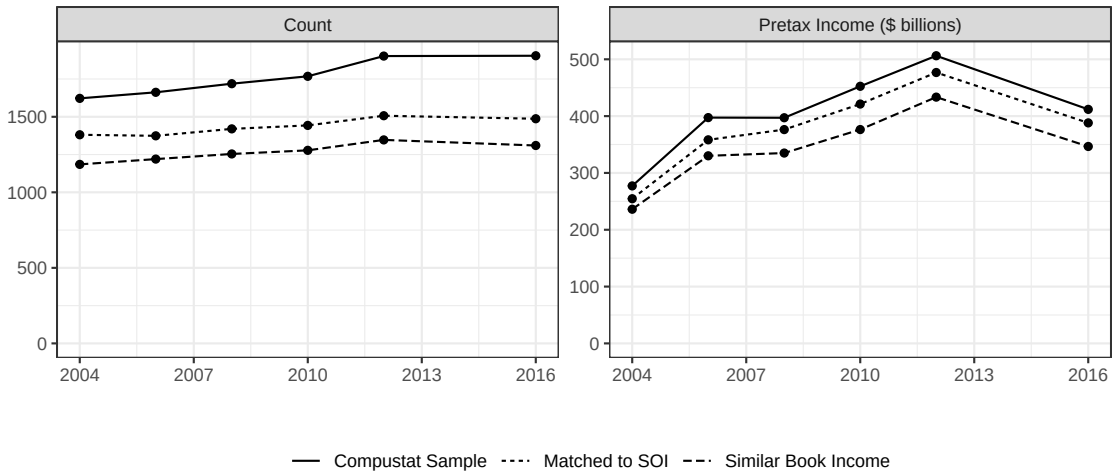

Figure 3.2: Compustat MNCs vs Matched Sample

Notes: This figure demonstrates the attrition associated with the sample construction process, following the same methodology as described in Figure 3.1 but starting with the Compustat sample rather than the SOI sample. Observations reporting missing pretax foreign income are removed from the initial Compustat sample. The left panel shows sample sizes in terms of number of MNCs, while the right panel shows sample sizes in terms of pretax income. These panels show that the majority of Compustat MNCs were able to be linked to an SOI MNC, and that the match quality of linked MNCs is high.

3.3.2 Comparison of Aggregate Foreign Income and Tax

Having constructed a sample of multinationals for which both public filings and tax filings are observed, I now examine aggregate foreign earnings and tax outcomes over time within the sample.

Figure 3.3 shows three aggregate measures of foreign income tax. The solid line plots aggregate foreign tax from Compustat. The dotted lines provide two different estimates from tax data — the series represented by short dashes are computed from Form 5471, Schedule C (Income Statement), and the series represented by the long dashes are computed from Form 5471, Schedule E (Income, War Profits, and Excess Profits Taxes Paid or Accrued). I use Schedule E as the preferred measure. The left panel contains aggregates computed from multinationals in all industries. The right panel excludes multinationals classified as operating in Financial, Utilities, Mining, Agriculture, or Oil & Gas.

Examining the full sample, it appears that Compustat tends to report much larger estimates of foreign income tax than what is contained in tax filings. This difference is markedly lower when excluding the industries selected above. One key reason for the discrepancy appears to be related to how extractive industries operate foreign concessions. These projects are often structured so that the foreign state receives a share of revenue or profits. Firms have discretion over whether to report these profit-sharing arrangements as income tax in their books. It does not appear that they have the same amount of discretion when they disclose information about foreign taxes in their returns.

Even after removing this set of industries, Compustat tends to report larger estimates of corporate income tax than what firms disclose in their tax returns. This suggests that aggregation error is likely not a significant problem when computing US multinationals' foreign tax bill from their 5471 filings.

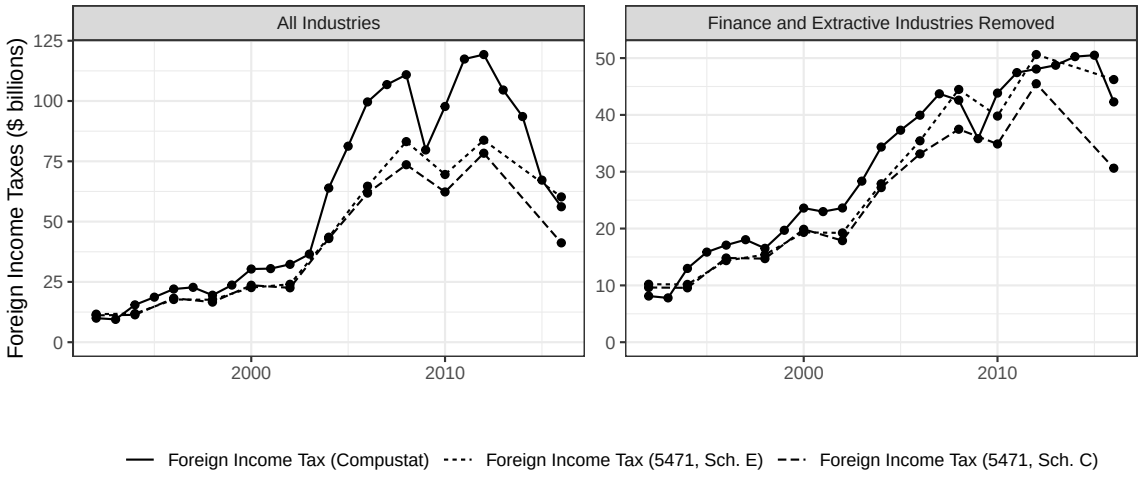

Figure 3.3: Comparison of Aggregate Foreign Tax Measures

Notes: This figure shows three measures of aggregate foreign income tax as respectively reported by Compustat; IRS Form 5471, Schedule C; and IRS Form 5471, Schedule E. The left panel contains aggregates computed from MNCs in all industries, while the right panel excludes MNCs operating in Financial, Utilities, Mining, Agriculture, or Oil & Gas. The panels show that Compustat tends to report larger estimates of foreign income tax than is contained in tax filings; however, this difference is lessened when excluding the industries mentioned above. We use Schedule E as the preferred measure.

Figure 3.4 shows three aggregate measures of foreign pretax earnings for the matched

sample. As before, the solid lines plot aggregate foreign pretax income from Compustat and the dotted lines provide two different estimates from tax data—the series represented by short dashes represents unadjusted earnings and profits (E&P) computed from Form 5471, Schedule H, and the series represented by the long dashes is adjusted to remove dividends received from related affiliates as reported on Schedule M. For the SOI data, I add back foreign taxes as reported on Schedule E so that both series represent pretax earnings. The left and right panels are inclusive and exclusive of the finance and extractive industries described above.

Notably, unadjusted E&P appear to diverge from the other measures over time. Compustat and adjusted SOI foreign income measures are more closely aligned, although SOI seems to report a slightly higher figure on average. This provides suggestive evidence that B&R's suggested correction for aggregation error is at least partially effective, removing the bulk of disagreement between data sources. Figure 3.5 provides a similar figure, but with aggregate post-tax earnings (instead of pretax), demonstrating a similar pattern.

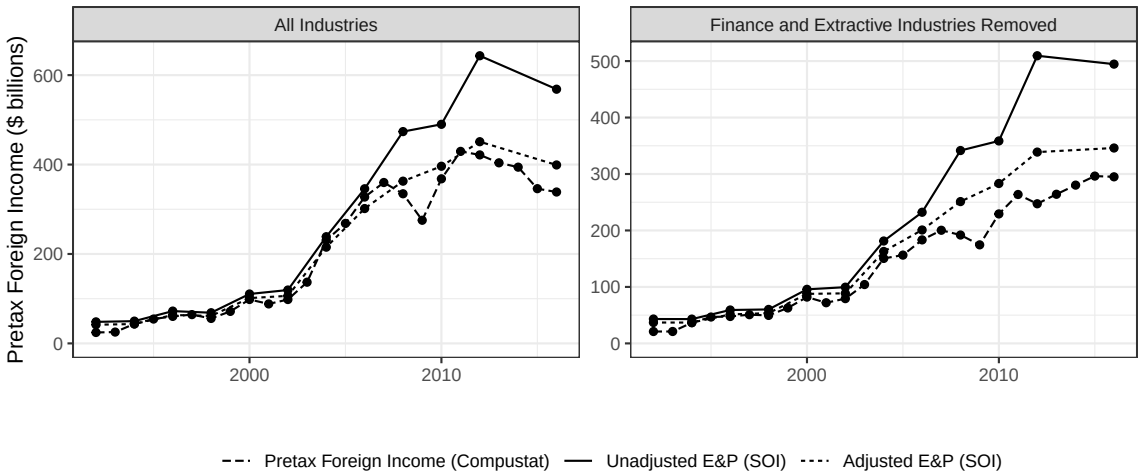

Figure 3.4: Comparison of Aggregate Foreign Income (Pre-Tax)

Notes: This figure shows three measures of aggregate foreign pretax income: aggregate foreign pretax income as reported by Compustat; unadjusted E&P computed from IRS Form 5471, Schedule H; and Schedule H E&P, adjusted to remove dividends received from related affiliates as reported on Schedule M. Both SOI measures include foreign taxes, as reported in Schedule E, added back to represent pretax earnings. As in Figure 3.3, the left panel contains aggregates computed from MNCs in all industries and the right panel excludes MNCs operating in Financial, Utilities, Mining, Agriculture, or Oil & Gas. The panels indicate that the E&P adjustment proposed by B&R is at least partially effective.

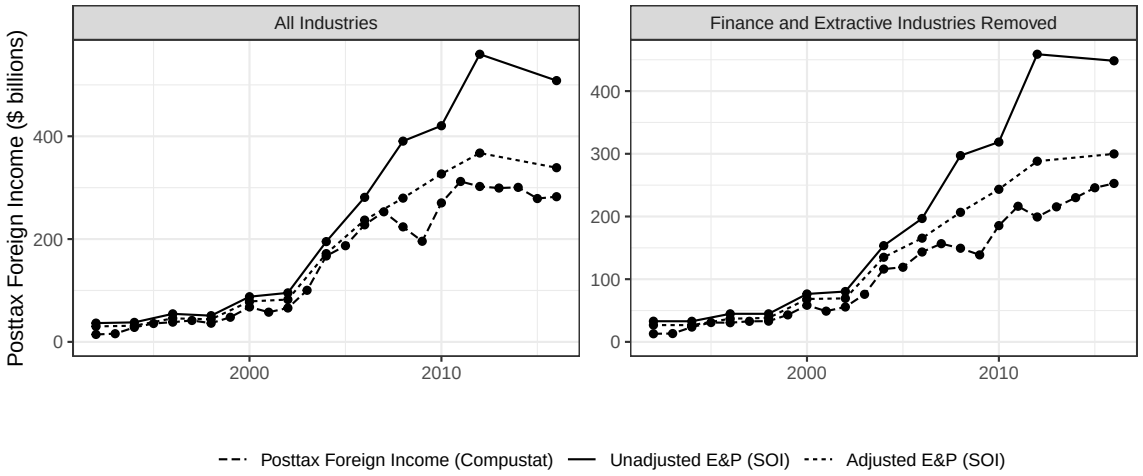

Figure 3.5: Comparison of Aggregate Foreign Income (Post-Tax)

Notes: This figure displays the same three aggregate foreign income measures as Figure 3.4 on a post-tax, rather than pretax, basis. It similarly underscores the effectiveness of B&R's adjustment for aggregation error in the SOI E&P data.

Finally, Figure 3.6 presents a comparison of average annual foreign effective tax rates from Compustat and SOI (one adjusted for aggregation error, and one unadjusted). Compared to the adjusted series, the average ETR computed from the unadjusted time series is downward-biased, as would be expected given that it overestimates the denominator. There is still considerable residual disagreement, however, when comparing adjusted ETRs from SOI with those computed from Compustat data. Compustat estimates are much higher throughout the sample period. This is a result of its higher tax estimates (Figure 3.3) and lower foreign earnings estimates (Figure 3.4).

Which source is more reliable? I do not have a full explanation for the remaining aggregate differences, but the Compustat series appear to be much noisier than aggregate series computed from tax data. Finance and extractive industries distort aggregate rates upward in both Compustat and SOI data, but the distortion is much more evident in Compustat data, generating what seem to be unrealistically high average rates in certain periods. This suggests that reporting of foreign taxes and income in public filings may be less consistent and comparable over time relative to tax data.

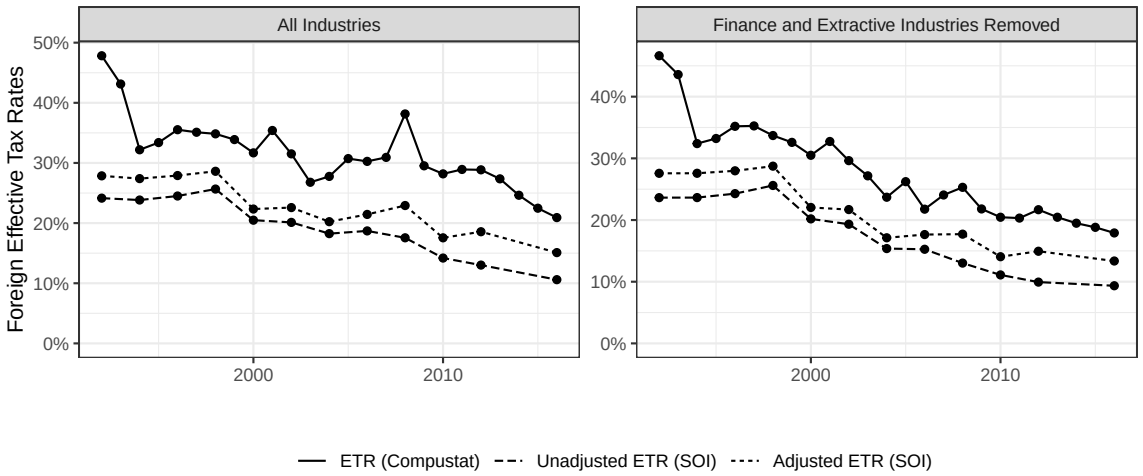

Figure 3.6: Comparison of Aggregate Foreign Effective Tax Rates

Notes: This figure presents a comparison of average foreign effective tax rates from Compustat, SOI (adjusted for aggregation error), and SOI (unadjusted). These tax rates are based on pretax income as reported in Figure 3.4 and foreign tax as reported in Figure 3.3. The left panel contains aggregates computed from MNCs in all industries and the right panel excludes MNCs operating in Financial, Utilities, Mining, Agriculture, or Oil & Gas. The panels demonstrate that lack of adjustment to the SOI data creates downward bias in the tax rate estimates. Alternately, they also show that comparatively higher taxes and lower foreign income in the Compustat data lead to larger Compustat ETR estimates than adjusted or unadjusted SOI.

3.3.3 Measurement Error and MNC Corporate Complexity

In this section, I show that aggregation error is closely related to the size of multinationals' foreign networks. More complex foreign networks create more potential for double-counting, as dividends may be reported on the books of a potentially large number of foreign affiliates. Once foreign earnings are adjusted to remove related dividends, however, there is no evidence that residual measurement error is systematically related to the complexity of a multinational's affiliate structure. Furthermore, I show that these adjustments result in residual measurement error that appears to be roughly constant over time. These results are encouraging for researchers, as they suggest that standard techniques to control for time-invariant heterogeneity and measurement error may be valid.

The aggregate trends shown above suggest that aggregation error primarily distorts estimates of foreign earning, not foreign taxes. Given that the primary mechanism behind aggregation error relates to double-counting of foreign earnings that are distributed up through a multinational's foreign ownership network, we would expect aggregation error to increase with the size of a multinational's foreign affiliate network. Figure 3.7 confirms that this is the case. The orange line and surrounding ribbon plot the estimates of a cublic spline regression. The dependent variable is the difference between log unadjusted foreign E&P as reported in SOI and log foreign income as reported in Compustat. The orange triangles are binscatter estimates produced for multinationals in 10 quantiles of their foreign affiliate network size, as measured by the number of opaque foreign entities (CFCs) in the tax data. Note that there is a clear upward trend, indicating that aggregation error is closely related to multinationals' foreign network size.

The black line and circles plot the same statistics computed for a dependent variable that uses adjusted SOI E&P. The spline is flat, suggesting that the residual differences compared with Compustat estimates are unrelated to the complexity of a firm's foreign network. If this adjustment were imperfect, and missed a portion of double-counted profits, I would expect this slope to perhaps be attenuated but not flat.

The black line does appear to be significantly positive, indicating that SOI foreign earnings are on average about 8% larger than the figures reported in Compustat. One reason for this may be that tax data include earnings by foreign affiliates that are associated with sales to US-based customers. These sales would be accounted for on the books of foreign affiliates, but might not be reported as foreign income on a firm's consolidated public filings.

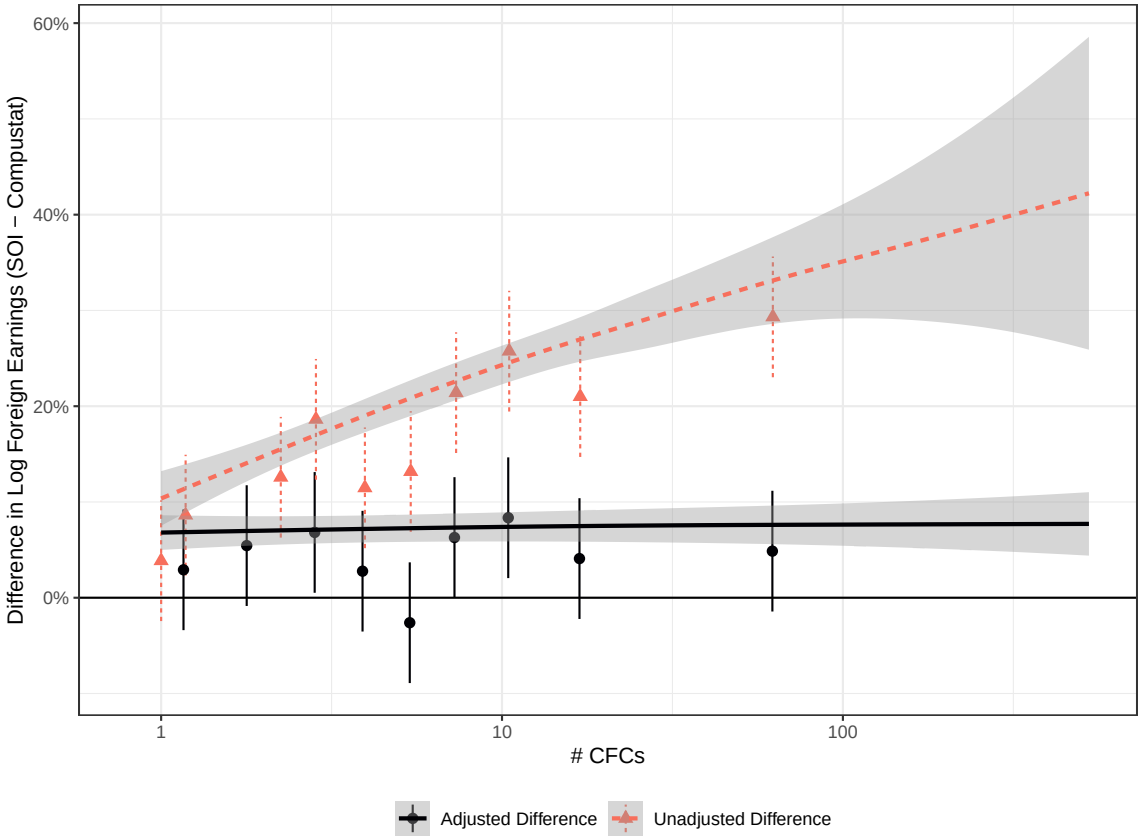

Figure 3.7: SOI/Compustat Foreign Earnings Differential

Notes: This figure presents the results of two cubic spline regressions where the dependent variable is the difference between log SOI foreign E&P and log Compustat foreign income. The orange line represents the estimates of a regression using unadjusted SOI E&P, while the black line uses adjusted SOI E&P. The orange triangles and black circles represent respective binscatter estimates produced for MNCs in 10 quantiles of foreign affiliate network size, as measured by the number of CFCs in the SOI data. The upward trend in the orange line and binscatter estimates indicates that aggregation error increases with the size of an MNC's foreign affiliate network. Alternately, the flat spline of the adjusted E&P regression indicates that residual differences between Compustat estimates are unrelated to a firm's foreign affiliate network size.

Figure 3.8 plots spline estimates for the same sample, but split into three distinct periods.

This confirms that aggregation error has been increasing over time, as reflected by the increasing slope of the orange splines. Adjusted E&P, however, remains flat and roughly constant at the same level. This again suggests that residual measurement error has not been increasing over time, and instead might be due to other reporting differences between tax and public filings.

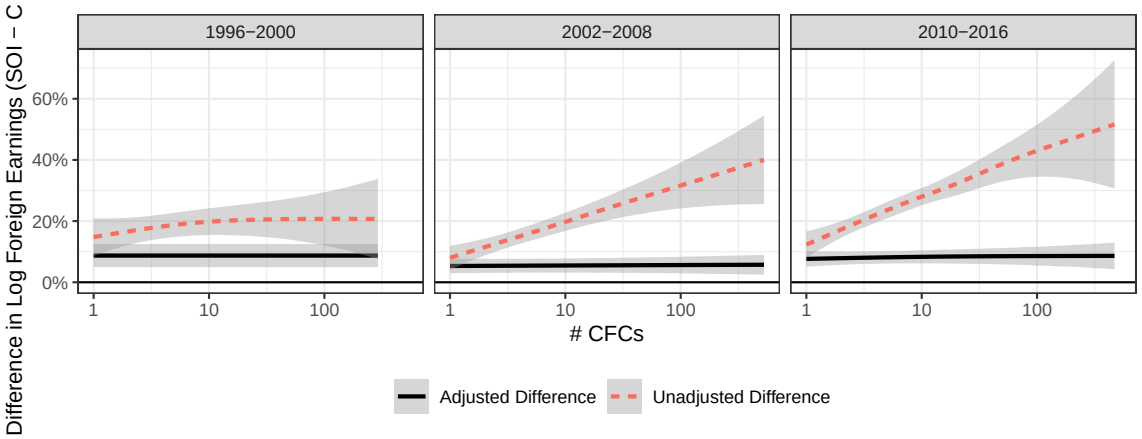

Figure 3.8: SOI/Compustat Foreign Earnings Differential by Period

Notes: This figure plots the same spline estimates as Figure 3.7 above, split into three time periods. The increase in the slope of the orange line over the three periods further confirms that aggregation error in the unadjusted SOI data has been increasing over time. However, the roughly constant black line indicates that residual measurement error has not been increasing.

Table 3.1 provides a regression formulation of the graphical analysis shown above. Again, the dependent variable is the difference between log foreign E&P as reported in SOI and log foreign income as reported in Compustat. Odd columns use unadjusted E&P and even columns introduce the double-counting adjustment. The independent variable is the log of the number of CFCs in a firm-year. For the unadjusted difference, the coefficients on this variable are positive, with large t-statistics. The adjusted coefficients, however, are close to zero with narrow confidence intervals. Columns (3), (4), (7), and (8) add year fixed effects, and columns (5) through (8) remove the industries excluded in the aggregate plots shown in the previous section.

Table 3.1: Measurement Error and MNC Foreign Affiliate Network Size

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(Intercept)	0.095***	0.082**			0.081***	0.079***		
	(0.027)	(0.027)			(0.016)	(0.016)		
Log # CFCs	0.043***	−0.001	0.043***	−0.012	0.047***	0.005	0.050***	0.000
	(0.009)	(0.009)	(0.009)	(0.009)	(0.005)	(0.005)	(0.006)	(0.005)
Num.Obs.	9993	9993	9993	9993	9040	9040	9040	9040
FE: year			X	X			X	X
Remove Outliers:	-	-	-	-	Yes	Yes	Yes	Yes
Adjusted Earnings	-	Yes	-	Yes	-	Yes	-	Yes

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table displays coefficients from regressions described above. The dependent variable is the difference between log SOI foreign E&P and log Compustat foreign income. Odd columns use unadjusted SOI E&P and even columns use adjusted SOI E&P. Columns (3), (4), (7), and (8) add year fixed effects, and Columns (5) through (8) remove MNCs operating in Financial, Utilities, Mining, Agriculture, or Oil & Gas.

3.4 Measurement Error and Tax Elasticities of Foreign Income

In the previous section, I show that introducing a simple correction to remove related dividends appears to remove systematic aggregation error in the measurement of foreign earnings in tax data. In this section, I recompute tax semi-elasticities of foreign earnings following the methodology in Dowd, Landefeld and Moore (2017, hereafter DLM). DLM use unadjusted earnings and a simple regression framework to examine how sensitive the allocation of multinational earnings are to average and statutory tax rates by country. First, I attempt to reproduce their results without any alterations using unadjusted earnings. DLM show that foreign earnings are more sensitive to tax rates in jurisdictions typically classified as tax havens. I construct a sample that is nearly identical to theirs, and obtain similar regression results with their aggregation process. I then modify their approach to utilize adjusted foreign earnings that remove aggregation error. Although this adjustment does result in noticeably different estimates, their general findings are robust to this measurement error.

Notably, my finding that DLM's results are robust to measurement error is different than the takeaway from B&R, who use BEA data to demonstrate that, once aggregation error is taken into account, multinationals do not appear to be more sensitive to rate changes in haven jurisdictions. One reason for this could be specification error—B&R run country-level regressions. I examine a similar specification that continue to show larger elasticities in haven jurisdictions. I discuss why this might be the case—the apparent contradiction can be resolved by understanding how foreign earnings are geographically classified in BEA vs. tax data.

3.4.1 Sample Construction

DLM provide a clear description of how they construct their sample. This allows me to nearly replicate their analysis sample with data from Form 5471. They start with a sample of opaque foreign affiliates (CFCs) with positive foreign earnings for multinationals that do

not operate in extractive industries, finance, insurance, or utilities. As shown in Section 3.3, this appears to be a reasonable selection criterion given that tax reporting can be quite different for these industries. DLM require that at least 10 CFCs operate in countries that are included in the analysis.[6] They also require that CFCs report positive amounts for wages and tangible assets.[7]

Table 3.2 contains the relevant sample statistics from DLM, Table 1. Compare these estimates with those contained in Table 3.3, which contains sample statistics for my reproduction. There are some small differences between the samples—mine is slightly smaller (90746 vs. 96959 in their sample), which may be due to differences in how they classify industries.[8] Overall, however, the summary statistics are very close.

Table 3.2: Unadjusted Sample Statistics from DLM, Table 1

	All	Nonhavens	Havens
Profits	14.41	9.01	57.19
Average Tax Rate	0.17	0.18	0.08
Statutory Tax Rate	0.29	0.30	0.18
Capital	27.37	24.40	50.91
Wages	7.75	7.18	12.32
GDP per capita	0.03	0.03	0.04
Population	150.24	168.54	5.14
2002	0.10	0.10	0.10
2004	0.16	0.16	0.15
2006	0.17	0.17	0.17
2008	0.18	0.18	0.18
2010	0.19	0.19	0.19
2012	0.20	0.20	0.20
Num obs.	96959	86099	10860

[6]This meets standard criteria for disclosure rules from the SOI.

[7]Assets are reported on the balance sheet disclosed on Form 5471, Schedule F. Tangible assets are calculated as the sum of end-year figures for buildings and other depreciable assets, depletable assets, and land (Sch. F, lines 8a, 9a, and 10 on the 12-2007 revision of the form). Although DLM indicate that depreciation is removed, I obtain similar summary statistics only when computing tangible assets before removing accumulated depreciation and depletion and proceed with the analysis accordingly. Wages are computed from the income statement disclosed in Schedule C (line 10 on the 12-2007 revision, "Compensation not deducted elsewhere").

[8]DLM do not provide the exact industry codes that they exclude in their analysis.

Table 3.3: Unadjusted Sample Statistics (reproduction)

	All	Nonhavens	Havens
Profits	15.25	9.42	76.31
Average Tax Rate	0.17	0.18	0.07
Statutory Tax Rate	0.29	0.30	0.18
Capital	27.15	23.25	68.05
Wages	8.08	7.37	15.49
GDP per capita	0.03	0.03	0.05
Population	154.50	168.82	4.60
2002	0.10	0.10	0.10
2004	0.16	0.16	0.15
2006	0.17	0.17	0.17
2008	0.18	0.18	0.19
2010	0.19	0.19	0.19
2012	0.20	0.20	0.20
Num obs.	90746	82831	7915

Next, I produce a similar sample using adjusted foreign earnings.[9] Table 3.4 contains sample statistics for this adjusted sampling procedure. Unsurprisingly, average profits are lower. The sample size is also somewhat smaller due to the fact that foreign earnings for a small part of the sample become negative once the adjustment is introduced.

[9]I make two adjustments to DLM's measurement of foreign earnings. They use pretax income disclosed in the income statement (Schedule C). I instead use foreign earnings and profits (E&P), disclosed on Schedule H, and remove related dividends disclosed on Schedule M. B&R suggest that E&P is reported in a more consistent manner than the earnings figure reported on Schedule C.

Table 3.4: Adjusted Sample

	All	Nonhavens	Havens
Profits	13.33	7.99	69.60
Average Tax Rate	0.20	0.21	0.08
Statutory Tax Rate	0.29	0.30	0.18
Capital	26.82	22.83	68.93
Wages	8.01	7.25	16.05
GDP per capita	0.03	0.03	0.05
Population	156.43	170.83	4.58
2002	0.10	0.11	0.10
2004	0.16	0.16	0.15
2006	0.17	0.17	0.17
2008	0.18	0.18	0.19
2010	0.19	0.19	0.19
2012	0.20	0.20	0.20
Num obs.	88801	81112	7689

3.4.2 Corrected Measurements of Tax Semi-Elasticities of Foreign Earnings

Adjusted Estimates from DLM

I use the samples described above to estimate the regression,

$$\log \pi_{ict} = \alpha + \beta_1 \log K_{ict} + \beta_2 \log L_{ict} + \beta_3(1 - \tau_{ct}) + \beta_4 X_{ct} + \mu_t + \psi_g + \epsilon_{it}, \qquad (3.1)$$

where π_{it} are profits, K_{it} are tangible assets, L_{it} is the wage bill, and τ_{ct} is either the average tax rate faced by CFCs incorporated in jurisdiction c or the statutory rate for country c in period t. X_{ct} is a vector of country-level controls that includes a second order polynomial of GDP per capita and population. μ_t is a year fixed effect and ψ_g is a multinational fixed effect. As in DLM, β_3 is the parameter of interest. This parameter captures how sensitive multinational profit allocations are to tax rates.

In practice, when using average tax rates, τ_{ct} is computed using a jackknife procedure to remove the effect of firm i on its own average rate. Figure 3.9 shows average tax rates for

the set of tax havens considered in this chapter.[10]

[10]I also include Netherlands in this figure given its importance as a domicile for the hybrid tax planning structures discussed in Chapter 2, although it is not classified as a tax haven for the DLM-style regressions for comparability.

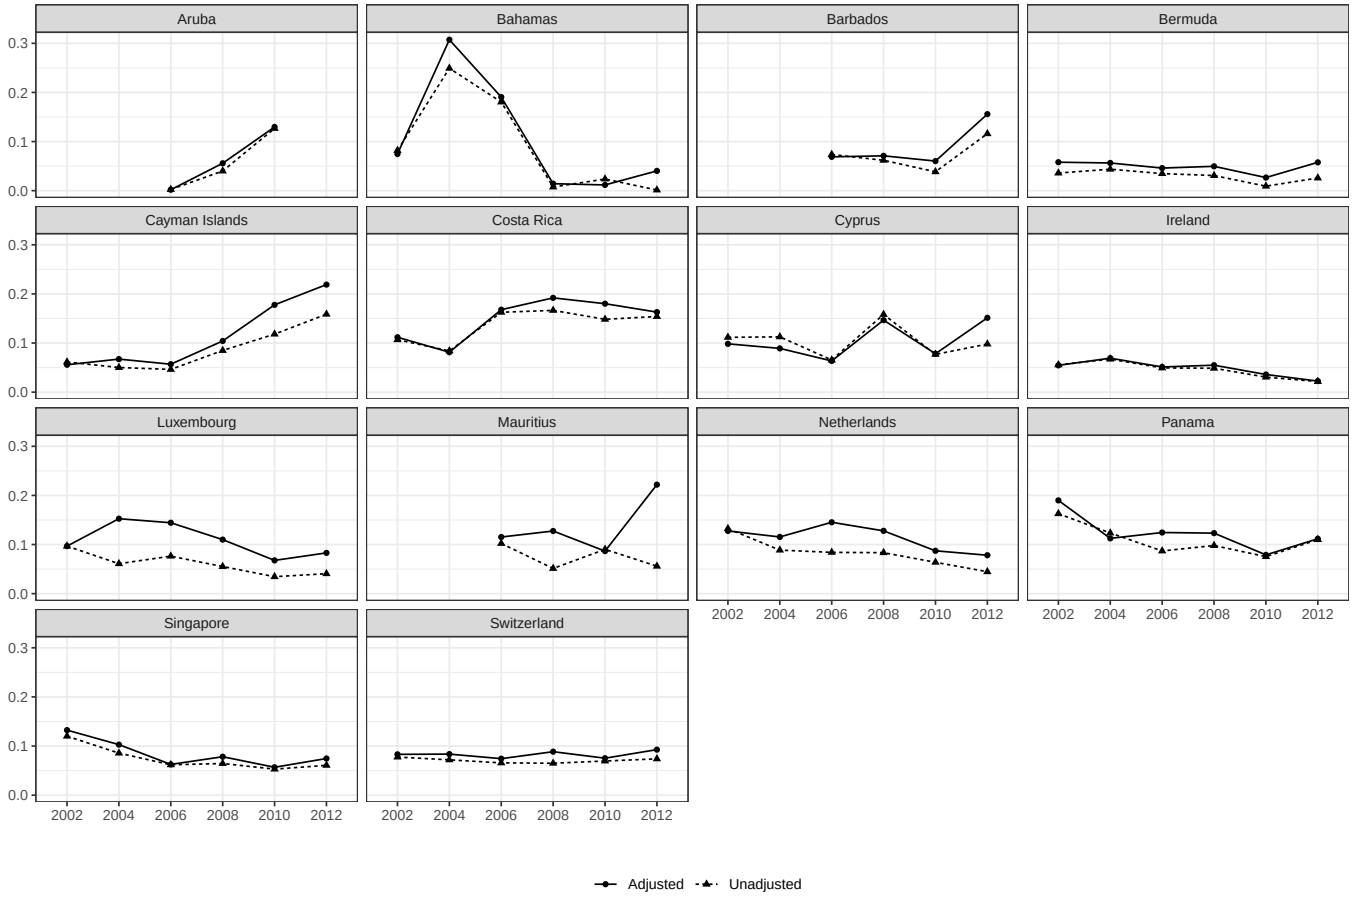

Figure 3.9: Adjusted and Unadjusted Average Tax Rates for Selected Haven Jurisdictions

Notes: This figure presents adjusted and unadjusted average tax rates over time for the set of tax haven countries we consider in our DLM-style regression analysis (Equation 3.1), as well as the Netherlands.

I also run two modified specifications that are also presented by DLM. The first allows the intercept and β_3 to vary according to whether country c is classified as a tax haven,

$$\log \pi_{ict} = \alpha_h + \beta_1 \log K_{ict} + \beta_2 \log L_{ict} + \beta_{3,h}(1 - \tau_{ct}) + \beta_4 X_{ct} + \mu_t + \psi_g + \epsilon_{it}. \quad (3.2)$$

Note that the only difference between equation Equation 3.1 and equation Equation 3.2 is the addition of h subscripts on α and β_3. The second allows for a second-order polynomial in the keep rate,

$$\log \pi_{ict} = \alpha + \beta_1 \log K_{ict} + \beta_2 \log L_{ict} + \beta_3(1 - \tau_{ct}) + \gamma_3(1 - \tau_{ct})^2 + \beta_4 X_{ct} + \mu_t + \psi_g + \epsilon_{it}. \quad (3.3)$$

Existing literature interprets β_3 as a semi-elasticity. Here, I do not emphasize a causal interpretation, which would require conditional exogeneity of τ_{ct}.[11] Instead, I focus on whether these parameter estimates are sensitive to measurement error.

Tables 3.5 and 3.6 show estimates for specification Equation 3.1 using unadjusted and adjusted earnings, respectively. Column (1) presents estimates using statutory rates and (2) presents estimates using average rates. The coefficient for the keep rate in the adjusted regression is slightly larger, but close in magnitude to the unadjusted estimates.

[11] One can imagine that countries who lower effective tax rates on foreign investment may also change other policy levers to attract capital, which would violate this assumption.

Table 3.5: DLM Table 2, Unadjusted

	(1)	(2)
1-Rate	1.293***	0.740***
	(0.150)	(0.086)
Log capital	0.321***	0.320***
	(0.009)	(0.009)
Log wages	0.394***	0.393***
	(0.012)	(0.012)
GDP per capita	4.866***	2.432+
	(1.329)	(1.313)
GDP per capita sq.	53.218***	72.183***
	(15.239)	(15.468)
Population	0.001***	0.001***
	(0.000)	(0.000)
Population sq.	0.000***	0.000**
	(0.000)	(0.000)
Num.Obs.	90 746	90 746
R2 Within	0.463	0.463
ETR	Statutory	Average

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.6: DLM Table 2, Adjusted

	(1)	(2)
1-Rate	1.308***	0.836***
	(0.155)	(0.102)
Log capital	0.317***	0.317***
	(0.009)	(0.009)
Log wages	0.391***	0.390***
	(0.013)	(0.013)
GDP per capita	4.159**	2.291+
	(1.384)	(1.388)
GDP per capita sq.	46.753**	61.697***
	(15.108)	(15.662)
Population	0.001***	0.001***
	(0.000)	(0.000)
Population sq.	0.000***	0.000**
	(0.000)	(0.000)
Num.Obs.	88 801	88 801
R2 Within	0.457	0.457
ETR	Statutory	Average

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Tables 3.7 and 3.8 present unadjusted and adjusted results for specification Equation 3.2. Columns (1) and (2) define haven countries in the same manner as DLM. These haven classifications follow Gravelle (2013), and include Aruba, the Bahamas, Barbados, Bermuda, the Cayman Isalnds, Costa Rica, Cyprus, Ireland, Luxembourg, Mauritius, Panama, Singapore, and Switzerland. Columns (3) and (4) define havens as countries in the smallest decile of the tax rate distribution by year. The coefficients on the interaction term are large and significant, which replicates the general findings of DLM. This suggests that foreign earnings are more sensitive to changes in tax rates within havens.

Table 3.7: DLM Table 4, Unadjusted

	(1)	(2)	(3)	(4)
Haven	−1.997***	−2.631**	−1.871***	−6.361***
	(0.410)	(0.971)	(0.400)	(1.849)
1-Rate	0.539**	0.578***	0.759***	0.597***
	(0.182)	(0.089)	(0.154)	(0.087)
Haven × 1 - Rate	2.657***	3.020**	2.582***	6.932***
	(0.519)	(1.054)	(0.517)	(1.961)
Log capital	0.319***	0.320***	0.319***	0.319***
	(0.009)	(0.009)	(0.009)	(0.009)
Log wages	0.395***	0.393***	0.395***	0.393***
	(0.012)	(0.012)	(0.012)	(0.012)
GDP per capita	3.566**	3.233*	4.185**	3.158*
	(1.318)	(1.318)	(1.318)	(1.319)
GDP per capita sq.	50.850***	48.653**	49.668***	56.189***
	(15.208)	(15.383)	(15.077)	(15.161)
Population	0.001***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Population sq.	0.000**	0.000***	0.000***	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)
Num.Obs.	90 746	90 746	90 746	90 746
R2 Within	0.464	0.463	0.464	0.464
ETR	Statutory	Average	Statutory	Average
Haven Def.	Gravelle (2013)	Gravelle (2013)	Lowest Decile	Lowest Decile

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.8: DLM Table 4, Adjusted

	(1)	(2)	(3)	(4)
Haven	−2.361***	−1.543*	−2.706***	−6.844***
	(0.397)	(0.728)	(0.416)	(1.801)
1-Rate	0.418*	0.605***	0.735***	0.644***
	(0.182)	(0.112)	(0.155)	(0.104)
Haven × 1 - Rate	3.144***	1.872*	3.518***	7.499***
	(0.502)	(0.804)	(0.523)	(1.933)
Log capital	0.316***	0.316***	0.316***	0.316***
	(0.009)	(0.009)	(0.009)	(0.009)
Log wages	0.392***	0.391***	0.392***	0.391***
	(0.013)	(0.013)	(0.013)	(0.013)
GDP per capita	2.626+	2.559+	3.700**	2.124
	(1.361)	(1.382)	(1.377)	(1.387)
GDP per capita sq.	44.019**	47.561**	39.629**	60.933***
	(14.896)	(15.555)	(14.897)	(15.460)
Population	0.001**	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Population sq.	0.000**	0.000***	0.000**	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)
Num.Obs.	88 801	88 801	88 801	88 801
R2 Within	0.458	0.457	0.458	0.457
ETR	Statutory	Average	Statutory	Average
Haven Def.	Gravelle (2013)	Gravelle (2013)	Lowest Decile	Lowest Decile

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Finally, Tables 3.9 and 3.10 present unadjusted and adjusted results for specification Equation 3.3. As in DLM, the second-order keep rate coefficient is significant, and actually increases in magnitude for the adjusted results.

Table 3.9: DLM Table 5, Unadjusted

	(1)	(2)
1-Rate	−8.847***	−0.542+
	(1.944)	(0.291)
1-Rate sq.	6.764***	0.885***
	(1.320)	(0.211)
Log capital	0.320***	0.320***
	(0.009)	(0.009)
Log wages	0.394***	0.393***
	(0.012)	(0.012)
GDP per capita	4.512***	2.514+
	(1.336)	(1.311)
GDP per capita sq.	47.706**	67.960***
	(15.173)	(15.392)
Population	0.001***	0.001***
	(0.000)	(0.000)
Population sq.	0.000**	0.000**
	(0.000)	(0.000)
Num.Obs.	90 746	90 746
R2 Within	0.464	0.463
ETR	Statutory	Average

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table 3.10: DLM Table 5, Adjusted

	(1)	(2)
1-Rate	−9.963***	−3.666**
	(1.943)	(1.150)
1-Rate sq.	7.521***	2.916***
	(1.316)	(0.751)
Log capital	0.317***	0.317***
	(0.009)	(0.009)
Log wages	0.391***	0.391***
	(0.013)	(0.013)
GDP per capita	3.755**	2.365+
	(1.388)	(1.389)
GDP per capita sq.	40.917**	57.132***
	(14.880)	(15.498)
Population	0.001***	0.001***
	(0.000)	(0.000)
Population sq.	0.000**	0.000**
	(0.000)	(0.000)
Num.Obs.	88 801	88 801
R2 Within	0.458	0.457
ETR	Statutory	Average

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Country-Level Regressions

The results shown above present a puzzle. B&R also estimate DLM-style regressions using corrected data from the BEA, and show that there is no evidence for heterogeneity once estimates of net income are corrected to adjust for aggregation error. One reason for this may be differences in specification. B&R estimate country-level regressions of the form,

$$\log \pi_{ct} = \alpha + \beta_1 \log K_{ct} + \beta_2 \log L_{ct} + \beta_3(1 - \tau_{ct}) + \beta_4 X_{ct} + \mu_t + \gamma_c + \epsilon_{ct}, \qquad (3.4)$$

where the dependent variable is aggregate foreign earnings for foreign affiliate located in country c in year t. They present result inclusive and exclusive of country fixed effects γ_c.

I aggregate the sample described in the previous section to the country level, and estimate similar regressions, presented in Tables 3.11 and 3.12 below. Even columns in Table 3.11 include the haven interaction term analagous to the specification Equation 3.2, and the same columns in Table 3.12 include second-order polynomials in the tax rate analagous to specification Equation 3.3. B&R present results with respect to the tax rate (as opposed to the keep rate in prior regressions). I do the same to faciliate comparison with their results.

In both tables, columns (1) through (4) present estimates using unadjusted foreign earnings, and columns (5) through (8) present estimates using adjusted foreign earnings. As in B&R, estimates are computed with and without country fixed effects. Table 3.11 demonstrates that, unlike B&R, the coefficient on the haven interaction remains large and significant, although slightly attenuated. Results in Table 3.12 are mixed—the adjusted estimates appear to weaken the non-linear term.

Table 3.11: Blouin & Robinson Table 2, Panel A

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ETR	−2.833**	−1.173+	−2.706*	−2.644**	−2.313***	−1.581*	−2.445*	−2.414*
	(0.480)	(0.575)	(0.678)	(0.631)	(0.286)	(0.408)	(0.641)	(0.675)
Haven × ETR		−12.099*		−11.802+		−9.054**		−8.290*
		(3.020)		(4.872)		(1.948)		(2.420)
Log capital	0.480***	0.437**	0.370*	0.341*	0.426***	0.390***	0.160	0.123
	(0.067)	(0.084)	(0.128)	(0.129)	(0.044)	(0.043)	(0.171)	(0.165)
Log wages	0.549**	0.554**	0.358	0.294	0.604***	0.606***	0.358	0.360
	(0.098)	(0.110)	(0.258)	(0.268)	(0.056)	(0.068)	(0.234)	(0.248)
GDP per capita	−0.177+	−0.167+	−0.009	0.010	−0.158+	−0.135+	−0.179	−0.128
	(0.085)	(0.070)	(0.069)	(0.085)	(0.068)	(0.053)	(0.130)	(0.104)
GDP per capita sq.	0.024*	0.014	0.008	0.003	0.025**	0.012*	0.021*	0.012+
	(0.009)	(0.008)	(0.006)	(0.007)	(0.006)	(0.004)	(0.008)	(0.005)
Population	0.137	0.105	−0.277	0.227	0.169	0.199	−3.250	−2.974
	(0.104)	(0.071)	(2.183)	(1.919)	(0.153)	(0.138)	(1.623)	(1.546)
Population sq.	−0.012	−0.008	0.016	0.001	−0.014	−0.016	0.131	0.122
	(0.007)	(0.005)	(0.100)	(0.091)	(0.011)	(0.010)	(0.070)	(0.068)
Num. Obs.	269	269	269	269	269	269	269	269
Year FEs	X	X	X	X	X	X	X	X
Country FEs			X	X			X	X
Adjusted Earnings	-	-	-	-	Yes	Yes	Yes	Yes

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table 3.12: Blouin & Robinson Table 2, Panel B

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ETR	−2.833**	−12.001**	−2.706*	−9.359+	−2.313***	−3.832*	−2.445*	−3.200+
	(0.480)	(2.781)	(0.678)	(4.519)	(0.286)	(1.360)	(0.641)	(1.470)
ETR Sq.		18.742*		12.321		1.760		0.759
		(4.817)		(7.329)		(1.689)		(1.159)
Log capital	0.480***	0.508***	0.370*	0.392*	0.426***	0.442***	0.160	0.153
	(0.067)	(0.071)	(0.128)	(0.136)	(0.044)	(0.041)	(0.171)	(0.175)
Log wages	0.549**	0.522**	0.358	0.278	0.604***	0.594***	0.358	0.368
	(0.098)	(0.096)	(0.258)	(0.271)	(0.056)	(0.054)	(0.234)	(0.245)
GDP per capita	−0.177+	−0.164+	−0.009	−0.084	−0.158+	−0.153+	−0.179	−0.188
	(0.085)	(0.079)	(0.069)	(0.069)	(0.068)	(0.067)	(0.130)	(0.119)
GDP per capita sq.	0.024*	0.019	0.008	0.009	0.025**	0.023*	0.021*	0.020*
	(0.009)	(0.010)	(0.006)	(0.005)	(0.006)	(0.006)	(0.008)	(0.007)
Population	0.137	0.079	−0.277	−0.278	0.169	0.194	−3.250	−3.442
	(0.104)	(0.094)	(2.183)	(1.913)	(0.153)	(0.151)	(1.623)	(1.767)
Population sq.	−0.012	−0.007	0.016	0.022	−0.014	−0.016	0.131	0.141
	(0.007)	(0.007)	(0.100)	(0.090)	(0.011)	(0.011)	(0.070)	(0.076)
Num. Obs.	269	269	269	269	269	269	269	269
Year FEs	X	X	X	X	X	X	X	X
Country FEs			X	X			X	X
Adjusted Earnings	-	-	-	-	Yes	Yes	Yes	Yes

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

3.5 Conclusion

Measurement error should be a primary concern for researchers who use accounting and tax data to study the activity of multinational firms. IRS tax data provide a wealth of information for tax researchers, but may generate significant measurement error if foreign earnings are improperly aggregated. I show that a simple proposed correction for this aggregation error performs well using a proxy of error based on a comparison with consolidated measures from public filings. I reproduce estimates from prior literature that examine the link between MNCs' profit allocations, tax rates, and heterogeneity in this relationship across haven and non-haven jurisdictions, generally showing that its conclusions are robust to this measurement error. These estimates may be different than those computed from BEA data given conceptual differences in how they define the geographical boundary of a foreign subsidiary — estimates from tax data that are based on average tax rates should be interpreted with the understanding that foreign affiliates may "hide" income that is generated by activity in other domiciles. Which of these estimates are "correct" depends on the estimand of the researcher.

4

Regulatory Uncertainty and Tax Avoidance: Evidence from Cost Sharing and Labor Costs of US Multinationals

4.1 Introduction

Tax experts have long recognized the potential for abuse of transfer pricing regulations to avoid taxation. Multinational firms have incentives to shift costs to high-tax jurisdictions and profits to low-tax jurisdictions to lower their overall tax bill. Transfers of intangible assets, in particular, have posed a difficult problem for regulators and tax authorities because their market value, particularly for unique assets, can be difficult to determine.

One way that the IRS regulates transfers of these intangible assets is through what are known as "Cost Sharing Agreements" (CSAs). CSAs allow a foreign affiliate to pay for a share of R&D costs of the US parent. Because these are foreign costs, they are not tax deductible and do not generate a tax shield from the US corporate tax. In return, however, the foreign affiliate can merchandise the IP to foreign markets, and (prior to 2017) could benefit from a US system that allows for deferral of tax on foreign income.

These firms face clear incentives to limit costs that are incurred by foreign affiliates while maintaining the benefits of tax deferral. They can do this either by (illegally) understating the true cost of their IP or by exploiting legal loopholes in the regulatory system. One such legal loophole emerged after new regulations governing CSAs were passed in the mid-1990s. The loophole allowed firms that had R&D labor costs to pay their employees with stock options instead of wages. A Tax Court ruling in 2005 determined that, under 1995 regulations, such costs did not have to be covered by a CSA. Although the IRS had attempted to close

this loophole with regulatory amendments passed in 2003, the decision was seen as a victory for certain MNCs, and the 2003 regulations continue to be litigated to the present.[1]

In this chapter, I examine the reaction of CSA-holding firms, as measured in IRS tax data, to the 2005 Tax Court ruling. In Section 4.4, I show that in the days following the ruling, the market value of these firms increased considerably. I then show that these exposed MNCs greatly increased their usage of stock compensation in the years following the court's decision, indicating that they were taking advantage of the loophole. These firms also expanded their R&D activity relative to other MNCs, and increased the share of total R&D for which they claimed an available tax credit from the United States. This suggests that these firms may have relabeled foreign R&D activity as dometic R&D activity for purposes of claiming the credit.

The chapter proceeds as follows. Section 4.2 discusses the institutional details behind the loophole, and the legal battles that emerged in the 2000s. I compile anecdotal evidence from experts and tax practitioners to argue that a period of regulatory uncertainty emerged after 2005 that strongly affected firm behavior. This evidence also suggests that there were large revenue concerns from the US Treasury related to this behavior. Section 4.3 describes the data we use to identify exposed firms and measure their responses. Section 4.4 shows that a legal victory for exposed firms increased their market value, and examines longer-term responses related to the underlying mechanism and other behavioral tax responses. Section 4.5 concludes.

4.2 Institutional Background

In 1986, as a response to concerns around the growing complexity and frequency of the transfer of intangible assets, combined with the potential for erosion of the corporate income tax base, Congress included the following language in Sec. 482 of the Tax Reform Act (TRA):

[1] In 2020, the Supreme Court upheld a favorable ruling for the IRS in the ninth circuit.

> In the case of any transfer (or license) of intangible property… the income with respect to such transfer or license shall be commensurate with the income attributable to the intangible. (TRA, 1986)

In other words, for assets that prove difficult to value, Congress suggested that transfer prices should be based on the income that those intangibles might generate as an alternative to finding comparable transactions that might not exist.[2] This law led to new regulations promulgated by the US Treasury in 1994 and 1995 that governed the tax treatment of transfers of intangible assets.

4.2.1 Cost Sharing Agreements

Consider a US multinational corporation that performs research and development to develop an intangible asset that it expects will be used to merchandise a product or service to be sold both domestically and abroad. Prior to 2017, the United States operated under a worldwide tax system, meaning that all of the firm's earnings would constitute taxable income under US law. The MNC typically has two options to develop its IP from a tax standpoint.

One option is for the associated R&D to be incurred by the US parent company. The IP that results from this R&D can then be licensed to a foreign affiliate, which will sell products or services to customers abroad. The foreign affiliate will then pay a license fee or royalty back to the US parent company. This income would be then be taxed by the United States. Notably, such agreements prevent the possibility for income deferral — the foreign affiliate must pay the royalty to the US parent company, where it would be immediately subject to US tax.

Another option that the firm can undertake prior to developing the IP is to arrange a cost-sharing agreement (CSA) with its foreign affiliate. Under a CSA, the foreign affiliate pays for a share of the R&D expense incurred by the multinational, with the domestic parent incurring the remainder. Under the 1995 regulations described above, these costs

[2]Note that this is not a departure from the arms-length standard insofar as the valuation mimics a similar exercise that might be undertaken by unrelated parties.

are typically split based on the income that the investment is expected to produce for each party. Under a CSA, the foreign affiliate will hold the rights to merchandise the IP abroad, and will not have to pay a royalty back to the US parent. In comparison to a royalty agreement, a CSA with a tax haven foreign affiliate will result in a short-term increase in the tax bill for the multinational — unlike a royalty agreement, R&D costs allocated to the foreign affiliate are not deductible for tax purposes in the United States, and therefore do not generate a tax shield. Once the IP is developed, however, the foreign affiliate can store any income it might generate from foreign sales abroad and defer US tax indefinitely.

4.2.2 CSAs, Xilinx, and the Allocation of Stock Compensation

The use of employee stock-options (ESOs) was uncommon in the 1980s and early 1990s. As a result, the 1995 regulations governing the transfer of intangible assets and the use of cost-sharing agreements did not specify how ESOs should be treated from a tax perspective. During this period, there was a poor understanding of whether they constituted real expenses and firms often did not report ESOs on their books.[3]

On August 30, 2005, the US Tax Court issued a judgment in favor of a company, Xilinx Inc., in its petition against the IRS. Xilinx had used ESOs to compensate employees that performed R&D subject to a cost-sharing agreement. The IRS audited the company and found that it had excluded ESO-related costs from its CSA with an Irish affiliate. Furthermore, Xilinx had claimed expense deductions as well as R&D credits related to the ESO expense.[4] If the ESOs were considered costs under the CSA, they would have been partially allocated to its foreign affiliate and Xilinx would not have been able to deduct this share or use it to claim an R&D tax credit; therefore, its US tax bill would have increased. Xilinx filed a petition in the Tax Court to challenge the results of the audit. In its 2005 decision, the Tax Court sided with Xilinx, holding that "unrelated parties" would not have agreed

[3]In fact, they are—ESOs dilute the ownership of existing shareholders. See Bodie, Kaplan and Merton (2003).
[4]See IRS (2010).

to share ESO costs.

After 2005, the Xilinx ruling continued to be litigated through an appeals process. The decision was initially reversed in May 2009 before a new opinion was issued in March 2010, once more siding with Xilinx.

4.2.3 2003 Amendments and Altera Corp. v. Commisioner

The 2005 Xilinx decision incentivized other MNCs to inflate domestic costs, despite the fact that the Treasury had attempted to address the regulatory ambiguity by clarifying the 1995 regulations in a set of amendments issued in 2003. These amendments explicitly classified the use of ESOs as costs that firms would have to make subject to CSAs. The Xilinx case was litigated under the 1995 regulations. There was uncertainty, however, over whether the 2003 amendments could withstand legal challenges. For example, at least one US multinational, Altera Corp., restructured its CSA immediately after the Xilinx ruling to omit ESO costs, and became involved in another dispute with the IRS. Prior to the Xilinx ruling, Altera had shared its ESO costs with its foreign affiliate under the CSA. This separate dispute was not resolved until June 2020, when the Supreme Court denied an appeal following a ruling that sided with the tax authorities.[5]

Between the 2005 Xilinx decision and the 2020 Supreme Court decision in Altera, US MNCs that used CSAs had clear incentives to attempt to restructure their wage bill. If firms were indifferent between paying their employees in cash and paying them with ESOs, this regulatory uncertainty would have shifted the balance in favor of stock compensation.

Concurrent legal analysis suggests that the Xilinx case had significant implications for tax authorities, and was closely monitored by other firms and practitioners. For example, after the May 2009 Tax Court reversal, a tax law professor at the University of Michigan wrote:

> On May 7 the Ninth Circuit decided Xilinx v. Commissioner. By a 2-1 majority,

[5]See Altera v. Commissioner (2019). See also Foley and Kuvadia (2021).

> the panel reversed the Tax Court and held that costs of employee stock options must be included in the pool of costs subject to a tax-sharing agreement. The Xilinx decision is important for three reasons. First, **cost sharing is probably the key element in current transfer pricing law** because it is the principal way in which profits from intangibles get shifted from the United States to low-tax jurisdictions. Moreover, informed observers agree that the allocation of income from intangibles is the most important problem in transfer pricing, and because **most intangible-intensive corporations rely heavily on employee stock options, the narrow issue decided in Xilinx has large revenue implications, especially for high-tech companies.** This is evidenced by the filing of two amicus briefs on behalf of **coalitions of high-tech companies siding with the taxpayer and by practitioners' reactions to the IRS victory.** (Avi-Yonah, 2009, emphasis added)

Finally, the Altera ruling also appears to have significant implications for the other US multinationals. Immediately after the Supreme Court's decision in June 2020, one tax law firm issued an announcement to clients:

> While Altera was pending, many taxpayers with related tax disputes agreed to be bound by the outcome of Altera, and many cost sharing arrangements were amended to include language requiring or not requiring cost sharing of stock-based compensation based on the outcome of Altera.[6]

While the Supreme Court's 2020 decision has resolved much of the regulatory uncertainty, the issue has not been completely put to rest: firms outside of the Ninth Circuit are not necessarily bound by the decision. Regardless, between 2005 and 2020, anecdotal evidence suggests that a large number of firms structured their foreign operations to take advantage of the ESO technicality.

[6] McWaters (2020).

4.3 Data

I combine information from several different datasets to examine multinationals' response to the period of regulatory uncertainty following the 2005 Xilinx decision. I use IRS tax data at the foreign affiliate level to identify multinationals that were using cost sharing agreements in the mid-2000s. IRS tax data also provide information about labor costs and claims for the domestic R&D tax credit. I merge this sample of multinationals with information on stock prices from the Center for Research in Security Prices (CRSP). I also obtain historical data from Compustat that contains information about stock compensation and global R&D expenses.

Table 4.1 provide summary statistics for the merged sample. I identify 168 multinationals with CSAs and 1,166 that do not have any such agreement. CSA users spend significantly more on R&D and stock compensation than non-adopters.

Table 4.1: Summary Statistics

	Cost Sharers		Other MNCs	
	Mean	SD	Mean	SD
Compustat R&D	488	1564	96	564
Stock Comp.	123	398	16	48
Num. Obs.	2404		12891	
Num. MNCs	168		1166	

Notes: This table contains summary statistics for MNCs with and without cost sharing agreements (CSAs). The sample was created by merging IRS tax data, which identifies MNCs using CSAs with their foreign affiliates in the mid-2000s, with stock price data from the Center for Research in Security Prices (CRSP) and stock compensation and global R&D expense data from Compustat.

4.4 Analysis of Multinational Response

Below, I provide two analyses that show how multinationals were affected by the positive shock from regulatory uncertainty created by the Xilinx decision. First, I show that over the short term, firm values for exposed multinationals increased following the positive news

(from the firms' perspective) that the Tax Court might prevent the IRS from requiring the inclusion of ESOs in CSAs. This suggests that the market assigned a significant value to the positive regulatory shock generated by the Tax Court decision. Next, I show that over the ensuing years, exposed firms increased their usage of stock-compensation and R&D activity compared to other multinationals. I also show that their claimed deductions under the domestic R&D tax credit increased relative to the consolidated R&D cost on their books.

4.4.1 Tax Court Ruling Effect on Firm Values

To determine whether the August 2005 Tax Court ruling positively affected the market value of CSA holders, I use a standard stock market event study methodology (MacKinlay, 1997), estimating the following regression of stock returns for a set of multinationals c with cost-sharing agreements at date t,

$$R_{ct} = X_t\theta_c + \gamma E_{ct}(k) + \epsilon_{ct}.$$

where R_{ct} are log raw daily returns for each exposed multinational. I examine two choices for controls X_t—one includes an intercept that captures average daily returns for the sample, the second adds a control for market returns, for which I use daily log returns from the S&P 500. I use the approach taken in Dube, Kaplan and Naidu (2011) to estimate the cumulative abnormal return between the event and k days after the event in a single step. This can be done by setting E_{ct} to be a dummy variable equal to one during the event window. γ captures average daily returns over the event window, and the cumulative abnormal return can then be computed as $\beta = \gamma k$.[7] The regression is estimated using a pre-period of 115 days prior to the event, and k days after the event. I also estimate placebo effects by computing cumulative abnormal returns for the 5 days prior to event. Firms are weighted by size.[8]

[7] In practice, I set E_{ct} equal to $1/k$ to recover β directly.
[8] I use total assets as reported in Compustat as a proxy for multinational size.

This specification also allows for E_{ct} to be a continuous measure of treatment exposure. Firms that generate a larger share of their aggregate profits from foreign sales would naturally be in a better position to take advantage of the tax planning opportunity created by the CSA regulations. Thus, I use firms' foreign share of profits as of 2004, computed with tax data, to serve as a continuous exposure measure and compare it to a binary treatment specification.

Figure 4.1 plots estimates produced from the raw returns model for the 10 trading days following the Tax Court ruling. The first trading day following the Tax Court ruling is August 31, and the first news article covering the ruling appears in the Factiva news database on September 2. The vertical lines encompass estimates of cumulative abnormal returns for the sample during this period. Placebo estimates are presented for the days preceding this event window, and longer-term cumulative abnormal returns are presented for subsequent trading days. As in Dube, Kaplan and Naidu (2011), I cluster standard errors both by date and by firm, taking the maximum across both specifications for each date to generate conservative confidence intervals.

Both the binary and continuous measures show considerable growth in firm value after the Tax Court ruling. This growth persisted over the following 10 days, with cumulative abnormal returns reaching close to 5% with a binary measure and 8% with a continuous measure.

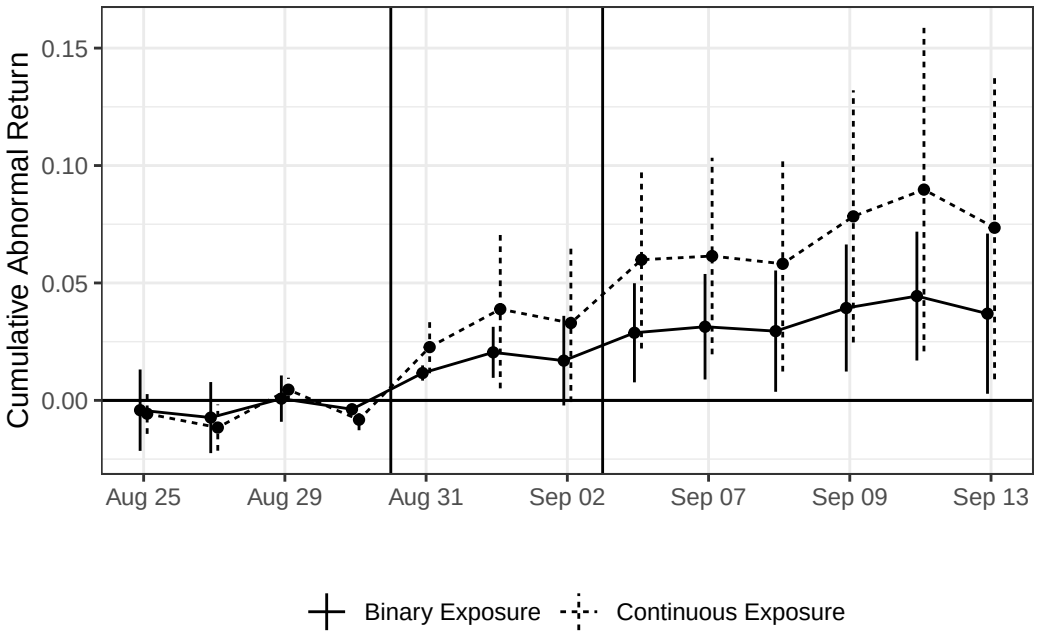

Figure 4.1: Firm Response to Tax Court Decision on Aug. 30, 2005

Notes: This figure presents cumulative abnormal return estimates, produced from raw returns, for the 10 trading days following the August 30, 2005 Tax Court ruling, from event study regressions. The vertical lines encompass estimates of cumulative abnormal returns for the sample during this period. Standard errors are clustered by both date and firm. Both exposures indicate sizeable firm value growth in the 10 trading days following the ruling.

Figure 4.2 shows estimates produced from the model that includes market returns. Compared with the raw returns model, these estimates have less power and show smaller effects (1.5% for the binary exposure measure and 3% for the continuous exposure measure). Placebo estimates, however, are more tightly centered around zero.

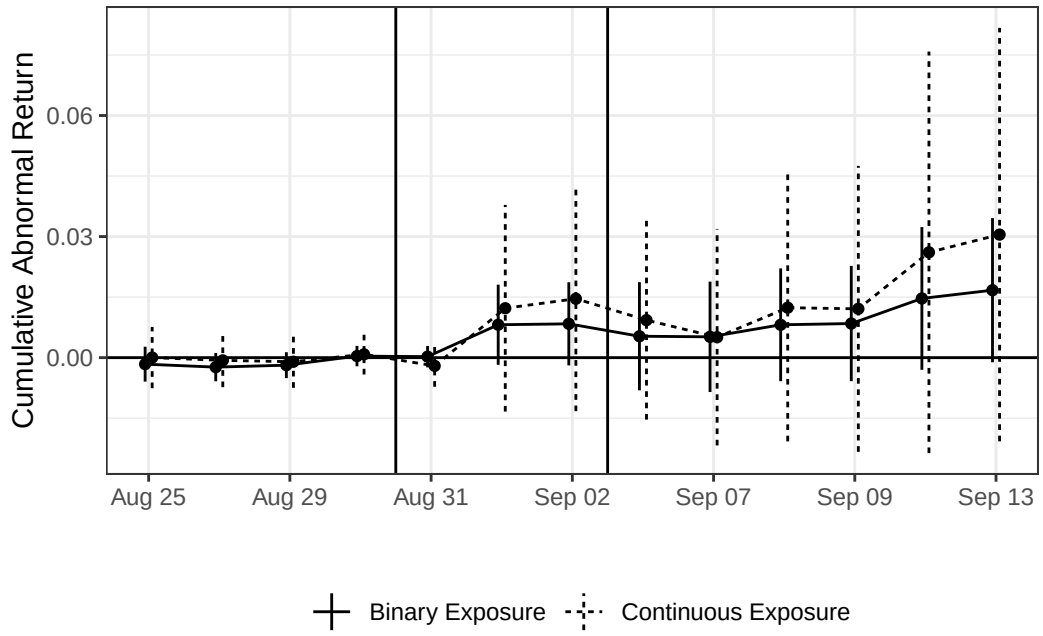

Figure 4.2: Firm Response to Tax Court Decision on Aug. 30, 2005 (Robustness)

Notes: This figure presents cumulative abnormal return estimates, produced from market returns, from event study regressions. The vertical lines encompass estimates of cumulative abnormal returns for the sample during this period. Standard errors are clustered by both date and firm. These exposures show a less significant effect than indicated by Figure 4.1.

4.4.2 Effects on Stock Compensation and Evidence for R&D Relabeling

Above, I demonstrate that multinationals who used cost-sharing agreements also received a positive shock from the Tax Court ruling, suggesting that the market believed these firms might benefit from its regulatory implications. Below, I show that these multinationals responded to this shock in the following years by increasing their use of ESOs and R&D. I also provide evidence that these firms increased claims for R&D tax credits as a share of total R&D expenditure. This indicates that relabeling of foreign expenses as domestic may be driving some of the observed increased R&D in the tax data.

Growth in Stock Compensation Among Exposed Firms

Using information about stock compensation expense provided in Compustat, I run a standard, unstaggered event study specification to estimate the growth in stock compensation,

$$Y_{ct} = \sum_{t>=2005} D_{ct} + \gamma_c + \sum_{g \in G} \mu_{gt} + \varepsilon_{it}, \tag{4.1}$$

where Y_{ct} is a measure of stock compensation for multinational c in year t, D_{ct} is an indicator equal to one in year t if multinational c is observed to have a cost-sharing agreement. γ_c is a multinational fixed effect and $\sum_{g \in G} \mu_{gt}$ are year-by-group fixed effects for a set of groups G that allow for rich time-varying heterogeneity. I examine a simple specification that only includes multinational and year fixed effects, as well as a richer specification that includes year-by-asset quartile fixed effects and year-by-sales quartile fixed effects.[9]

Figure 4.3 shows the resulting estimates where Y_{ct} is the log of stock compensation as reported in Compustat. After 2005, relative to multinationals without CSAs, exposed multinationals reported large increases in stock compensation. Table 4.2 contains point estimates from pooled regressions.

[9] I use consolidated asset and sales figures from Compustat to compute these quartiles.

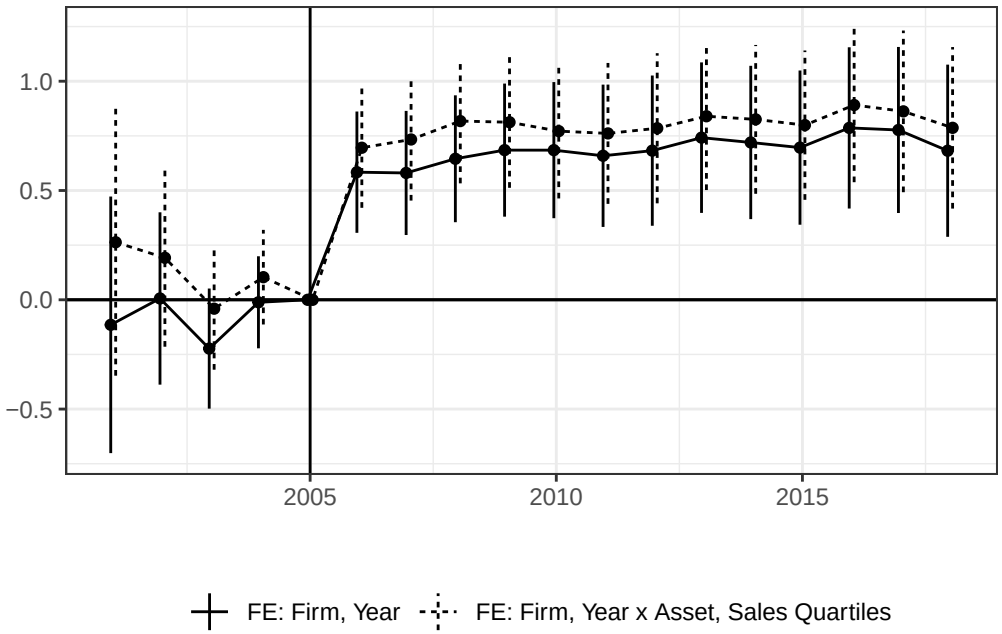

Figure 4.3: Effect on Stock Compensation

Notes: This table displays estimates from regressions under Equation 4.1. The dependent variable is log stock compensation for MNC c in year t as reported in Compustat. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects, year-by-group fixed effects, and year-by-sales quartile fixed effects.

Table 4.2: Effect on Log Stock Compensation

	(1)	(2)
ATT	0.731***	0.716***
	(0.128)	(0.127)
Num.Obs.	15 295	15 295
R2	0.812	0.818
R2 Adj.	0.794	0.798
R2 Within	0.014	0.013

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: This table presents point estimates from pooled regressions under Equation 4.1. The dependent variable is log stock compensation for MNC c in year t as reported in Compustat.

Figure 4.4 and Table 4.3 show dynamic and pooled estimates from a similar specification where Y_{ct} is the ratio of stock compensation scaled by R&D expenses incurred in 2005. I trim this dataset to remove obsevations where this ratio is not contained in $[0, 1]$. This alternative specification also shows large increases in stock compensation for exposed multinationals on the order of 10% of 2004 R&D spending.

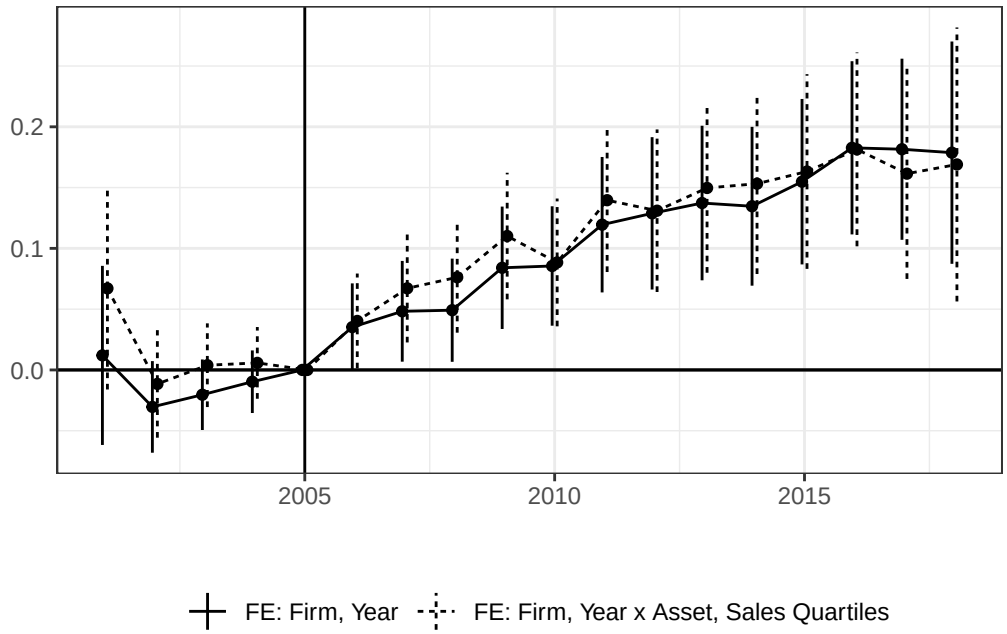

Figure 4.4: Effect on Ratio of Stock Compensation to R&D

Notes: This table presents estimates from pooled regressions under Equation 4.1. The dependent variable is the ratio of stock compensation for MNC c in year t to R&D expenses incurred in 2005. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects, year-by-group fixed effects, and year-by-sales quartile fixed effects.

Table 4.3: Effect on Ratio of Stock Compensation to R&D

	(1)	(2)
ATT	0.104***	0.096***
	(0.020)	(0.022)
Num.Obs.	7163	7163
R2	0.611	0.626
R2 Adj.	0.566	0.575
R2 Within	0.015	0.011

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: This table presents estimates from pooled regressions under Equation 4.1. The dependent variable is the ratio of stock compensation for MNC c in year t to R&D expenses incurred in 2005.

Shifts in the Wage Structure of Exposed Firms

The event studies shown above demonstrate that exposed multinationals increased their use of stock compensation after the Tax Court ruling in 2005. It could be the case, however, that exposed firms were simply growing faster than other multinationals. Figure 4.5 and Table 4.4 show dynamic and pooled estimates from a similar specification to Equation 4.1 where Y_{ct} is the ratio of stock compensation, as measured in Compustat, to the total wage bill, as observed in IRS tax data. As before, I trim this dataset to remove observations where this ratio is not contained in $[0, 1]$. The results show that, relative to other multinationals, exposed multinationals shifted their wage structure toward stock compensation.

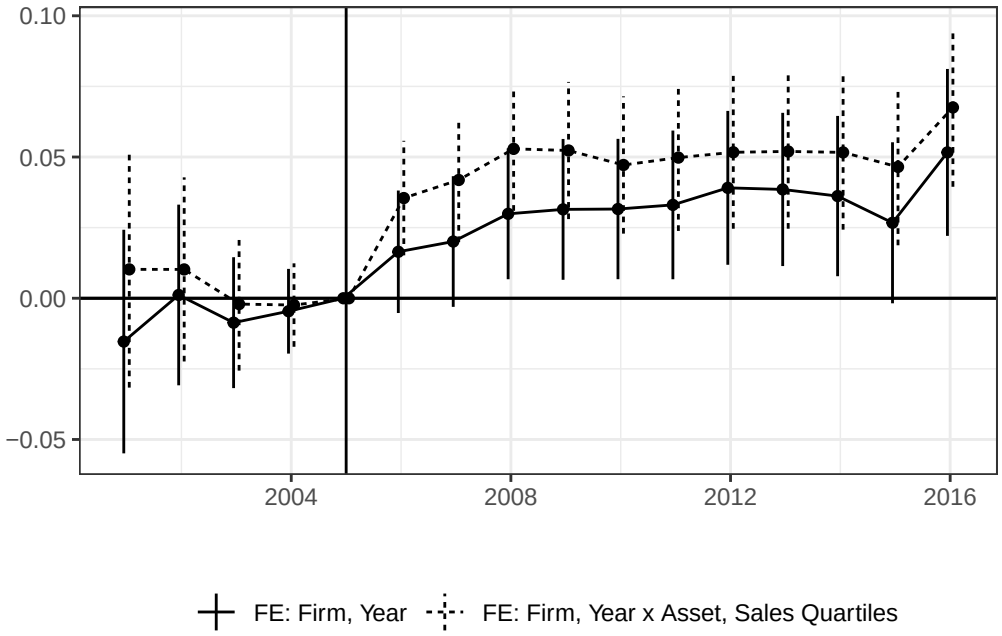

Figure 4.5: Effect on Ratio of Stock Compensation to Total Labor Compensation

Notes: This table presents point estimates from pooled regressions under Equation 4.1. The dependent variable is the ratio of stock compensation for MNC c in year t, as measured in Compustat, to the total wage bill as measured in IRS data. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects, year-by-group fixed effects, and year-by-sales quartile fixed effects.

Table 4.4: Effect on Ratio of Stock Compensation to Total Labor Compensation

	(1)	(2)
ATT	0.035***	0.047***
	(0.010)	(0.010)
Num.Obs.	11 453	11 453
R2	0.667	0.676
R2 Adj.	0.628	0.635
R2 Within	0.003	0.006

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table presents point estimates from pooled regressions under Equation 4.1. The dependent variable is the ratio of stock compensation for MNC c in year t, as measured in Compustat, to the total wage bill as measured in IRS data.

Evidence for R&D Relabeling by Exposed Firms

It is possible that MNCs could also use their shifted labor costs to claim additional domestic R&D tax credits—the Tax Court briefing indicates that Altera did exactly this. Figure 4.6 and Table 4.5 show dynamic and pooled estimates from a similar specification to Equation 4.1 where Y_{ct} is the ratio of domestic R&D expenses claimed under the R&D tax credit to the firm's total R&D expense as reported in Compustat.[10] Note that, if the results shown above were confounded by international expansion of the group of exposed firms, one would expect this ratio to be declining with their domestic share of R&D. Instead, this ratio is **increasing**, indicating that exposed multinationals performed domestic R&D relatively more intensely than other multinationals. To the extent that this increase was driven by reclassification of foreign expenses, it does not represent real R&D growth for these firms, but rather a behavioral response to the regulatory uncertainty surrounding the classification of ESOs.

Note that, when examining raw R&D growth, exposed firms generated much larger increases in R&D compared to unexposed firms as shown in Figure 4.10 and Table 4.9.

[10] As with the other ratios, I trim this dataset to remove observations where this ratio is not contained in $[0, 1]$.

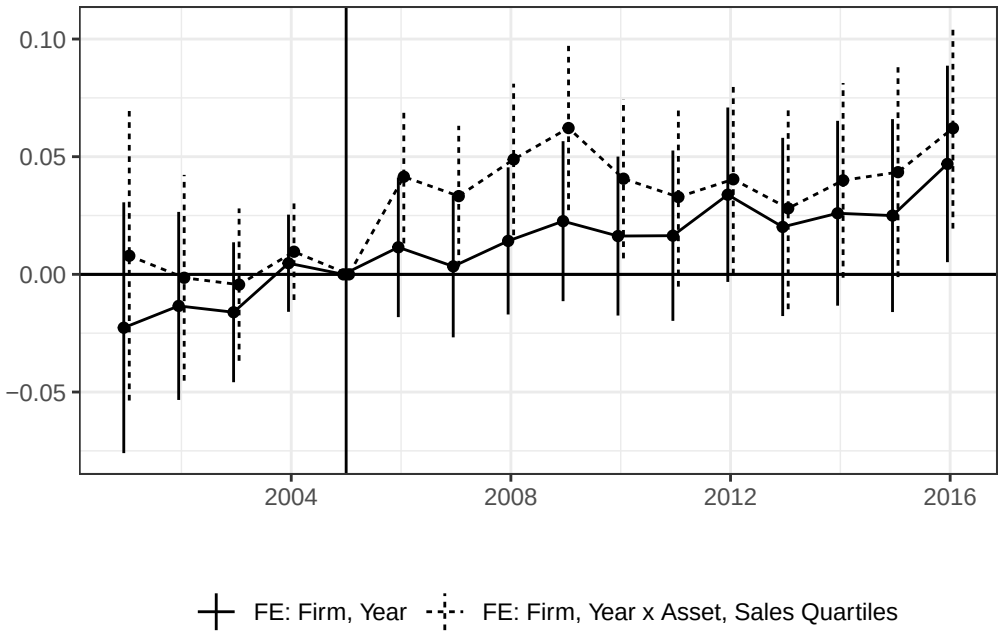

Figure 4.6: Effect on Ratio of Claimed R&D Expenses to Compustat R&D

Notes: This table presents point estimates from pooled regressions of the form

$$Y_{ct} = \sum_{t>=2005} D_{ct} + \gamma_c + \sum_{g \in G} \mu_{gt} + \varepsilon_{it}.$$

The dependent variable is the ratio of domestic R&D expenses claimed under the R&D tax credit to a firm's total R&D expense as reported in Compustat. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects, year-by-group fixed effects, and year-by-sales quartile fixed effects.

Table 4.5: Effect on Ratio of Claimed R&D Expenses to Compustat R&D

	(1)	(2)
ATT	0.025+	0.041**
	(0.013)	(0.015)
Num.Obs.	6258	6258
R2	0.636	0.651
R2 Adj.	0.587	0.596
R2 Within	0.002	0.004

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: This table presents point estimates from pooled regressions under Equation 4.1. The dependent variable is the ratio of domestic R&D expenses claimed under the R&D tax credit to a firm's total R&D expense as reported in Compustat.

Threats to Identification and Robustness Checks

The results above provide compelling evidence that exposed multinationals strongly reacted to the period of regulatory uncertainty starting in 2005 by shifting employee compensation toward ESOs and by reclassifying foreign R&D expenses as domestic expenses.

One potential concern with these results is that measurement error may be introducing spurious effects. This is particularly a concern with regard to estimates that rely on stock compensation figures reported in Compustat. There were significant changes to accounting standards that also affected reporting of stock compensation in 2005. After these changes, starting in 2006, Compustat's coverage of ESOs became much more widespread. Figures 4.7 through 4.9 and Tables 4.6 through 4.8 provide alternative estimates that use 2006 as the reference year for the policy shock. While these estimates also show positive effects, they are attenuated relative to the baseline estimates presented above. Note that, for measurement error to have a biased effect on the baseline estimates, it would have to differentially affect exposed multinationals—homogeneous forms of measurement error in the pre-period would result in parallel trends, satisfying the difference-in-differences identifying assumption. Results that are limited to R&D variables (Figure 4.6 and Figure 4.10) should not be affected

by this potential measurement issue.

4.5 Conclusion

Multinational firms face incentives to engage in behavior that lowers their aggregate tax cost. One way that they can do this is by reclassifying costs that occur in a low-tax jurisdiction to occur in a high-tax jurisdiction, where it can generate a larger tax shield. In 1995, the Treasury passed regulations that failed to prevent firms from reclassifying stock-based compensation. Although it attempted to fix these regulations in 2003 to address this shortcoming, litigation generated a long period of regulatory uncertainty. A favorable ruling increased the value of exposed firms, and these firms responded by structuring their R&D costs to take advantage of the potential tax advantage, increasing their use of stock-based compensation. Evidence from tax data suggests they also may have used the to relabel foreign R&D costs as domestic R&D costs for the purposes of claiming an R&D credit, further increasing their tax shield. This contributes to the larger literature on corporate tax and income shifting by providing clear measures of the behavioral response of firms to the corporate income tax. Further work might clarify the extent of the tax losses that may have been generated by this activity, and whether legal rulings in favor of the regulatory authorities result in their recapture.

4.6 Appendix

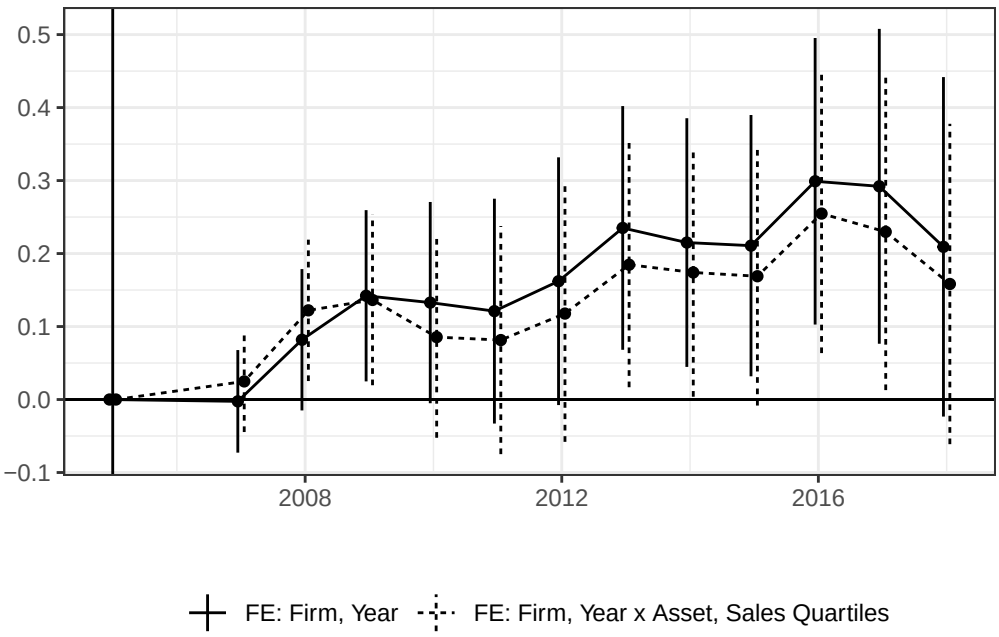

Figure 4.7: Effect on Stock Compensation (Robustness, Ref Year = 2006)

Notes: This figure presents estimates, using an alternate reference year of 2006, from regressions under Equation 4.1. The dependent variable is log stock compensation for MNC c in year t as reported in Compustat—comparable to Figure 4.3. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects, year-by-group fixed effects, and year-by-sales quartile fixed effects.

Table 4.6: Effect on Log Stock Compensation (Robustness, Ref Year = 2006)

	(1)	(2)
ATT	0.147*	0.127*
	(0.060)	(0.061)
Num.Obs.	10 504	10 504
R2	0.909	0.912
R2 Adj.	0.895	0.897
R2 Within	0.001	0.001

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: This table presents estimates, using an alternate reference year of 2006, from regressions under Equation 4.1. The dependent variable is log stock compensation for MNC c in year t as reported in Compustat—comparable to Table 4.2.

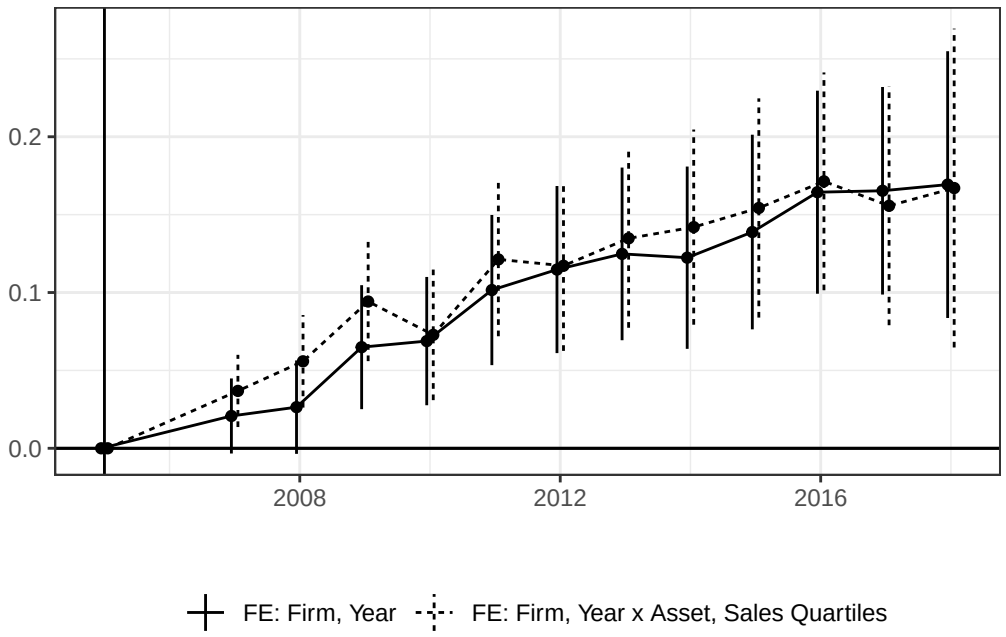

Figure 4.8: Effect on Ratio of Stock Compensation to R&D (Robustness, Ref Year = 2006)

Notes: This figure presents estimates, using an alternate reference year of 2006, from regressions under Equation 4.1. The dependent variable is the ratio of stock compensation for MNC c in year t to R&D expenses incurred in 2005—comparable to Table 4.3.

Table 4.7: Effect on Ratio of Stock Compensation to R&D (Robustness, Ref Year = 2006)

	(1)	(2)
ATT	0.080***	0.092***
	(0.017)	(0.016)
Num.Obs.	4630	4630
R2	0.702	0.715
R2 Adj.	0.649	0.656
R2 Within	0.006	0.007

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: This table presents estimates, using an alternate reference year of 2006, from regressions under Equation 4.1. The dependent variable is the ratio of stock compensation for MNC c in year t to R&D expenses incurred in 2005—comparable to Table 4.3.

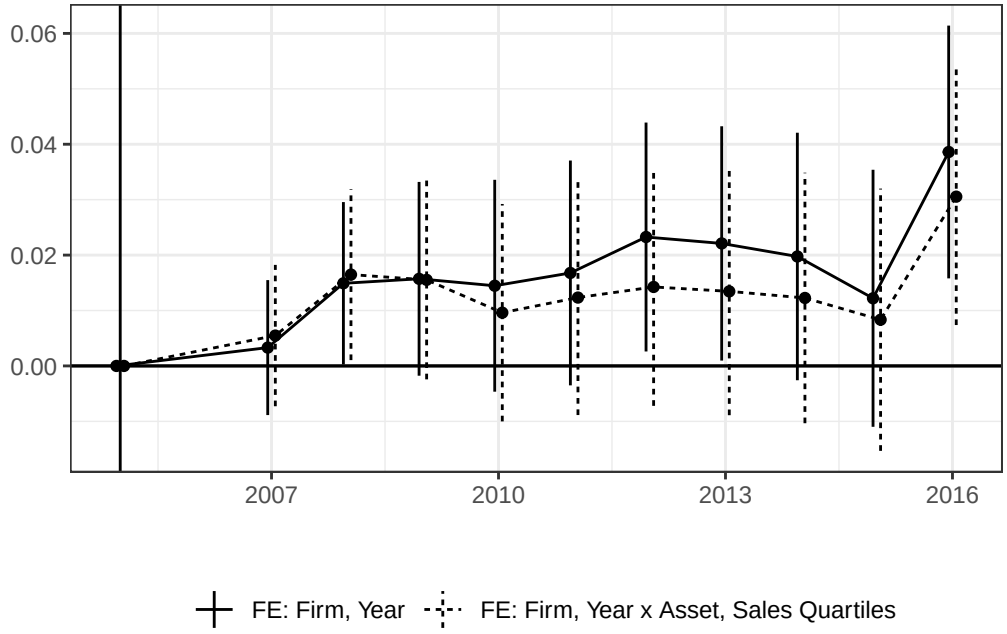

Figure 4.9: Effect on Ratio of Stock Compensation to Total Labor Compensation (Robustness, Ref Year = 2006)

Notes: This figure presents estimates, using an alternate reference year of 2006, from regressions under Equation 4.1. The dependent variable is the ratio of stock compensation for MNC c in year t to R&D expenses incurred in 2005—comparable to Figure 4.5. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects, year-by-group fixed effects, and year-by-sales quartile fixed effects.

Table 4.8: Effect on Ratio of Stock Compensation to Total Labor Compensation (Robustness, Ref Year = 2006)

	(1)	(2)
ATT	0.016*	0.013+
	(0.007)	(0.008)
Num.Obs.	7684	7684
R2	0.765	0.770
R2 Adj.	0.726	0.728
R2 Within	0.001	0.000

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: This table presents estimates, using an alternate reference year of 2006, from regressions under Equation 4.1. The dependent variable is the ratio of stock compensation for MNC c in year t, as measured in Compustat, to the total wage bill as measured in IRS data—comparable to Table 4.4.

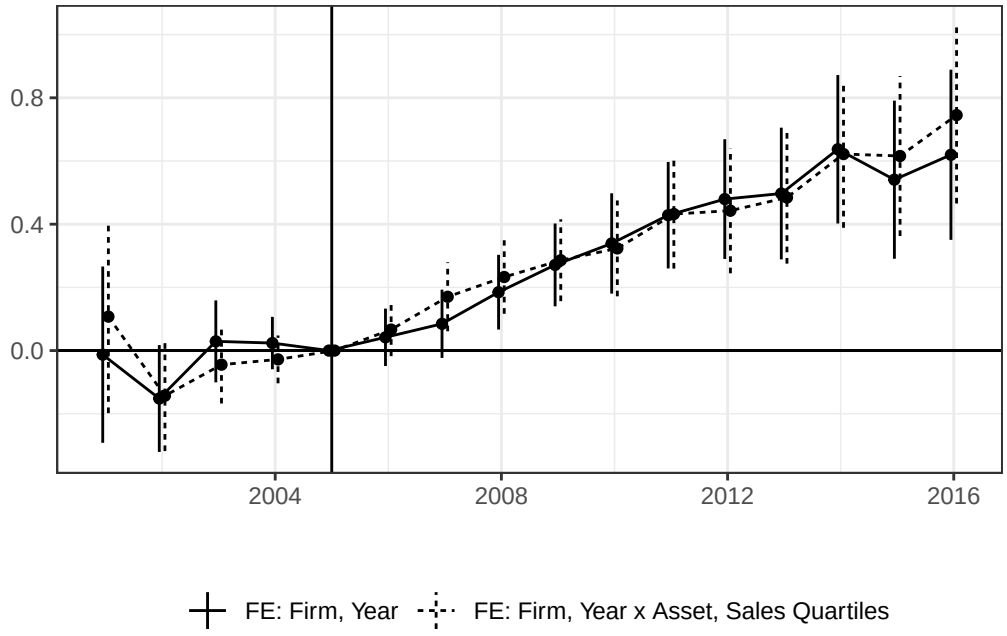

Figure 4.10: Effect on Log R&D

Notes: This figure presents estimates, using an alternate reference year of 2006, from regressions under Equation 4.1. The dependent variable is the ratio of stock compensation for MNC c in year t, as measured in Compustat, to the total wage bill as measured in IRS data—comparable to Table 4.4.

Table 4.9: Effect on Log R&D

	(1)	(2)
ATT	0.363***	0.410***
	(0.079)	(0.081)
Num.Obs.	8340	8340
R2	0.918	0.921
R2 Adj.	0.909	0.911
R2 Within	0.009	0.010

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: This table presents estimates, using an alternate reference year of 2006, from regressions under Equation 4.1. The dependent variable is log R&D expenses incurred in 2005.

5

Conclusion

This dissertation examines how multinational corporations based in the United States react to incentives created by the global system of corporate income taxation. These firms have the ability to operate across countries that offer a menu of different tax systems. As a result, they face incentives to shift capital and earnings in part to minimize their tax burden. The lack of harmonization in the global tax system, during the period examined in the preceding analyses, violates basic tenets of optimal taxation by encouraging corporations to engage in behavior that is unrelated to production efficiency.

Multinationals respond to the global tax system in myriad complex ways. In the second chapter, I demonstrate how multinationals structure their foreign affiliate networks to lower their foreign effective tax rates. I show that, over the period examined, affiliate structures in Ireland, the Netherlands, and Luxembourg became increasingly widespread. Adoption of these structures coincided with stark decreases in multinationals' foreign effective tax rates, and large shifts in their economic activity.

Researchers who study multinationals often use disaggregated affiliate-level data. This data is key to understanding how these firms allocate capital and earnings across different jurisdictions. Researchers, however, must be careful with how they use this data. The third chapter of this dissertation examines the extent to which certain types of measurement error may appear when researchers aggregate tax data. I examine the performance of a corrected aggregation method to measure foreign earnings, and show that it yields promising results.

The fourth chapter examines how the global tax system incentivizes changes in the domestic activity of US multinationals. I examine a period of regulatory uncertainty around the treatment of multinationals' IP transfers, and show that exposed firms strongly reacted to tax minimization opportunities. This chapter offers an alternative explanation for growth

in the use of stock compensation and demonstrates that tax incentives have shifted R&D activity for a particular set of multinationals.

This dissertation examines firm behavior during a particular global tax regime that existed between 1992 and 2016. During this period, the global economy became increasingly integrated. At the same time, there was a general lack of global coordination in the tax treatment of multinational firms. Since 2016, there have been deep changes to the global tax system. In the United States, the Tax Cuts and Jobs Act of 2017 shifted the tax regime away from a worldwide system toward a territorial system. Simultaneously, there have been recent attempts to coordinate international tax policy to minimize the distortions caused by the previous system and increase the global corporate tax base. Future research should shed light on the success of these efforts.

Bibliography

Albertus, James F. 2019. "Does Foreign Tax Arbitrage Promote Innovation?" SSRN.

Altera Corporation & Subsidiaries v. Commissioner. 2019. "Altera Corporation & Subsidiaries v. Commissioner." 926 F.3d. 1061.

Altshuler, Rosanne, and Harry Grubert. 2006*a*. "Governments and Multinational Corporations in the Race to the Bottom." *Tax Notes*, 110(8).

Altshuler, Rosanne, and Harry Grubert. 2006*b*. "The Three Parties in the Race to the Bottom: Host Governments, Home Governments and Multinational Corporations." *Florida Tax Review*, 7(3): 137–209.

Avi-Yonah, Reuven S. 2009. "Xilinx and the Arm's-Length Standard." *Tax Notes*, 123(10): 1231–1236.

Bilicka, Katarzyna Anna. 2019. "Comparing UK Tax Returns of Foreign Multinationals to Matched Domestic Firms." *American Economic Review*, 109(8): 2921–53.

Bilicka, Katarzyna Anna, Yaxuan Qi, and Jing Xing. 2019. "Real Responses to Anti-tax Avoidance Policies: Evidence from UK's Worldwide Debt Cap." *Working Paper*.

Blouin, Jennifer L., and Leslie Robinson. 2020. "Double Counting Accounting: How Much Profit of Multinational Enterprises Is Really in Tax Havens?" SSRN September 1, 2020.

Blouin, Jennifer L., and Linda K Krull. 2014. "Does Organizational Form Affect Firms' Foreign Operations? The Role of "Check-the-Box" on Multinational Tax Planning." Vol. 107, 1–31, JSTOR.

Bodie, Zvi, Robert S. Kaplan, and Robert C. Merton. 2003. "For the Last Time: Stock Options Are an Expense." Harvard Business Review Web Site `https://hbr.org/2003/03/for-the-last-time-stock-options-are-an-expense` (Accessed: July 4, 2023).

Bound, John, and Alan B Krueger. 1991. "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" *Journal of labor economics*, 9(1): 1–24.

Bound, John, Charles Brown, Greg J Duncan, and Willard L Rodgers. 1994. "Evidence on the Validity of Cross-sectional and Longitudinal Labor Market Data." *Journal of Labor Economics*, 12(3): 345–368.

Clausing, Kimberly A. 2016. "The Effect of Profit Shifting on the Corporate Tax Base in the United States and Beyond." *National Tax Journal*, 69(4): 905–934.

de Mooij, Ruud A., and Li Liu. 2018. "At A Cost: the Real Effect of Transfer Pricing Regulations on Multinational Investment."

Diamond, Peter A., and James A. Mirrlees. 1971*a*. "Optimal Taxation and Public Production I: Production Efficiency." *American Economic Review*, 61(1): 8–27.

Diamond, Peter A., and James A. Mirrlees. 1971*b*. "Optimal Taxation and Public Production II: Tax Rules." *American Economic Review*, 261–278.

Dowd, Tim, Paul Landefeld, and Anne Moore. 2017. "Profit Shifting of U.S. Multinationals." *Journal of Public Economics*, 148: 1–13.

Drucker, Jesse. 2010. "Google 2.4% Rate Shows How $60 Billion Is Lost to Tax Loopholes."

Dube, Arindrajit, Ethan Kaplan, and Suresh Naidu. 2011. "Coups, Corporations, and Classified Information." *Quarterly Journal of Economics*, 126(3): 1375–1409.

Duhigg, Charles, and David Kocieniewski. 2012. "How Apple Sidesteps Billions in Taxes."

Dyreng, Scott, and Michelle Hanlon. 2021. "Tax Avoidance and Multinational Firm Behavior."

Fajgelbaum, Pablo D., Eduardo Morales, Juan Carlos Suárez Serrato, and Owen M. Zidar. 2018. "State Taxes and Spatial Misallocation." *Review of Economic Studies*.

Faulkender, Michael W, Kristine W Hankins, and Mitchell A Petersen. 2019. "Understanding the Rise in Corporate Cash: Precautionary Savings or Foreign Taxes." *The Review of Financial Studies*, 32(9): 3299–3334.

Foley, Sean, and Hardi Kuvadia. 2021. "Stock-based Compensation Cost-sharing Regs. Following Altera." The Tax Adviser Web Site `https://www.thetaxadviser.com/issues/2021/jun/stock-based-compensation-cost-sharing-regs-altera.html` (Accessed: July 4, 2023).

Gravelle, Jane G. 2013. "Tax Havens: International Tax Avoidance and Evasion." Congressional Research Service Web Site `https://crsreports.congress.gov/product/details?prodcode=R40623` (Accessed: July 4, 2023).

Grubert, Harry. 2012. "Foreign Taxes and the Growing Share of US Multinational Company Income Abroad: Profits, Not Sales, Are Being Globalized." *National Tax Journal*, 65(2): 247–281.

Grubert, Harry, and Joel Slemrod. 1998. "The Effect of Taxes on Investment and Income Shifting to Puerto Rico." *Review of Economics and Statistics*, 80(3): 365–373.

Guardian. 2018. "From Seattle to Luxembourg: How Tax Schemes Shaped Amazon."

Guardian. 2019. "EU Investigates Nike's Tax Status in the Netherlands."

Hardeck, Inga, and Patrick U. Wittenstein. 2018. "Assessing the Tax Benefits of Hybrid Arrangements - Evidence from the Luxembourg Leaks." *National Tax Journal*, 71(2): 295–334.

Hines, James R., and Eric M. Rice. 1994. "Fiscal Paradise: Foreign Tax Havens and American Business." *The Quarterly Journal of Economics*, 109(1): 149–182.

IRS. 2011. "SOI Sampling Methodology and Data Limitations." U.S. Internal Revenue Service `https://www.irs.gov/pub/irs-prior/p946--2002.pdf` (Accessed: February 7, 2022).

IRS. 2022*a*. "SOI Tax Stats - Controlled Foreign Corporations ." U.S. Internal Revenue Service Web Site `https://www.irs.gov/statistics/soi-tax-stats-controlled-foreign-corporations` (Accessed: February 7, 2022).

IRS. 2022*b*. "SOI Tax Stats - Corporation Complete Report." U.S. Internal Revenue Service Web Site `https://www.irs.gov/statistics/soi-tax-stats-corporation-complete-report` Accessed: February 7, 2022).

IRS Acquiesces in Result, But Not Reasoning, of Xilinx Transfer Pricing Case. 2010. "IRS Acquiesces in Result, But Not Reasoning, of Xilinx Transfer Pricing Case." Internal Revenue Service `https://www.taxnotes.com/research/federal/irs-guidance/actions-on-decisions/irs-acquiesces-in-result%2C-but-not-reasoning%2C-of-xilinx-transfer/1fchf` (Accessed: July 4, 2023).

Kanbur, Ravi, and Michael Keen. 1993. "Jeux Sans Frontières: Tax Competition and Tax Coordination When Countries Differ in Size." *The American Economic Review*, 83(4): 877–892.

Keen, Michael, and Kai A. Konrad. 2013. "The Theory of International Tax Competition and Coordination." In *Handbook of Public Economics*. Vol. 5, 257–328. Elsevier.

Kleinbard, Edward D. 2013. "Through a Latte, Darkly: Starbucks's Stateless Income Planning." *Tax Notes*, 1515–1535.

MacKinlay, A. Craig. 1997. "Event Studies in Economics and Finance." *Journal of Economic Literature*, 35(1): 13–39.

McWaters, Eleanor Moran. 2020. "Supreme Court Declines Certiorari in Altera, Stock-Based Compensation Cost Sharing Regulations Stand." Crowell & Moring LLP Web Site `https://www.crowell.com/en/insights/client-alerts/supreme-court-declines-certiorari-in-altera-stock-based-compensation-cost-sharing-regulations-stand` (Accessed: July 4, 2023).

Mutti, John, and Harry Grubert. 2009. "The Effect of Taxes on Royalties and the Migration of Intangible Assets Abroad." In *International trade in services and intangibles in the era of globalization*. 111–137. University of Chicago Press.

Samarakoon, Navodhya. 2022. "The Effect of the Closure of the Double Irish Loophole on the Location of U.S. Multinational Companies' Profits." SSRN November 23, 2022.

Suárez Serrato, Juan Carlos. 2018. "Unintended Consequences of Eliminating Tax Havens." Duke University Working Paper.

Suárez Serrato, Juan Carlos, and Owen Zidar. 2016. "Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms." *American Economic Review*.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics*, 225(2): 175–199.

Tax Reform Act of 1986. 1986. "Tax Reform Act of 1986." 99th United States Congress.

Tørsløv, Thomas R, Ludvig S Wier, and Gabriel Zucman. 2018. "The Missing Profits of Nations." National Bureau of Economic Research Working Paper 24701.

Yagan, Danny. 2015. "Capital Tax Reform and the Real Economy: The Effects of the 2003 Dividend Tax Cut." *The American Economic Review*, 105(12): 3531–3563.

Zodrow, George R, and Peter Mieszkowski. 1986. "Pigou, Tiebout, Property Taxation, and the Underprovision of Local Public Goods." *Journal of urban economics*, 19(3): 356–370.

Zucman, Gabriel. 2014. "Taxing Across Borders: Tracking Personal Wealth and Corporate Profits." *Journal of Economic Perspectives*, 28(4): 121–48.

Zucman, Gabriel. 2015. *The Hidden Wealth of Nations.* The University of Chicago Press.

Zwick, Eric, and James Mahon. 2017. "Tax Policy and Heterogeneous Investment Behavior." *American Economic Review*, 107(1): 217–48.

Biography

Lysle Boller is a Ph.D. candidate in the Department of Economics at Duke University. His research interests lie at the intersection of industrial organization and public economics, with a focus on how firms respond to the incentives created by the global tax system. He holds a B.A. in mathematics and economics from the University of Virginia and an M.Sc. in economics from the London School of Economics.