Tax Planning and Multinational Behavior*

Rosanne Altshuler
Rutgers University

Lysle Boller
Penn Wharton Budget Model

Juan Carlos Suárez Serrato
Stanford GSB & NBER

February 2024

Abstract

We study the adoption and use of a specific form of tax planning by US multinational corporations (MNCs). Using IRS data, we identify "hybrid" tax planning structures (HTPs) which can be used to avoid corporate income tax by targeting mismatches between US and Irish, Dutch, and Luxembourgish tax law. By 2016, more than 35% of the foreign profits of US MNCs were linked to HTPs. Difference-in-differences models comparing adopting and non-adopting MNCs reveal that after HTP adoption, MNCs intensify behaviors related to profit shifting, significantly increasing related-party loans, foreign intangible assets, and profits held abroad. These changes result in stark reductions in foreign effective tax rates. Adopting MNCs also experience larger increases in foreign tangible assets and in global R&D, payroll, and investment.

Keywords: international taxation, profit shifting, Double Irish, Reverse Hybrid Mismatch
JEL Codes: D22, H25, H26, H32

This research was 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.

*We are very grateful for comments from David Agrawal, Hunt Allcott, Alan Auerbach, Michael Best, Jeff Clemens, Julie Cullen, Dhammika Dharmapala, Rebecca Diamond, Tim Dowd, Clare Doyle, Jesse Drucker, Naomi Feldman, Roger Gordon, Mindy Herzfeld, Jim Hines, Chad Jones, Wojciech Kopczuk, Wouter Leenders, Rebecca Lester, Antoine Levy, Li Liu, Clément Malgouyres, Peter Merrill, Mathilde Muñoz, Karthik Muralidharan, Paul Oosterhuis, Mathieu Parenti, Emmanuel Saez, Navodhya Samarakoon, Shoshanna Vasserman, Daniel Xu, Danny Yagan, Gabriel Zucman, and from seminar participants at CESifo, Columbia University, the Federal Reserve Bank of San Francisco, the National Tax Association Annual Meetings, the NBER Public Economics Meetings, the Office of Tax Analysis of the US Treasury Department, the Oxford University Centre for Business Taxation, the Paris School of Economics, the Penn Wharton Budget Model, Rutgers University, Stanford University, UC Berkeley, UC San Diego, the University of Illinois at Urbana-Champaign, the University of South Carolina, and the Statistics of Income unit of the Internal Revenue Service. We are grateful to Arnold Ventures, the International Tax Policy Forum, and the NBER for their support of this project. Special thanks to Chloe Gagin, Nuria McGrath, Melanie Patrick, Ralph Rector, Michael Weber, and Jason Wenrich for assistance with data access. Clare Doyle, Jason Harrison, and Kevin Roberts provided excellent research assistance. All errors remain our own.

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 leveraging mismatches in tax laws across countries. While media attention following these revelations and dissatisfaction with the current system have motivated important international tax policy changes and multilateral projects to address tax avoidance, 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 tax data from the Internal Revenue Service (IRS) to answer these questions by analyzing the adoption and use of a set of complex tax planning strategies that target mismatches between US and 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 leverage mismatches in tax laws across these countries. Data from the tax returns of US MNCs is crucial for the purposes of identifying the adoption of these tax planning structures and understanding how MNCs use them to shift profits across countries. This paper is the first to systematically uncover these tax planning structures and to study how their adoption is related to changes in tax avoidance and real economic activity.

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 were gradually adopted by MNCs following the 1997 regulations known as "Check the Box" (CTB), which facilitated this form of tax planning. By 2016, 17.5% 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 all MNCs in our sample are linked to these structures. MNCs that use these strategies are also responsible for significant shares of domestic economic activity,

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

Figure 1: Comparison of Foreign Effective Tax Rates

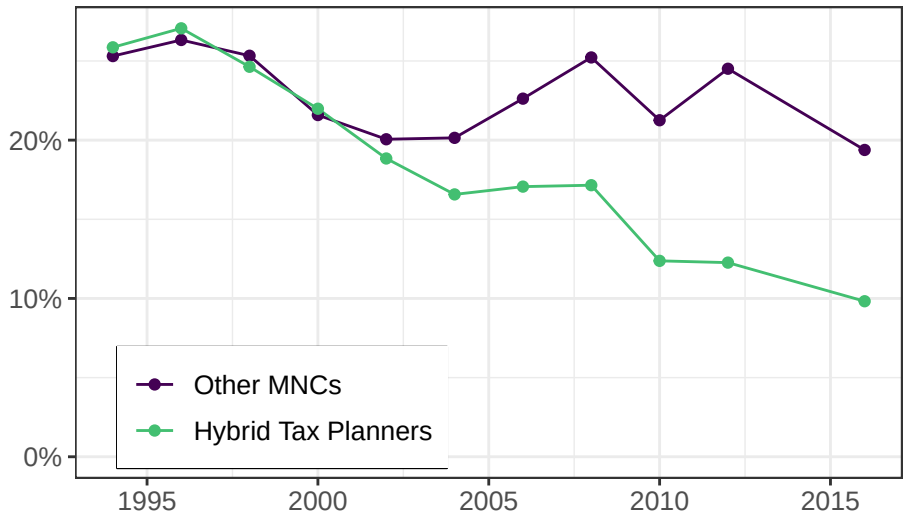

Notes: This figure compares the aggregate annual foreign ETR for two groups of US MNCs. The light green line shows the ETR for a group of MNCs that eventually adopt at least one of the hybrid tax structures described in Section 2. The dark blue line shows the ETR for MNCs that did not adopt any of these structures during the sample window. The combined sample includes most large US C corporations as described in Section 3.

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 1, over the period during which these structures are gradually put in place, the foreign ETRs of adopting MNCs experience a dramatic decline. By 2016, 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. The models measure changes in firm outcomes surrounding the adoption of hybrid tax planning 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 and that these mechanisms of tax avoidance lead to significant declines in foreign ETRs, as in Figure 1.

Concurrent with declining foreign ETRs, we estimate that adopting US MNCs have larger increases in foreign investment and accumulate 40% ($p < 0.01$) 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 20% ($p < 0.01$) increase in domestic payroll and a 50% ($p < 0.01$) increase in expenditures on research and development (R&D).

We develop our results in three 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, and domestic investment. Parents also file an information form (Form 5471) for each of their controlled foreign corporations (CFCs), 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 considered 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 leverage 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 route foreign profits to low-tax foreign affiliates. 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 mismatch arrangements.

Using these unique indicators of tax planning, we measure the growth, prevalence, and importance of these hybrid tax planning (HTP) 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 magnitude.

[2]Use of hybrid tax planning structures by US MCNs has been previously reported by Drucker (2010), Duhigg and Kocieniewski (2012), Guardian (2019), Guardian (2018), and Kleinbard (2013).

In our third and final step, we compare the foreign and domestic operations of US MNCs that adopt one of these strategies to those of MNCs that do not. In the early 1990s, these two groups of firms experience similar trajectories with regard to their foreign ETRs and measures of domestic and foreign activity. After the 1997 CTB regulations are put in place, foreign affiliates of MNCs that engage in hybrid tax planning experience significantly larger declines in ETRs; more rapid foreign and domestic growth; and sharp increases across several proxies for profit shifting and tax deferral, including related-party loans, intangible assets, payments related to cost sharing agreements, royalty income, and cash held abroad.

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 using a "stacked" difference-in-differences estimator (e.g., as in Cengiz, Dube, Lindner and Zipperer, 2019) provide 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 fixed effects, and when we flexibly control for differences in intangible assets across firms. We also obtain similar results when we additionally use inverse propensity score weights to attain more similar firm-to-firm comparisons, when we use a two-way fixed effects (TWFE) specification, and when we use an alternative staggered estimator from Sun and Abraham (2021).

Estimates of our staggered difference-in-differences models have a causal interpretation under the assumption that the outcomes of MNCs that adopted 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

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

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. Interpreting these results through the lens of an economic model could help decompose firm-level changes into two components, one that captures selection on gains from tax planning and another that captures the effects of the HTP tax advantage.

Overall, our results provide the first systematic documentation of the prevalence of hybrid tax planning structures among US MNCs. We show that, by 2016, these structures accounted for about a third of their foreign profits and that these MNCs represent a large fraction of domestic corporate activity. 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 MNCs. 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 understanding 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 the suggestions in Blouin and Robinson (2020), 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. The characterization of the structure of the foreign activities of US MNCs prior to the passage of the Tax Cuts and Jobs Act of 2017 (TCJA) is an important contribution of this paper, as the response of US MNCs to the many changes and new incentives in the TCJA likely depends on these preexisting structures.[4]

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) 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 transferred intangible assets abroad. Blouin and Krull (2014) show that MNCs had more tax haven affiliates and longer ownership chains after the enactment of CTB. Faulkender, Hankins and Petersen (2019) study the 440% 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 the bulk of the decrease in foreign ETRs and increase in cash held abroad over the sample period

[4]The TCJA lowered the corporate tax rate to 21% and made significant changes to US taxation of international income. It also introduced four new provisions. First, due to the transition to territorial taxation, US MNCs now can deduct dividends received from foreign affiliates from their US taxable income, thereby eliminating any repatriation tax. Second, the TCJA introduced a new tax on Global Intangible Low-Taxed Income (GILTI)—defined as income that includes low-tax foreign income exceeding 10% of an MNC's tangible foreign capital investment (adjusted for depreciation). Third, a new category of income—Foreign-Derived Intangible Income (FDII)—is subject to a reduced tax rate. FDII encompasses income derived from intellectual property held in the United States that generates foreign sales. Finally, the TCJA introduced the Base Erosion and Anti-Abuse Tax (BEAT) to curb the erosion of the tax base by both US and foreign MNCs.

[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.

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 experienced a larger decline in average rates and increased 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 the types of firms that benefited 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 discusses how CTB facilitated the creation of foreign disregarded entities and describes the three hybrid tax planning structures we examine. Section 3 provides an overview of the data. Section 4 discusses how the creation of foreign disregarded entities and hybrid tax planning structures has grown over time. Section 5 estimates firm-level changes in foreign tax rates and economic activity of US MNCs following the adoption of tax planning strategies. Section 6 concludes. We conduct additional analyses in the appendices. Appendix A provides robustness checks for our main analyses. Appendix B discusses the measurement of foreign earnings and taxes. Appendix C studies how MNCs can structure cost sharing agreements to shift profits from the US parent to foreign affiliates.

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.

[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) show that a worldwide debt cap that limited interest stripping as a form of profit shifting also impacts the investment decisions of UK MNCs.

2.1 Check The Box Regulations

During the period we study, the United States imposed a 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 US 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 Corporation 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 related CFC in another jurisdiction.

In 1996, the US 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 leverage 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 leverage mismatches in tax laws across countries to lower foreign tax bills.

2.2 Hybrid Tax Planning Strategies

CTB facilitates tax planning by allowing MNCs to easily create "hybrid" entities. 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 disregard a foreign corporation, thereby creating an FDE, a type of pass-through entity. If an entity is disregarded, the transactions with its entity parent and with

other FDEs owned by the same parent become transparent to the US Treasury, as they are all viewed as 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." Though the interest is deductible abroad against taxable income, it remains 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 tax purposes. Since 1997, however, 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 CFC is an unincorporated branch of the tax haven FDE; the interest payment is thus transparent to the US Treasury, which regards the combined tax haven/high-tax operation as one consolidated corporation. The interest payment therefore avoids Subpart F taxes and the company can defer US income tax by holding profits in the tax haven.

Panel A in Figure 2 depicts this simple hybrid tax planning structure using a tax haven affiliate. 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 transparent, thereby avoiding any current US tax on interest.

While this structure avoids Subpart F tax and defers US income tax, the MNC would still be subject to corporate income tax in the tax haven (if it exists) and potentially to foreign withholding tax 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.

[7] Although it was possible for MNCs to create foreign disregarded entities prior to 1997, 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.

[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.

2.2.1 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. This structure has the same foundation as in Panel A of Figure 2, but replaces the equity injection with a transfer of IP (via a cost sharing agreement) and uses royalty payments instead of interest to shift profits.

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

Even if IP is not underpriced, MNCs have historically attempted to strategically allocate allowable costs to generate tax savings through cost sharing agreements. We discuss these cost allocation strategies as well as their legal challenges in Section 5.3 and Appendix C.

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, including the US.

[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.

Figure 2: Diagrams of Hybrid Tax Planning Structures

(A) Simple CTB Structure

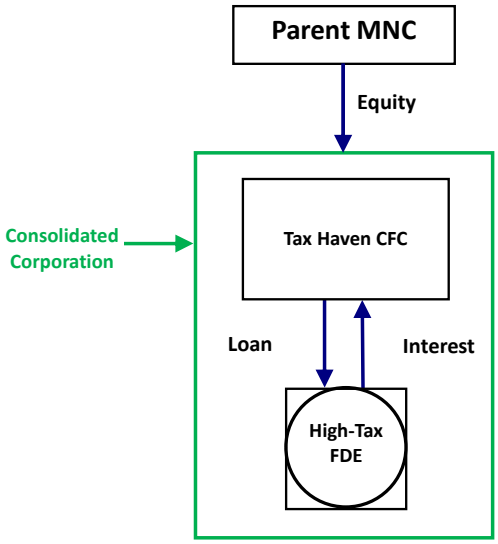

(B) Double Irish CTB with Cost Sharing

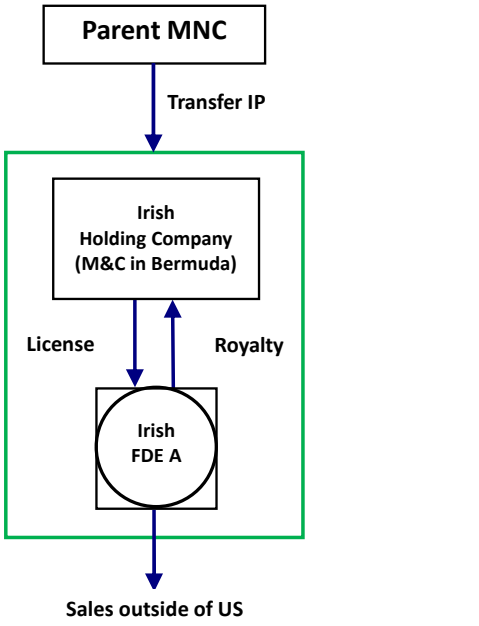

(C) Double Irish with Dutch Sandwich

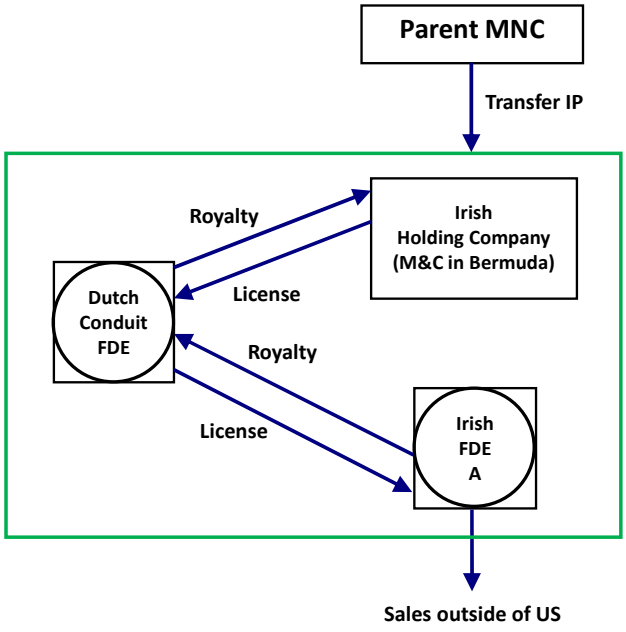

(D) Reverse Hybrid Mismatch: Dutch CV-BV

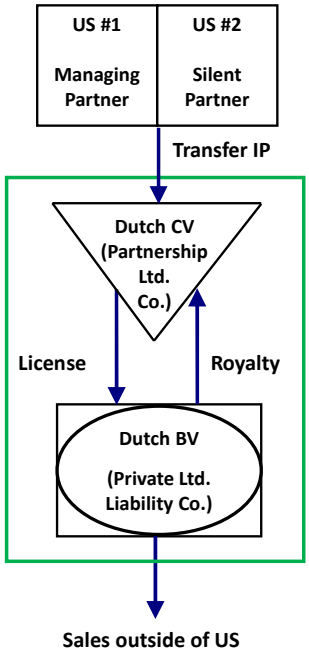

Notes: Panel A of Figure 2 depicts a hypothetical financing CTB structure; Panel B describes a Double Irish cost sharing structure with CTB; Panel C illustrates the Double Irish with a Dutch Sandwich; and Panel D describes a Reverse Hybrid Mismatch structure, otherwise known as a CV-BV (or SCS-SARL in the case of Luxembourgish entities). In each of these diagrams, the green rectangle depicts the combined structures as perceived by the IRS, and squares denote corporations. Squares with circles inside denote hybrid entities, which are corporations in the local country but disregarded for US purposes. In Panel D, the CV (or SCS if Luxembourg entities are used) 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.

2.2.2 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 a tax haven country such as Bermuda. To motivate this structure, 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 holding company in a tax haven using a cost sharing arrangement. The tax haven holding company then licenses the IP to an Irish CFC (CFC A), which pays royalties back to the holding company from the sales revenue it receives selling the IP abroad.

This initial cost sharing structure 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 Irish CFC A to the tax haven holding company. Second, taxes will be due in Ireland on any profits that remain in CFC A after royalties are paid to the haven holding company. Finally, the parent will owe Irish withholding taxes on the royalty transfers to the haven.

The first two of these tax problems can be solved using CTB and a Double Irish tax planning structure, as summarized in Panel B of Figure 2. In this structure, the parent transfers the IP to a holding company managed and controlled in a tax haven (e.g., Bermuda) but legally incorporated in Ireland. Though the US considers this an Irish holding company, under Irish tax law, the holding company is a Bermuda company and therefore not subject to Irish tax. The parent also checks the box on Irish CFC A to avoid current US tax on the royalties: the CFC becomes an FDE and is therefore fiscally transparent to the US Treasury. This eliminates Irish tax on any profits remaining in Ireland (i.e., the holding company) after royalties are paid, as well as current US taxes on the royalties. However, the transfer from the Irish FDE to the IP-holding company still generates Irish withholding tax.

MNCs can eliminate this withholding tax, thus solving the final tax problem, by inserting a Dutch conduit—a "Dutch Sandwich"—between the Irish affiliates. 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 holding company because no withholding tax is imposed on these transfers under Dutch law. To avoid Subpart F taxes on these 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, as shown in Panel C of Figure 2, solves all three tax problems we identified above.

2.2.3 Reverse Hybrid Mismatch

The final two tax planning structures we examine use a strategy known as a Reverse Hybrid Mismatch. While this strategy can be employed using affiliates in different countries, we 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 CV owns a Dutch 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 the IP to foreign companies, and pays royalties to the CV. Panel D of Figure 2 depicts this structure.

The Reverse Hybrid Mismatch allows the MNC to avoid tax on foreign income by solving three tax problems. First, tax may be due in the Netherlands. By Dutch tax law, the CV is a pass-through entity, so corporate tax is not levied in the Netherlands. Second, payments from BV to CV can generate Subpart F tax. The (reverse hybrid) CV is a corporation from the US perspective, and if the parent "checks the box" to disregard the (hybrid) BV, the US sees the two entities as a consolidated operation. Thus, no Subpart F tax will be due on the royalties. Finally, payments from the BV to the CV could trigger Dutch withholding tax. However, during our period of analysis, a 2005 decree by the Dutch Finance Ministry exempted US-based CV-BVs from withholding tax. With this Reverse Hybrid Mismatch structure in place, profits from US-developed IP sold abroad were not subject to corporate tax in the Netherlands and enjoyed indefinite deferral from US tax (under pre-TCJA law). The CV-BV structure is effectively a "sink" for foreign profits.

A Reverse Hybrid Mismatch structure can be set up through other countries. In particular, a combination of two types of Luxembourgish companies, known respectively as SCS and SARL, yields a structure similar to the Dutch CV-BV.[10] We study both CV-BV and SCS-SARL structures for the purpose of our analyses.

[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.

While the description of these structures emphasizes their 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 to engage in transactions to form the structures. In addition, company executives differ in their perceived cost of adopting tax-aggressive positions. To the extent that MNCs incur these costs to avoid paying taxes, tax planning is distortionary from an economic perspective.

3 Data and Sample Construction

3.1 IRS Business Tax Data

We rely primarily on several IRS datasets for our analysis. These administrative datasets provide parent and (both CFC- and FDE-level) affiliate-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 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, and 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 contains a subset of C corporations from the SOI Corporate Sample that file Form 5471 or Form 8858: Form 5471 provides financial information and activity of CFCs, and Form 8858 provides similar information for FDEs, the entity type enabled by CTB. We refer to this sample as the "SOI International Business Tax Sample." Unlike the SOI Corporate Sample, which is provided annually, CFC data is collected only in even years. FDE data is collected for four of

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

Table 1: Data Sources 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	Country of Incorporation Foreign Assets Foreign E&P Foreign Taxes Transactions Between CFCs Transactions Between US Parent and CFCs
8858	FDE Information Return	Country of Incorporation Date Disregarded Foreign Assets Foreign E&P Pass-through Owners Tax Owner

Compustat Data

Description	Selected Outcomes
Consolidated Public MNC Data	Deferred Foreign Taxes (txdfo) Foreign Taxes (txfo) Net Income (ni) Pretax Foreign Income (pifo) R&D Expense (xrd) Revenue (sale) Total Assets (at)

the years in our sample period (2006, 2008, 2012, and 2016). Prior to 2004, SOI statistics only included information related to CFCs for large MNCs with more than $500 million in assets. Starting in 2004, the sampling procedure became much broader.

In some instances, we also supplement the IRS datasets with financial statements data on

Table 2: Sample Sizes

	(1) Int'l. Business Sample	(2) SOI Corp. Sample	(3) Stable Sample
MNC Count	23,222	20,029	3,635
CFC Count	333,438	322,538	43,941
FDE Count	58,690	57,685	53,141

Notes: This table provides the size of three different samples of US MNCs, along with their related controlled foreign corporations (CFCs) and 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 were not C corporations. 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.

public companies from Compustat. Table 1 summarizes the tax forms described above along with selected outcomes that we use in our analysis, both from SOI data and from Compustat.

3.2 Sample Construction

Table 2 shows the size of several different samples of MNCs, along with their foreign affiliates (CFCs and FDEs). We consider a firm to be an 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 Sample. Column (2) shows the number of MNCs from Column (1) that are C corporations contained in the SOI Corporate Sample. 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. Our analysis primarily uses the sample shown in Column (3) to study firm-level outcomes. Summary statistics for this stable sample are provided in Table A.1.

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 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 issued as a dividend to CFC A. CFC A 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.[12] To calculate pretax foreign earnings and profits, we obtain pretax 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).

In Appendix B, we examine the performance of this correction. We construct a proxy for aggregation error by generating a link between IRS data, which provide disaggregated information about foreign affiliates, and Compustat, which provides data from public disclosures of consolidated MNCs. By comparing the disaggregated data to the consolidated figures, we 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.[13] Applying the correction proposed

[12]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.

[13]Extractive industries often operate under contracts with foreign governments that include forms of revenue sharing, which can be misreported as a corporate income tax.

by Blouin and Robinson (2020) yields a 30% reduction in the magnitude of foreign earnings as measured in tax data in 2016 and significantly reduces book-tax differences. Furthermore, as shown in Figures B.7 and B.8, 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.[14]

3.4 Measuring Changes in International Corporate Structures

Both the SOI Corporate Sample and the International Business Tax Sample have been used to study domestic and international business taxation. Relatively little work, however, has utilized the wealth of information regarding FDEs.[15] Although data from Form 8858 are collected less frequently than other samples, they allow us to observe two important features of US MNC structures. First, they reveal the date when an entity was first disregarded by an MNC, which allows us to measure adoption of CTB among MNCs. Second, they allow 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 and FDEs that have particular structures associated with the tax planning strategies described in Section 2.

3.4.1 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

[14] In Appendix B.3, we also provide corrected estimates of the elasticity of foreign earnings with respect to foreign tax rates following the methodology of Dowd, Landefeld and Moore (2017). We show that their general finding that earnings are more sensitive to rates in haven jurisdictions is robust to this correction.

[15] 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.

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.

3.4.2 Detecting Reverse Hybrids

The other type of structures we consider are Reverse Hybrid Mismatch arrangements common in the Netherlands and Luxembourg. This arrangement, also described in Section 2, involves a top-level entity that is classified as a partnership and a bottom-level entity that is classified as a 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.

4 Adoption and Prevalence of CTB and Hybrid Tax Planning Structures

Below, we describe how US MNCs used CTB starting 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. 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 generated a majority of foreign earnings among firms in our sample and within these MNCs, a majority of foreign earnings were connected to these structures. We show that these MNCs also represent a large share of domestic corporate

Figure 3: Adoption of Foreign Disregarded Entities

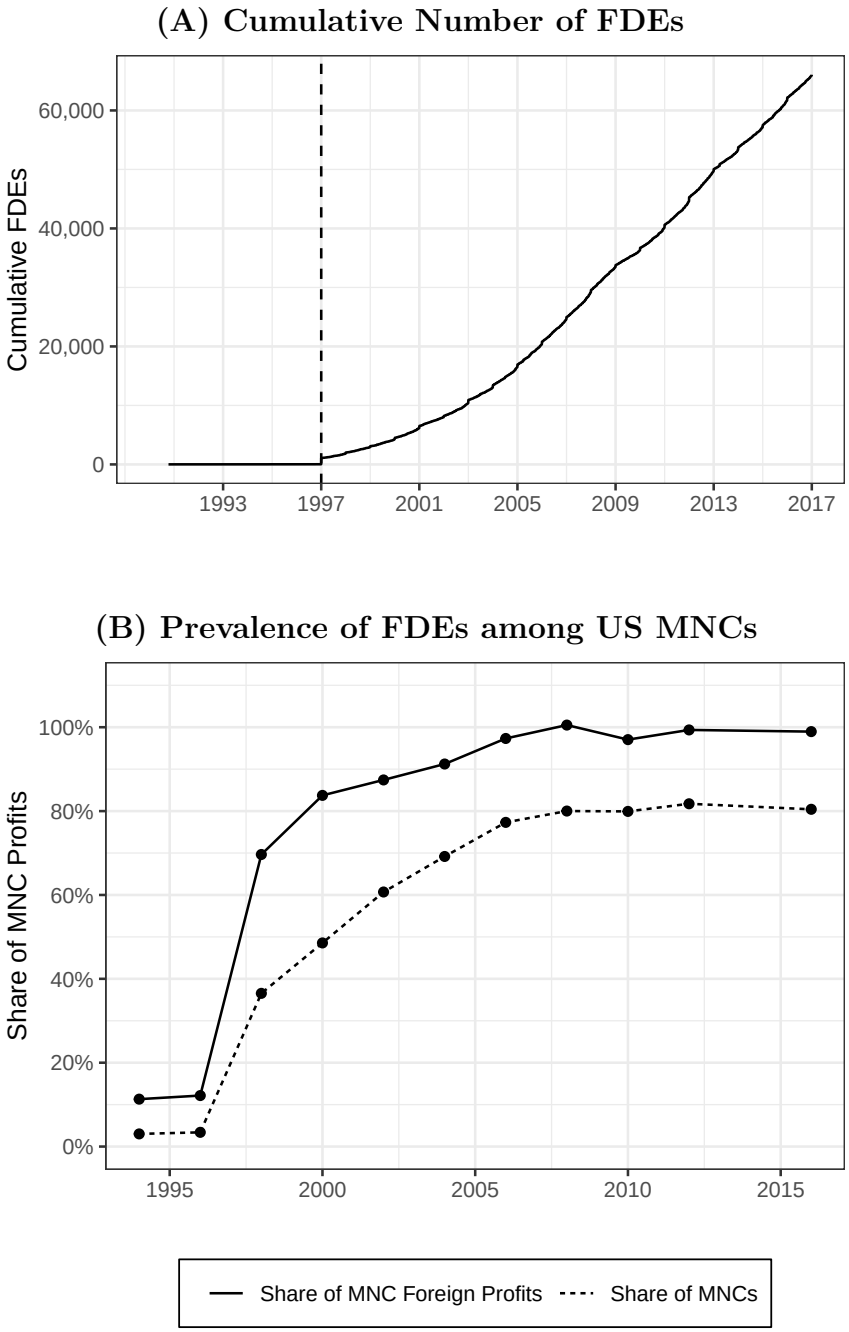

Notes: Panel A of Figure 3 plots the cumulative number of foreign disregarded entities (FDEs). There were fewer than 50 FDEs prior to 1997 and this number grew rapidly following the adoption of CTB regulations. Panel B plots the fraction of US MNCs with an FDE, as well as the share of foreign profits 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 3: Disregarded Entities by Country of Incorporation

Country Name	Unadj. E&P (billions)	Num. FDEs
Ireland	224	4,844
Netherlands	136	12,236
United Kingdom	130	26,982
Switzerland	82	2,138
Cayman Islands	71	5,277
Singapore	68	3,458
Luxembourg	53	4,592
Bermuda	51	2,126
Canada	41	7,427
Australia	34	6,780

Notes: This table shows the largest ten countries by total foreign earnings generated by foreign disregarded entities. Column (2) provides aggregate unadjusted E&P, generated by FDEs in 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.

activity, generating 20% of domestic payroll and holding 15% of domestic capital assets by 2016 among US C corporations.

4.1 Adoption of Check the Box

Panel A of Figure 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. 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 3 shows the largest ten countries according to total earnings generated by FDEs. The Netherlands and Ireland are some of 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.

4.2 Adoption of Hybrid Tax Planning Structures

Panel A of Figure 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 17.5% 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 linked to HTP structures relative to aggregate foreign E&P for MNCs in our sample—by 2016, more than 35% of all foreign E&P of US MNCs was routed through an HTP.

Figure 5 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 accounted 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 decades after adoption.

Figure 4: Adoption of Hybrid Tax Planning (HTP) Structures

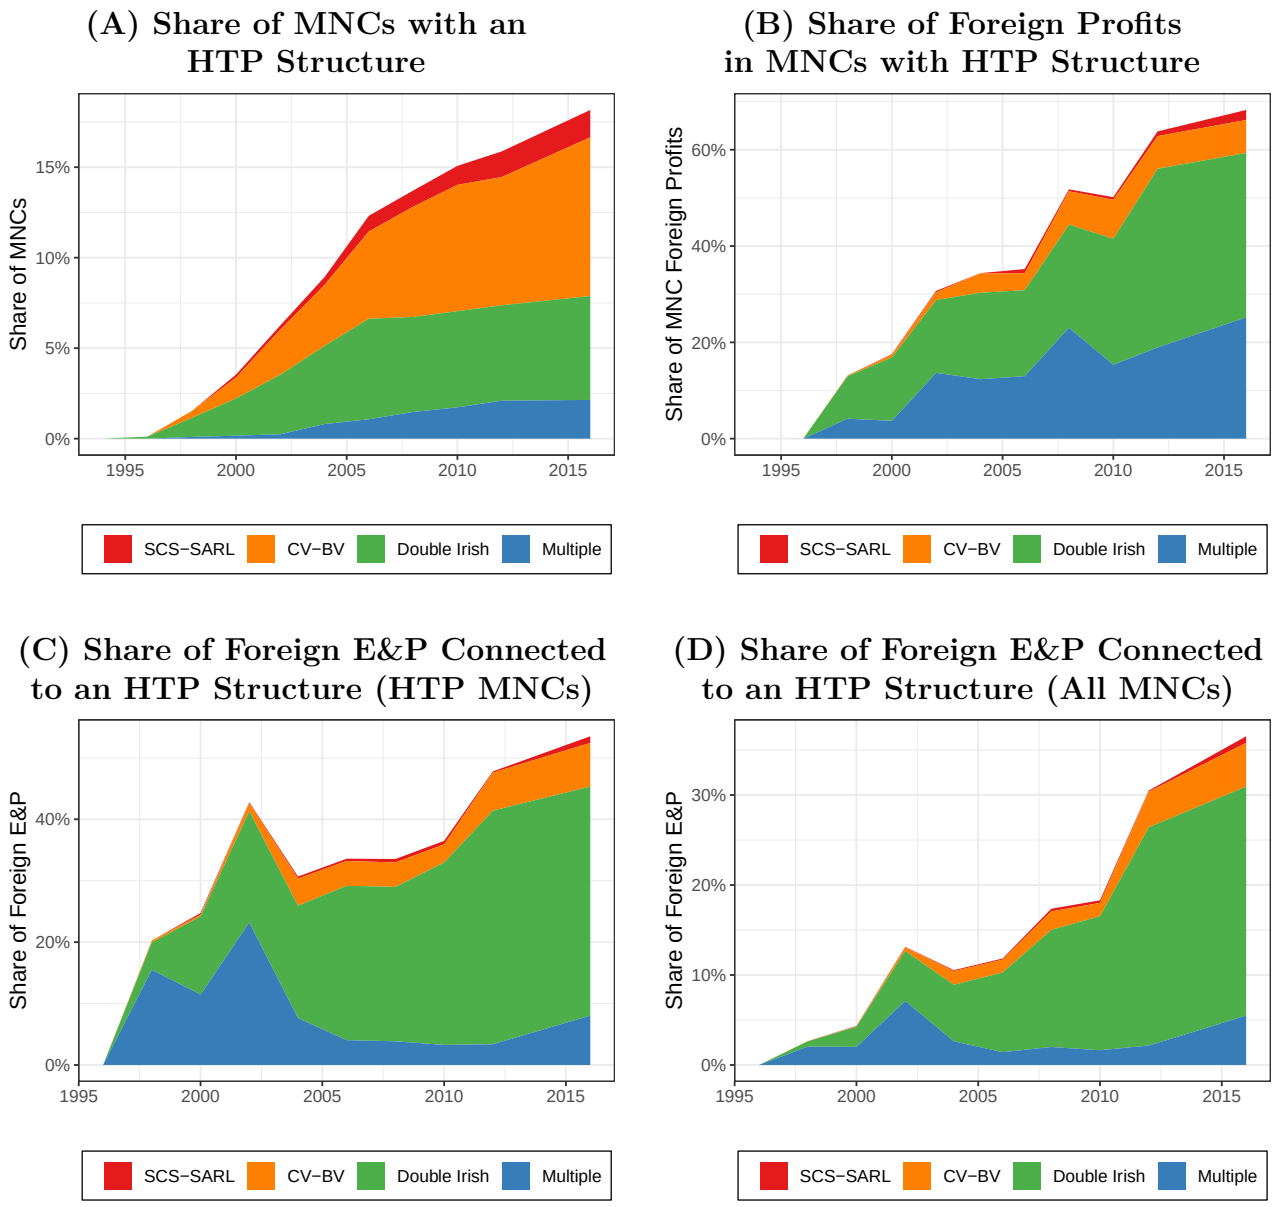

Notes: These figures show that a growing share of US MNCs have adopted hybrid tax structures over time (Panel A) and that these MNCs are responsible for a large share of the overall foreign E&P of US MNCs (Panel B). Panel C shows that a large share of foreign E&P within adopting MNCs is connected to 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 5: Hybrid Tax Planning Structures and Domestic Economic Activity

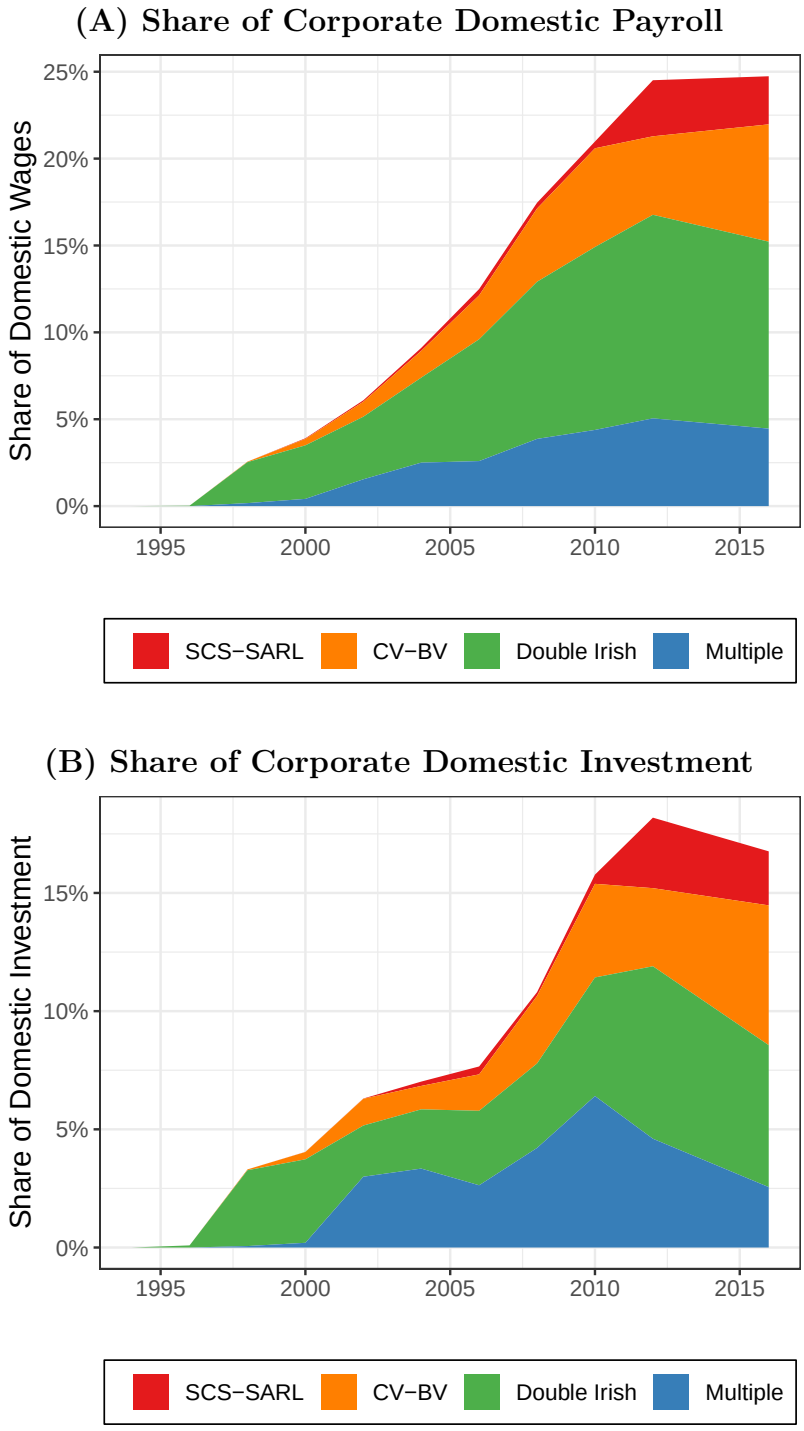

Notes: These figures show the share of domestic wages (Panel A) and domestic capital investment (Panel B) paid by US MNCs that adopt one of the hybrid tax structures described in Section 2. This share is computed as a fraction of all domestic wages and capital investment among C corporations in the IRS Statistics of Income Corporate Sample. See Section 3 and Table 2, Column 2 for a description of the sample.

Figure 6: Predicted Probability of HTP Adoption by Industry

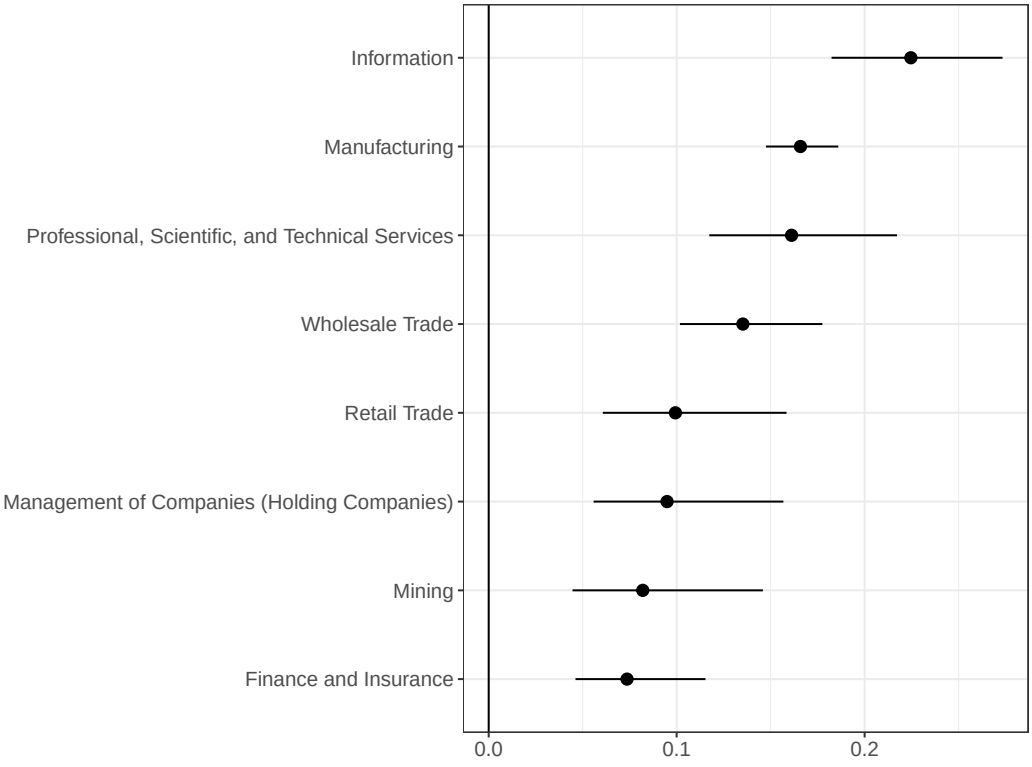

Notes: This figure reports the predicted probability of HTP adoption by sector according to a simple logit model. We remove industries with fewer than ten MNCs in either group.

4.3 Which Firms Adopt Hybrid Tax Planning Structures?

To examine characteristics of MNCs that adopt hybrid structures, we estimate a series of simple logistic regressions that predict HTP adoption and estimate industry shares within adopting and non-adopting groups. These analyses support anecdotes that aggressive tax planning MNCs tend to be larger firms that operate in industries with large amounts of IP. We also examine whether a set of additional characteristics are predictive of adoption and discuss patterns that emerge from this analysis.

We start by examining industry variation in the adoption of hybrid tax planning by sector.[16] Figure 6 reports the predicted probability of adoption by sector using estimates from a simple logit model with industry dummies. There is considerable variation in adoption across different sectors. For example, MNCs that are classified within the Information sector are more than twice as likely to adopt HTP structures compared to MNCs classified under Finance and Insurance.

[16]We use IRS industry classifications that are analogous to 2-digit SIC codes. For consistency, we use the most recently observed industry classification for each MNC across all years.

Figure 7: Sector Shares, HTPs vs. Other MNCs

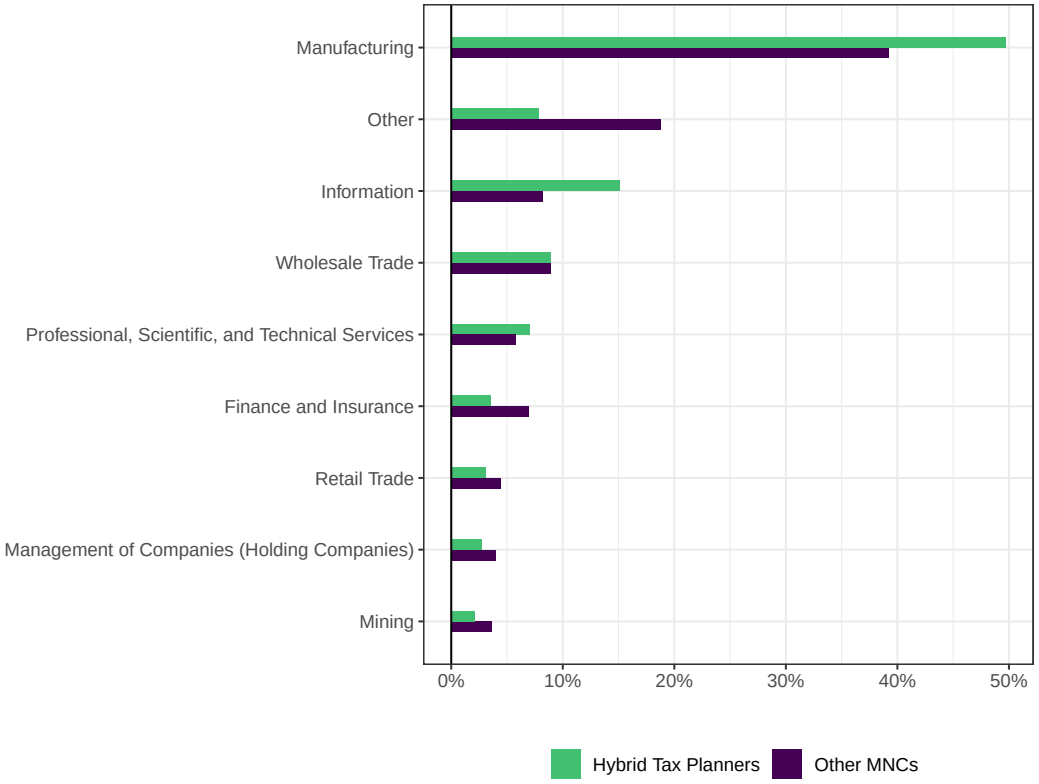

Notes: This figure reports industry shares by sector. Sectors with fewer than 10 firms in either group are collected into the "Other" category.

Figure 7 shows industry shares for MNCs that adopt HTPs and for those that do not. A plurality of MNCs are classifed within the Manufacturing sector. MNCs that adopt HTP structures disproportionately come from the Information and Manufacturing sectors.

Figure 8 provides predicted probability estimates from a logit model that is analagous to Figure 6 but for subindustries.[17] Adoption is stronger within subindustries that contain tech and pharmaceutical firms (e.g., data processing and chemical manufacturing). Other IP-intensive industries, such as publishing, also adopt HTP structures at relatively high rates. Figure 9 displays subindustry shares for adopting and non-adopting MNC groups, providing a similar takeaway.

[17]These IRS subindustry classifications are roughly equivalent to 3-digit SIC codes.

Figure 8: Predicted Probability of HTP Adoption by Subindustry

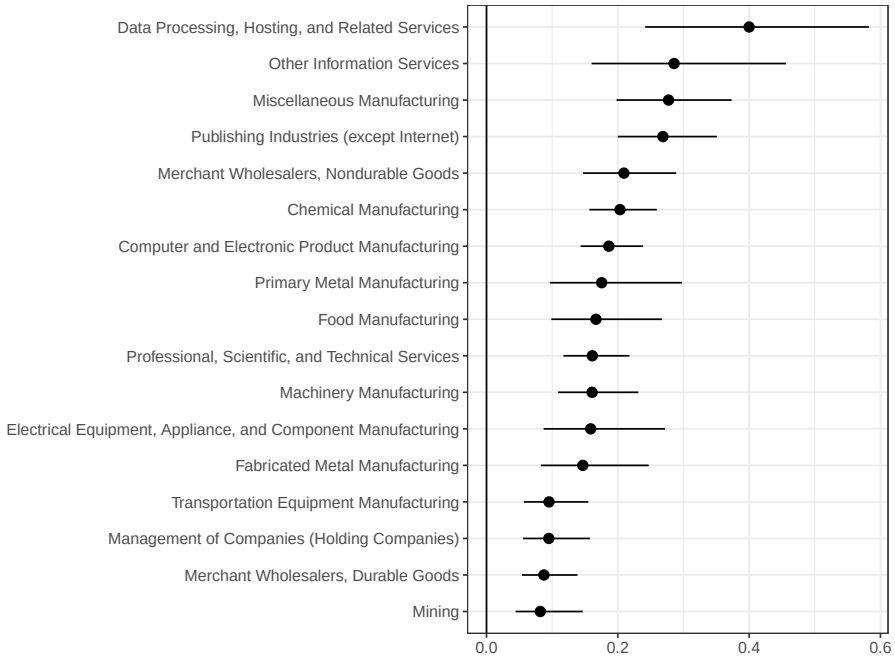

Notes: This figure reports the predicted probability of adoption by major group according to a simple logit model. We remove groupings with fewer than ten MNCs in either group.

Figure 9: Major Shares, HTPs vs. Other MNCs

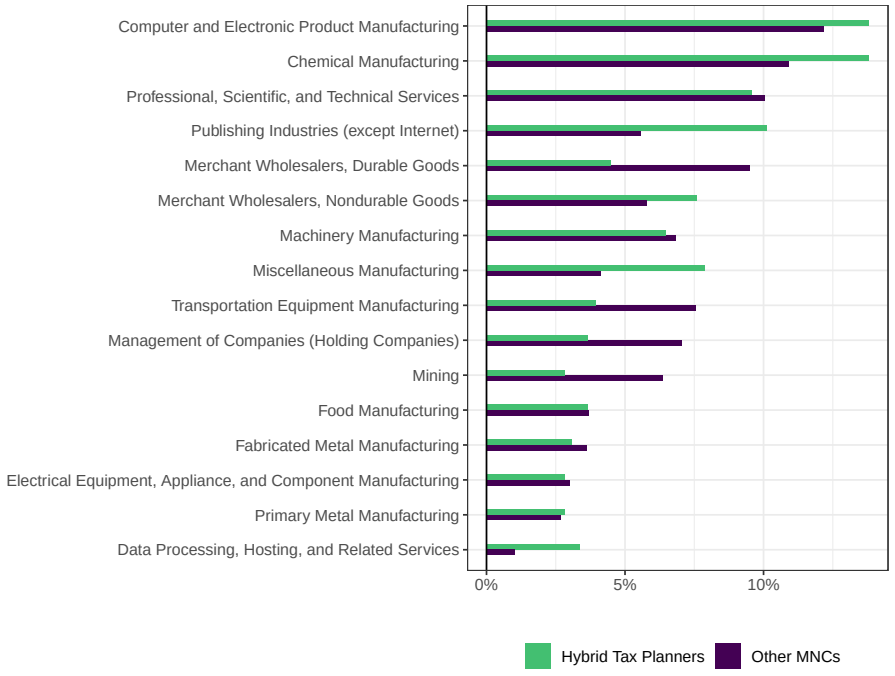

Notes: This figure reports industry shares by major group. We remove groupings with fewer than ten MNCs in either group.

Figure 10: Predicted Probability of HTP Adoption by Firm Size

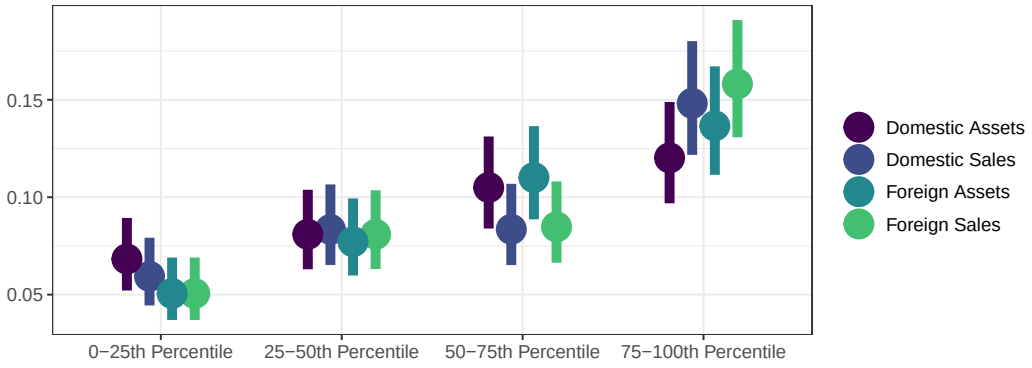

Notes: This figure reports the predicted probability of adoption by size according to four simple logit models that examine four different measures of firm size. We remove industry groupings with fewer than ten MNCs in either group.

Figure 10 provides estimates of the predicted probability of HTP adoption according to several different measures of MNC size. Each color provides estimates computed from a separate logit model corresponding to a different measure of size (domestic assets, domestic sales, foreign assets, and foreign sales). We bin the MNCs into size quartiles. All of these measures indicate that larger MNCs tend to adopt HTP structures at higher rates.

In addition to industry and firm size, we examine whether a set of other observable characteristics is predictive of HTP adoption. Table A.4 reports coefficient estimates for a battery of logit regressions that predict HTP adoption based on whether an MNC claims a tax credit for R&D; its age (binned by quartile with the youngest firms set as the reference category); the average statutory foreign ETR that it faces (and its share of foreign sales in jurisdictions with unobserved statutory rates); a measure of its geographic exposure to Check the Box; whether it operated in in Ireland, the Netherlands, or Luxembourg prior to adoption; whether it has negative domestic earnings; and its advertising to sales ratio (a proxy for intangibles used in, e.g., Grubert and Slemrod, 1998).[18] These models use observations from adopting MNCs in the period prior to

[18]To compute a firm-level measure of exposure to Check the Box, we first compute a country-level measure of exposure $\gamma_c = \pi_c^p/\pi_c$, where π_c^p are aggregate foreign earnings generated by pass-through foreign affiliates (FDEs) for country c, and π_c are aggregate foreign earnings for all foreign affiliates in country c for the years that we observe FDE earnings (2006, 2008, 2012, and 2016). Next, for each MNC i we compute the share of foreign sales by country s_{ic} in the period prior to adoption. Finally, for each firm, we compute the exposure measure as a weighted average of the country-level exposure measures, $\gamma_i = \sum_{c \in C_i} \gamma_c s_{ic}$ for the set of countries C_i where MNC i has positive sales in the period prior to adoption, where the weights s_{ic} are the firm-level country shares calculated in the second step. For firms that never adopt, we compute this exposure measure in every year.

adoption, and include all observations for never-adopters. Table A.5 combines the variables from these regressions into a single logit regression. Both tables provide specifications with year fixed effects only and with interactions between industry, sales quartile bins, asset quartile bins, and year fixed effects.[19] Finally, Table A.6 provides logit estimates from a set of MNCs that also appear in Compustat to examine whether the identity of firms' auditors plays a role in adoption of hybrid structures.[20]

Of the characteristics listed above, only a few appear to be predictive of HTP adoption. Unsurprisingly, MNCs that previously operated in the jurisdictions where we detect HTP structures (Ireland, the Netherlands, and Luxembourg) are more likely to adopt them. More generally, the geographic distribution of MNC activity is predictive of adoption—MNCs that operate in countries where Check the Box is more heavily used may be able to shift income more easily, which may explain why this measure of exposure is predictive of adoption. Finally, MNCs that are in a domestic loss position are less likely to adopt HTP structures—this also has an intuitive explanation as these firms may have tax credits that offset income tax levied on repatriated foreign earnings. Confidence intervals for these estimates are generally wider in the combined regression, and when more granular fixed effects are included. These tables show that, even though some characteristics are predictive of adopting an HTP, most variables are not statistically significant after the inclusion of industry and size fixed effects. While some variables statistically correlate with adoption, it is hard to predict which specific firms will adopt an HTP, even conditional on a rich set of covariates.

Overall, the results in this section confirm the conventional wisdom that larger firms in industries that rely on intellectual property are more likely to adopt HTPs. At the same time, there remains considerable unexplained variation in HTP adoption even after controlling for a large set of firm characteristics that have been used as proxies for tax planning, indicating the existence of idiosyncratic costs and benefits of adoption that vary across firms.

[19]Note that the number of observations goes down in the second Column (from 8,608 to 4,920). This is because to be included, each bin must have both HTP adopters and non-adopters.

[20]We divide the firms into three groups, multinationals with a Big 4 auditor (Ernst & Young, PWC, Deloitte, and KPMG), medium size auditors, and small auditors. These results suggest that auditors do not play a significant role in the adoption of HTP structures.

5 Hybrid Tax Planning Structures and Multinational Activity

Below, we examine changes in domestic and foreign economic activity of adopting MNCs. We start with a set of descriptive facts that compare select aggregate outcomes for MNCs that utilize HTPs to other MNCs that do not rely on these arrangements. We then estimate staggered difference-in-differences models that control for firm characteristics in order to compare outcomes for MNCs that do and do not adopt hybrid tax structures. These models allow us to tie changes in firm activity to the timing of HTP adoption.

5.1 Comparison of Aggregate Trends

While the structures we study have been suspected of being used for profit shifting, lack of tax data prevented prior researchers from confirming this role in a systematic manner. We therefore start by examining whether hybrid tax planning firms engage in the kinds of transactions that are associated with profit shifting. Figure 11 compares hybrid tax planners to other MNCs along a number of these dimensions. For a given outcome, this and related figures plot selected outcomes as a percentage of 1996 levels. The secondary y-axis to the right of each graph indicates the level values for the group of hybrid tax planning firms. In Panel A, we first document that hybrid tax planning firms generate a much larger aggregate loan balance between their related CFCs when compared to other MNCs. These balances may be related to interest stripping strategies, as discussed in Panel A of Figure 2.

Panel B of Figure 11 shows that hybrid tax planning MNCs experience a much faster rise in the book value of foreign intangible assets when compared to non-HTP MNCs. This growth is consistent with the use of HTP structures to shift income generated by intangible assets. Panel C shows that these MNCs also increased compensation for services paid by CFCs to parent companies—which includes cost sharing payments that are used to transfer intangible assets from the US to foreign affiliates. Starting in 2008, Schedule G of Form 5471 allows us to 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 1 compares average foreign ETRs of hybrid tax planners with that of 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

Figure 11: Mechanisms for Profit Shifting

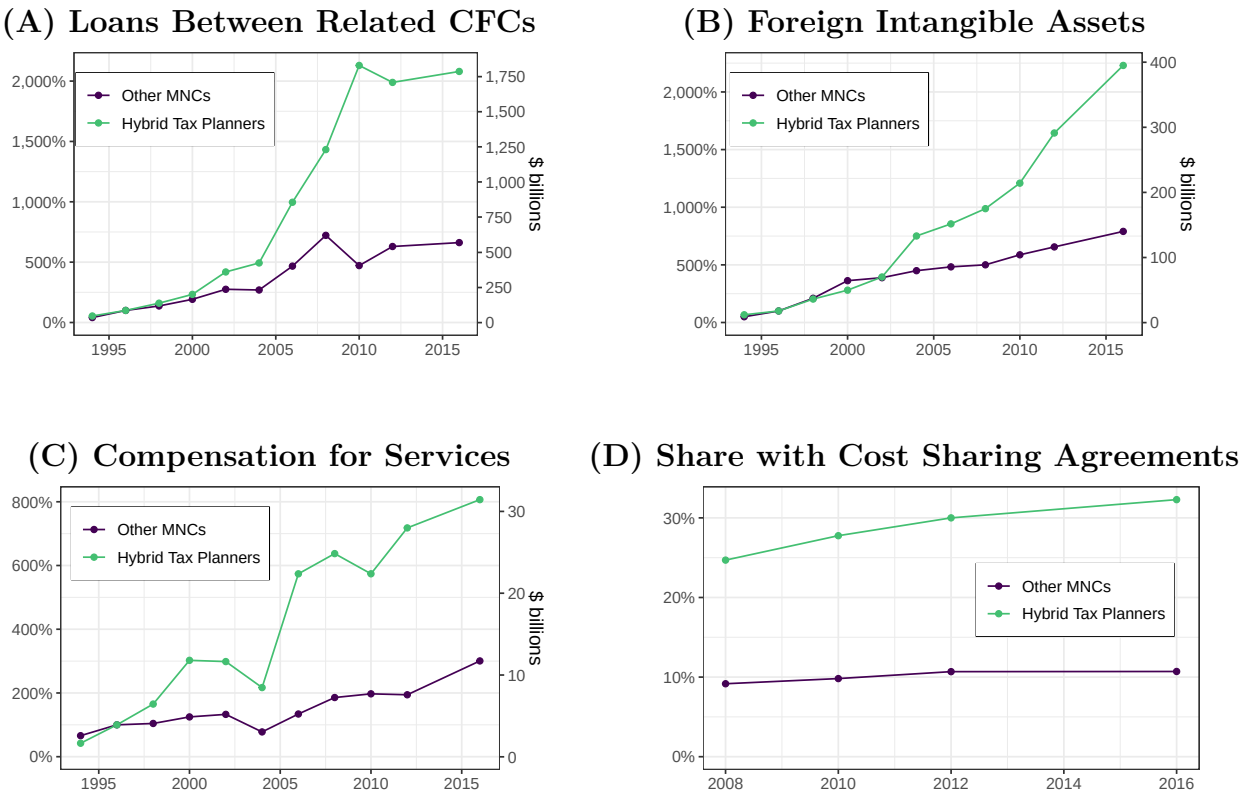

Notes: These figures show the evolution of aggregate loans between related CFCs (Panel A), foreign intangible assets (Panel B), payments from CFCs to US parent companies for technical services (Panel C), and the share of MNCs in each group with active cost sharing agreements with a CFC (Panel D). For 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. The dark blue line shows aggregate values for MNCs that did not adopt any of these structures during the sample window. For most outcomes, the right-hand axis displays dollar value in billions, relative to the aggregate 1996 dollar value for HTP MNCs.

each group's ETR. By 2016, hybrid tax planners faced a foreign ETR that was about half of that incurred by other MNCs. Figure 11 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 11 shows that MNCs with the largest reduction in foreign ETRs were also engaging in behavior that has been linked to aggressive tax planning. Indeed, as we see in Figure 1, firms that did not adopt HTP structures experienced a much smaller decline in their foreign ETR 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 A of Figure 12 shows that, relative to foreign E&P, hybrid tax planning firms saw faster declines in royalty

Figure 12: Evidence of Deferral

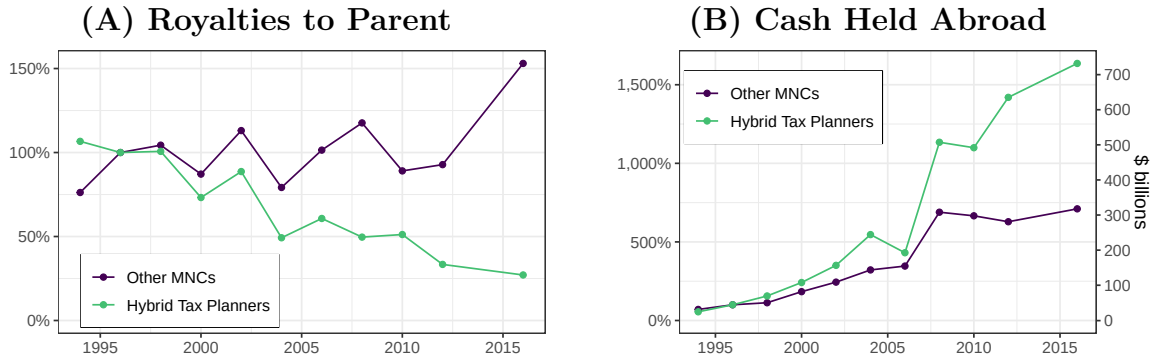

Notes: These figures show the evolution of royalty payments from CFCs to domestic parent entities as a share of foreign E&P (Panel A) and of aggregate foreign cash balances (Panel B) for two groups of US MNCs. For comparability, aggregate values for both groups are normalized to 100% as of 1996 in Panel 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. The dark blue line shows aggregate values for MNCs that do not adopt any of these structures during the sample window. For most outcomes, the right-hand axis displays dollar value in billions, relative to the aggregate 1996 dollar value for HTP MNCs.

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 this growth occurred among the 300 firms that we observe with hybrid tax planning agreements through Ireland, Netherlands, and Luxembourg.

The results in Figures 11 and 12 provide *prima facie* evidence of the specific mechanisms through which hybrid tax planning strategies operate. The ability to shift profits to lower-tax countries allows MNCs to avoid foreign income tax and defer US income tax.

These figures also highlight the value of using tax data, because they allow us to (i) identify the adoption of specific tax planning structures (Figure 4), (ii) link the adoption of HTP structures to specific profit shifting mechanisms (Figures 11 and 12), and (iii) measure the associated impact on ETRs (Figure 1).

We now examine whether firms that benefited from tax planning also experienced differential evolution in their real operations. Figure 13 shows that, while hybrid tax planning firms and non-HTP MNCs had similar patterns of economic activity prior to 1997, their economic activities diverged over the same time period that hybrid tax planning strategies were adopted. Panel A shows that hybrid tax planners had larger increases in domestic capital assets; Panel B shows

Figure 13: Hybrid Tax Planning and Real Economic Activity

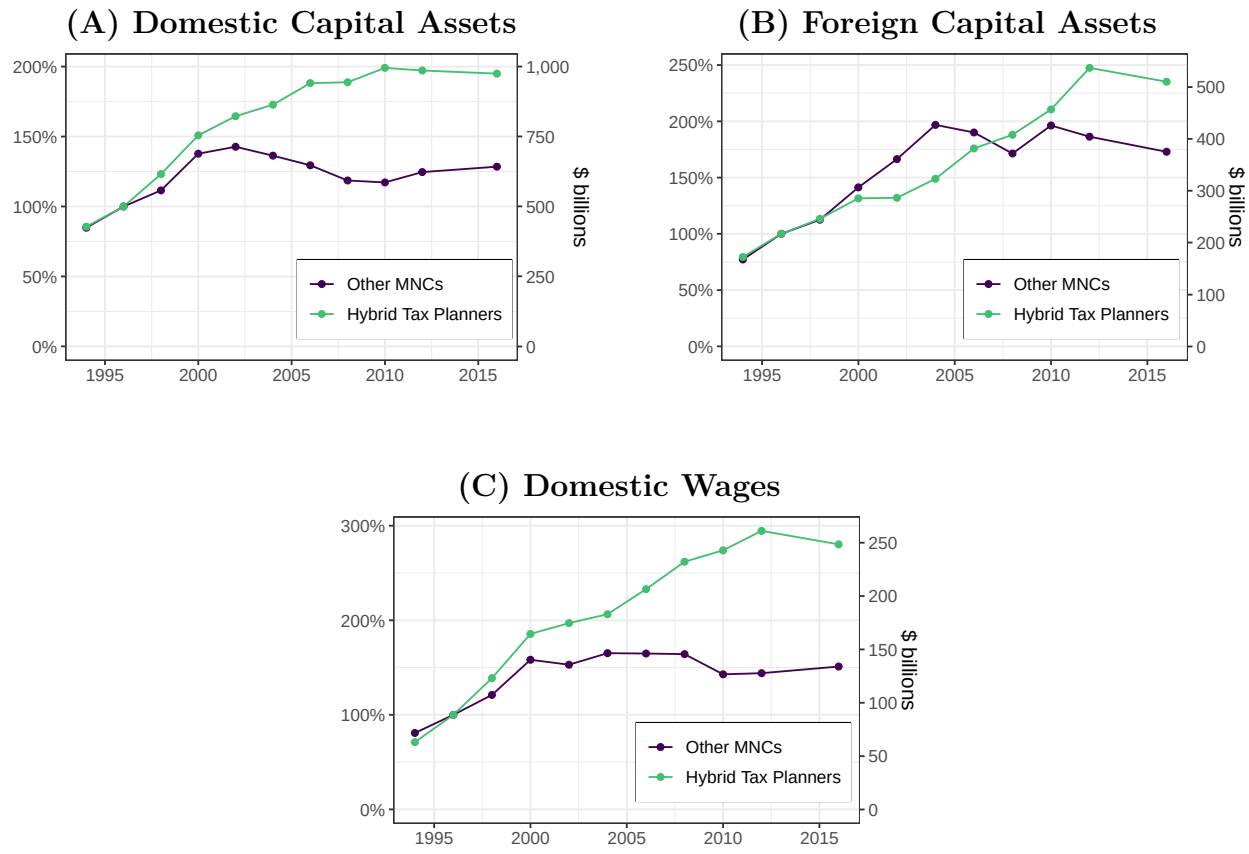

Notes: These figures show the evolution of aggregate domestic capital assets (Panel A), foreign capital assets (Panel B), and domestic wages (Panel C) for two groups of US MNCs. For comparability, aggregate values for both groups are normalized to 100% as of 1996. The light green line shows aggregate values for a group of MNCs that eventually adopt at least one of the hybrid tax structures described in Section 2. The dark blue line shows aggregate values for MNCs that do not adopt any of these structures during the sample window. The right-hand axis displays dollar value in billions relative to the aggregate 1996 dollar value for HTP MNCs.

that these firms also accumulated more foreign capital assets; and Panel C shows larger increases in domestic payroll. Across all of these measures, declines in foreign ETRs were accompanied by increases in foreign and domestic economic activity.

5.2 Estimating Staggered Difference-in-Differences Models

We now show that the changes in firm outcomes described above are closely tied to the adoption of HTPs. Recent literature has provided several alternative models that researchers may use to produce difference-in-differences estimates in staggered contexts. We provide estimates for three of these models. Our main specification relies on the "stacked" design from Cengiz, Dube, Lindner and Zipperer (2019). We choose this estimator as our main specification because it lends

itself easily to the addition of propensity score weights, which we describe later in this section.

The stacked design creates a data set for each cohort. This dataset includes MNCs that adopt a hybrid structure, as well as MNCs that do not adopt a hybrid structure in the six years before and after the adoption year c. Non-adopting MNCs may therefore be repeated in the regression dataset as comparison units for different cohorts. Formally, the stacked design estimates the regression equation,

$$Y_{ict} = \alpha_{ic} + \lambda_{ct} + \sum_{\ell} \mu_{\ell} \mathbf{1}\left\{t - c = \ell\right\} + v_{ict}, \tag{1}$$

where i indexes MNCs, c is the year in which a particular cohort first adopts a hybrid structure, and t indexes years. ℓ is an indicator for the relative number of periods after MNC i adopts a foreign tax planning structure.[21] α_{ic} and λ_{ct} are MNC-by-cohort and year-by-cohort fixed effects. We estimate this regression for various outcomes Y_{ict}. In all specifications, we cluster standard errors at the MNC level.

We interpret results of Equation 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-adoption 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.[22] As a robustness check, we also estimate effects using the stacked specification with inverse probability weights (IPW) as well as an alternative specification proposed by Sun and Abraham (2021) and a standard TWFE estimator.

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 estimates of Equation 1 can be interpreted as causal effects of HTPs. The usual assumptions for a causal interpretation

[21] For comparison, a standard two-way fixed effects specification does not use repeated comparison units and instead estimates the regression equation

$$Y_{it} = \alpha_i + \lambda_t + \sum_{\ell} \mu_{\ell} \mathbf{1}\left\{t - c = \ell\right\} + v_{it}. \tag{2}$$

[22] To be precise, this implies an augmented version of Equation 1 where λ_{ct} is replaced by $\sum_{g \in G} \lambda_{gct}$, where G is a set of groups for which we include group-by-cohort-by-year fixed effects.

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. MNCs would therefore have to adjust behavior two years in advance for this form of bias to be present. Because the parallel trends assumption is inherently untestable and that HTP adoption is an endogenous choice of the firm, it is important to consider that MNCs may select into HTPs because they have more to gain from tax planning.

5.2.1 Estimates of Changes in Financial and Tax Outcomes

Figure 14 plots estimates of Equation 1 for 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-cohort-by-industry and year-by-cohort-by-group fixed effects, where groups include domestic and foreign sales quartiles and domestic and foreign intangible asset quartiles. Quartiles are computed using the period prior to adoption for each cohort.

Panel A provides estimates for the log of the balance of loans between CFCs, Panel B for the log of intangibles held abroad, and Panel C for the log of cash held abroad. Across these three outcomes, we observe similar trends for HTP-adopting and non-adopting MNCs prior to the adoption of a hybrid structure followed by relative increases for HTP-adopting MNCs after the period of adoption. Consistent with these mechanisms and the result of Figure 1, Panel D shows that the foreign ETRs of HTP-adopting MNCs gradually declined relative to non-adopting MNCs following adoption. Six years after adoption, MNCs experience a reduction in their foreign ETR of between three and four percentage points. For all of these outcomes, we find that inclusion of size-bin-by-cohort-by-year fixed effects and industry-by-cohort-by-year fixed effects does not significantly impact the estimates.

Table 4 provides aggregated estimates using a similar specification that replaces relative time dummies with pre and post dummies. The estimates in Column (1) indicate that loans between CFCs increased by 40%, foreign intangible assets increased by 68%, and cash held abroad increased by 56% on average relative to the period prior to adoption. We also estimate

Figure 14: Profit Shifting Mechanisms, Deferral, and Foreign ETRs

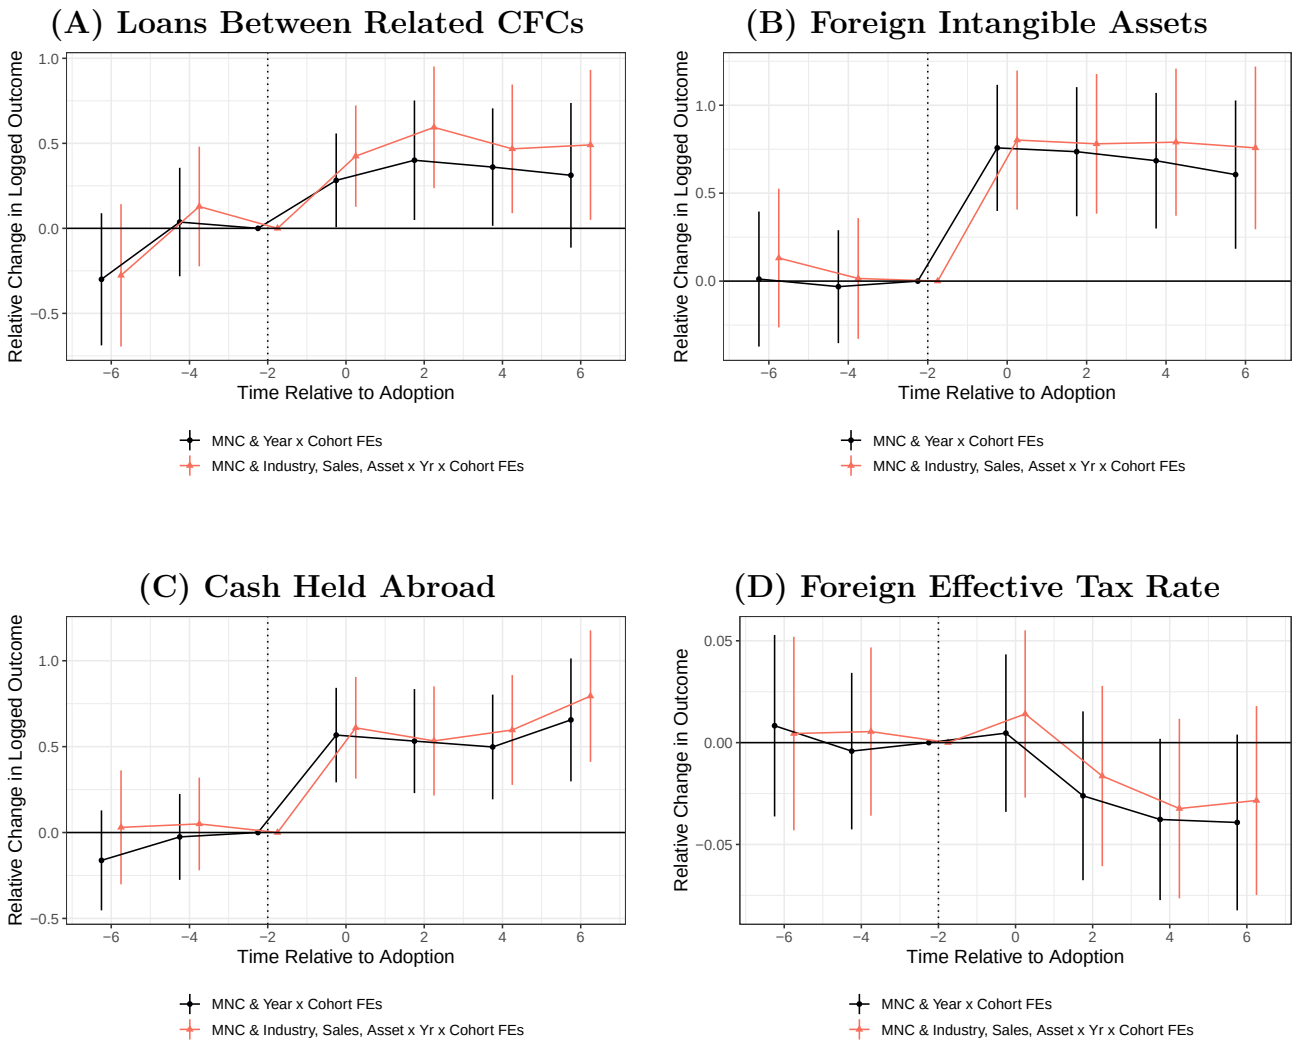

Notes: This figure provides estimates of μ_ℓ from Equation 1 for the corresponding foreign outcome listed in each panel, using 2-year pooled average data from Compustat and dropping odd years to match the IRS SOI International Business Tax Sample. R&D Intensity (Panel A) is calculated as the ratio of annual R&D to the MNC's most recent sales value pre-adoption, and is restricted to be less than 1. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-cohort-by-industry fixed effects and year-by-cohort-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-adoption values for each cohort.

an average decline in foreign ETRs of 3.9 percentage points in the period after adoption. Column (2) shows that these estimates are stable across specifications that include interactions between bins of firm size and year fixed effects, bins of intangible assets interacted with year fixed effects, and industry-by-year fixed effects, suggesting that our results 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 4: Profit Shifting Mechanisms and Foreign ETRs

	(1)	(2)	(3)	(4)	(5)
Panel A					
Foreign Intangibles	0.681***	0.735***	0.637***	0.590***	0.709***
	(0.172)	(0.177)	(0.193)	(0.171)	(0.175)
Num. Treated	248	248	208	250	229
Num. Control	1524	1490	1221	1532	1757
Panel B					
Foreign Cash	0.563***	0.555***	0.469**	0.539**	0.472**
	(0.133)	(0.136)	(0.149)	(0.173)	(0.144)
Num. Treated	257	257	214	257	252
Num. Control	2054	2000	1518	2023	2037
Panel C					
Rltd. CFC Loans	0.401*	0.524**	0.505**	0.344+	0.491**
	(0.157)	(0.161)	(0.174)	(0.176)	(0.163)
Num. Treated	233	233	197	238	205
Num. Control	1263	1240	1041	1294	1597
Panel D					
Foreign ETR	-0.039**	-0.032*	-0.024	-0.037*	-0.042**
	(0.015)	(0.016)	(0.018)	(0.017)	(0.015)
Num. Treated	257	257	214	257	252
Num. Control	2100	2043	1529	2046	2042
MNC & Year x Cohort FEs	Yes	-	-	-	-
MNC & Ind., Size x Yr FEs	-	-	-	Yes	Yes
MNC & Ind., Size x Yr x Cohort FEs	-	Yes	Yes	-	-
Inverse Prob. Weights	-	-	Yes	-	-
Model	Stacked	Stacked	Stacked	SA	TWFE

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

Notes: This table provides estimates of the difference-in-differences model discussed in Section 5.2 for the corresponding foreign outcome listed in each panel, where Columns (1) - (3) use Equation 1 and Columns (4) and (5) use Equation 2. Column (1) does not include additional controls. Columns (2) - (5) include year-by-cohort-by-industry and year-by-cohort-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-adoption values. Column (3) uses inverse probability-weighted data. Column (4) estimates an alternative specification from Sun and Abraham (2021). Column (5) estimates a standard TWFE specification.

We conduct an additional exercise to provide an alternative summary of the average change in the foreign effective tax rate to reduce potential measurement error. In contrast to the estimates above, which use year-by-year foreign ETRs, Table 5 shows pooled estimates where ETRs are

Table 5: Pooled Foreign ETRs

	(1)	(2)
Unweighted Foreign ETR	−0.035+	−0.032*
	(0.021)	(0.016)
Weighted Foreign ETR	−0.092**	−0.061*
	(0.034)	(0.028)
Num. Treated	155	83
Num. Control	1101	308
Sample	SOI	Compustat
Outcomes	SOI	Compustat
Years Included	Even	All

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

Notes: This table estimates a difference-in-differences model using a two-period specification. Foreign effective tax rates are computed using aggregate taxes and earnings from each period to reduce measurement error relative to an annual specification. Column (1) provides estimates using measures from SOI data. Column (2) provides estimates using measures from Compustat. Results are provided for an unweighted specifciation as well as a weighted specification where firms are weighted by aggregate pre-period earnings. Firms are only included if they have positive earnings in both pre and post periods.

computed at the firm level for all pre and post periods. We include firms that have positive aggregate earnings in both periods.[23] This reduces year-to-year volatility in the firm-level ETR that can be generated by losses and tax credits. We winsorize these rates to ensure they do not exceed 100% so that outliers do not skew the average. Finally, we compute the average foreign ETR using two different measures. The first measure, which is also used for Figure 14 above, uses SOI tax data to compute the foreign ETR following the methodology described in Section 3.3. The second measure uses foreign taxes and pretax income as reported in Compustat following the method used to calculate foreign effective tax rates in Dyreng, Hanlon, Maydew and Thornock (2017). There are advantages and disadvantages to both measures—SOI data provide a larger sample of multinationals, but only even years are present in our data. Compustat has smaller coverage, excluding private firms, but these firms are observed annually. Finally, we include two specifications for each measure—the first is an unweighted regression and the second weights firms by aggregate pre-period foreign income. Results are shown in Table 5. The unweighted pooled results are broadly similar to the event study estimates, showing declines of between 3.2 and

[23]We use the stacked panel in this analysis in order to create distinct pre and post periods for the comparison group of firms that do not adopt hybrid structures.

3.5 percentage points among firms that adopt hybrid structures. Weighted results demonstrate much stronger reductions in foreign ETRs after adoption of hybrid structures, between 6.1 and 9.2 percentage points depending on the measure used. These results suggest that although firms generally experienced reductions in foreign tax rates after adopting hybrid structures, these reductions were concentrated in the largest adopting firms.

5.2.2 Estimates of Changes in Real Economic Activity

To measure changes in real activity, we first estimate Equation 1 for a set of outcomes that includes foreign assets, domestic wages, and domestic capital assets. Foreign assets are reported on IRS Form 5471, domestic wages and capital assets are reported on Form 1120, and domestic investment is measured as the sum of reported assets placed into service on Form 4562. Figure 15 reports the results of the event study analyses for these outcomes. As with the financial and tax outcomes, estimates are not significantly different from zero in the pre-adoption period for 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 and domestic wages.

Table 6 summarizes these estimates using an aggregated specification. The specification with more granular controls (Column (2)) reports an increase in domestic investment of 27% and a corresponding increase in domestic capital assets of 16%. Panel C reports a 40% increase in foreign capital assets, while Panel D reports a 20% increase in domestic payroll.

Figure 15: Event Studies: Hybrid Tax Planning and Real Economic Activity

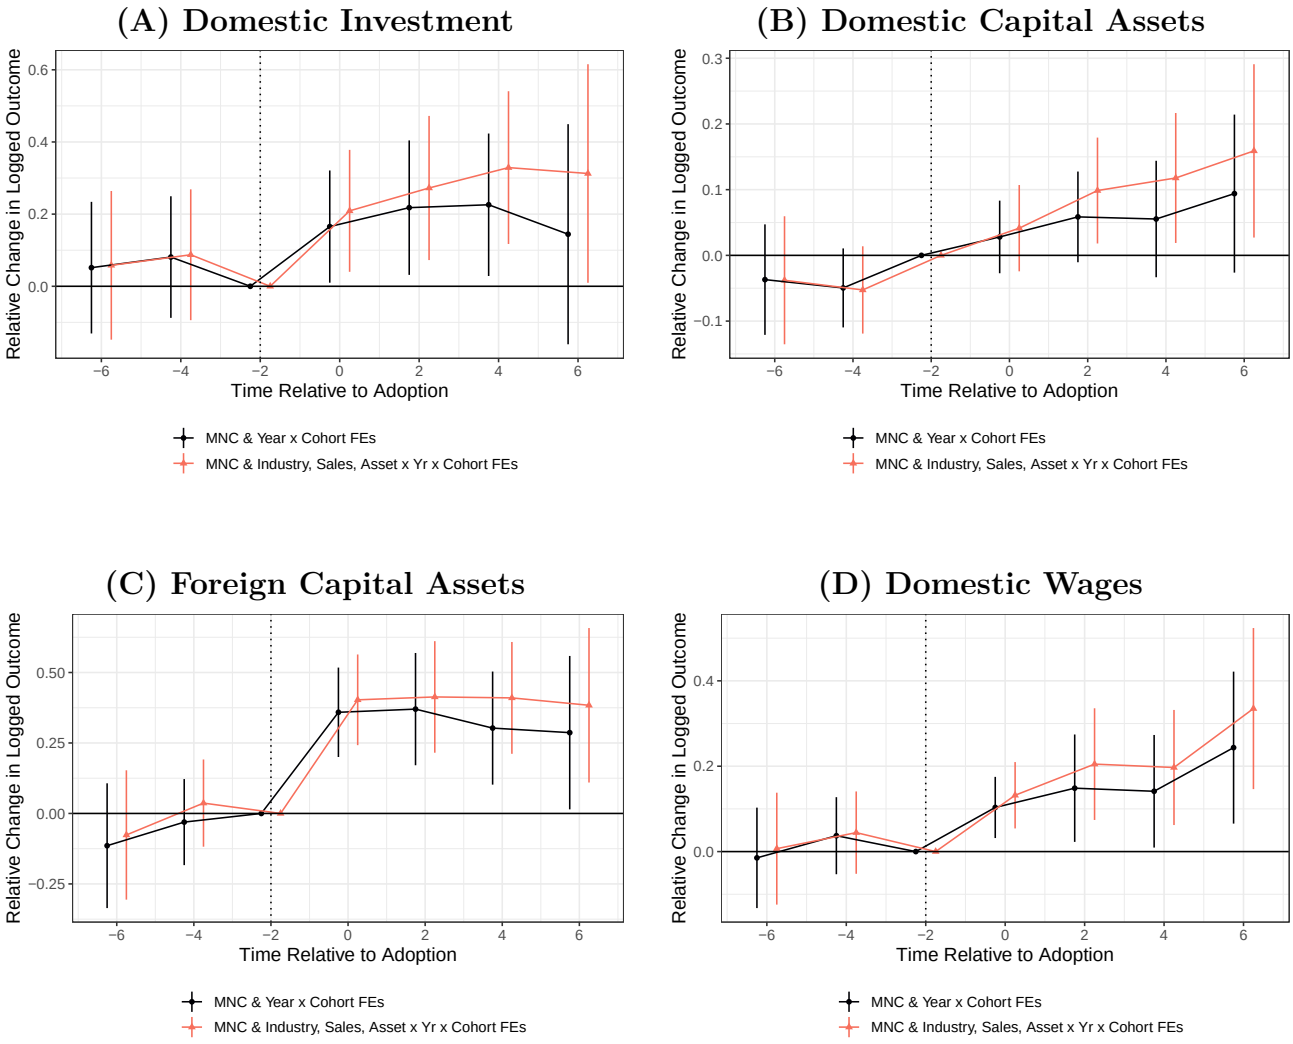

Notes: This figure provides estimates of μ_ℓ from Equation 1 for the corresponding outcome listed in each panel. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-cohort-by-industry fixed effects and year-by-cohort-by-group fixed effects, where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

Table 6: Hybrid Tax Planning and Real Economic Activity

	(1)	(2)	(3)	(4)	(5)
Panel A					
Domestic Investment	0.170*	0.267**	0.307**	0.285**	0.206*
	(0.079)	(0.088)	(0.109)	(0.104)	(0.086)
Num. Treated	241	241	209	242	240
Num. Control	1802	1756	1472	1798	1919
Panel B					
Domestic Capital	0.098*	0.156**	0.147**	0.076	0.105*
	(0.046)	(0.051)	(0.056)	(0.048)	(0.049)
Num. Treated	250	250	214	252	247
Num. Control	1882	1828	1528	1856	1964
Panel C					
Foreign Capital	0.351***	0.404***	0.377***	0.340**	0.366***
	(0.101)	(0.100)	(0.106)	(0.111)	(0.105)
Num. Treated	257	257	214	257	247
Num. Control	1906	1852	1492	1866	1978
Panel D					
Domestic Wages	0.150**	0.204***	0.182**	0.186*	0.148*
	(0.055)	(0.059)	(0.062)	(0.078)	(0.060)
Num. Treated	250	250	214	252	247
Num. Control	1872	1819	1519	1850	1961
MNC & Year x Cohort FEs	Yes	-	-	-	-
MNC & Ind., Size x Yr FEs	-	-	-	Yes	Yes
MNC & Ind., Size x Yr x Cohort FEs	-	Yes	Yes	-	-
Inverse Prob. Weights	-	-	Yes	-	-
Model	Stacked	Stacked	Stacked	SA	TWFE

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

Notes: This table provides estimates of the difference-in-differences model discussed in Section 5.2 for the corresponding outcome listed in each panel, where Columns (1) - (3) use Equation 1 and Columns (4) and (5) use Equation 2. Column (1) does not include additional controls. Columns (2) - (5) include year-by-cohort-by-industry and year-by-cohort-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-adoption values. Column (3) uses inverse probability-weighted data. Column (4) estimates an alternative specification from Sun and Abraham (2021). Column (5) estimates a standard TWFE specification.

Figure 16: Event Studies: Hybrid Tax Planning and R&D Data from Compustat

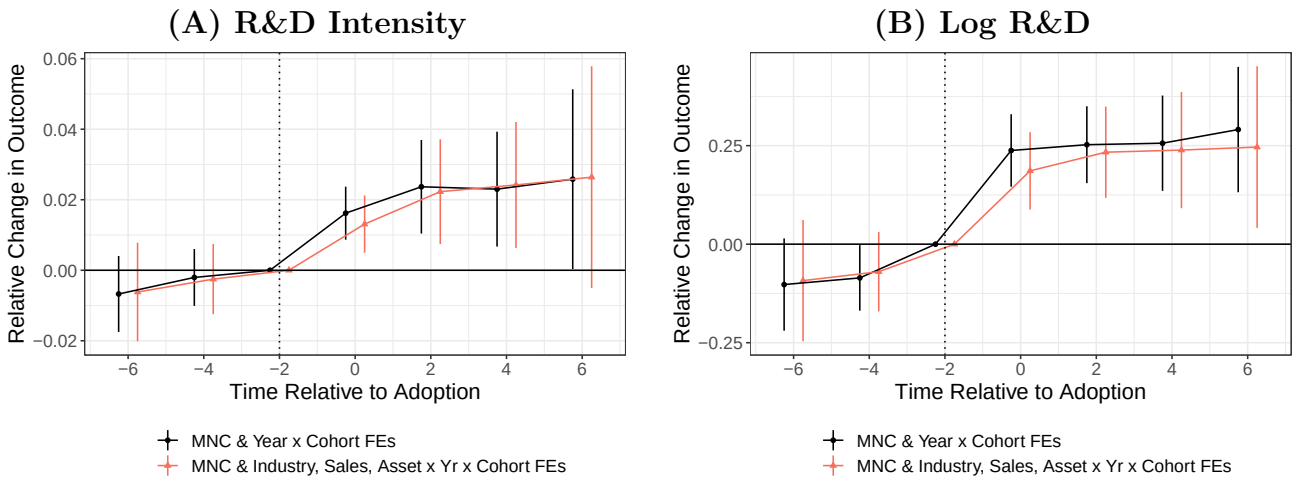

Notes: This figure provides estimates of μ_ℓ from Equation 1 for the corresponding R&D outcome listed in each panel. R&D Intensity (Panel A) is calculated as the ratio of annual R&D to the MNC's most recent sales value pre-adoption, and is restricted to be less than 1. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-cohort-by-industry fixed effects and year-by-cohort-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-adoption values for each cohort.

5.2.3 R&D Activity

To examine R&D activity, we merge in data from Compustat, which gathers R&D expenditures from MNCs' consolidated financial statements. Because this procedure introduces sampling attrition, as not all MNCs in the tax data are public, we separate this analysis from the estimates contained above that rely on administrative tax data.

In order to match the biannual frequency of observations in the SOI sample, we create a two-year pooled average for each Compustat variable of interest and focus the estimation on the same years available in the tax data. We study two measures of R&D using the same event study model discussed above. The first measure, R&D intensity, is computed as the ratio of R&D to revenue. We fix the denominator to the period prior to HTP adoption. A valuable feature of this measure is that it does not exclude firm-years in which zero R&D expense is reported. The second measure is the log of R&D. This measure leads to further sample attrition because of the large number of firm-years that report zero R&D expenditures.

The results of these estimations are shown in Figure 16. As with other measures of real economic activity, this figure shows that adopting firms experience a significant increase in both R&D intensity and log R&D following HTP adoption. These results are summarized in Table 7.

Table 7: Hybrid Tax Planning and R&D Data from Compustat

	(1)	(2)	(3)	(4)	(5)
Panel A					
R&D Intensity	0.026**	0.026**	0.032**	0.033***	0.034***
	(0.008)	(0.010)	(0.010)	(0.009)	(0.009)
Num. Treated	136	124	113	124	124
Num. Control	452	385	335	380	380
Avg. R&D Intensity (All Firms)	0.051	0.051	0.047	0.051	0.051
Avg. R&D Intensity (Treated Firms)	0.054	0.054	0.054	0.054	0.054
Panel B					
Log R&D	0.319***	0.285***	0.362***	0.241***	0.268**
	(0.068)	(0.078)	(0.077)	(0.069)	(0.084)
Num. Treated	89	79	73	79	79
Num. Control	298	252	235	266	266
MNC & Year x Cohort FEs	Yes	-	-	-	-
MNC & Ind., Size x Yr FEs	-	-	-	Yes	Yes
MNC & Ind., Size x Yr x Cohort FEs	-	Yes	Yes	-	-
Inverse Prob. Weights	-	-	Yes	-	-
Model	Stacked	Stacked	Stacked	SA	TWFE

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

Notes: This table provides estimates of the difference-in-differences model discussed in Section 5.2 for the corresponding R&D outcome listed in each panel, where Columns (1) - (3) use Equation 1 and Columns (4) and (5) use Equation 2. Estimations use 2-year pooled average data from Compustat, where odd years are dropped to match the IRS SOI International Business Tax Sample. R&D Intensity (Panel A) is calculated as the ratio of annual R&D to the MNC's most recent sales value pre-adoption, and is restricted to be less than 1. Column (1) does not include additional controls. Columns (2) - (5) include year-by-cohort-by-industry and year-by-cohort-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-adoption values. Column (3) uses inverse probability-weighted data. Column (4) estimates an alternative specification from Sun and Abraham (2021). Column (5) estimates a standard TWFE specification.

Column (2) reports that R&D intensity increased by 0.026, which corresponds to a 48% increase relative to the average R&D intensity of 0.054 for HTP-adopting firms. Panel B shows that log R&D increased by 0.285 following adoption of an HTP. The larger estimate for R&D intensity suggests that incorporating extensive-margin responses is important in this setting and leads to larger estimates of the change in R&D activity following the adoption of an HTP structure.

5.2.4 Further Robustness Checks

To ensure that we are comparing similar firms in our regressions, we extend the estimation of Equation 1 by applying inverse probability weighting to the sample of MNCs. A standard approach to this form of weighting in an unstaggered difference-in-differences setting estimates a single propensity score for each unit. These scores are canonically estimated using data in periods prior to HTP adoption. In staggered designs, however, control units may serve as comparisons for multiple treated cohorts, and there is not a well defined pre-period. The advantage of the stacked estimator is that it can be thought of as combining a set of unstaggered difference-in-differences datasets, one for each treated cohort.

As shown in Section 4.3, we find evidence for considerable variation in rates of HTP adoption by industry and by firm size, which motivates our usage of industry and firm-size fixed effects in the preceding difference-in-differences estimates to control for time-varying heterogeneity along these dimensions. Section 4.3 also shows that several other observable characteristics predict MNC adoption of hybrid tax structures (see Tables A.4 and A.5). Propensity score weighting offers a convenient way to control for potential selection bias that may be introduced by a lack of balance along these observable dimensions between adopting and non-adopting MNCs.

We approach the computation of propensity scores with a similar logic to the design of the stacked estimator by computing a set of propensity scores for each treated cohort and never-treated observations that are used as cohort-level comparisons. We use the same predictor variables examined in Section 4.3 and Tables A.4 and A.5.[24] This gives us a set of propensity scores for never-treated units that are used as comparison units for multiple cohorts. These scores increase comparability within cohorts across adopting and non-adopting MNCs for other observable characteristics. The score incorporates pre-adoption measures of R&D activity; firm age; average foreign statutory rates; exposure to geographies where it may have been easier to use Check the Box; prior activity in Ireland, the Netherlands, and Luxembourg; whether or not the MNC had negative domestic income; and the firm's advertising intensity relative to sales. Because the stacked estimator also explicitly duplicates never-treated units when they are used as comparisons for multiple cohorts, we can apply these sets of propensity scores directly in

[24] In contrast to the logit estimates shown in Section 4.3, we use a random forest model to compute propensity scores. In our context, random forest models outperform logit models from a predictive standpoint, and can flexibly account for non-linearities. They are also theoretically invariant to transformations such as the natural log, which would require dropping observations with non-positive values.

the stacked design in a similar way to an unstaggered propensity score weighted difference-in-differences model.

Columns (3) of Tables 4, 6, and 7 show that we obtain similar estimates of changes in firm outcomes when we expand the model in Equation 1 by including inverse probability weights. See Figure A.1 and A.2 for corresponding event studies that also use IPW.

We also explore the robustness of our results to using alternative estimators. Columns (4) of Tables 4, 6, and 7 show that we generally obtain similar estimates when using the estimator of Sun and Abraham (2021). Columns (5) of Tables 4, 6, and 7 show that we also find similar effects when using a simpler two-way fixed effects estimator.[25]

Overall, the results in this section show that the adoption of HTP structures precedes large increases in tax avoidance behavior as well as increases in real economic activity, both at home and abroad.

5.3 Domestic to Foreign Profit Shifting and the Role of Cost Sharing Agreements

In Section 2.2.1, we briefly discuss how MNCs can use a contract known as a cost sharing agreement (CSA) to transfer IP to foreign affiliates. In Section 5.1, we show that HTPs utilize CSAs at a much higher rate than non-HTP MNCs. One underlying reason for this is that MNCs with hybrid tax planning structures may be able to use CSAs to shift profits from the US to their low-tax hybrid entities.

Note that this behavior may in some instances be permitted under US tax law. One such legal approach to profit shifting, which is described in publicly available legal briefs, was used by MNCs in the early 2000s. During this period, MNCs were able to exclude certain R&D costs that the IRS argued were subject to a CSA (and therefore, would have been allocated to an MNC's foreign affiliates). Specifically, 1995 regulations issued by the IRS did not explicitly require stock-based compensation to be covered by CSAs.

After extensive litigation, a technology firm won a ruling in the US Tax Court that permitted the exclusion of these costs. Although the US Treasury attempted to clarify its regulations to disincentivize this type of tax planning, other MNCs continued to engage in similar strategies during a long period of regulatory uncertainty. Appendix C provides additional institutional

[25]See Figure A.3 and A.4 for event studies using the estimator of Sun and Abraham (2021) and Figure A.5 and A.6 for event studies using the TWFE estimator.

detail about this strategy, along with empirical analysis that demonstrates a substantial behavioral shift by MNCs that held CSAs during the mid-2000s. We find evidence that MNCs exposed to this regulatory uncertainty expanded their use of stock-based compensation, both in absolute terms and as a fraction of their overall wage bill. We also find evidence that these MNCs increased their overall R&D activity and reported larger amounts of their wage expenses as eligible R&D costs for a domestic tax credit.

6 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 help to fill this 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 two forms of Reverse Hybrid Mismatch arrangement, one through the Netherlands, the other through 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. Although only 17.5% of US MNCs adopt these structures, those that do obtain a significant tax advantage over other MNCs. Remarkably, this small fraction of companies generates a majority of foreign earnings by US MNCs and is responsible for the bulk of the increase in cash held abroad over the period we study.

Our analyses use this tax data to reveal that the adoption of hybrid tax planning structures is accompanied by significant changes in real economic activity. We find that firms adopting these structures also have experience significant foreign and domestic growth.

While our analysis sheds light on multinational tax planning behavior over the last several decades, a large number of policy changes have reshaped the incentives and feasibility of tax planning in recent years. As new data becomes available, researchers should be able to determine how firms have reacted to the new landscape and the extent to which current policy attempts to curb profit shifting and tax avoidance have been successful.

References

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. "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.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The effect of minimum wages on low-wage jobs." *The Quarterly Journal of Economics*, 134(3): 1405–1454.

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."

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."

Dyreng, Scott, Michelle Hanlon, Edward L Maydew, and Jacob R Thornock. 2017. "Changes in Corporate Effective Tax Rates Over the Past Twenty-Five Years." *Journal of Financial Economics*.

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, 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).

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.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*, 225(2): 175–199. Themed Issue: Treatment Effect 1.

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.

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.

Appendices

These appendices include supplemental information and additional analyses. Appendix A provides additional tables and figures related to the main analyses. Appendix B provides an overview of measurement issues in several different sources of multinational data. Appendix C describes how multinationals with cost sharing agreements may have used a legal strategy to reclassify foreign costs as domestic costs for tax planning purposes.

A Additional Tables and Figures

Table A.1: MNC Summary Statistics

Domestic Outcomes	All MNCs				Hybrid Tax Planners				Other MNCs			
	Mean	P25	P75	SD	Mean	P25	P75	SD	Mean	P25	P75	SD
Tangible Assets	3024	226	1981	9977	4407	341	2894	14023	2659	201	1762	8557
Wages	583	58	447	1636	1125	114	928	2440	440	49	366	1308
Investment	267	10	141	1619	356	17	216	1242	243	9	125	1704
R&D	73	0	38	343	179	0	106	644	44	0	27	184

Foreign Outcomes	All MNCs				Hybrid Tax Planners				Other MNCs			
	Mean	P25	P75	SD	Mean	P25	P75	SD	Mean	P25	P75	SD
Tangible Assets	886	16	409	3874	1706	56	811	6598	668	11	313	2692
Intangible Assets	355	0	136	1856	898	6	462	3467	211	0	91	1039
Pretax E&P	279	2	105	1607	771	12	323	3062	148	1	74	840
Income Taxes	54	0	22	284	117	4	59	378	38	0	16	250
Cash	496	5	125	4579	1429	21	430	8368	250	4	90	2775

Foreign Affiliate Counts	All MNCs				Hybrid Tax Planners				Other MNCs			
	Mean	P25	P75	SD	Mean	P25	P75	SD	Mean	P25	P75	SD
FDE Count	38	4	34	103	86	16	81	188	23	4	24	48
CFC Count	13	2	12	41	34	4	33	86	10	2	10	26

Sample Sizes	3635	512	3123

Notes: This table contains summary statistics for our stable sample as described in Section 3.2. This sample was created by identifying MNCs with coverage in both the SOI Corporate and SOI International Business Tax samples, not including MNCs that did not have at least $500 million in domestic assets as well as at least one CFC with $50 million in foreign assets. To ensure we do not disclose information about individual MNCs, the values listed for P25 and P75 are the means of ten observations surrounding a given percentile.

Figure A.1: Event Studies: Profit Shifting Mechanisms and Foreign ETRs (IPW)

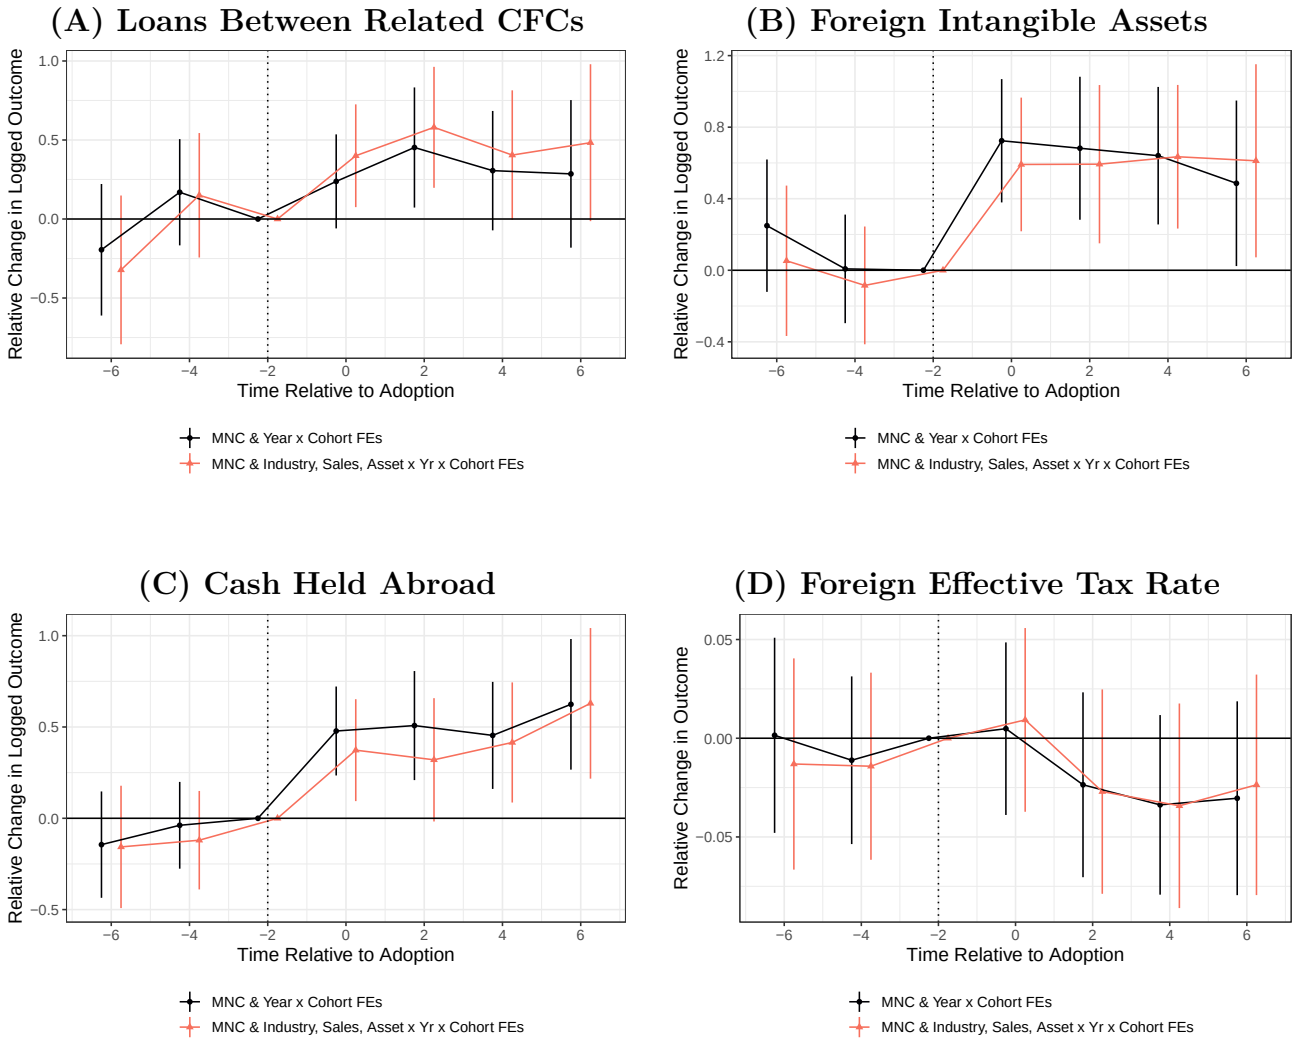

Notes: This figure provides estimates of μ_ℓ from Equation 1 for the corresponding foreign outcome listed in each panel, using inverse probability-weighted data. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-cohort-by-industry fixed effects and year-by-cohort-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-adoption values for each cohort.

Figure A.2: Event Studies: Hybrid Tax Planning and Real Economic Activity (IPW)

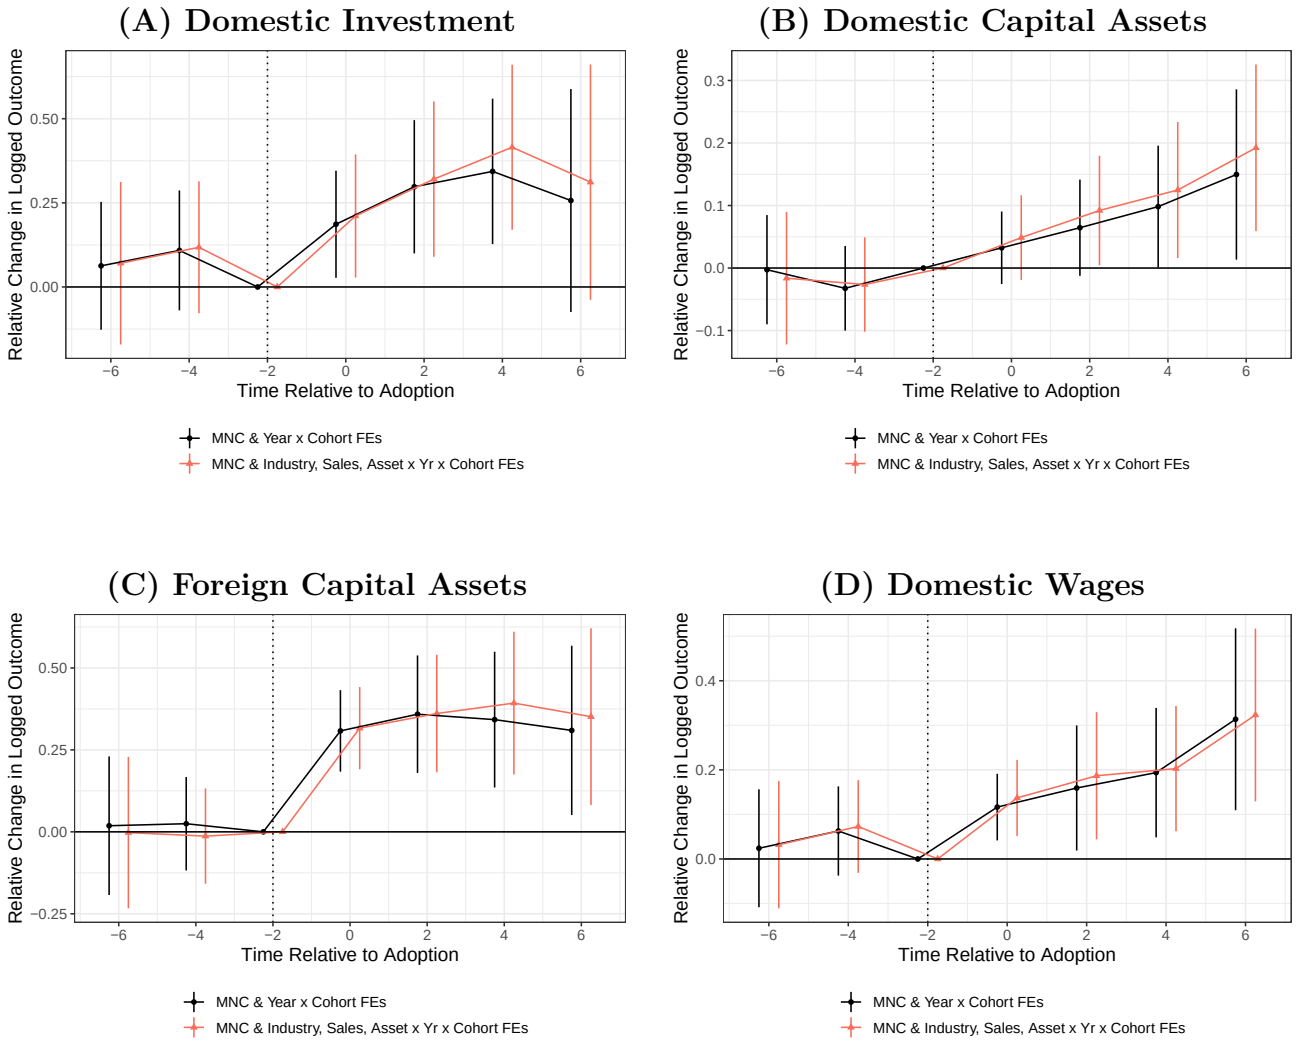

Notes: This figure provides estimates of μ_ℓ from Equation 1 for the corresponding outcome listed in each panel, using inverse probability-weighted data. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-cohort-by-industry fixed effects and year-by-cohort-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-adoption values for each cohort.

Figure A.3: Event Studies: Profit Shifting Mechanisms and Foreign ETRs (Sun Abraham)

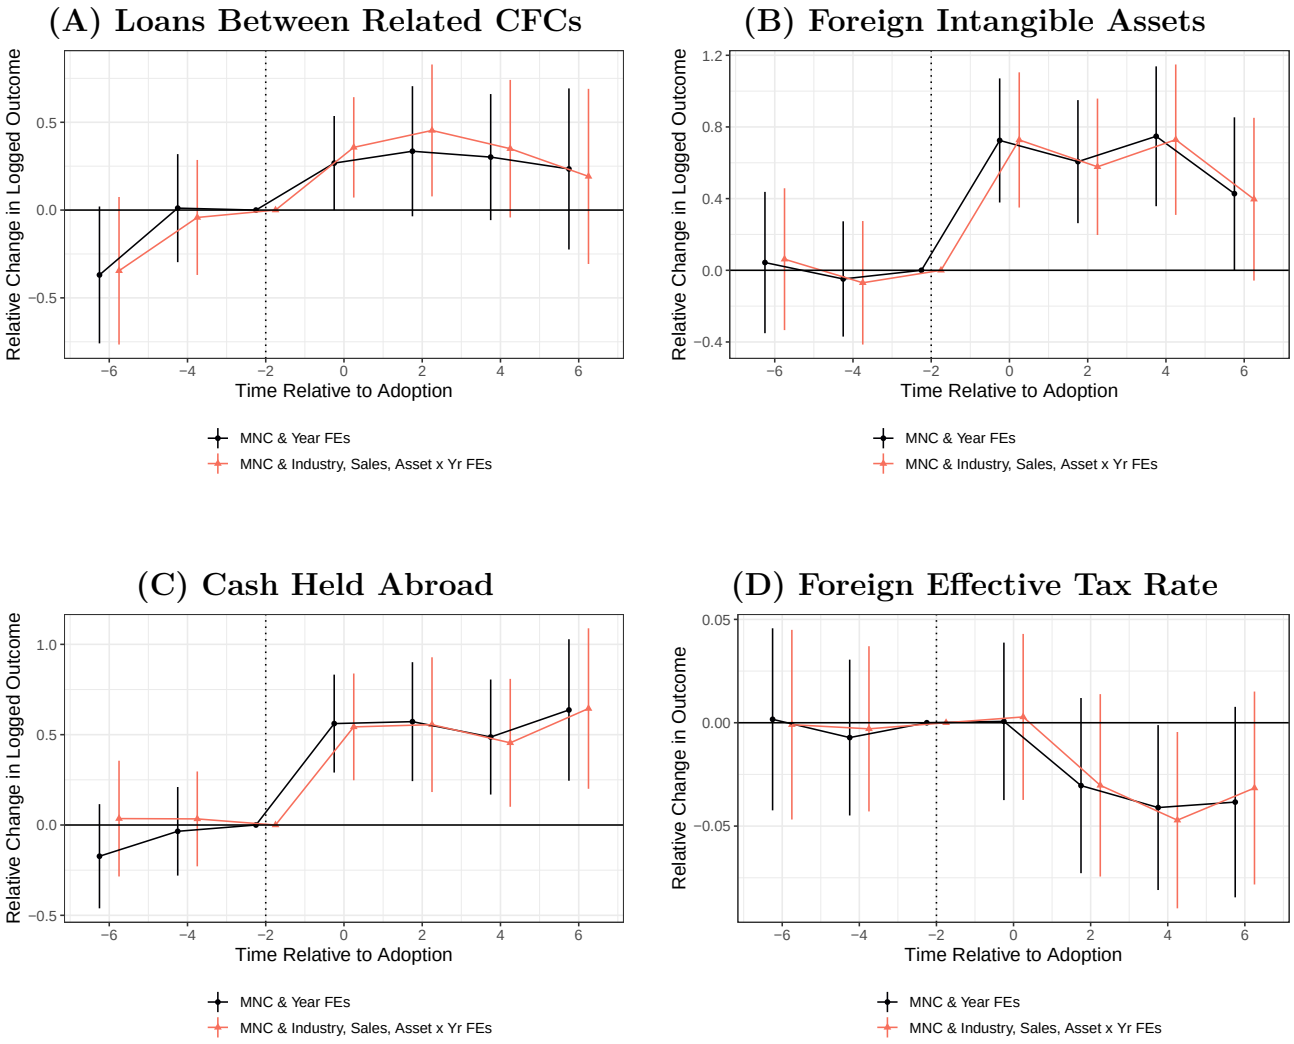

Notes: This figure provides estimates of μ_ℓ from Equation 2 for the corresponding foreign outcome listed in each panel, using the staggered estimator from Sun and Abraham (2021). Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-cohort-by-industry fixed effects and year-by-cohort-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-adoption values for each cohort.

Figure A.4: Event Studies: Hybrid Tax Planning and Real Economic Activity (Sun Abraham)

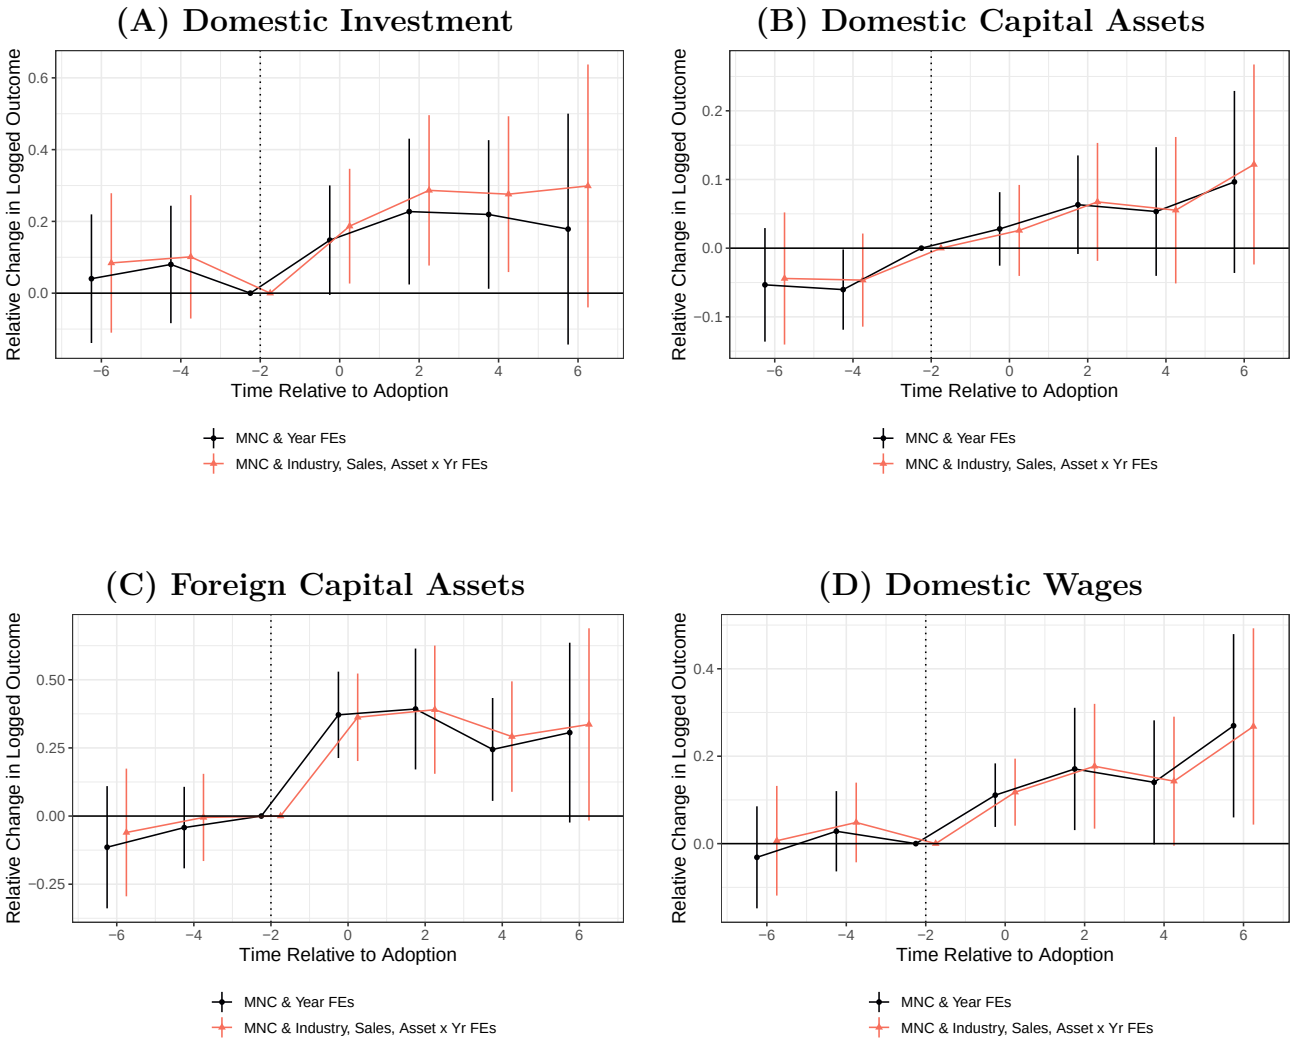

Notes: This figure provides estimates of μ_ℓ from Equation 2 for the corresponding outcome listed in each panel, using the staggered estimator from Sun and Abraham (2021). Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-cohort-by-industry fixed effects and year-by-cohort-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-adoption values for each cohort.

Figure A.5: Event Studies: Profit Shifting Mechanisms and Foreign ETRs (TWFE)

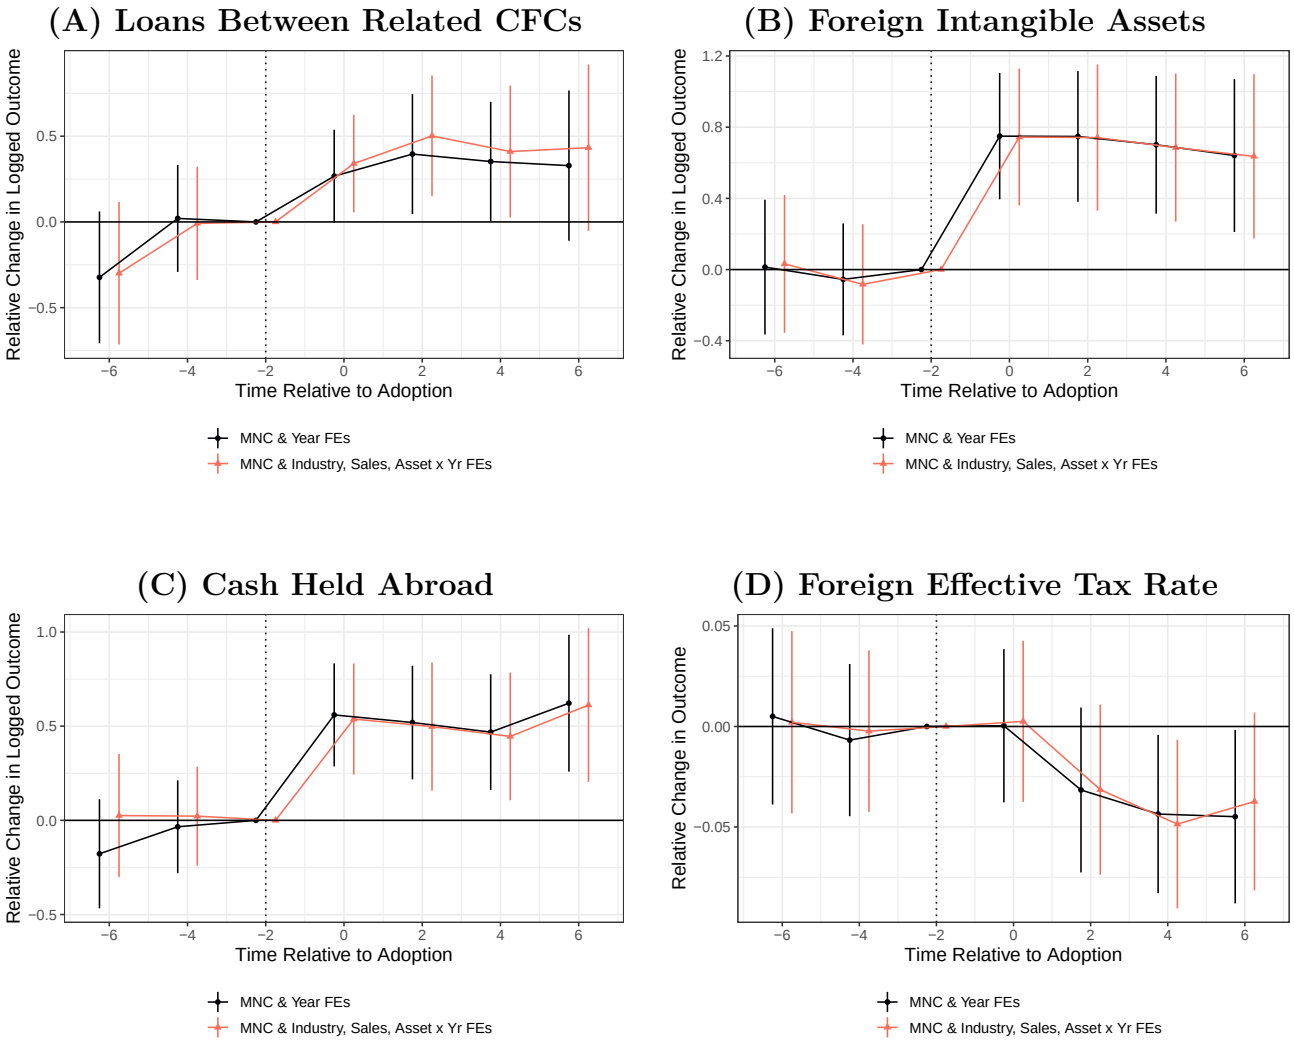

Notes: This figure provides estimates of μ_ℓ from Equation 2 for the corresponding foreign outcome listed in each panel. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-cohort-by-industry fixed effects and year-by-cohort-by-group fixed effects,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

Figure A.6: Event Studies: Hybrid Tax Planning and Real Economic Activity (TWFE)

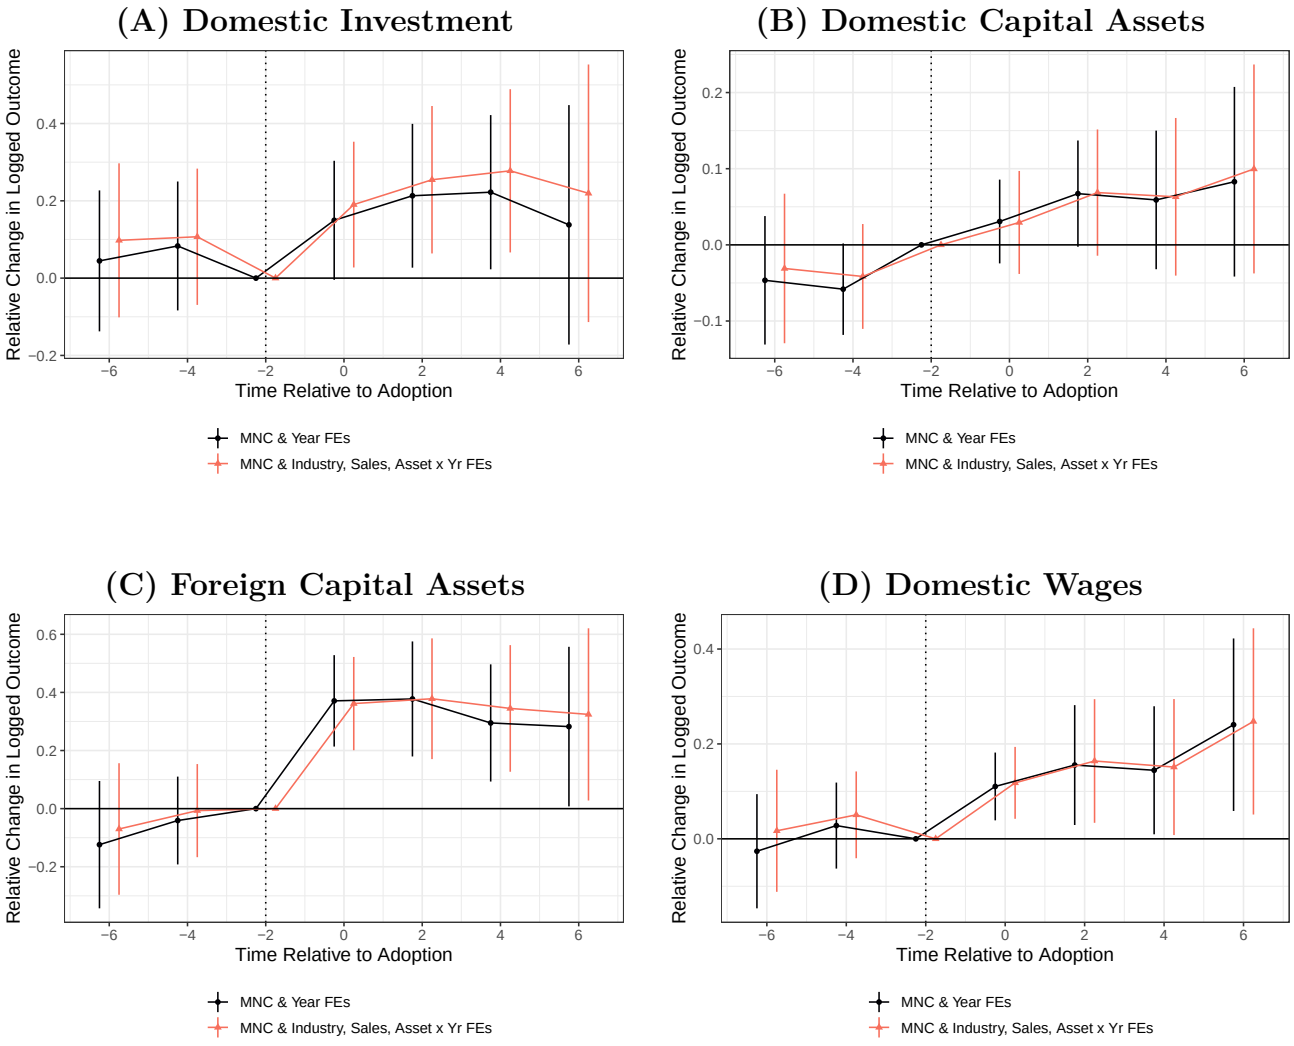

Notes: This figure provides estimates of μ_ℓ from Equation 2 for the corresponding outcome listed in each panel. Specification 1 (in black) does not include additional controls. Specification 2 (in orange) includes year-by-cohort-by-industry fixed effects and year-by-cohort-by-group fixed effects, where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

Figure A.7: Event Studies: Compustat R&D (Robustness)

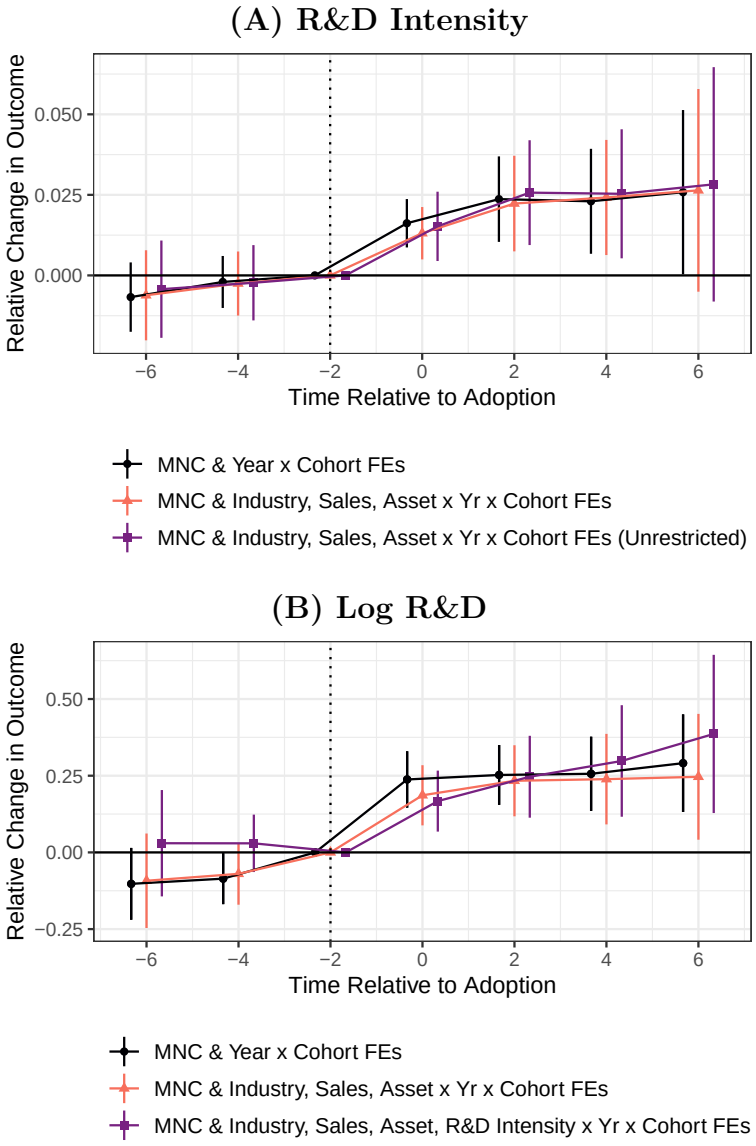

Notes: This figure provides estimates of μ_ℓ from Equation 1 for the corresponding R&D outcome listed in each panel, using 2-year pooled average data from Compustat and dropping odd years to match the IRS SOI International Business Tax Sample. Specifications 1 and 2 (in black and orange, respectively) are identical to the results shown in Figure 16. For R&D Intensity (Panel A), Specification 3 (in purple) does not restrict R&D intensity to be less than 1. For Log R&D (Panel B), Specification 3 includes an additional R&D intensity quartile fixed effect in its year-by-cohort-by-group fixed effects, where R&D intensity is not restricted to be less than 1.

Table A.2: Hybrid Tax Planning and R&D Data from Compustat (Robustness)

	(1)	(2)	(3)	(4)
Panel A				
R&D Intensity (Unrestricted)	0.028**	0.034***	0.033**	0.034**
	(0.010)	(0.010)	(0.010)	(0.011)
Num. Treated	125	114	125	125
Num. Control	388	337	387	387
Avg. R&D Intensity (All Firms)	0.060	0.050	0.059	0.059
Avg. R&D Intensity (Treated Firms)	0.056	0.055	0.056	0.056
R&D Intensity x Yr FEs	-	-	-	-
R&D Intensity x Yr x Cohort FEs	-	-	-	-
Panel B				
Log R&D	0.285***	0.362***	0.241***	0.268**
	(0.078)	(0.077)	(0.069)	(0.084)
Num. Treated	79	73	79	79
Num. Control	252	235	266	266
R&D Intensity x Yr FEs	-	-	Yes	Yes
R&D Intensity x Yr x Cohort FEs	Yes	Yes	-	-
MNC & Ind., Sales, Asset x Yr FEs	-	-	Yes	Yes
MNC & Ind., Sales, Asset x Yr x Cohort FEs	Yes	Yes	-	-
Inverse Prob. Weights	-	Yes	-	-
Model	Stacked	Stacked	SA	TWFE

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

Notes: This table provides estimates of the difference-in-differences model discussed in Section 5.2 for the corresponding R&D outcome listed in each panel, where Columns (1) - (3) use Equation 1 and Columns (4) and (5) use Equation 2. Estimations use 2-year pooled average data from Compustat, where odd years are dropped to match the IRS SOI International Business Tax Sample. R&D Intensity (Panel A) is calculated as the ratio of annual R&D to the MNC's most recent sales value pre-adoption. All specifications include year-by-cohort-by-industry and year-by-cohort-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-adoption values. Log R&D (Panel B) also includes an additional R&D intensity quartile fixed effect in its year-by-cohort-by-group fixed effects. Column (2) uses inverse probability-weighted data. Column (3) estimates an alternative specification from Sun and Abraham (2021). Column (4) estimates a standard TWFE specification.

Table A.3: HTP Uptake by Size

	Dom. Sales	For. Sales	Dom. Assets	For. Assets
Quartile 2	0.350+	0.483*	0.193	0.442*
	(0.198)	(0.207)	(0.193)	(0.210)
Quartile 3	0.350+	0.504*	0.482**	0.803***
	(0.199)	(0.207)	(0.187)	(0.198)
Quartile 4	0.949***	1.147***	0.623***	1.018***
	(0.184)	(0.193)	(0.184)	(0.195)
Num. Obs.	8809	8809	8809	8809
R2 Pseudo	0.019	0.024	0.012	0.020
Year FEs	Yes	Yes	Yes	Yes

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

Notes: This table displays coefficients from logit regressions, as discussed in Section 4.3, that predict HTP adoption based on 4 determinations of firm size: domestic sales, foreign sales, domestic assets, and foreign assets. The model uses observations from adopting MNCs in the period prior to adoption and includes all observations from never-adopters. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Table A.4: Selection into Hybrid Tax Planning

	(1)	(2)
Panel A		
Claim Research Credit	0.568***	0.215
	(0.128)	(0.170)
Num.Obs.	8809	5030
Panel B		
Age Quartile 2	0.499**	0.446*
	(0.178)	(0.200)
Age Quartile 3	0.205	0.071
	(0.191)	(0.221)
Age Quartile 4	0.452*	0.128
	(0.185)	(0.222)
Num.Obs.	8809	5030
Panel C		
Avg. Statutory Foreign ETR	0.045	-1.653
	(1.047)	(1.148)
Share of Foreign Sales with Unobserved Statutory Rate	-0.212	-0.529
	(0.344)	(0.364)
Num.Obs.	7832	4528
Panel D		
Exposure to CTB	2.411***	1.930*
	(0.596)	(0.873)
Num.Obs.	8280	4796
Panel E		
Presence in Ireland	1.460***	1.356***
	(0.183)	(0.223)
Presence in Netherlands	0.708***	0.440**
	(0.139)	(0.168)
Presence in Luxembourg	0.023	-0.105
	(0.220)	(0.249)
Num.Obs.	8809	5030
Panel F		
Domestic Loss	-0.861***	-0.824***
	(0.172)	(0.200)
Num.Obs.	8809	5030
Panel G		
Advertising to Sales Ratio	-0.007	0.561
	(0.010)	(1.714)
Num.Obs.	8608	4920
Year FEs	Yes	-
Industry, Sales, Asset Quartiles x Year FEs	-	Yes

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

Notes: This table displays coefficients from a series of logit regressions, as discussed in Section 4.3, that predict HTP adoption based on various MNC characteristics (Panels A-G). The model uses observations from adopting MNCs in the period prior to adoption and includes all observations from never-adopters. Column (1) includes year fixed effects; Column (2) includes year-by-industry, sales, and asset quartile fixed effects. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Table A.5: Selection into Hybrid Tax Planning (Combined)

	(1)	(2)
Claim Research Credit	0.105	−0.045
	(0.144)	(0.178)
Firm Age Quartile 2	0.417*	0.399+
	(0.188)	(0.212)
Firm Age Quartile 3	0.000	−0.015
	(0.209)	(0.241)
Firm Age Quartile 4	0.010	−0.041
	(0.206)	(0.238)
Avg. Statutory Foreign ETR	0.588	0.041
	(0.993)	(1.146)
Share of Foreign Sales with Unobserved Statutory Rate	0.117	0.035
	(0.320)	(0.352)
Exposure to CTB	1.608*	1.605
	(0.795)	(1.007)
Advertising to Sales Ratio	−0.147	−1.919
	(0.500)	(2.111)
Domestic Loss	−0.793***	−0.910***
	(0.184)	(0.216)
Presence in Ireland	1.371***	1.331***
	(0.199)	(0.244)
Presence in Netherlands	0.578***	0.446*
	(0.150)	(0.182)
Presence in Luxembourg	0.027	−0.077
	(0.229)	(0.265)
Num.Obs.	7624	4483
Year FEs	Yes	-
Industry, Sales, Asset Quartiles x Yr FEs	-	Yes

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

Notes: This table displays coefficients from logit regressions, as discussed in Section 4.3, that predict HTP adoption based on various MNC characteristics. This specification combines all characteristics tested in Table A.4 into a single logit regression. Column (1) includes year fixed effects; Column (2) includes year-by-industry, sales, and asset quartile fixed effects. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Table A.6: HTP Uptake by Auditor

	(1)	(2)	(3)	(4)
Big 4	0.805	0.378	0.342	0.044
	(1.019)	(1.082)	(0.293)	(0.348)
Medium	0.511	0.370		
	(1.054)	(1.119)		
Num.Obs.	4610	2491	4610	2491
Year FEs	Yes	-	Yes	-
Industry, Sales, Asset Quartiles x Year FEs	-	Yes	-	Yes

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

Notes: This table displays coefficients from logit regressions, as discussed in Section 4.3, that predict HTP adoption based on the identity of an MNC's auditors, determined using Compustat data. Columns (1) and (2) include dummy variables for a "Big 4" auditor as well as a medium-sized auditor; Columns (3) and (4) include only the Big 4 dummy. Odd columns include year fixed effects and even columns include year-by-industry, sales, and asset quartile fixed effects. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

B Measurement of Foreign Earnings and Taxes

In this appendix, we examine the potential for aggregation error in the measurement of foreign earnings and taxes when using IRS data. The potential for measurement error is clearly summarized in Blouin and Robinson (2020, henceforth B&R). Appendix B.1 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 MNCs. Appendix B.2 describes how we construct a matched sample of MNCs which appear in both tax data and public filings compiled by Compustat. We 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 Appendix B.3, we 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. We also provide corrected estimates of average foreign effective tax rates by jurisdiction.

B.1 Overview of Data Sources for Measuring Tax Outcomes of US Multinationals

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

B.1.1 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.

B.1.2 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 MNCs. Research that examines foreign earnings by US MNCs 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 tax 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 MNCs. These data are 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 affil-

iates, 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.

B.1.3 Administrative Tax Data

A full description of the SOI data is provided in Section 3. The discussion below focuses on the potential for aggregation error in this data.

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 AMNE data is that, as alluded to above, foreign affiliates may be classified for tax purposes as opaque entities (CFCs) or pass-through 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 MNC 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 pass-through (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, because information covering pass-through 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 pass-throughs 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 overstate foreign earnings, estimates from research that examine foreign earnings in tax data may not be reliable. Worse, to the extent this error is growing over time, which one might suspect if MNCs' 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).

B.2 Assessing Aggregation Error in Foreign Earnings and Taxes

In this section, we 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, we describe how we create a matched sample of US MNCs that appear both in SOI and Compustat. Compustat reports foreign income and foreign tax figures on a consolidated basis for a set of MNCs.[26] The consolidated nature of Compustat's

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

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 in B&R.[27]

B.2.1 Sample Construction

Figure B.1 shows attrition associated with the sample construction process. The top line in the left panel shows the annual sample size for US MNCs contained in the SOI sample, which ranges between 7,597 and 9,157. US MNCs are defined as US companies that file a corporate tax return and that own at least one opaque foreign affiliate (CFC) 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 B.1, ranges between 1,845 and 2,334. 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, we use information from Form 1120, Schedule M-3, which must be filed by large US MNCs 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, we 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 1,680 to 2,023, is shown on the bottom line of the left panel of Figure B.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 relatively

[27] 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.

small US MNCs.

Figure B.1: SOI MNCs vs. Matched Sample

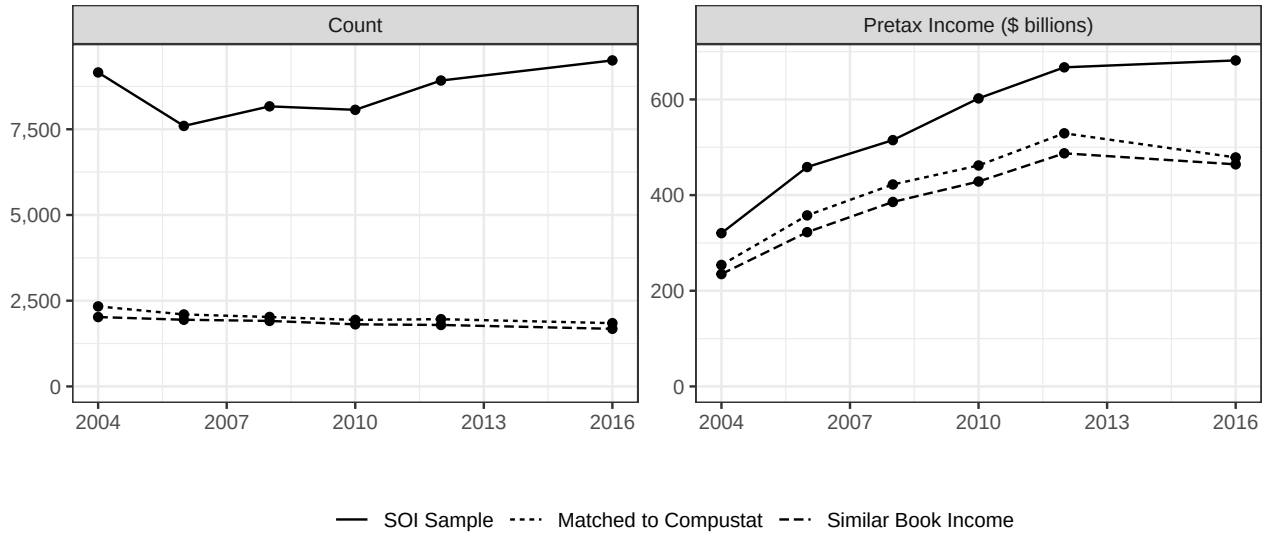

Notes: This figure demonstrates the attrition associated with the sample construction process. The "SOI Sample" contains all US MNCs from the SOI data sample. The "Matched to Compustat" sample contains SOI Sample MNCs matched to Compustat via unique firm identifier (EIN). The final linked sample, "Similar 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 B.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 MNCs are dropped in the match, which is unsurprising given that SOI is the more comprehensive database. In 2016, for example, the matched sample includes 68.8% of Compustat MNCs. In the right panel, attrition is presented in terms of aggregate foreign earnings as reported in Compustat, with the matched sample including 84% of the aggregate foreign earnings reported in Compustat in 2016.

Figure B.2: Compustat MNCs vs. Matched Sample

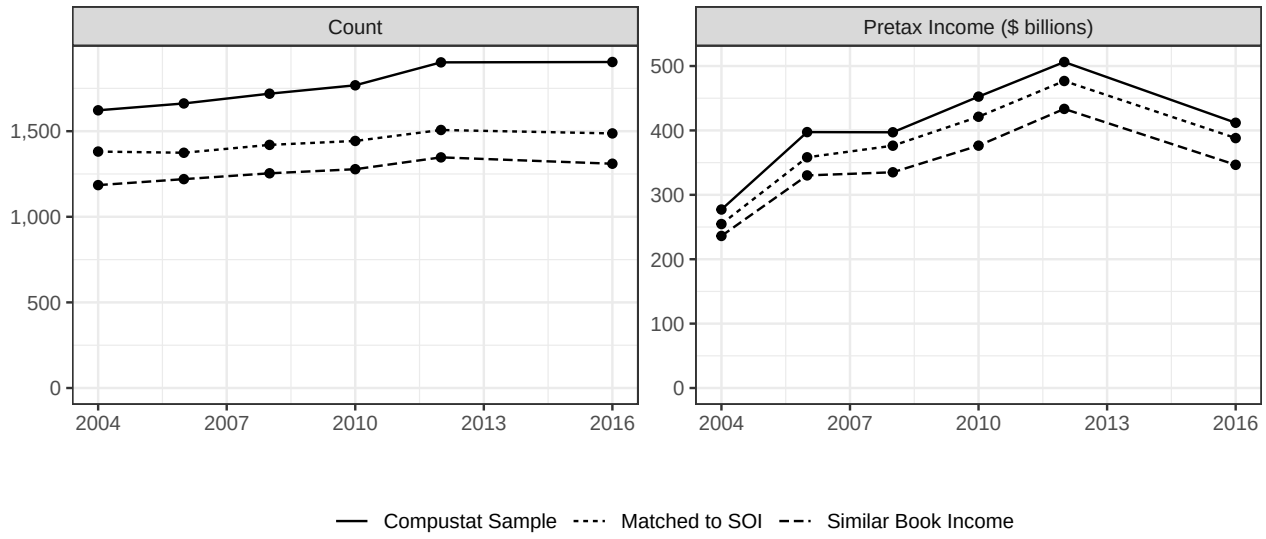

Notes: This figure demonstrates the attrition associated with the sample construction process, following the same methodology as described in Figure B.1 but starting with the Compustat sample rather than the SOI sample. Observations reporting missing pretax foreign income are removed from the initial 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.

B.2.2 Comparison of Aggregate Foreign Income and Tax

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

Figure B.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). We use Schedule E as the preferred measure. The left panel contains aggregates computed from MNCs in all industries. The right panel excludes MNCs 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 this discrepancy appears to be related to how extractive industries operate foreign concessions. These projects are often

Figure B.3: Comparison of Aggregate Foreign Tax Measures

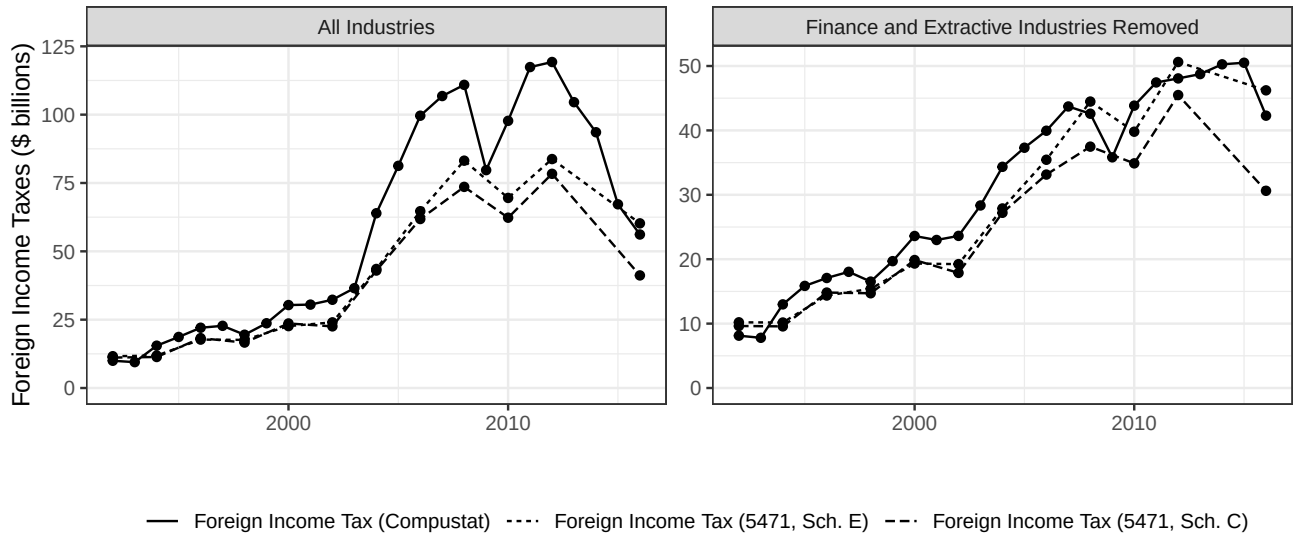

Notes: This figure shows three measures of aggregate foreign income tax as respectively reported by Compustat; IRS Form 5471, Schedule C; and IRS Form 5471, Schedule E. The left panel contains aggregates computed from MNCs in all industries, while the right panel excludes MNCs operating in Financial, 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.

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 a US MNC's foreign tax bill from its 5471 filings.

Figure B.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 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, we 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

Figure B.4: Comparison of Aggregate Foreign Income (Pretax)

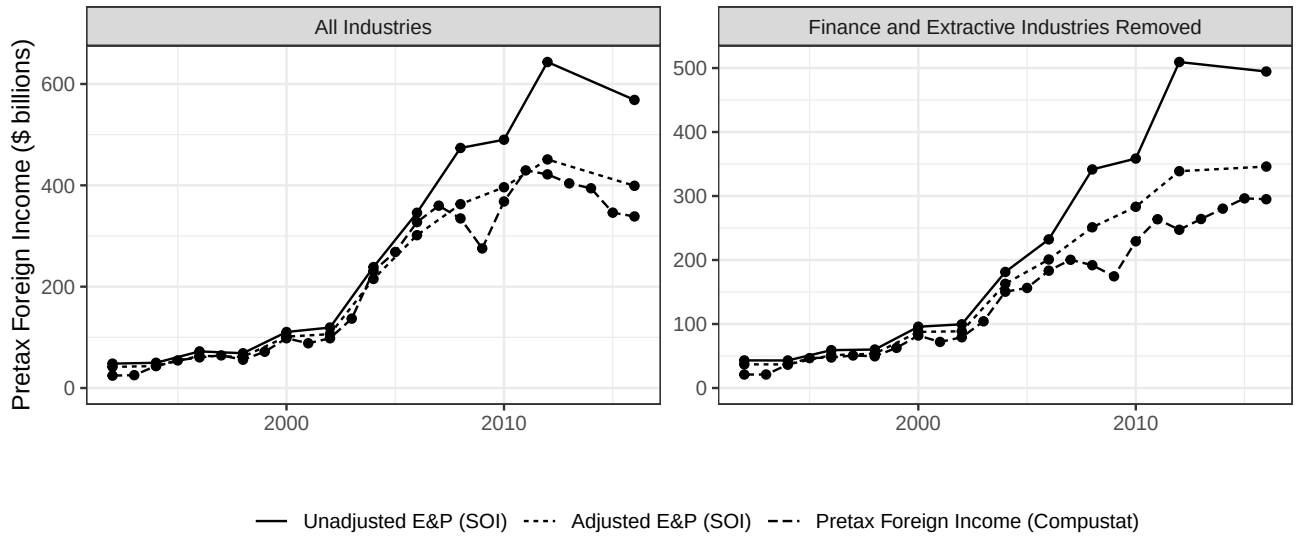

Notes: This figure shows three measures of aggregate foreign pretax income: aggregate foreign pretax income as reported by Compustat; unadjusted E&P computed from IRS Form 5471, Schedule H; and Schedule H E&P, adjusted to remove dividends received from related affiliates as reported on Schedule M. Both SOI measures include foreign taxes, as reported in Schedule E, added back to represent pretax earnings. As in Figure B.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.

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 B.5 provides a similar figure, but with aggregate post-tax earnings (instead of pretax), demonstrating a similar pattern.

Figure B.5: Comparison of Aggregate Foreign Income (Post-Tax)

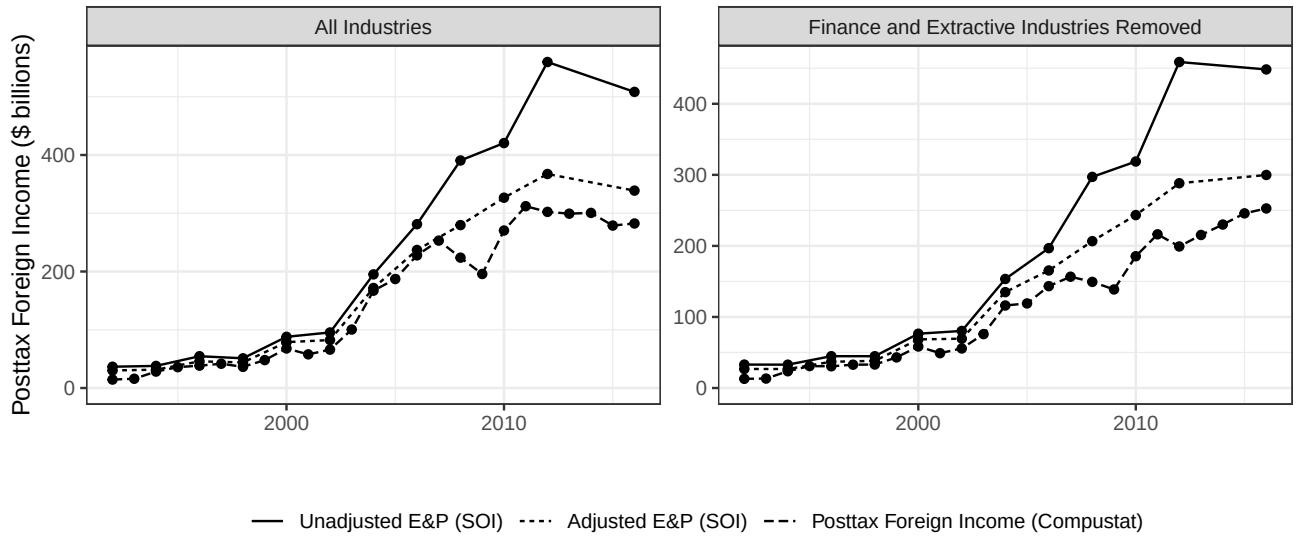

Notes: This figure displays the same three aggregate foreign income measures as Figure B.4 on a post-tax, rather than pretax, basis. It similarly underscores the effectiveness of B&R's adjustment for aggregation error in the SOI E&P data.

Finally, Figure B.6 presents a comparison of average annual foreign effective tax rates from Compustat and SOI (one adjusted for aggregation error, 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 B.3) and lower foreign earnings estimates (Figure B.4).

Figure B.6: Comparison of Aggregate Foreign Effective Tax Rates

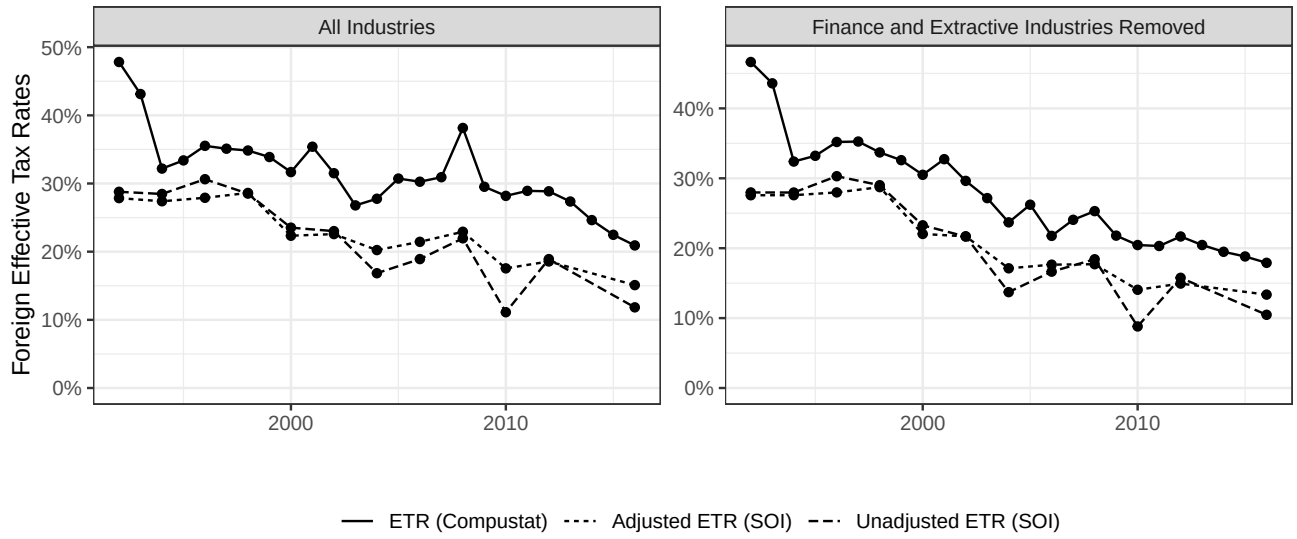

Notes: This figure presents a comparison of average foreign effective tax rates from Compustat, SOI (adjusted for aggregation error), and SOI (unadjusted). These tax rates are based on pretax income as reported in Figure B.4 and foreign tax as reported in Figure B.3. The left panel contains aggregates 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.

B.2.3 Measurement Error and MNC Corporate Complexity

In this section, we show that aggregation error is closely related to the size of MNCs' 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 an MNC's affiliate structure. Furthermore, we 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 earnings, not foreign taxes. Given that the primary mechanism behind aggregation error relates to the double-counting of foreign earnings that are distributed up through an MNC's foreign ownership network, we would expect aggregation error to increase with the size of an

Figure B.7: SOI/Compustat Foreign Earnings Differential

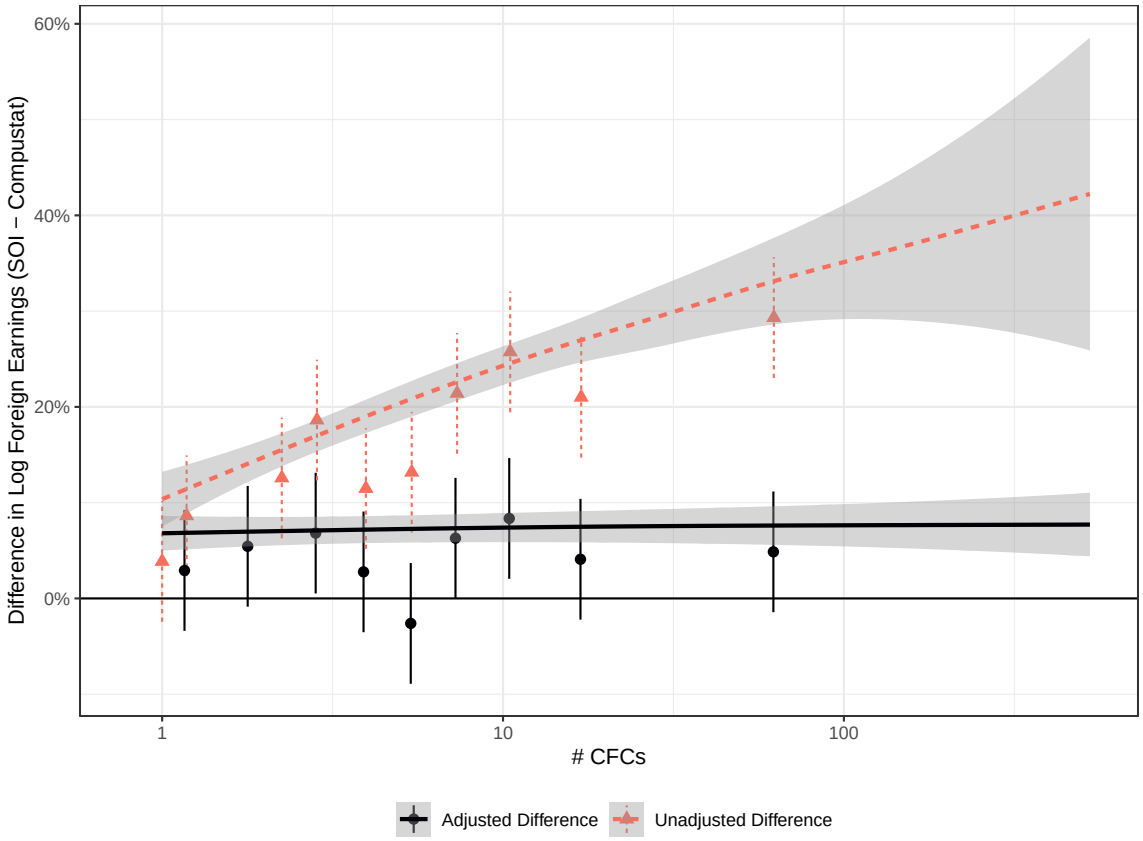

Notes: This figure presents the results of two cubic spline regressions where the dependent variable is the difference between log SOI foreign E&P and log Compustat foreign income. The orange line represents the estimates of a regression using unadjusted SOI E&P, while the black line uses adjusted SOI 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.

MNC's foreign affiliate network. Figure B.7 confirms that this is the case. The orange line and surrounding ribbon plot the estimates of a cubic 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 MNCs in 10 quantiles of 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 MNCs' 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 between Compustat

Figure B.8: SOI/Compustat Foreign Earnings Differential by Period

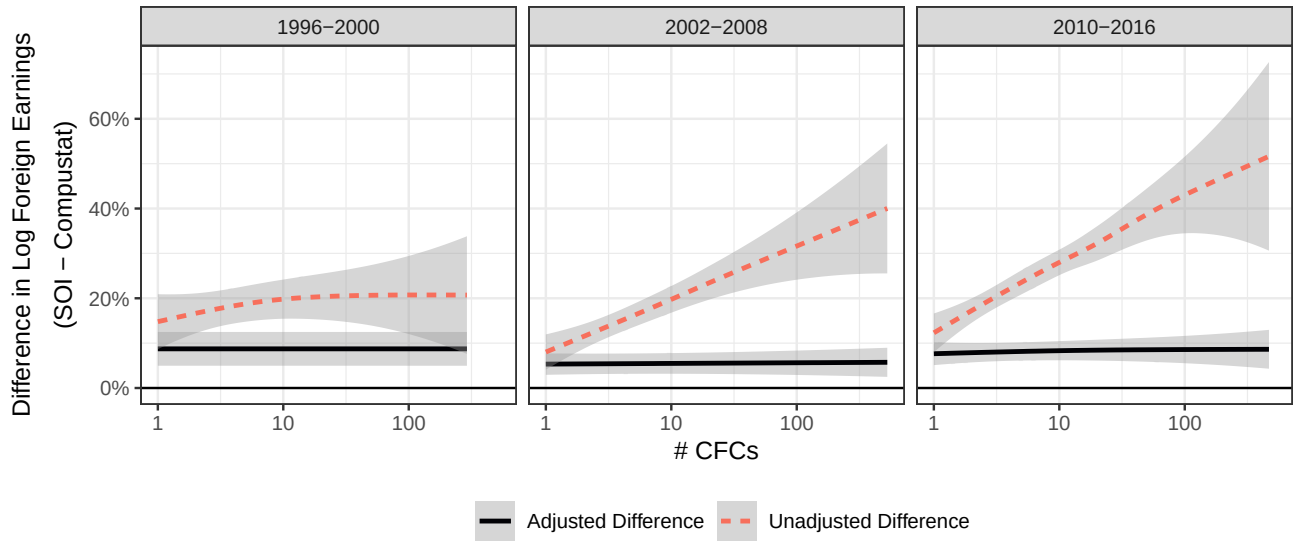

Notes: This figure plots the same spline estimates as Figure B.7 above, split into three time periods. The increase in the slope of the orange line over the three periods further confirms that aggregation error in the unadjusted SOI data has been increasing over time. However, the roughly constant black line indicates that residual measurement error has not been increasing.

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, we 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 B.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.

Table B.1 provides a regression formulation of the graphical analysis shown above that estimates the following specification,

$$\log(\Pi_{it}^{f,\text{SOI}}) - \log(\Pi_{it}^{f,\text{Compustat}}) = \beta \log(N_{it}^{\text{CFC}}) + \mu_t + u_i. \tag{B.1}$$

Table B.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 of the form B.1. 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. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

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) include the year fixed effects μ_t, and Columns (5) through (8) remove the industries excluded in the aggregate plots shown in the previous section.

B.3 Measurement Error and Tax Elasticities of Foreign Income

In the previous section, we 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, we recompute tax semi-elasticities of foreign earnings following the methodology in Dowd, Landefeld and Moore (2017, henceforth DLM). DLM use unadjusted earnings and a simple regression framework to examine how sensitive the allocation of MNC earnings are to average and statutory tax rates by country. First, we 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. We construct a

sample that is nearly identical to theirs, and obtain similar regression results with their aggregation process. We 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, our 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. We examine a similar specification that continues to show larger elasticities in tax haven jurisdictions. We 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.

B.3.1 Sample Construction

DLM provide a clear description of how they construct their sample. This allows us 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 MNCs that do not operate in extractive industries, finance, insurance, or utilities. As shown in Appendix B.2, 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.[28] They also require that CFCs report positive amounts for wages and tangible assets.[29]

Table B.2 contains the relevant sample statistics from DLM, Table 1. We compare these estimates with those contained in Table B.3, which contains sample statistics for our reproduction. There are some small differences between the samples—ours is slightly smaller (90,746 vs. 96,959 in their sample), which may be due to differences in how they classify industries.[30] Overall, however, the summary statistics are very close.

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

[29] 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, we 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").

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

Table B.2: Unadjusted Sample Statistics from Dowd, Landefeld and Moore (2017), 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

Notes: This table contains relevant unadjusted sample statistics from DLM Table 1. DLM use a sample of CFCs with positive foreign earnings for MNCs that report positive amounts for wages and tangible assets and do not operate in extractive industries, finance, insurance, or utilities. Statistics are shown for the entire sample, for CFCs located in non-haven countries, and for CFCs located in tax havens.

Table B.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

Notes: This table contains our reproduction of the DLM unadjusted sample statistics presented in Table B.2. While our sample is slightly smaller, the summary statistics are very similar between DLM's original and our reproduction.

Table B.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

Notes: This table contains sample statistics following the same form as DLM (Table B.2) but using adjusted foreign earnings data. Whereas DLM use pretax income disclosed in IRS Form 5471, Schedule C, we instead use foreign E&P disclosed in Schedule H and remove related dividends disclosed in Schedule M. This results in lower average profits and a slightly smaller sample size post-adjustment.

Next, we produce a similar sample using adjusted foreign earnings.[31] Table B.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.

B.3.2 Corrected Measurements of Tax Semi-Elasticities of Foreign Earnings

Adjusted Estimates from DLM

We use the samples described above to estimate the regression,

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

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

[31] We make two adjustments to DLM's measurement of foreign earnings. They use pretax income disclosed in the income statement (Schedule C). We 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.

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 MNC 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 B.9 shows average tax rates for the set of tax havens considered in this appendix.[32]

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

Figure B.9: Adjusted and Unadjusted Average Tax Rates for Selected Haven Jurisdictions

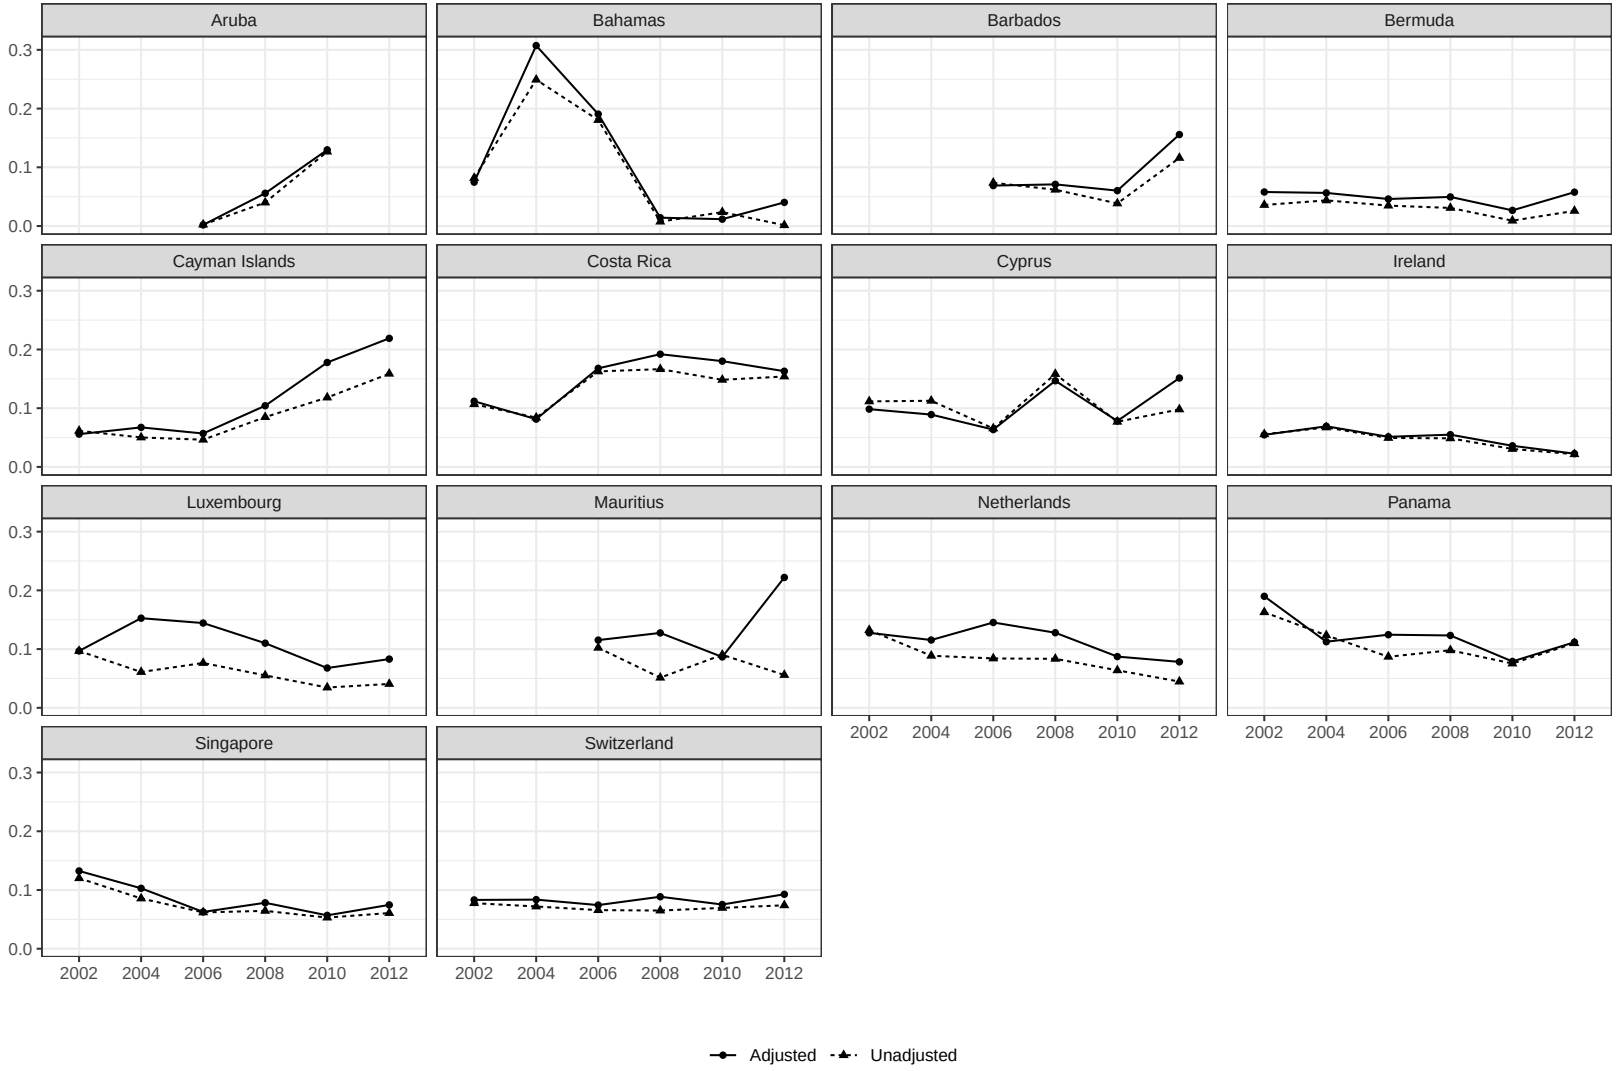

Notes: This figure presents adjusted and unadjusted average tax rates over time for the set of tax haven countries we consider in our DLM-style regression analysis (Equation B.2), as well as the Netherlands.

Table B.5: Dowd, Landefeld and Moore (2017) 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

Notes: This table displays coefficients from regressions under Equation B.2. The dependent variable is log profits for MNC i in country c in period t, using unadjusted SOI foreign E&P. The parameter of interest, 1-Rate, captures the sensitivity of MNC profit allocation to changes in tax rates. Column (1) presents estimates using statutory tax rates and Column (2) presents estimates using average rates. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

We 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,

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

Note that the only difference between equation Equation B.2 and Equation B.3 is the addition of h subscripts on α and β_3. The second allows for a second-order polynomial in the keep rate,

$$\log \pi_{ict} = \alpha + \beta_1 \log K_{ict} + \beta_2 \log L_{ict} + \beta_3(1 - \tau_{ct}) + \gamma_3(1 - \tau_{ct})^2 + \beta_4 X_{ct} + \mu_t + \psi_g + \epsilon_{it}. \tag{B.4}$$

Existing literature interprets β_3 as a semi-elasticity. Here, we do not emphasize a causal

Table B.6: Dowd, Landefeld and Moore (2017) 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

Notes: This table displays coefficients from regressions under Equation B.2. The dependent variable is log profits for MNC i in country c in period t, using adjusted SOI foreign E&P. The parameter of interest, 1-Rate, captures the sensitivity of MNC profit allocation to changes in tax rates. Column (1) presents estimates using statutory tax rates and Column (2) presents estimates using average rates. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

interpretation, which would require conditional exogeneity of τ_{ct}.[33] Instead, we focus on whether these parameter estimates are sensitive to measurement error.

Tables B.5 and B.6 show estimates for specification Equation B.2 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.

Tables B.7 and B.8 present unadjusted and adjusted results for specification Equation B.3. 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

[33]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 B.7: Dowd, Landefeld and Moore (2017) 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

Notes: This table displays coefficients from regressions under Equation B.3. The dependent variable is log profits for MNC i in country c in period t, using unadjusted SOI foreign E&P. The parameter of interest is an interaction term between 1-Rate, which captures the sensitivity of MNC profit allocation to changes in tax rates, and a dummy variable equal to one if country c is classified as a tax haven. Columns (1) and (2) classify tax havens as defined in Gravelle (2013), the same methodology used by DLM (Figure B.9). Columns (3) and (4) define havens as the countries in the smallest decile of tax rate distribution by year. Odd columns present estimates using statutory tax rates and even columns present estimates using average rates. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

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.

Finally, Tables B.9 and B.10 present unadjusted and adjusted results for specification Equa-

Table B.8: Dowd, Landefeld and Moore (2017) 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

Notes: This table displays coefficients from regressions under Equation B.3. The dependent variable is log profits for MNC i in country c in period t, using adjusted SOI foreign E&P. The parameter of interest is an interaction term between 1-Rate, which captures the sensitivity of MNC profit allocation to changes in tax rates, and a dummy variable equal to one if country c is classified as a tax haven. Columns (1) and (2) classify tax havens as defined in Gravelle (2013), the same methodology used by DLM (Figure B.9). Columns (3) and (4) define havens as the countries in the smallest decile of tax rate distribution by year. Odd columns present estimates using statutory tax rates and even columns present estimates using average rates. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

tion B.4. As in DLM, the second-order keep rate coefficient is significant, and actually increases in magnitude for the adjusted results.

Table B.9: Dowd, Landefeld and Moore (2017) 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

Notes: This table displays coefficients from regressions under Equation B.4. The dependent variable is log profits for MNC i in country c in period t, using unadjusted SOI foreign E&P. The parameter of interest, 1-Rate, is the semi-elasticity of sensitivity of MNC profit allocation to changes in tax rates. Column (1) presents estimates using statutory tax rates and Column (2) presents estimates using average rates. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Country-Level Regressions

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,

$$\log \pi_{ct} = \alpha + \beta_1 \log K_{ct} + \beta_2 \log L_{ct} + \beta_3(1 - \tau_{ct}) + \beta_4 X_{ct} + \mu_t + \gamma_c + \epsilon_{ct}, \tag{B.5}$$

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

Table B.10: Dowd, Landefeld and Moore (2017) 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

Notes: This table displays coefficients from regressions under Equation B.4. The dependent variable is log profits for MNC i in country c in period t, using adjusted SOI foreign E&P. The parameter of interest, 1-Rate, is the semi-elasticity of sensitivity of MNC profit allocation to changes in tax rates. Column (1) presents estimates using statutory tax rates and Column (2) presents estimates using average rates. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

We aggregate the sample described in the previous section to the country level, and estimate similar regressions, presented in Tables B.11 and B.12 below. Even columns in Table B.11 include the haven interaction term analogous to the specification Equation B.3, and the same columns in Table B.12 include second-order polynomials in the tax rate analogous to specification Equation B.4.[34] B&R present results with respect to the tax rate (as opposed to the keep rate in prior regressions). We do the same to facilitate comparison with their results.

In both tables, Column (1) through (4) present estimates using unadjusted foreign earnings, and Column (5) through (8) present estimates using adjusted foreign earnings. As in B&R,

[34]We include the same tax havens used by B&R—Bermuda, the Cayman Islands, Ireland, Luxembourg, the Netherlands, Singapore, and Switzerland.

estimates are computed with and without country fixed effects. Table B.11 demonstrates that, unlike B&R, the coefficient on the haven interaction remains large and significant, although slightly attenuated. Results in Table B.12 are mixed—the adjusted estimates appear to weaken the non-linear term.

Table B.11: Replication of Blouin and Robinson (2020) 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

Notes: This table displays coefficients from regressions under Equation B.5. The dependent variable is aggregate foreign earnings for a foreign affiliate located in country c in year t. The parameter of interest is effective foreign tax rate. Even columns include the same tax haven interaction term as presented in Equation B.3 (Tables B.7 and B.8). Columns (1) through (4) use unadjusted SOI foreign E&P; Columns (5) through (8) use adjusted E&P. Columns (3), (4), (7), and (8) also include country fixed effects.

Table B.12: Replication of Blouin and Robinson (2020) 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

Notes: This table displays coefficients from regressions under Equation B.5. The dependent variable is aggregate foreign earnings for a foreign affiliate located in country c in year t. The parameter of interest is effective foreign tax rate. Even columns include the same tax haven interaction term as presented in Equation B.4 (Tables B.9 and B.10). Columns (1) through (4) use unadjusted SOI foreign E&P; Columns (5) through (8) use adjusted E&P. Columns (3), (4), (7), and (8) also include country fixed effects.

C Cost Sharing Agreements and Cost Relabeling

In this appendix, we provide evidence that some US MNCs use cost sharing agreements to relabel foreign costs as domestic costs for tax purposes, which is a form of domestic-to-foreign profit shifting. As shown in Figure 11, CSAs are commonly used by MNCs that adopt HTPs.

The appendix is structured as follows. Appendix C.1 discusses the institutional details behind a regulatory ambiguity that allowed MNCs to exclude stock-based compensation from their CSAs, and the legal battles that emerged in the 2000s. We 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. Appendix C.2 describes the data we use to identify exposed firms and measure their responses. Appendix C.3 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.

C.1 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):

> 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.[35] 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.

[35]Note that this is not a departure from the arm's length standard insofar as the valuation mimics a similar exercise that might be undertaken by unrelated parties.

C.1.1 Cost Sharing Agreements

Consider a US MNC 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 costs 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 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 is immediately subject to US tax.

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

C.1.2 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.[36]

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.[37] 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 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.

C.1.3 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 MNC, 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.[38]

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

[36] In fact, they are—ESOs dilute the ownership of existing shareholders. See Bodie, Kaplan and Merton (2003).
[37] See IRS (2010).
[38] See Altera v. Commissioner (2019). See also Foley and Kuvadia (2021).

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 leading international tax scholar at the University of Michigan wrote:

> On May 7 the Ninth Circuit decided Xilinx v. Commissioner. By a 2-1 majority, 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 MNCs. 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 (McWaters, 2020).

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

Table C.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 our merged sample of MNCs with and without 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.

that a large number of firms structured their foreign operations to take advantage of the ESO technicality.

C.2 Data

We combine information from several different datasets to examine MNCs' response to the period of regulatory uncertainty following the 2005 Xilinx decision. We use IRS tax data at the foreign affiliate level to identify MNCs 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. We merge this sample of MNCs with information on stock prices from the Center for Research in Security Prices (CRSP). We also obtain historical data from Compustat that contain information about stock compensation and global R&D expenses.

Table C.1 provides summary statistics for the merged sample. We identify 168 MNCs 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.

C.3 Analysis of Multinational Response

Below, we provide two analyses that show how MNCs were affected by the positive shock from regulatory uncertainty created by the Xilinx decision. First, we show that over the short term, firm values for exposed MNCs 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, we show that over the ensuing years, exposed firms increased their usage of stock-compensation and R&D activity compared to other MNCs. We also show that their claimed deductions under the domestic R&D tax credit increased relative to the consolidated R&D cost on their books.

C.3.1 Tax Court Ruling Effect on Firm Values

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

$$R_{ct} = X_t\theta_c + \gamma E_{ct}(k) + \epsilon_{ct}, \tag{C.1}$$

where R_{ct} are log raw daily returns for each exposed MNC. We 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 we use daily log returns from the S&P 500. We 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 $\beta = \gamma k$.[39] The regression is estimated using a pre-period of 115 days prior to the event, and k days after the event. We also estimate placebo effects by computing cumulative abnormal returns for the 5 days prior to the event. Firms are weighted by size.[40]

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, we 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.

[39] In practice, we set E_{ct} equal to $1/k$ to recover β directly.
[40] We use total assets as reported in Compustat as a proxy for MNC size.

Figure C.1: Firm Response to Tax Court Decision on Aug. 30, 2005

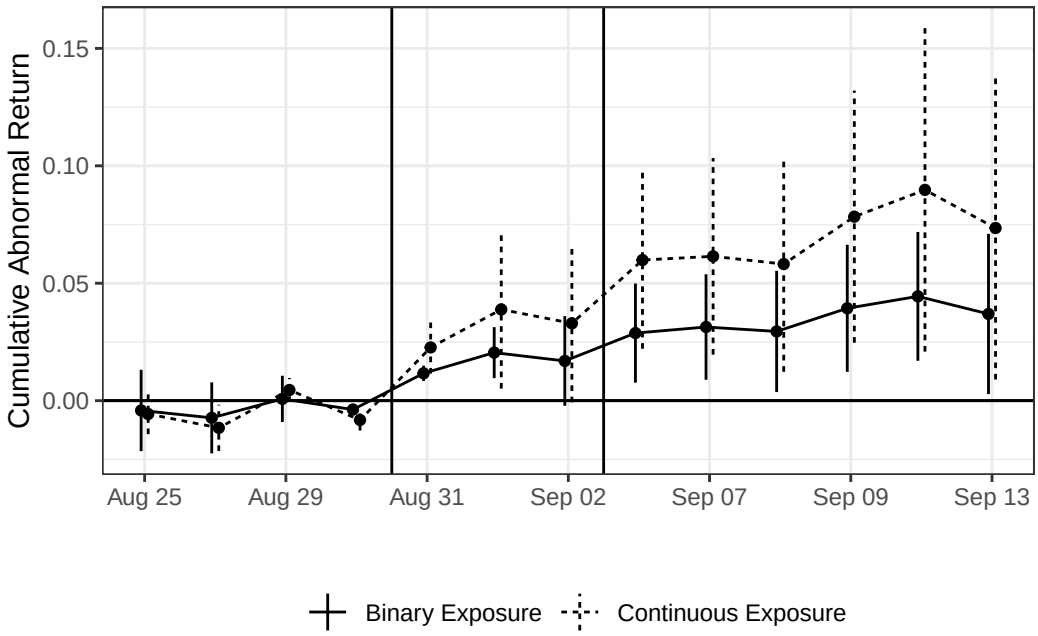

Notes: This figure presents cumulative abnormal return estimates, produced from raw returns, for the 10 trading days following the August 30, 2005 Tax Court ruling, from event study regressions of the form in Equation C.1. The vertical lines encompass estimates of cumulative abnormal returns for the sample during this period. Standard errors are clustered by both date and firm with the widest confidence interval shown for both types of clustering. Both exposures indicate sizeable firm value growth in the 10 trading days following the ruling.

Figure C.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), we 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 C.2 shows estimates produced from the model that includes market returns. Compared

Figure C.2: Firm Response to Tax Court Decision on Aug. 30, 2005 (Robustness)

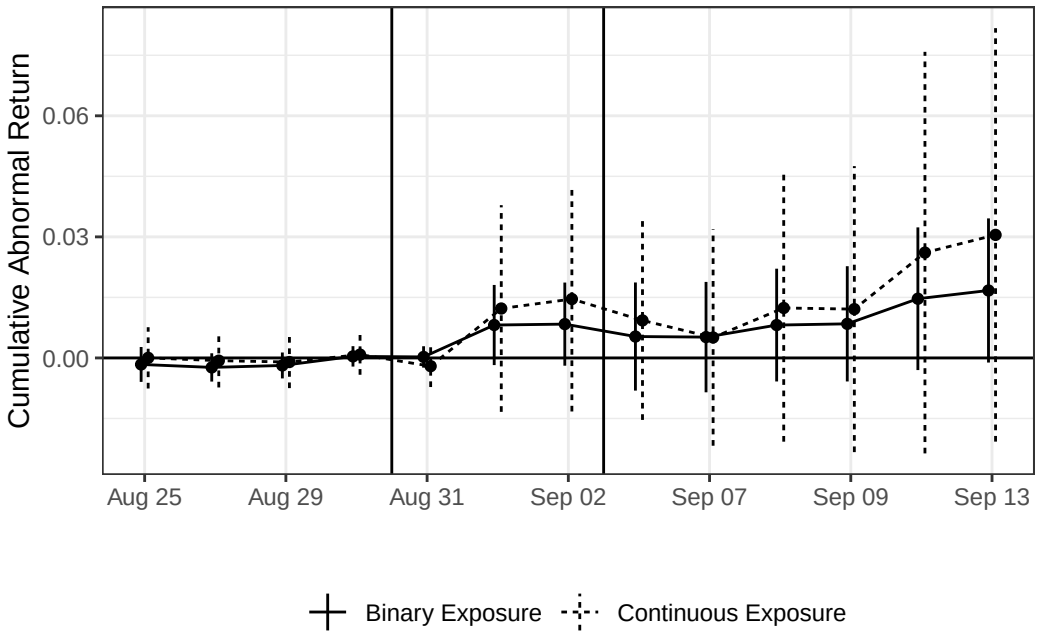

Notes: This figure presents cumulative abnormal return estimates, produced from market returns, from event study regressions of the form in Equation C.1. The vertical lines encompass estimates of cumulative abnormal returns for the sample during this period. Standard errors are clustered by both date and firm with the widest confidence interval shown for both types of clustering. These exposures show a less significant effect than indicated by Figure C.1.

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.

C.3.2 Long-Term Effects on Stock Compensation and Evidence for R&D Relabeling

Above, we demonstrate that MNCs who used cost sharing agreements received a positive shock from the Tax Court ruling, suggesting that the market believed these firms might benefit from its regulatory implications. Below, we show that these MNCs responded to this shock in the following years by increasing their use of ESOs and R&D. We also provide evidence that these firms increased claims for R&D tax credits as a share of total R&D expenditure. This indicates that the 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 expenses provided in Compustat, we run a standard, unstaggered event study specification to estimate the growth in stock compensation,

$$Y_{ct} = \sum_{t>=2005} D_{ct} + \gamma_c + \sum_{g \in G} \mu_{gt} + \varepsilon_{it}, \tag{C.2}$$

where Y_{ct} is a measure of stock compensation for MNC c in year t, and D_{ct} is an indicator equal to one in year t if MNC c is observed to have a cost sharing agreement. γ_c is a MNC 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. We examine a simple specification that only includes MNC 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.[41]

Figure C.3 shows the resulting estimates where Y_{ct} is the log of stock compensation as reported in Compustat. After 2005, relative to MNCs without CSAs, exposed MNCs reported large increases in stock compensation. Table C.2 contains point estimates from pooled regressions.

[41] We use consolidated asset and sales figures from Compustat to compute these quartiles.

Figure C.3: Effect on Log Stock Compensation

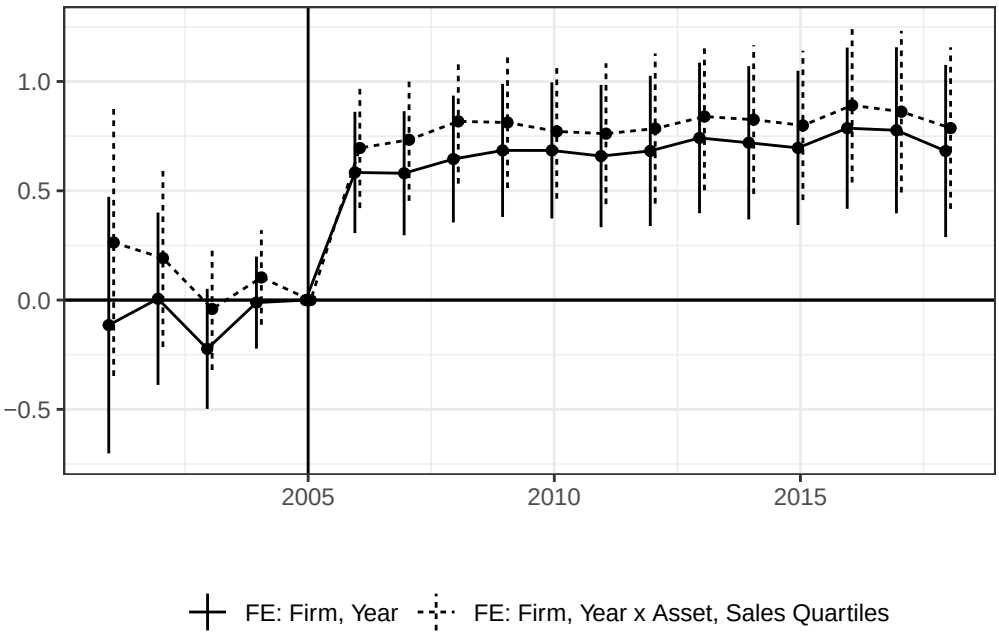

Notes: This table displays estimates from regressions under Equation C.2. The dependent variable is log stock compensation for MNC c in year t as reported in Compustat. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004.

Table C.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 C.2. The dependent variable is log stock compensation for MNC c in year t as reported in Compustat. Column (1) includes firm and year fixed effects and Column (2) includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Figure C.4: Effect on Ratio of Stock Compensation to R&D

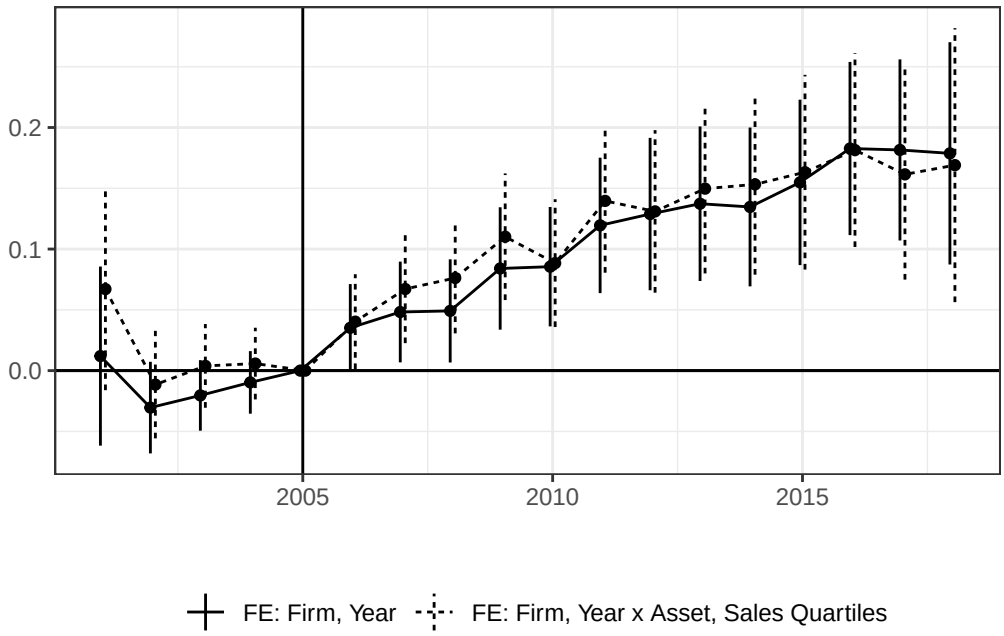

Notes: This table presents estimates from pooled regressions under Equation C.2. The dependent variable is the ratio of stock compensation for MNC c in year t to R&D expenses incurred in 2005. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004.

Figure C.4 and Table C.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. We trim this dataset to remove observations where this ratio is not contained in $[0, 1]$. This alternative specification also shows large increases in stock compensation for exposed MNCs on the order of 10% of 2004 R&D spending.

Table C.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 C.2. The dependent variable is the ratio of stock compensation for MNC c in year t to R&D expenses incurred in 2005. Column (1) includes firm and year fixed effects and Column (2) includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Shifts Toward Stock Compensation in the Wage Structure of Exposed Firms

The event studies shown above demonstrate that exposed MNCs 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 MNCs. Figure C.5 and Table C.4 show dynamic and pooled estimates from a similar specification to Equation C.2 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, we trim this dataset to remove observations where this ratio is not contained in $[0, 1]$. The results show that, relative to other MNCs, exposed MNCs shifted their wage structure toward stock compensation.

Figure C.5: Effect on Ratio of Stock Compensation to Total Labor Compensation

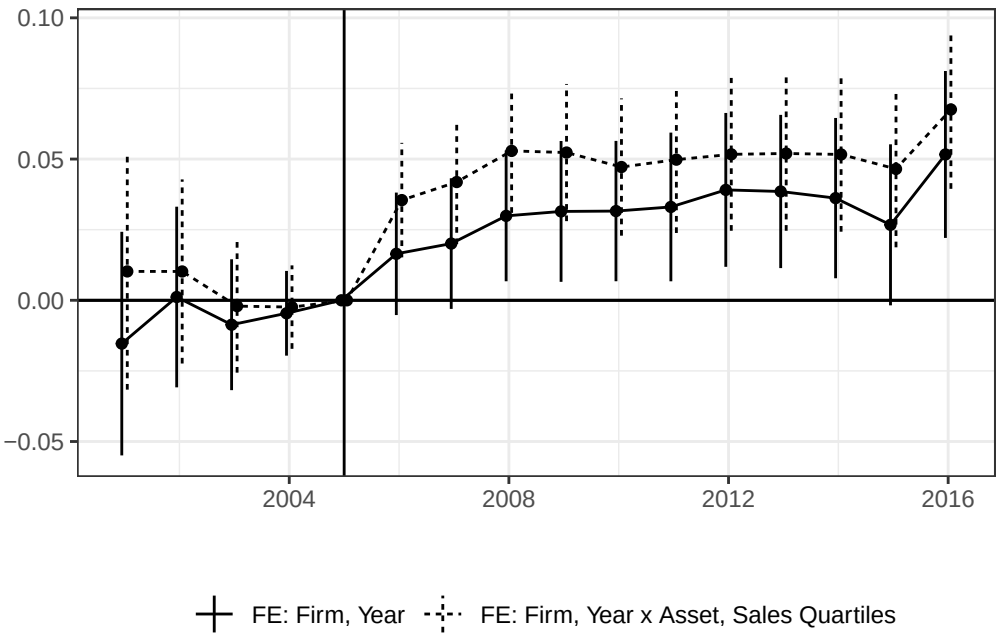

Notes: This table presents point estimates from pooled regressions under Equation C.2. The dependent variable is the ratio of stock compensation for MNC c in year t, as measured in Compustat, to the total wage bill as measured in IRS data. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004.

Table C.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 C.2. 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. Column (1) includes firm and year fixed effects and Column (2) includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Table C.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 C.2. 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. Column (1) includes firm and year fixed effects and Column (2) includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

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 C.6 and Table C.5 show dynamic and pooled estimates from a similar specification to Equation C.2 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.[42] Note that, if the results shown above were confounded by international expansion of the group of exposed firms, one would expect this ratio to decline as their domestic share of R&D declines. Instead, this ratio is **increasing**, indicating that exposed MNCs performed domestic R&D relatively more intensely than other MNCs. 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 C.10 and Table C.9.

Threats to Identification and Robustness Checks

The results above provide compelling evidence that exposed MNCs reacted strongly to the period

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

Figure C.6: Effect on Ratio of Claimed R&D Expenses to Compustat R&D

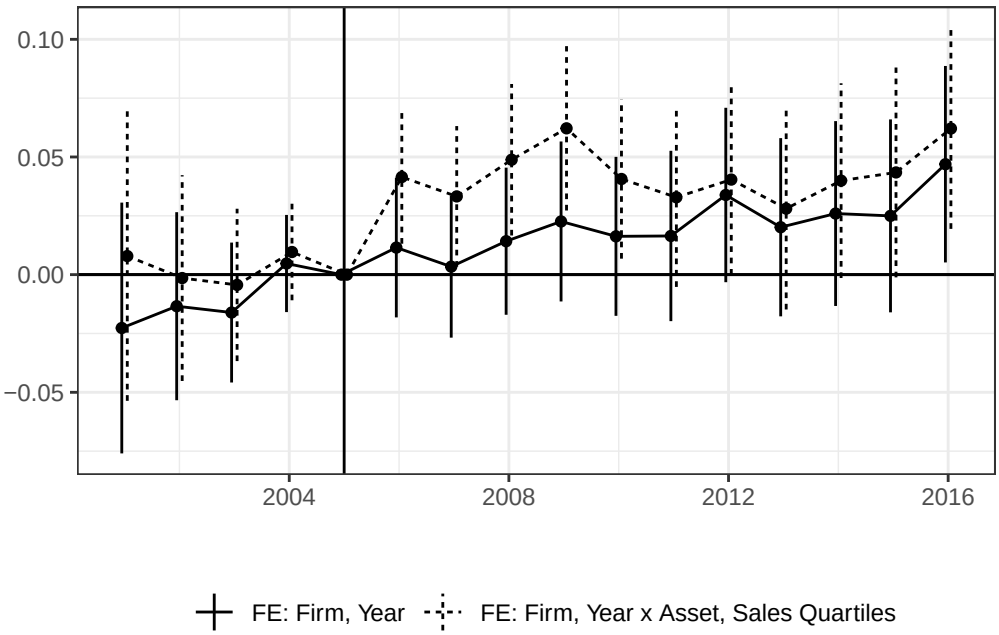

Notes: This table presents point estimates from pooled regressions of the form

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

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 and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004.

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 C.7 through C.9 and Tables C.6 through C.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 MNCs—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 C.6 and Figure C.10) should not be affected by this potential measurement issue.

C.4 Additional Figures and Tables

Figure C.7: Effect on Log Stock Compensation (Robustness, Ref Year = 2006)

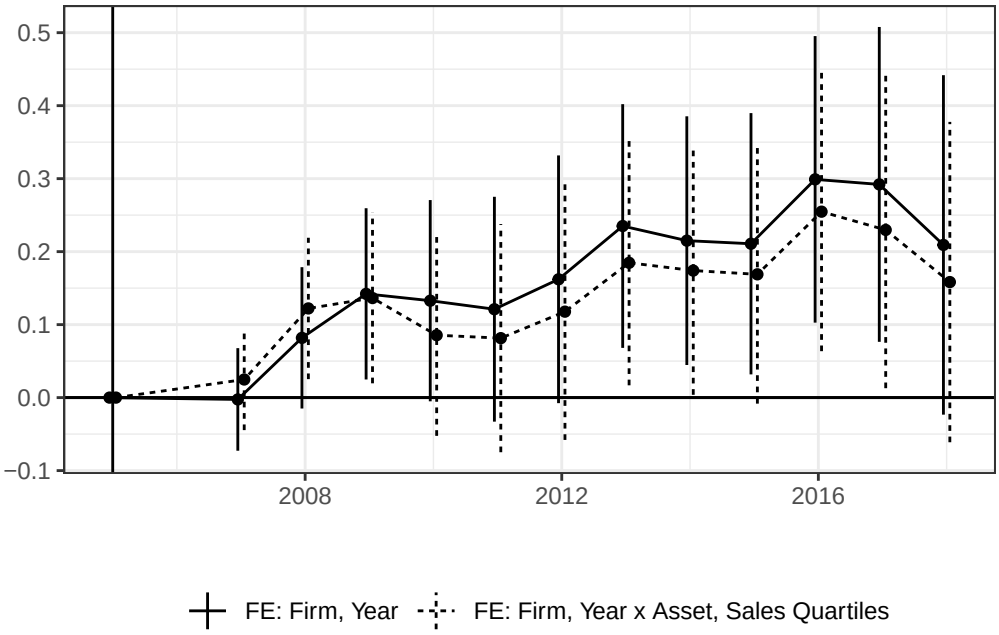

Notes: This table presents estimates, using an alternate reference year of 2006, from regressions under Equation C.2. The dependent variable is log stock compensation for MNC c in year t as reported in Compustat—comparable to Figure C.3. The figure displays estimates under two different sets of fixed effects: the solid line includes only firm and year fixed effects, while the dashed line includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004.

Table C.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 C.2. The dependent variable is log stock compensation for MNC c in year t as reported in Compustat—comparable to Table C.2. Column (1) includes firm and year fixed effects and Column (2) includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Figure C.8: Effect on Ratio of Stock Compensation to R&D (Robustness, Ref Year = 2006)

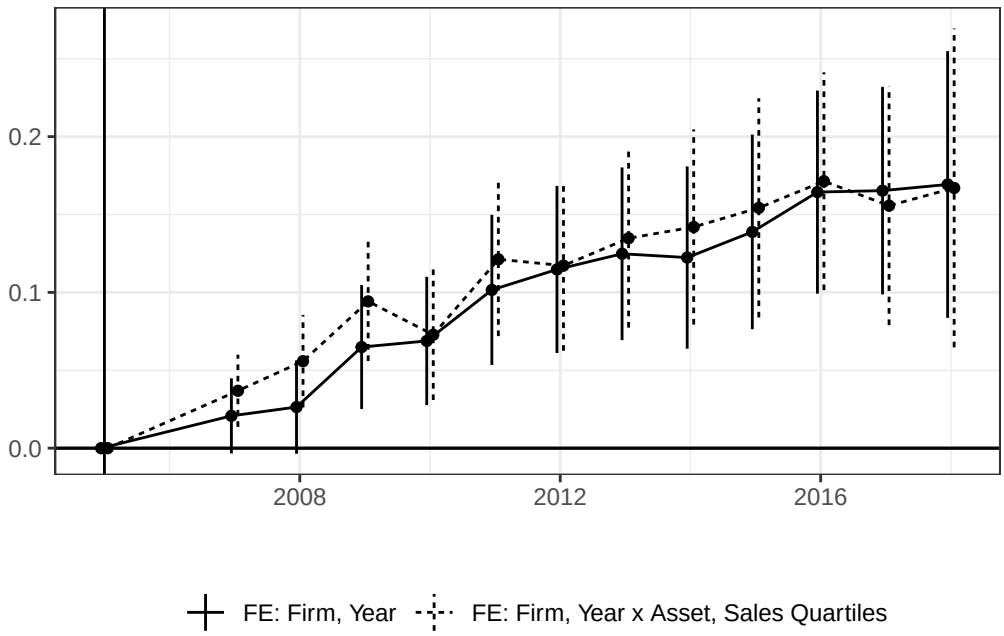

Notes: This table presents estimates, using an alternate reference year of 2006, from regressions under Equation C.2. The dependent variable is the ratio of stock compensation for MNC c in year t to R&D expenses incurred in 2005—comparable to Table C.3. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Table C.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 C.2. 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 C.3. Column (1) includes firm and year fixed effects and Column (2) includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Figure C.9: Effect on Ratio of Stock Compensation to Total Labor Compensation (Robustness, Ref Year = 2006)

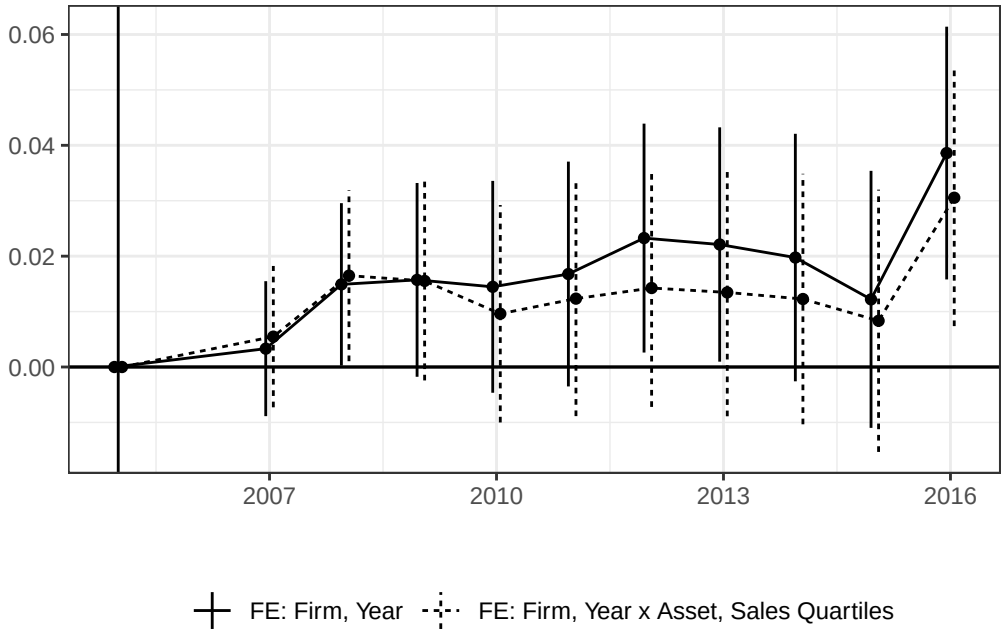

Notes: This table presents estimates, using an alternate reference year of 2006, from regressions under Equation C.2. The dependent variable is the ratio of stock compensation for MNC c in year t to R&D expenses incurred in 2005—comparable to Figure C.5. The figure displays estimates under two different 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 C.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 point estimates, using an alternate reference year of 2006, from pooled regressions under Equation C.2. 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. Column (1) includes firm and year fixed effects and Column (2) includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Figure C.10: Effect on Log R&D

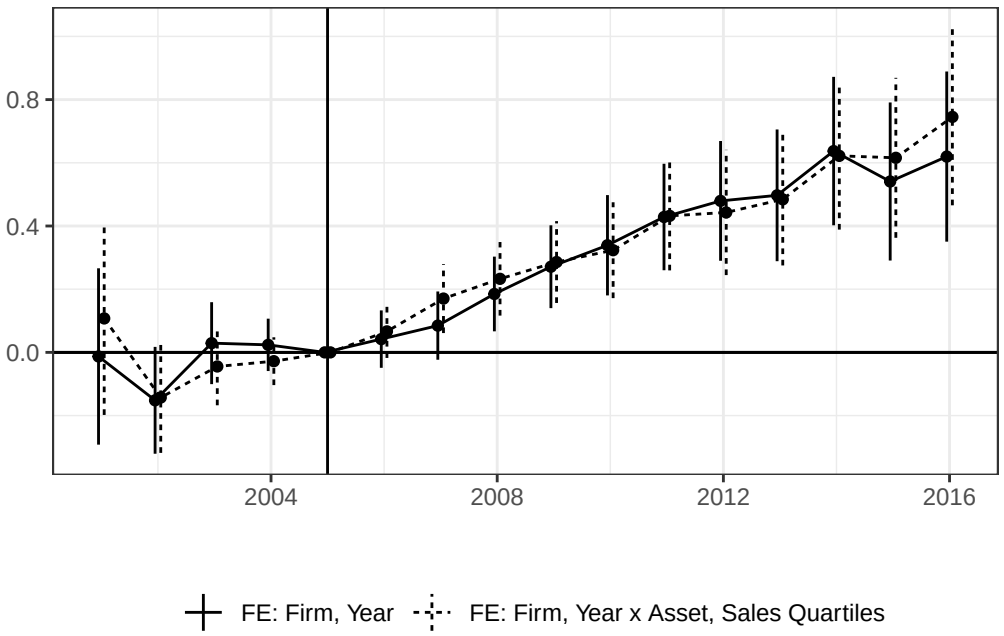

Notes: This table presents estimates, using an alternate reference year of 2006, from regressions under Equation C.2. The dependent variable is the ratio of stock compensation for MNC c in year t, as measured in Compustat, to the total wage bill as measured in IRS data—comparable to Table C.4. *, **, and *** denote statistical significance at the 10, 5, and 1% level.

Table C.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 C.2. The dependent variable is log R&D expenses incurred in 2005. Column (1) includes firm and year fixed effects and Column (2) includes firm fixed effects and year fixed effects interacted with indicators for quartile bins of assets and sales in 2004. *, **, and *** denote statistical significance at the 10, 5, and 1% level.