Substance-ial Investment Shifting: The Role of Substance in Anti-Tax Avoidance Rules*

Ruby Doeleman†

This Version: 2 August 2026

Work in progress.

Abstract

This paper analyzes how multinational enterprises structure their corporate groups under anti-tax avoidance rules with substance exemptions. Substance exemptions stipulate that if an MNE has enough investment in a low-tax country, it can continue shifting its profits there. Multinational enterprises weigh the tax benefits of shifting profits to tax-haven affiliates against the investment required under substance rules. The model predicts two effects: an extensive margin effect where substance rules reduce the number of tax-haven affiliates, and an intensive margin effect where remaining tax-haven affiliates receive greater investment. I confirm this pattern in a staggered difference-in-differences design using affiliate-level data from 2010–2023.

Keywords: Real Investment, Profit Shifting, Anti-Tax-Avoidance Rules

JEL Classification: H25, H26, F23

*I thank Harald Amberger, Kurt Brekke, Antonia Hohmann, Dominika Langenmayr, Kevin Markle, Johannes Scheurer, Dirk Schindler, participants at the 80th IIPF Conference and the 2024 RSIT Conference, and seminar participants at LMU, NHH, the Stanford Tax Lab, and WU Vienna for their valuable comments and suggestions. I acknowledge support by the Austrian Science Fund (FWF) (Doc 92-G).

†WU Vienna. E-mail: ruby.doeleman@wu.ac.at.

1 Introduction

Multinational enterprises (MNEs) can reduce their tax burden by shifting profits from affiliates in high-tax countries to affiliates in tax havens (Merlo and Wamser, 2024). Most MNEs have multiple tax-haven affiliates: Among the hundred largest MNEs worldwide, the average corporate group structure includes 549 affiliates, of which 68 are located in offshore financial centres (United Nations Conference on Trade and Development, 2016). These tax-haven affiliates are geographically dispersed and tend to be located near the group's productive affiliates, rather than in a single global hub. Why an MNE operates *multiple* tax-haven affiliates, and why it places them close to where it produces, is not yet well understood.

This paper studies the organizational structure of MNEs in response to anti-tax avoidance rules. I argue that MNEs trade off the benefit of shifting profit to a nearby tax-haven affiliate against the cost of establishing and maintaining each one, and that the cost of shifting profits increases with distance. Shifting profit to a close-by tax-haven affiliate is cheaper than shifting to a far-away one, because proximity lowers the cost of familiarization with local legal institutions and establishing the effective place of management. This gravity-like structure is well documented in the profit-shifting literature (Desai et al., 2004, 2006; Laffitte, 2024; Ferrari et al., 2023; Tørsløv et al., 2023).

To identify how MNEs make locational choices under anti-tax avoidance rules, I analyze anti-tax avoidance rules that affect MNE group structures. To align taxation with real economic activities in low-tax affiliates, many anti-tax avoidance rules provide carve-outs in the form of substance requirements (OECD, 2015). If the MNE is successful in arguing it has sufficient substance in low-tax affiliates, the anti-tax avoidance rule is not applied, and the MNE can continue shifting profits.[1]

In this paper, I assess the tension between the goal of curbing profit shifting from high-tax countries to tax havens, versus the goal of fostering economic activity in high-tax countries. Substance rules could imply a shift from paper profit shifting toward shifting of real activity to satisfy substance requirements. Therefore, I analyze whether substance rules incentivize MNEs to increase employment and fixed assets in tax-haven affiliates, and whether this affects investment in high-tax affiliates.

To analyze the decision of an MNE to set up tax-haven affiliates under substance exemptions, I model an MNE's group structure using a Salop (1979) framework, based

[1] See, e.g., Art. 7(2)a of Council Directive (EU) 2016/1164 of 12 July 2016 laying down rules against tax avoidance practices that directly affect the functioning of the internal market.

on the literature on regional tax havens (Desai et al., 2004, 2006; Laffitte, 2024; Ferrari et al., 2023; Tørsløv et al., 2023) and previous analytical work that uses a Salop framework to model MNE behavior (Brekke et al., 2017; Gresik and Konrad, 2017). Making use of a close-by tax-haven affiliate is cheaper than a far-away tax-haven affiliate for two main reasons. First, tax advisors in the high-tax country might be more familiar with the legal institutions in a close-by low-tax country. Second, to operate a tax-haven affiliate, the effective place of management requires flying in directors to the tax-haven affiliate. This explains the notion of adjustment costs created by the *legal* distance between productive affiliates and tax-haven affiliates. The model optimally trades off the benefits of shifting profit to a close-by tax-haven affiliate against the costs of setting up a tax-haven affiliate.

The model is solved in three stages by backward induction. In the first step, each productive affiliate chooses the optimal amount of profit to shift to its nearest tax-haven affiliate. The affiliate effectively equates the marginal benefit of the tax rate differential with the marginal costs of profit shifting, which increase in the legal distance to the tax-haven affiliate. In the second step, the MNE chooses the optimal level of investment in each tax-haven affiliate, trading off the marginal benefit of reducing adjustment costs against the marginal cost of investment required by the substance rule. In the third step, the MNE determines the optimal number of tax-haven affiliates in the group structure, balancing the benefits of having nearby tax-haven affiliates for profit shifting against the costs of investing in each one.

Substance rules change the group structure of MNEs along two margins. The *extensive margin* is the number of tax-haven affiliates the group holds. The *intensive margin* is the amount of real investment—fixed assets and employment—in each tax-haven affiliate. A substance rule increases the fixed cost of each tax-haven affiliate. On the extensive margin, this incentivizes the MNE to hold *fewer* tax-haven affiliates. On the intensive margin, each remaining tax-haven affiliate now serves a larger share of the group's productive affiliates, so the MNE invests *more* in each one. In short, the group consolidates its shifting into fewer affiliates with larger substance.

To empirically assess the intensive and extensive margin effects of substance rules, I analyze whether MNEs increase real activity in the form of tangible fixed assets, number of employees, and costs of employees after becoming subject to a substance rule. I implement a staggered difference-in-differences design using data from Bureau van Dijk's Historical Orbis over the years 2010–2023 covering affiliates in 68 different countries. In addition to data on substance rules from the OECD (2024), I hand-collect data from national CFC legislation. I exploit variation in whether, when, and how

substance rules are implemented and essentially compare affiliates in the same country that differ in where their shareholders are located. The treatment group contains MNE affiliates that face a tax rate below the CFC tax rate threshold and become subject to a CFC rule that incorporates a substance rule. The control group consists of affiliates that face a CFC rule that does not contain a substance rule and affiliates that face a tax rate above the CFC threshold.

For the extensive margin, I find that the share of substance affiliates in the MNE group decreases significantly after the MNE becomes subject to a substance rule. On the intensive margin, I find that tangible fixed assets, the costs of employees, and the number of employees increase after an affiliate becomes subject to a substance rule. This pattern is in line with the model: the substance rule increases the fixed cost of investment for each tax-haven affiliate, so the MNE reduces the number of tax-haven affiliates and concentrates substance in the remaining tax-haven affiliates. These results are in line with the notion that MNEs exploit substance rules in order to forego application of anti-tax avoidance rules.

The interesting question is whether the increase in substance in tax-haven affiliates is due to an investment-shifting effect, where MNEs move real activity from high-tax affiliates to tax-haven affiliates. To analyze this, I estimate whether substance rules lead to a relocation of fixed assets and employment from high-tax affiliates to tax-haven affiliates. I do not find significant evidence for such investment shifting behavior, which is in contrast with Suárez Serrato (2019), Williams (2018), and Drake et al. (2022). The absence of investment shifting suggests that real activity in low-tax and high-tax affiliates is complementary rather than substituted: meeting substance requirements in tax-haven affiliates does not appear to come at the expense of investment in high-tax affiliates, consistent with the view that profit shifting can lower the cost of capital for the MNE as a whole.

The paper's primary contribution is to analyze how MNEs structure their corporate groups in response to anti-tax avoidance rules. This sets the paper apart from other real-effects literature. Whereas the existing literature estimates how MNEs' real investment responds to anti-avoidance rules, I model the organizational choice itself—how many tax-haven affiliates a group holds and where it locates them—and identify how substance carve-outs in CFC rules reshape that structure along an extensive and an intensive margin. This paper contributes to the literature on the real effects of anti-avoidance rules (De Simone and Olbert, 2022; Suárez Serrato, 2019; Williams, 2018; Drake et al., 2022; Buettner et al., 2018; De Mooij and Liu, 2020) and to the theoretical literature on CFC rules and tax competition (Weichenrieder, 1996; Haufler et al.,

2018; Pieretti and Pulina, 2020; Clifford, 2019). This paper also relates to the literature that poses that profit shifting to low-tax countries reduces the cost of capital in high-tax countries (Hong and Smart, 2010; Becker and Riedel, 2012). Previous literature shows that taxation affects location decisions (Voget, 2011; Beer et al., 2020), both when it comes to setting up affiliates and allocating investment within the MNE group. The real effects of substance rules depend on whether investment in high-tax and tax-haven affiliates is complementary: If substance rules lead to shifting of fixed assets and employment, this may reduce activities in high-tax countries (Suárez Serrato, 2019; Williams, 2018; Drake et al., 2022).

By analyzing whether substance rules distort real investment, I aim to provide policy implications for the use of substance rules to curb profit shifting. The alignment with substance is important in assessing transfer pricing: The OECD Transfer Pricing Guidelines stipulate that dealing at arm's length with affiliates requires a functional analysis of economically significant activities and responsibilities (OECD, 2022, p. 46). Substance therefore allows allocating profits to affiliates. The OECD's global minimum tax also contains a substance-based income exclusion, putting substance at the forefront of current tax policy.

In the following, I discuss the institutional background and stylized facts to motivate the model in Section 2. Subsequently, in Section 3, I review related literature. Section 4 contains the model and hypotheses, and Section 5 describes the data. Section 6 analyzes the extensive margin, Section 7 analyzes the intensive margin, and Section 8 analyzes spillover effects. Section 9 concludes.

2 Institutional Background & Stylized Facts

One of the most common types of legislation to counter profit shifting are Controlled Foreign Company (CFC) rules.[2] In Action 3 of its Base Erosion and Profit Shifting project, the OECD recommends CFC rules to combat tax avoidance (OECD, 2015). Since 2019, the EU requires its Member States to implement CFC rules in the Anti-Tax Avoidance Directive 2016/1164. Because some countries implement CFC rules with substance requirements, this is a useful setting to analyze the effects of substance rules on real activities of MNEs.

[2]CFC rules were not originally developed to combat tax avoidance, however. In 1962, the U.S. Kennedy administration introduced a CFC system in Subpart F to mitigate the post-war balance of payments deficit. The focus on passive income in the CFC rule was a compromise reached after Congress opposed including active income, see Kofler et al. (2020, Chapter 1).

In short, CFC rules are a cross-border corporate tax on foreign affiliates of domestic MNEs. Under a CFC rule, if a parent company controls a foreign affiliate, certain types of passive income of that subsidiary may be attributed to the parent entity and taxed accordingly. Thus, CFC rules essentially create an exception to the tax exemption principle in territorial tax systems. Take the Dutch CFC rule for example: It stipulates that if a Dutch taxpayer has more than 50% ownership in a foreign entity that faces a statutory tax rate below 9%, the Dutch taxpayer must include in its tax base the CFC's passive income, proportional to the ownership that the Dutch taxpayer holds in the CFC.[3]

CFC rules have several scoping characteristics, including the definition of controlling ownership, the tax rate threshold, and types of income to be attributed. First, controlling ownership means having decisive influence, often defined as a threshold of more than 50% of ownership (OECD, 2015, p.21). Second, the tax rate threshold determines the tax rate below which a CFC becomes subject to CFC rules (OECD, 2015, p.33). This threshold can be absolute or relative to the domestic tax rate, and can be based on statutory or effective tax rates. Instead of a tax rate threshold, some countries apply black lists and white lists that list countries that are in or out of scope of the CFC rule. Third, types of income targeted by CFC rules usually include passive income, such as dividend, interest, royalties, rents, capital gains and other income that is deemed passive (OECD, 2015, p.43).

Many countries have also implemented exemptions to the application of CFC rules. These exemptions can be categorized as being in one (or all) of three forms: active income exemptions (transactional), business reason exemptions (motive), and active business exemptions (functional).[4] First, active income exemptions usually apply a safe harbor ratio of passive income to total income.[5] Second, business reason exemptions require that a certain arrangement has not been put in place for the essential purpose of obtaining a tax advantage.[6] Third, there are active business exemptions, so-called substance rules.[7] These substance rules essentially require taxpayers to demonstrate that the CFC carries on substantive economic activity supported by staff, equipment, assets, and premises.

[3]The rule is more detailed. See Art. 13ab of the Dutch Corporate Income Tax Act 1969.

[4]I categorize these exemptions based on data from OECD (2024) and hand-collected data from national tax legislation. Examples can be found in EU Directive 2016/1164 (ATAD1).

[5]For a transactional exemption, see Art. 7(4)b ATAD1.

[6]For a motive exemption, see Art. 7(2)b ATAD1.

[7]For a functional exemption, see Art. 7(2)a ATAD1. See OECD (2015, p. 47) for an overview of the considerations on using substance analysis.

Because profit shifting provides tax benefits, MNEs are incentivized to circumvent the application of CFC rules by manipulating the three scoping thresholds above (Clifford, 2019). First, MNEs can manipulate the control requirement, but countries have implemented anti-abuse provisions that will test whether the MNE actually controls the CFC. Second, MNEs can make use of passive income thresholds.[8] Third, MNEs can manipulate tax rate thresholds, primarily if these are based on the effective tax rate.[9] Fourth, MNEs can exploit substance rules.

Substance rules are central to this paper. These rules provide an incentive to shift real investment, because they exempt an affiliate from CFC rules if the CFC has sufficient employees and assets, i.e., a substantial physical presence. I obtain data on all substance rules worldwide from OECD (2024) and I hand-collect data from national legislation on the details of the substance requirement. Table A1 in the Appendix gives an overview of countries' substance legislation. For illustrative purposes, Table 1 shows a sample of countries that have implemented substance rules.

Some countries apply substance rules in the form of *safe harbors*, implying that a CFC must have a minimum measurable amount of substance. For example, until 2020, the Netherlands applied a safe harbor. In short, this required a CFC to have at least EUR 100,000 in employee expenses and a permanent office. As Table 1 shows, Canada's substance rule also contains a partial safe harbor, requiring that the CFC transacts principally with arm's length persons and has more than five full-time employees. Another example is South Africa, which lists types of businesses that qualify as real economic activity.

Substance rules also come in the form of *subjective tests*. These tests require a case-by-case evaluation of all relevant facts and circumstances that provide evidence that the CFC mainly carries out an economic activity. This often involves testing whether a CFC has enough fixed assets and employees, see for example Austria and Germany in Table 1. The majority of countries apply such subjective tests. For example, Table 1 shows that Norway and Sweden require that the CFC is established for business reasons and that it engages in real economic activities. Similarly, Belgium (shown in Table 1), the U.K. and Ireland require that the CFC carries out significant people functions.

[8] Desai and Dharmapala (2009) and Gunny (2010) provide general evidence that MNEs manipulate income to shift profits.

[9] In addition to incentives for MNEs, low-tax countries have an incentive to facilitate MNEs in shifting their profits. Low-tax countries can increase their statutory corporate tax rate to increase the MNE's absolute and effective tax rate.

Table 1: Selected Examples of Substance Rules in National Legislation

Country	Introduction	Substance Rule	Burden of Proof	Legislation
Austria	2019	Significant economic activity in terms of personnel, equipment, assets and premises	Taxpayer	Art. 10a KStG 1988
Belgium	2019	Owner of assets and assume risks related to this asset – comparison whether Belgian taxpayer carries out key functions	Tax Authorities	Art. 185/2 WIB 1992
Canada	1976	More than five full-time employees in carrying out an active business	Taxpayer	Art. 95(1) ITA 1985
Germany	1972	Material resources and sufficiently qualified personnel to independently carry out real economic activity while assuming full risk, in EU CFC state only	Taxpayer	Art. 8(2) AStG 1972
Hungary	2017	Personnel, facilities, equipment and premises to carry out actual business activities; safe harbor if at least 50% of total revenue comes from production, manufacturing, agriculture, services, investment, or trading carried out with own assets and employees	Taxpayer	Art. 8(1)f Act LXXXI 1996
Italy	2001	Effective economic activity through the use of personnel, equipment, assets, and premises; economic integration in local market; safe harbor if more than 50% of costs and revenue are sourced in CFC country	Taxpayer	Art. 167 TUIR 1086
Norway	1994	In fact established in an EEA state or treaty state and carries out genuine economic activity there	Taxpayer	Art. 10-64 TA 1999
Poland	2015	Established in EU; Local, qualified, commensurate personnel and equipment within scope of activity; Economic reasons for arrangements; Independent performance of economic functions, including local management	Taxpayer	Art. 24a CITA 1992
Sweden	2008	Established in EU for business reasons and engaged in real economic activities, supported by staff and risk-taking	Taxpayer	Section 39a ITA 1999

Note: Data from Kofler et al. (2020) and national legislation.

Some countries apply a mix between a safe harbor and subjective tests. For example, Table 1 shows that Italy and Hungary apply a safe harbor combined with a subjective test, namely if at least 50% of total revenue comes from real activities carried out with own assets and employees in the CFC country.

The 2006 Cadbury-Schweppes ruling (C-196/04) by the Court of Justice of the European Union (CJEU) had an important influence on the use of substance rules within the EU (Whitehead, 2007). The ruling assessed whether the UK's CFC rules were in line with EU law. Cadbury-Schweppes had two Irish subsidiaries which were subject to CFC rules. Had the subsidiaries been established in the U.K., they would not have been subject to CFC rules. This could violate the right of establishment within the EU. However, only if the CFC is a 'wholly artificial arrangement aimed at circumventing national law', imposing a CFC rule would be allowed to prevent tax avoidance (par. 108 of C-194/04). The ruling therefore effectively introduced substance rules within the EEA, with the risk of violating the EU fundamental freedoms when failing to do so correctly.

Because of the Cadbury-Schweppes ruling, some EU countries only allow substance exemptions within the European Economic Area (EEA) and do not provide substance rules for non-EEA CFCs.[10] These countries include Germany, Greece, Norway, Poland, Sweden, and Romania. Alternatively, some countries apply a stricter substance rule outside of the EEA. Examples are France and Italy, which apply a stricter burden of proof for non-EEA CFCs. Another example is Finland, which requires that Finland and the CFC country have an Exchange of Information Agreement for applying the substance rule.

The investment effects of substance rules depend on their design. Countries define substance differently and apply different standards to assess whether the substance requirements are fulfilled. From the subjective tests described, it seems that substance is defined as being more than having an employee in an office picking up the phone. This suggests that substance rules could have consequences for the shifting of real investment, not just substance on paper or "phantom investment".

2.1 Stylized Facts on MNE Group Structures

To motivate the model, I provide three stylized facts about MNE group structures. These motivate the two building blocks of the model: that MNEs hold *several* tax-haven

[10]For empirical research on the effects of these rules on investment and taxable income, see Ruf and Weichenrieder (2013) and Schenkelberg (2020).

affiliates, and that they locate them *close* to the productive affiliates. Throughout, I look at ownership structures of MNE groups as recorded in Bureau van Dijk's Historical Orbis Data. The MNE group is the set of affiliates sharing a Corporate Global Ultimate Owner, and I restrict to groups with affiliates in at least two countries. Statistics are for 2019; I provide a complete description of the sample construction in Section 5. In this section, I classify an affiliate as a tax-haven affiliate if it is located in one of the jurisdictions listed in Dharmapala and Hines (2009); all remaining affiliates are productive affiliates.

Stylized Fact 1: MNEs hold multiple tax-haven affiliates. Many MNEs have several tax-haven affiliates. Of the 276,514 MNE groups in the sample, 48.1% hold at least one tax haven affiliate. The average MNE group with at least one tax-haven affiliate holds 3.7 haven affiliates. Figure 1 plots the distribution of the number of haven affiliates per group among haven-holding MNEs. 58% of MNEs has a single tax-haven affiliate, but 25% hold three or more. 6% hold ten or more. Holding multiple tax-haven affiliates is not exclusive to the largest MNEs: among groups with 5-9 affiliates the median group holds 2 haven affiliates, increasing to 5 among groups with 25-99 affiliates.[11] This is in line with Merlo and Wamser (2024), who find that even small MNEs hold tax-haven affiliates.

Stylized Fact 2: Tax-haven affiliates are geographically dispersed and located near productive affiliates. MNEs do not hold tax-haven affiliates in a single location. Instead, haven affiliates are spread across regions, each near a cluster of the group's productive affiliates. Figure 2 maps the haven and productive affiliates of a large, globally active multinational (Panel A) and, in aggregate, of all sample groups (Panel B), illustrating that haven affiliates are geographically close to production. The illustrative group in Panel A shows its European tax-haven affiliates (the Netherlands, Belgium, Ireland, and Switzerland) close to European operations and its Asian haven affiliates (Hong Kong and Singapore) close to its Asian operations. In aggregate (Panel B), haven affiliates equally concentrate in Europe and Hong Kong.

[11] The top 20% of MNEs holds 37 affiliates on average.

Figure 1: Distribution of the Number of Tax-Haven Affiliates in MNEs

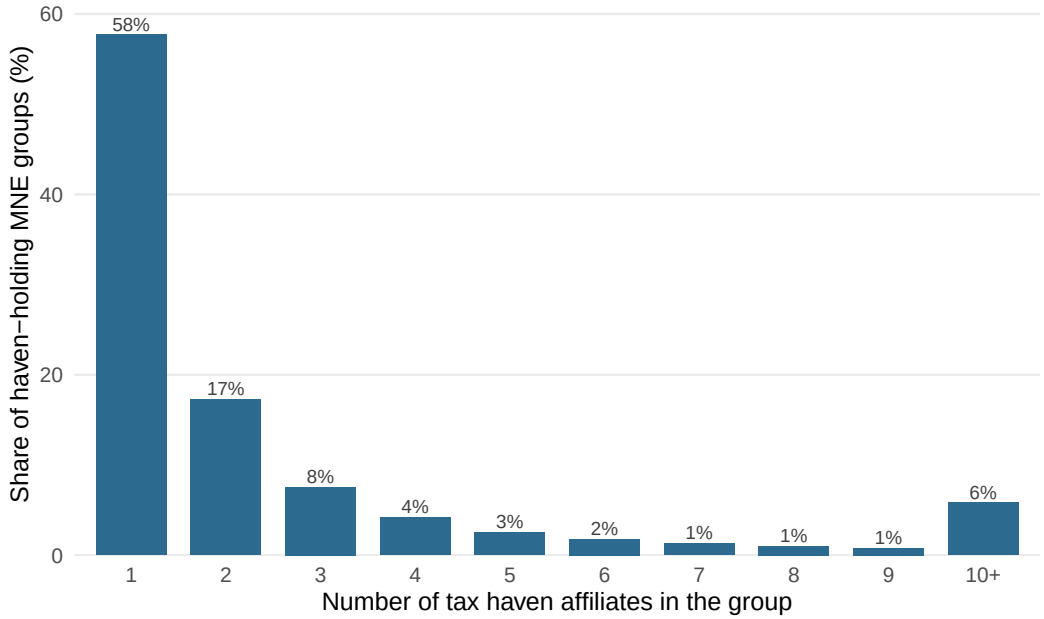

Note: Data from Orbis, 2019. The figure plots how many tax-haven affiliates MNEs hold. Groups are defined by the Corporate Global Ultimate Owner. Tax havens defined as in Dharmapala and Hines (2009).

Stylized Fact 3: Productive affiliates are close to a tax-haven affiliate, and closer when the group holds more of them. On average, a productive affiliate's nearest haven affiliate is 2,584 km away, whereas it would be 7,464 km for a tax haven drawn at random from all haven locations. Moreover, this distance falls as the group holds more haven affiliates: Figure 3 shows a negative relationship between the number of tax-haven affiliates in a group and the average distance from its productive affiliates to the nearest one.

Together, these three facts—multiplicity, geographic dispersion, and proximity—motivate the model: MNEs hold several haven affiliates and locate them close to the production they serve.

Figure 2: Geographic Spread of MNE Affiliates

Note: Data from Orbis, 2019. Panel A shows the productive affiliates (blue circles) and tax-haven affiliates (red diamonds) of a single large multinational, with line segments linking each productive location to its nearest within-group haven. Panel B aggregates all sample MNE groups; marker size is proportional to the number of affiliates. Groups are defined by the Corporate Global Ultimate Owner; tax havens follow Dharmapala and Hines (2009). Affiliate locations are capital cities.

Figure 3: Distance to the Nearest Tax-Haven Affiliate

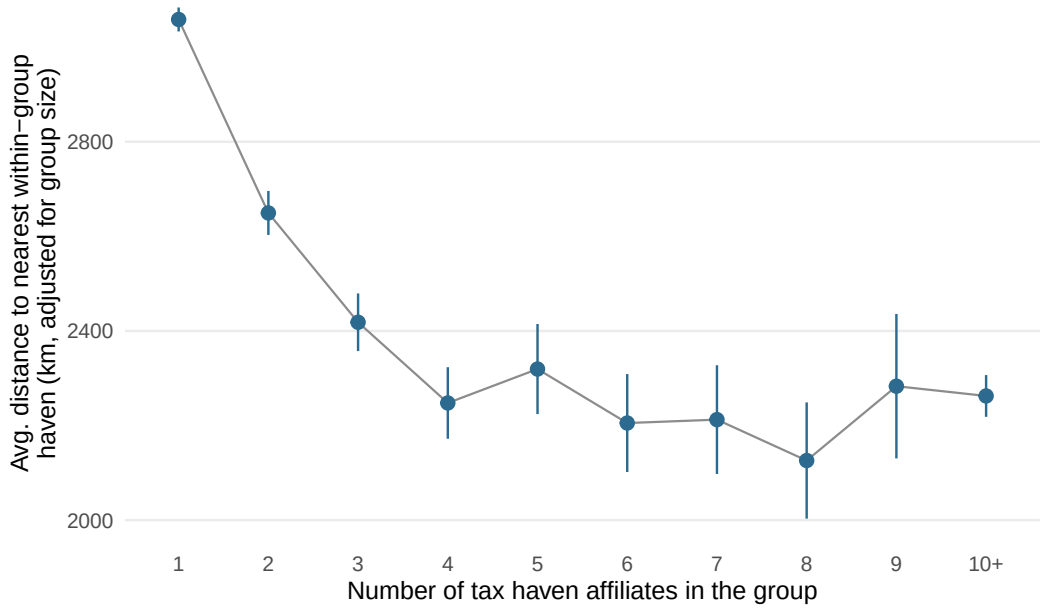

Note: Data from Orbis, 2019. Average distance with 95% confidence intervals from a group's productive affiliates to the nearest within-group haven affiliate after adjusting for total group size. Affiliate locations are capital cities. Groups are defined by the Corporate Global Ultimate Owner; tax havens follow Dharmapala and Hines (2009).

3 Literature

This paper investigates whether substance exemptions generate real effects. Therefore, this paper builds on the literature around the real effects of anti-tax abuse rules and the potentially beneficial role of profit shifting in investment (Alstadsæter et al., 2024). Hong and Smart (2010) show that tax avoidance may be beneficial because it reduces the tax burden on mobile capital and thereby increases investment. Suárez Serrato (2019) finds that shutting down tax havens might have negative consequences for real investment and employment in high-tax countries. Becker and Riedel (2012) find that there are within-MNE spillovers of taxation, where an increase in the parent company's corporate tax rate decreases investment in affiliates. CFC rules also place importance on the location of parent entities. Voget (2011) shows that an increase in repatriation taxes and the introduction of CFC rules increase the number of headquarter relocations. Buettner et al. (2018) and De Mooij and Liu (2020) find that the introduction of transfer pricing regulation decreases real investment. These papers do not find that total investment by MNEs decreases, which suggests that investment is instead shifted abroad. I add to this literature by specifically looking into the real effects of exemptions in anti-profit shifting rules that are meant to foster investment.

The literature around CFC rules has mostly focused on the use of CFC rules to curb profit shifting. Overesch et al. (2024) provide an extensive overview of the existing research on CFC rules. Weichenrieder (1996) formally analyzes the impact of anti-tax avoidance legislation on the size and growth of MNEs' low-tax affiliates. The model in his paper provides intuition for why MNEs want to increase substance, since substance allows to continue shifting profits. Pieretti and Pulina (2020) provide theoretical analysis of the incentives of anti-profit shifting rules on the relocation of economic activity to low-tax countries. They conclude that eradicating paper profit shifting only increases corporate tax revenue in high-tax countries if tax havens cannot efficiently provide infrastructure to attract real investment.

There is also empirical literature that considers CFC rules. Gschossmann and Pfrang (2024) analyze the effects of CFC rules on European MNEs' location, financial, and economic activity decisions. They focus on extensive and intensive margin effects. On the extensive margin, the share of CFCs decreases, but the financial income of the remaining affiliates remains unchanged. Their results show that the costs of employees are increased in CFCs, which suggests that substance rules are a driver of CFC activity. Egger and Wamser (2015) use MiDi data in a Regression Discontinuity Design over the period 1996-2008 to study the impact of CFC rules on real FDI. They find that CFC rules decrease real investment in fixed assets in low-tax affiliates. On the extensive investment margin, Egger and Wamser (2015) find that unconditionally, exit of treated affiliates is about seven percentage points more likely and entry is five percentage points less likely. However, this effect is not robust to controlling for the forcing variables, suggesting that the overall effect on investment in fixed assets is not distorted by extensive margin investment decisions. Clifford (2019) uses data from Bureau van Dijk's Orbis to analyze CFC rules from 2003-2013. Using a Regression Discontinuity Design, she shows that affiliates just below the tax rate threshold experience a 13% decrease in financial profits. She also shows that increases in group-level CFC exposure increases financial profit within higher-tax affiliates. This indicates a spillover effect of CFC rules for high-tax countries, implying that CFC rules can protect both the global and the domestic tax base. This spillover effect is confirmed by Hansen et al. (2023). On the extensive investment margin, Clifford (2019) finds that CFC rules increase (decrease) the propensity to open a new affiliate just above (below) the tax rate threshold.

4 Model and Hypotheses

4.1 Model Set-Up

To formulate the hypotheses, I model the decision of an MNE to set up tax-haven affiliates for profit shifting.[12] I focus on an MNE that has both productive affiliates which are used to generate real profits, and tax-haven affiliates that are used for profit shifting. Figure 4 depicts this type of ownership structure.[13]

In line with what we observe empirically, MNEs can have several tax-haven affiliates. Merlo and Wamser (2024) show that even small MNEs are likely to have several tax-haven affiliates, providing several shifting opportunities for different productive affiliates. For large MNEs, the UNCTAD concludes in its World Investment Report 2016 that based on Orbis data, the largest hundred MNEs worldwide have on average 549 affiliates, of which on average 68 are located in Offshore Financial Centre, see United Nations Conference on Trade and Development (2016).

In the model, MNEs choose their tax-haven affiliates to be close to productive affiliates. An important reason to locate tax-haven affiliates nearby is that this allows the MNE to benefit from the familiarity of legal institutions in the tax haven. Dharmapala and Hines (2009) show that governance quality has a statistically significant and quantitatively large association with the probability of being a tax haven. Dharmapala (2020) discusses that nontax factors are relevant for location decisions. Examples of these nontax factors are the availability of local legal and business expertise relevant to managing the holding structure. This could reflect the specialization of legal and accounting firms with respect to nearby tax havens (Dharmapala, 2020).

Another reason for a distance effect is the need to locate the effective place of management in the low-tax country.[14] The place of effective management requirement forces management decisions to actually be made in the tax-haven affiliate's jurisdiction in order to be treated as tax resident. The requirement of physical presence entails travel costs that increase in distance.

[12] I talk about tax-haven affiliates, but the model applies to all affiliates used for profit shifting. Some countries apply relatively high tax rate thresholds, targeting not only tax-haven affiliates but also medium- to high-tax affiliates.

[13] For the empirics, which anti-avoidance rule constrains the affiliate is determined at the shareholder level: any (intermediate) parent country can apply a CFC rule that targets the affiliate. It is therefore the (intermediate) parent's country, not the country of a productive affiliate, that governs whether the group falls under a CFC rule. Under the EU Anti-Tax Avoidance Directive, for example, the CFC rule applies in the Member State of the controlling parent (Art. 7–8 of Directive 2016/1164).

[14] This concept is primarily a legal concept and has not been studied directly in the empirical literature.

Figure 4: Ownership Structure of the Multinational Group

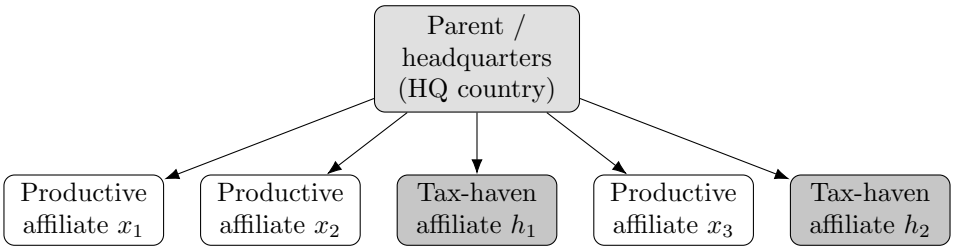

Note: The MNE consists of a parent (headquarters) in the HQ country that owns productive affiliates (x) and tax-haven affiliates (h) located in different countries. Ownership links run from the parent to each subsidiary. The headquarters' country determines which anti-avoidance and substance rule applies to the group. Figure 5 rearranges the same affiliates on a Salop circle to represent the geographic and legal distance between them.

Empirically, this distance relationship also exists: Desai et al. (2004) and Desai et al. (2006) show that tax-haven affiliates may enhance investment in especially nearby non-haven affiliates. Laffitte (2024), Ferrari et al. (2023), and Tørsløv et al. (2023) also provide evidence of a gravity-like profit shifting relationship for the distance between between productive affiliates and tax-haven affiliates. In particular, Dharmapala (2020) discusses that Caribbean havens seem particularly important for US MNEs, while European havens are important for European MNEs. Therefore, MNEs seem to face affiliate-level profit-shifting costs that increase in the distance to the tax-haven affiliate.

To analyze the notion that regional tax-haven affiliates serve close-by high-tax affiliates and that MNEs optimize the number of tax-haven affiliates in their group structure, I focus on a simple Salop (1979) framework, inspired by the model in Brekke et al. (2017).[15] The Salop circle in Figure 5 represents the MNE's structure. The productive affiliates are assumed to be uniformly distributed on a Salop circle of circumference 1, represented by the x in Figure 5. There are $h \in 1, ..., j$ identical tax-haven affiliates with tax rates t_j. I assume they are identical, because tax-haven affiliates serve as profit-shifting destinations and generally have disproportionately low real economic activity in comparison with profit levels (Fuest et al., 2022).[16]

The tax-haven affiliates are located equidistant from each other at $z_j \in [0, 1]$ on

[15]For another application of a Salop framework in studying MNE behavior, see Gresik and Konrad (2017).

[16]I treat tax-haven affiliates as pure shifting vehicles. Allowing them to also have real production would enrich the distinction between productive and non-productive havens, but it would not overturn the extensive- and intensive-margin mechanisms, which operate through the fixed cost and distance of each tax-haven affiliate rather than through its own productive output.

Figure 5: Multinational Group Structure

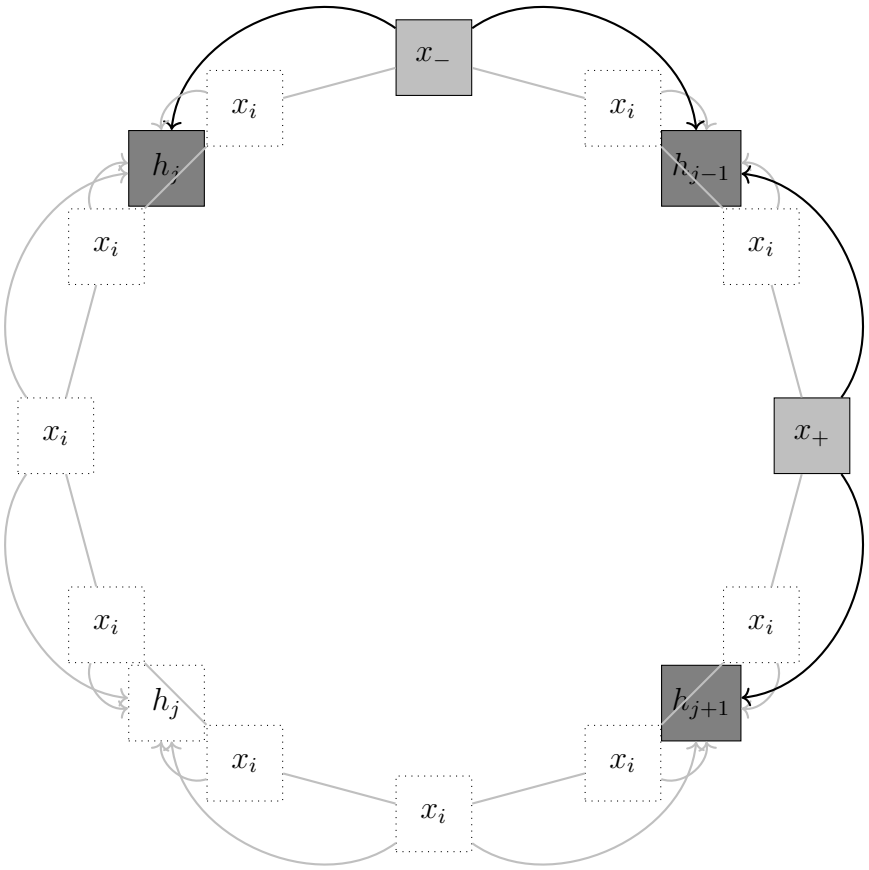

Note: The Salop circle represents the geographic and legal distance between the affiliates of the MNE shown in Figure 4; the parent is not pictured. Productive affiliates (x) are uniformly distributed on the circle and tax-haven affiliates (h) sit between them. Arrows show the direction of profit shifting: each productive affiliate shifts to its nearest tax-haven affiliate. x_- and x_+ represent affiliates that are just indifferent between shifting to two close-by tax-haven affiliates.

the Salop circle in Figure 5. Each productive affiliate shifts profit to the closest tax-haven affiliate. The two shaded productive affiliates labelled x_- and x_+ in Figure 5 are the marginal productive affiliates that are indifferent between two closest-by tax-haven affiliates.

The productive affiliate generates real profit m_i in its host country i. If the MNE does not establish a tax-haven affiliate, the productive affiliate pays $t_i m_i$ in corporate taxes in its host country. However, to decrease its tax burden, the productive affiliate can shift part of its profit $P_{i,j}$ to the nearest-by tax-haven affiliate j located at z_j. Shifting profits is costly, however. In particular, there are two types of variable costs, borne by a productive affiliate when it shifts; the third type of cost is a fixed cost, borne by the MNE as a whole for each tax-haven affiliate it operates.

Each productive affiliate faces two types of *variable costs of shifting*. First, in line with the literature on regional tax havens, the cost of profit shifting increases in the distance to the tax-haven affiliate. Thus, the MNE faces costs that are similar to transport costs in a traditional Salop framework. I refer to these costs as adjustment costs, defined as $\alpha(k_j)d_{i,j} = \alpha(k_j)\,|z_j - x_i|$. The adjustment costs increase in the productive affiliate's distance to the tax-haven affiliate $|z_j - x_i|$ with $x_i \in [0, 1]$. This distance cost is the central variable cost of the model: it makes a *nearby* tax-haven affiliate more valuable than a distant one, which is what gives the MNE an incentive to operate several, regionally placed tax-haven affiliates.[17] The adjustment cost parameter $\alpha(k_j) = \frac{\bar{\alpha}}{k_j}$ captures the proximity of legal institutions and ease of obtaining tax advice, but also anti-avoidance rules. Note that adjustment costs decrease in investment k_j in the tax-haven affiliate, because greater substance in tax-haven affiliates facilitates transfer pricing documentation and reduces the risk of disputes by the tax authorities.

Second, profit shifting involves convex tax-planning costs $\frac{P_{i,j}^2}{2}$, which increase in the amount of profit shifted $P_{i,j}$.[18] These are standard convex tax planning costs for hiring tax advisors to set up transfer pricing documentation and dealing with the productive affiliate's tax authorities in case of disputes.

The tax-haven affiliate faces a proportional corporate income tax t_j. This tax rate determines whether affiliates established in the low-tax country will be subject to CFC rules—and substance requirements. Having the tax-haven affiliate allows the productive

[17] Because adjustment costs increase in distance, the average adjustment cost a group faces decreases as it operates more tax-haven affiliates, since additional affiliates shorten the distance to the nearest one. Modelling the shifting cost as directly decreasing in the number of tax-haven affiliates would be an alternative that delivers the same extensive-margin comparative statics.

[18] The quadratic form is a convenient convex specification that guarantees an interior shifting optimum; the qualitative results hold for any strictly convex tax-planning cost.

affiliate to shift profits and avoid paying taxes in its host country. Therefore, if the productive affiliate decides to shift part of its profits $P_{i,j}$ to a tax-haven affiliate, it gains net benefits $(t_i - t_j)P_{i,j}$.

The *fixed cost of a tax-haven affiliate* is the investment the MNE must make for it to have enough substance. The MNE invests in fixed assets and local employees at cost $f(k_j, \phi) = \frac{\phi k_j^2}{2}$, with investment $k_j \geq 0$ and diminishing marginal returns. This is a fixed cost in the sense that the MNE incurs it once per tax-haven affiliate, independently of how much any individual productive affiliate shifts to it. The increases ϕ indicates the substance rule: a stricter rule raises the investment the affiliate needs to satisfy the substance requirements, so $\frac{\partial f(k_j,\phi)}{\partial \phi} > 0$.[19] Investment k_j therefore enters both cost categories with opposite signs: On the one hand, increasing investment decreases the variable adjustment cost of shifting because it facilitates proving to tax authorities that there are real economic activities in the tax haven. On the other hand, higher investment increases the fixed cost of the affiliate, which is intensified by the substance requirement. This trade-off between a lower variable adjustment cost and a higher fixed substance cost is the core tension of the model.

Given these three types of costs, the MNE decides whether to make use of tax-haven affiliates for profit shifting and how many tax-haven affiliates it wants to establish. This problem can be solved in three stages. First, the MNE decides on the number of tax-haven affiliates h. Second, the MNE chooses the level of investment k_j in each tax-haven affiliate, taking into account substance rules ϕ.[20] Third, the MNE maximizes after-tax profits by choosing the optimal amount of profits shifted $P_{i,j}$ from each productive affiliate.

The model is solved by backward induction: Section 4.2 solves the third-stage shifting decision, Section 4.3 the second-stage investment decision, and Section 4.4 the first-stage choice of the number of tax-haven affiliates.

[19] Here k_j stands for substance broadly—local staff, tangible assets, and the functions they perform—in line with the subjective substance tests described in Section 2.

[20] Investment therefore takes place in the second stage—after the group structure is chosen but before profits are shifted. This ordering also answers why the fixed cost does not appear in the third-stage shifting problem: by the time a productive affiliate shifts $P_{i,j}$, both h and k_j are already determined, so $f(k_j, \phi)$ is sunk and does not affect the marginal shifting decision.

4.2 How Much Profit Does the MNE Shift?

Starting with the profit-shifting decision, the MNE maximizes the sum of profit over all of its affiliates:

$$\Pi = \sum_{i=1} \pi_{i,j} - hf(k_j, \phi),$$

which consists of the after-tax and after-shifting profits in each productive affiliate i minus the investment in the h tax-haven affiliates, conditional on the MNE group structure and investment in tax-haven affiliates.[21]

The after-tax profit-maximization problem for each symmetric affiliate is given by

$$\max_{P_{i,j}}: \quad \pi_{i,j} = (1 - t_i)\, m_i + (t_i - t_j)\, P_{i,j} - \tfrac{\bar{\alpha}}{k_j} d_{i,j} P_{i,j} - \tfrac{P_{i,j}^2}{2}.$$

The MNE decides how many profits $P_{i,j}$ to shift from each productive affiliate i to tax-haven affiliate j. Taking the first-order condition yields

$$\frac{\partial \Pi_{i,j}}{\partial P_{i,j}} = (t_i - t_j) - \frac{\bar{\alpha}}{k_j} d_{i,j} - P_{i,j} = 0.$$

Rewriting gives the optimal amount of profits shifted

$$P_{i,j}^* = (t_i - t_j) - \frac{\bar{\alpha}}{k_j} d_{i,j}. \tag{1}$$

The optimal amount of profit shifted increases in the affiliate's tax rate t_i, since this makes profit shifting more attractive. It decreases in the tax-haven affiliate's tax rate t_j and the adjustment costs $\frac{\bar{\alpha}}{k_j} d_{i,j}$, because these costs reduce the net benefit of profit shifting. Thus, each affiliate optimally equalizes the marginal benefits of shifting against the marginal costs of shifting to maximize after-tax profit.

Substituting the optimal amount of profits shifted in equation (1) into the affiliate's after-tax profit, maximized after-tax profits in each productive affiliate are given by

$$\pi_{i,j}^* = (1 - t_i)\, m_i + \frac{1}{2} \left(t_i - t_j - \frac{\bar{\alpha}}{k_j} d_{i,j} \right)^2, \tag{2}$$

which consists of the after-tax profit from real activities $(1 - t_i)\, m_i$ as well as the net benefits from profit shifting $\frac{1}{2} \left(t_i - t_j - \frac{\bar{\alpha}}{k_j} d_{i,j} \right)^2$. These net benefits increase in the

[21] The MNE maximization problem is additively separable across productive affiliates, so the results do not rely on investments in different affiliates being complements or substitutes. The extensive and intensive margins arise from the fixed cost and distance of each tax-haven affiliate, not from cross-affiliate spillovers.

productive affiliate's tax rate t_i and decrease in the tax-haven affiliate's tax rate t_j and the adjustment costs $\frac{\bar{\alpha}}{k_j} d_{i,j}$. The further away the productive affiliate is from the tax-haven affiliate, the higher the adjustment costs. This introduces the incentive to set up nearby tax-haven affiliates, based on which the MNE decides on its optimal group structure.

4.3 How Much Does the MNE Invest in each Tax-Haven Affiliate?

The MNE chooses how much to invest in each tax-haven affiliate. Substituting the optimal shifting decision from equation (2), each tax-haven affiliate serves the productive affiliates located closest to it on the Salop circle. I assume the tax-haven affiliates are symmetric, so that $t_j = t_{j+1} = t_{j-1}$. The net profits shifted *aggregated* across the productive affiliates that each tax-haven affiliate serves are

$$D_{i,j} = \frac{(t_i - t_j)^2}{2h} - \frac{\bar{\alpha}(t_i - t_j)}{4h^2 k_j} + \frac{\bar{\alpha}^2}{24h^3 k_j^2}. \tag{3}$$

The aggregated net profits shifted per tax-haven affiliate are obtained by finding the productive affiliates that are just indifferent between two close-by tax-haven affiliates and integrating their net shifting benefits over the interval that each tax-haven affiliate covers on the Salop circle; Appendix A provides the full derivation.

The three terms of equation (3) can be interpreted as follows. The first term, $\frac{(t_i - t_j)^2}{2h}$, is the net shifting benefit the tax-haven affiliate would deliver *absent* adjustment costs: the tax saving from the rate differential $(t_i - t_j)$.[22] The second and third terms capture how adjustment costs decrease this shifting benefit. The second term, $-\frac{\bar{\alpha}(t_i - t_j)}{4h^2 k_j}$, is the first-order reduction in profits shifted, because more distant productive affiliates shift less. The third term, $+\frac{\bar{\alpha}^2}{24h^3 k_j^2}$, is a smaller second-order term of the opposite sign. Because both involve the adjustment-cost parameter $\bar{\alpha}$ divided by the investment k_j, investment in the tax-haven affiliate decreases the importance of the adjustment costs, which provides the intensive-margin incentive to invest. The aggregate shifting benefit decreases in the number of tax-haven affiliates h, since each then serves a lower number of productive affiliates, which incentivizes the extensive margin later on.

[22] The tax-rate differential is squared because the amount optimally shifted is itself proportional to the differential. This squares when integrating over the productive affiliates the tax-haven affiliate serves.

Because the affiliate mass is normalized to one[23] and there are h tax-haven affiliates, the MNE's after-tax profit can be written as

$$\Pi = (1 - t_i)m_i + h\left[D_{i,j} - f(k_j, \phi)\right], \tag{4}$$

where the first term shows after-tax profits from production in the productive affiliates. The second term shows, for each of the h tax-haven affiliates, the net benefits from shifting profits to nearby tax-haven affiliates, taking into account the fixed costs of setting up h tax-haven affiliates.

In the second stage, the MNE chooses investment k_j to maximize (4), taking the number of tax-haven affiliates h as given:

$$\max_{k_j}: \quad \Pi = (1 - t_i)m_i + h\left[D_{i,j} - f(k_j, \phi)\right].$$

The first-order condition is

$$\frac{\partial \Pi}{\partial k_j} = \frac{\bar{\alpha}(t_i - t_j)}{4h^2k_j^2} - \frac{\bar{\alpha}^2}{12h^3k_j^3} - \phi k_j = 0, \tag{5}$$

where the first two terms capture the marginal benefit of investment: higher investment lowers adjustment costs, which increases the net benefit of profit shifting and thereby after-tax profit. The third term captures the marginal cost of investment under the substance rule, because the substance rule requires higher investment in each tax-haven affiliate, reducing after-tax profit. In the optimum, the MNE equates the marginal benefit of reducing adjustment costs against the marginal cost of investment in each tax-haven affiliate.

4.4 How Many Tax-Haven Affiliates Does the MNE Establish?

Given optimal profits shifted and optimal investment, the MNE chooses the number of tax-haven affiliates h to maximize equation (4):

$$\max_{h}: \quad \Pi = (1 - t_i)m_i + h\left[D_{i,j} - f(k_j, \phi)\right].$$

[23] Normalizing the mass of productive affiliates to one abstracts from group-size or scale effects. The model speaks to how a group of a given size restructures under substance rules, not to how the absolute number of affiliates scales with firm size.

The corresponding first-order condition is

$$\frac{\partial \Pi}{\partial h} = \frac{\bar{\alpha}(t_i - t_j)}{4h^2 k_j} - \frac{\bar{\alpha}^2}{12h^3 k_j^2} - \frac{\phi k_j^2}{2} = 0, \tag{6}$$

which contains three arguments that capture the trade-off that the MNE faces: it wants to shift to nearby tax-haven affiliates but faces fixed costs of investing in each tax-haven affiliate. The first term shows that the optimal number of tax havens affiliates depends on the benefit of profit shifting, as indicated by the tax rate differential. The first and second term together show that the optimal number of tax-haven affiliates also depends on the adjustment costs. The third term indicates that the optimal number of tax-haven affiliates depends on the fixed costs of investment in each tax-haven affiliate, which is influenced by substance rules. This forms the starting point for my hypotheses.

4.5 Hypotheses

Based on the first-order conditions for optimal investment and the optimal number of tax-haven affiliates in equations (5) and (6), I derive the effects of changes in the adjustment costs, the tax rates in both the productive affiliates and tax havens, and the costs of investment.[24]

Extensive Margin. The comparative statics of the first-order condition (6) deliver two distinct extensive-margin predictions: one for the introduction of a CFC rule, which compresses the tax-rate differential, and one for the introduction of a substance requirement, which raises the fixed cost of each tax-haven affiliate. I state them as two separate hypotheses.

The tax-rate effect. The optimal number of tax-haven affiliates increases in the tax-rate differential $(t_i - t_j)$: It increases in the productive affiliate's tax rate t_i[25] and decreases in the tax-haven affiliate's tax rate t_j.[26] This mechanism is visible in the first term

[24] To derive the comparative statics, first define the FOC for the optimal number of tax-haven affiliates as $g(\bar{\alpha}, t_i, t_j, k_j, \phi, h) = \frac{\bar{\alpha}(t_i - t_j)}{4h^2 k_j} - \frac{\bar{\alpha}^2}{12h^3 k_j^2} - \frac{\phi k_j^2}{2}$. Furthermore, I assume the second-order conditions are negative: $SOC : \frac{\partial^2 \Pi}{\partial h^2} = -\frac{\bar{\alpha}(t_i - t_j)}{2h^3 k_j} + \frac{\bar{\alpha}^2}{4h^4 k_j^2} < 0$, which is required to maximize profits. Conversely, the FOC for the optimal amount of investment in each tax-haven affiliate is defined as $p(\alpha, t_i, t_j, k_j, \phi, h) = \frac{\bar{\alpha}(t_i - t_j)}{4h^2 k_j^2} - \frac{\bar{\alpha}^2}{12h^3 k_j^3} - \phi k_j$. The second-order condition is negative again: $SOC : \frac{\partial^2 \Pi}{\partial k_j^2} = -\frac{\bar{\alpha}(t_i - t_j)}{2h^2 k_j^3} - \frac{\bar{\alpha}^2}{4h^3 k_j^4} - \phi < 0$

[25] $\frac{dh}{dt_i} = -\frac{\bar{\alpha}/(4h^2 k_j)}{SOC} > 0.$

[26] $\frac{dh}{dt_j} = \frac{\bar{\alpha}/(4h^2 k_j)}{SOC} < 0.$

of equation (6), $\frac{\bar{\alpha}(t_i - t_j)}{4h^2 k_j}$: a larger tax-rate differential increases the marginal benefit of shifting, so the MNE optimally has more tax-haven affiliates. CFC rules essentially equalize the tax haven's tax rate to the productive affiliate's tax rate ($t_i = t_{CFC} = t_j$), which eliminates the profit shifting benefit in the first term of equation 6 for an MNE whose parent country operates a CFC rule. Indeed, Clifford (2019) finds that MNEs open affiliates just above the CFC threshold. This suggests that the optimal number of tax-haven affiliates decreases because the tax rate differential becomes smaller.

Hypothesis 1a *After the MNE group becomes subject to a CFC rule, the number of tax-haven affiliates in the group decreases.*

The substance effect. A substance requirement adds a second, distinct extensive-margin force. It raises the investment each tax-haven affiliate must have to qualify for the exemption, and hence the fixed cost $f(k_j, \phi)$ in the third term of (6). Formally, a stricter rule lowers the optimal number of tax-haven affiliates.[27]

Under a substance rule, providing sufficient evidence that there is substance in a low-tax country facilitates profit shifting because the substance rule provides an exemption from the application of CFC rules. Before the substance rule, the MNE holds the number of tax-haven affiliates at which the marginal affiliate—the one serving the smallest, most distant arc of productive affiliates—is just indifferent between two tax-haven affiliates. At that point, the aggregate net shifting benefit $D_{i,j}$ equals the fixed cost of investment in each tax-haven affiliate $f(k_j, \phi)$. With fewer tax-haven affiliates, each remaining tax-haven affiliate serves more productive affiliates—a share $1/h$ of the Salop circle—which increases the aggregate profits shifted to each tax-haven affiliate. Substance rules thus consolidate profit shifting into fewer, larger tax-haven affiliates.

In their working paper, Ruf and Weichenrieder (2013) suggest that the number of tax haven affiliates does not decrease after becoming subject to German CFC rules, but Gschossmann and Pfrang (2024) show the opposite. Concentration could imply that within-affiliate investment in tax havens is increased, whereas other tax-haven affiliates are closed to curb the fixed costs of investment within the MNE group. The model in this paper provides an explanation for the geographical dispersion of haven affiliates and explains why multinationals might have several haven affiliates.[28]

[27] $\frac{dh}{d\phi} = \frac{\partial f(k_j, \phi)/\partial \phi}{SOC} < 0$.

[28] Anecdotal evidence even suggests that MNE group structures often have several affiliates in tax havens, for example to selectively apply the U.S. Check-the-Box Regime or benefit from specific Tax Treaty provisions through particular routes of ownership. This suggests that the concentration effect occurs.

Hypothesis 1b *After the MNE group becomes subject to a substance rule, the number of tax-haven affiliates in the group decreases.*

Intensive Margin Effect. Under a substance rule, providing sufficient evidence that there is substance in a low-tax country facilitates profit shifting because the substance rule provides an exemption from the application of CFC rules. This incentivizes the MNE to invest more in its tax-haven affiliates.

Taking the total derivative of k_j with respect to the substance requirement ϕ in equation (5) yields:

$$\frac{dk_j}{d\phi} = \frac{\partial k_j}{\partial \phi} + \frac{\partial k_j}{\partial h} \cdot \frac{dh}{d\phi} > 0.$$

The first term is a *direct effect*, which captures how a stricter substance rule raises the marginal cost of investment, which decreases investment.[29] The second term is a *concentration effect*, which captures the indirect effect of substance rules on investment through the extensive margin. From Hypothesis 1b, a stricter substance rule decreases the optimal number of tax-haven affiliates. As the optimal number of tax-haven affiliates decreases, each remaining tax-haven affiliate serves a larger share of the MNE's productive affiliates, increasing the marginal return to investment in that tax haven. Formally, from equation (5), a decrease in h increases the first two terms, which increases optimal investment in each tax-haven affiliate.[30] This prediction is consistent with Gschossmann and Pfrang (2024), who find that employee costs rise in CFCs after the ATAD reform. Therefore, the intensive-margin hypothesis is:

Hypothesis 2 *After an MNE affiliate becomes subject to a substance rule, fixed assets and employment in that affiliate are increased.*

Investment Shifting Effect. An interesting policy question is whether substance rules in tax-haven affiliates prevent loss of real investment in high-tax affiliates. This could be the result of an investment shifting effect, which implies a re-routing of investment from high-tax to low-tax countries (Buettner et al., 2018; Suárez Serrato, 2019; Drake et al., 2022; Williams, 2018). Investment shifting behavior might provide a double hit for high-tax countries: They lose both tax base through profit shifting, and economic activities as MNEs increase real investment in the tax-haven affiliate to avoid applying CFC rules. However, it may also be the case that investment in tax-haven

[29] $\frac{\partial k_j}{\partial \phi} < 0$.

[30] Since $\frac{dh}{d\phi} < 0$ and $\frac{\partial k_j}{\partial h} < 0$, the concentration effect is positive.

affiliates is complementary to investment in high-tax affiliates (Hong and Smart, 2010). Therefore, the fourth hypothesis is stated in the null form.

Hypothesis 3 *Real investment in fixed assets and employment in unaffected affiliates of the MNE remains unchanged as another affiliate in the MNE group becomes subject to a substance rule.*

Summarizing, as a result of substance rules within a CFC rule, the optimal number of tax-haven affiliates decreases—*extensive margin effect*. The remaining number of tax-haven affiliates receive a higher level of investment—*intensive margin effect*. It is a priori unclear how the substance rule affects investment in non-substance affiliates and whether an *investment shifting effect* occurs.

5 Data

I obtain data on substance rules from OECD (2024). In addition, I hand-collect data from national CFC legislation on the details of the substance rule. Table A1 in the Appendix shows the shareholder countries included in the sample which have implemented CFC rules. An asterisk (*) indicates that the sample observes affiliates in this country.

To test the hypotheses, I use unbalanced panel data for affiliates that have at least one foreign majority-shareholder from Bureau van Dijk's Historical Orbis over the period 2010–2023. To define treatment based on the tax rate threshold, I collect data from the Tax Foundation on yearly statutory corporate tax rates in percentages, which include both state and local taxes. I exclude affiliates with negative values in any year for total assets, tangible fixed assets, long-term debt, turnover, profit, and the number and costs of employees.

The sample contains 53 parent countries that have a CFC rule in place at some point during the sample period, see Table A1. 30 of these countries operate a substance requirement, which the affiliate must satisfy for the CFC rule not to apply. 3 of them provide a safe harbor rule; the other 27 apply a subjective test, which requires evaluation of all relevant facts and circumstances to assess whether there is substantial economic substance. The sample includes affiliates in 68 different countries. Table A2 lists the countries of those affiliates in countries that do not have a CFC rule in place, so that combined with Table A1, it lists all countries with observed affiliates.

Table 2 shows descriptive statistics in four panels. Panel A explains the sample construction. Restricting to affiliates with a foreign majority shareholder, and censor-

ing those whose ultimate owner changes, leaves 8.4 million affiliate-years covering 1.2 million affiliates in 472,565 groups. The three effects from the model draw on this sample differently, which is why they report different totals: the intensive margin uses the 2,655,900 affiliate-years belonging to treated affiliates and their two control groups, the spillover margin uses the 6,076,048 affiliate-years of affiliates that are never themselves within scope, and the extensive margin aggregates to 2,586,731 group $\times$ year $\times$ affiliate-type cells across 242,591 groups.

Panel B reports affiliate characteristics *before* treatment, separately for affiliates that eventually fall within the scope of a substance rule, affiliates whose parent operates a CFC rule without one, and affiliates whose parent operates neither. Substance affiliates face a mean statutory rate of 18.9 percent against 22.7 percent for CFC-only affiliates. Substance affiliates are also larger than untreated affiliates—a median of 12 employees against 5—and are part of substantially larger MNE groups, with a median of 11 affiliates against 8 in CFC-only groups and 2 in untreated groups.

Panel C reports when treatment occurs. 29,892 affiliates fall within the scope of a substance rule at some point, spread over 13 cohorts. 12 parent countries first operate a substance rule in 2019, when the EU Anti-Tax Avoidance Directive required member states to implement CFC rules, and that cohort accounts for 21.5 percent of all treated affiliates. Only 15 of the 30 substance countries adopt within the sample period; the remainder already operated a requirement before 2010, and contribute affiliates that obtain treatment when their own statutory rate crosses the threshold rather than when the rule is introduced.

Panel D describes the two group-level samples. The extensive margin contains 242,591 groups with a median of 2 affiliates in the first year the group is observed ('base year'), of which 79.1 percent are exposed to a substance rule at some point. The spillover sample contains 122,441 groups with a median of 3 affiliates, of which 5.7 percent ever contain a treated affiliate.

Table 2: Descriptives

Panel A: Sample Construction

	Affiliate-years	Affiliates	Groups
Orbis affiliate-years, 2010–2023	147,759,136	36,270,315	31,485,867
with an ownership group	144,780,434	35,692,831	31,485,867
with a foreign majority shareholder	8,604,115	1,239,483	473,736
after removing ownership changes	8,426,906	1,233,272	472,565
Extensive-margin sample (H1)	2,586,731	—	242,591
Intensive-margin sample (H2)	2,655,900	421,878	214,329
matched control group	145,283	17,526	10,899
Spillover sample (H3)	6,076,048	862,115	122,441
matched control group	588,333	77,367	6,140

Panel B: Pre-Treatment Affiliate Characteristics

	Substance	CFC-only	Never-rule
Total Assets (EUR Tsd)	70,035.6	79,669.5	58,240.5
	(859,326.7)	(5,008,564.3)	(7,447,815.4)
Median	1,349.8	1,662.1	703.8
Tangible Fixed Assets (EUR Tsd)	1,824.2	1,998.8	1,242.4
	(4,371.9)	(4,645.9)	(3,652.9)
Median	56.8	32.9	9.9
Number of Employees	52.1	46.7	31.4
	(87.2)	(83.3)	(66.6)
Median	12.0	9.0	5.0
Costs of Employees (EUR Tsd)	2,084.2	1,837.5	1,049.4
	(3,721.0)	(3,353.4)	(2,581.3)
Median	412.0	462.7	128.0
ln(Tangible Fixed Assets)	11.98	12.37	11.59
	(2.88)	(2.96)	(3.05)
ln(Number of Employees)	2.67	2.64	2.09
	(1.76)	(1.78)	(1.71)
ln(Costs of Employees)	12.70	12.75	11.61
	(2.44)	(2.31)	(2.49)
Statutory Tax Rate (%)	18.93	22.67	21.10
	(6.35)	(4.59)	(6.18)
Affiliates in Group	84.0	72.1	41.0
	(205.2)	(185.6)	(147.7)
Affiliate-years	133,367	1,151,758	970,893
Affiliates	29,892	173,869	191,368

Note: Affiliate-level data from Bureau van Dijk's Historical Orbis for the years 2010–2023. Panel A describes the sample construction. I keep multinational affiliates with a foreign majority shareholder, and drop affiliates from the group if their ownership changes. Panel B reports variable means with standard deviations in parentheses, and medians. For substance affiliates, I show years before the affiliate first becomes subject to a substance requirement; for CFC-only affiliates, I show years before it first becomes subject to a CFC rule; for never-rule affiliates, I show all years. Logarithmic variables are winsorised at the 1st and 99th percentiles. Affiliate-level statutory tax rates are obtained from the Tax Foundation.

Table 2: Descriptives (continued)

Panel C: Timing of Substance-Rule Treatment

	Affiliates newly treated	Parent countries newly adopting	Cumulative share (%)
2011	1,675	0	5.6
2012	1,349	0	10.1
2013	3,519	1	21.9
2014	2,231	1	29.4
2015	1,788	0	35.3
2016	1,701	0	41.0
2017	3,413	1	52.4
2018	1,853	0	58.6
2019	6,424	12	80.1
2020	2,175	0	87.4
2021	2,146	0	94.6
2022	1,385	0	99.2
2023	233	0	100.0
Total	29,892	15	

Panel D: Number of Affiliates in each MNE Group

	Groups	Mean size	Median size	Ever exposed (%)
Extensive margin (H1)	242,591	6.5	2	79.1
Spillover margin (H3)	122,441	7.2	3	5.7

Note: Affiliate-level data from Bureau van Dijk's Historical Orbis for the years 2010–2023. Panel C reports the year in which affiliates first fall within the scope of a substance requirement, based on the tax-rate threshold or the introduction of a substance requirement. Panel D reports the number of affiliates in the in the two group-level samples, measured.

6 The Extensive Margin

6.1 Measuring the Number of Tax-Haven Affiliates

Hypotheses 1a and 1b predict that after the MNE becomes subject to a CFC rule, and separately to a substance rule, the number of tax-haven affiliates in the group should decrease. Using Orbis Historical Data, I link every affiliate to an MNE group in every year using the direct ownership links, tracing each entity to its global ultimate owner using the fifty-percent control threshold. The group structure is measured using ownership snapshots from each current year or earlier, so no assignment uses future information. This matters because group membership is not stable: 4.6 percent of affiliates change ultimate owner during the sample period, which is why I update the ownership structure once a change is recorded in Orbis.

I define tax-haven affiliates based on Dharmapala and Hines (2009), excluding the countries that can be considered conduit countries: Belgium, Cyprus, Hong Kong, Ireland, Luxembourg, Malta, the Netherlands, Singapore and Switzerland. These are relatively large economies usually used in conduit structures because they have good treaty networks. To focus on pure tax havens, I use the remaining thirty-five jurisdictions as the baseline definition and report the full list as a robustness check.[31] It should be noted that Orbis records entities that file publicly, and the most secretive jurisdictions do not require public filing. For example, the British Virgin Islands contribute 268 affiliates to the ownership universe and Jersey and Guernsey fewer than thirty. The estimates below therefore should be read as applying to the part of the MNE that shows up in Orbis, and they are likely to understate the response of the unreported part.

Treatment is defined as follows: An MNE group is considered subject to a CFC rule when its affiliates fall within the rule's scope. Variation in treatment comes from three sources: a shareholder country adopts a CFC rule, an affiliate country's statutory rate crosses below the threshold, or a parent country's own rate changes and moves the threshold with it. Treatment requires that a controlling shareholder is resident in a country operating a CFC rule and that the affiliate's statutory tax rate lies below the

[31] Andorra, Anguilla, Antigua and Barbuda, Aruba, the Bahamas, Bahrain, Barbados, Belize, Bermuda, the British Virgin Islands, the Caribbean Netherlands, the Cayman Islands, Curacao, Gibraltar, Grenada, Guernsey, the Isle of Man, Jersey, Lebanon, Liechtenstein, Macao, the Marshall Islands, Mauritius, Monaco, the Netherlands Antilles, Panama, Puerto Rico, Samoa, the Seychelles, Sint Maarten, St Kitts and Nevis, St Lucia, St Vincent and the Grenadines, the Turks and Caicos Islands and Vanuatu. The Netherlands Antilles is carried alongside its post-2010 successor codes so that pre- and post-dissolution Orbis vintages both resolve.

shareholding country's threshold. Having a shareholder in a CFC-rule country applies to 91 percent of MNEs by 2019, and of 70 percent of affiliate-years, mostly because the United States, Germany, Japan, France and the United Kingdom have long operated CFC regimes and the EU Anti-Tax Avoidance Directive extended them to every EU member state in 2019. Being *within scope* is less prevalent: 6 percent of affiliate-years for CFC rules and 1 percent for substance requirements.

6.2 Identification

I take the first year in which I observe an MNE group in the sample and hold the composition of the group fixed ('base-year'). Orbis coverage of ownership links improves over time, and it is unclear whether these are *new* affiliates, or improvements in data quality. I therefore follow the affiliates in the group from the first time I observe the MNE group in the data. The outcome is their *survival*: the share of the base-year affiliates of a given type (CFC/substance versus non-CFC/substance) that still belong to the group in year t. Entities that join the MNE group later enter neither the numerator nor the denominator so expansion of the group—either through group expansion or data improvement—does not enter the measure. This set-up allows including group $\times$ year fixed effects and identify the effect from the difference between (CFC/substance versus non-CFC/substance *within the same group and year*. I estimate

$$Survival_{g,t,c} = \beta \left(Exposure_{g,t,c} \times TaxHaven_c \right) + \mu_{g,t} + \lambda_c + \varepsilon_{g,t,c}, \tag{7}$$

with g for the MNE group, t for the year and c for the affiliate type. $Survival_{g,t,c}$ is the share of group g's base-year affiliates of type c still belonging to the group in t; $Exposure_{g,t,c}$ is the continuous share of those affiliates within the scope of a CFC rule (Hypothesis 1a) or a substance requirement (Hypothesis 1b). $TaxHaven_c$ indicates tax-haven affiliates in the MNE group. The group $\times$ year fixed effects $\mu_{g,t}$ absorb year-varying group trends, including its growth, its underlying propensity to restructure, and other anti-avoidance legislation such as the Anti-Tax Avoidance Directive. λ_c absorbs baseline differencse in survival between tax-haven and other affiliates. Standard errors are clustered at the group level.

The estimated coefficient of interest β compares a group's own tax-haven affiliates with its other affiliates as its exposure to CFC rules and substance rules increases. Hypotheses 1a and 1b predict $\beta < 0$: the share of tax-haven affiliates should decrease when a CFC rule or substance rule kicks in. Because the comparison is within-MNE, the design does not rely on finding a comparable untreated MNE.

Because the outcome is a mean of binary indicators, observations are weighted by the number of base-year affiliates of the relevant type.[32] I require a group to be observed for at least 5 years and to hold at least 2 affiliates in its base year. [33]

6.3 Extensive Margin: Results

Figure 6 plots event-study coefficients for equation (7). The pre-treatment estimates are close to zero and individually insignificant. The CFC-rule estimates show that first-increased exposure to a CFC rule decreases the likelihood of survival of a specific affiliate significantly. The same holds for the substance rule: With increased exposure, the share of tax-haven affiliates decreases significantly.[34]

Table 3 reports estimates for equation (7). Exposure to a CFC rule varies within more groups than exposure to a substance requirement—782 versus 568. Once an affiliate becomes subject to a CFC rule, its likelihood of survival decreases by -0.017 with a standard error of 0.016: the point estimate has the predicted sign but cannot be distinguished from zero. The results thus do not provide robust evidence for Hypothesis 1a.

Exposure to a substance requirement decreases the survival of a group's substance-treated affiliates by 4.6 percentage points relative to its other affiliates, against a mean survival rate of 0.85, and the estimate is significant at the five percent level. Overall, the results support the extensive margin effect for substance affiliates in Hypothesis 1b: Multinational groups seem to reduce their number of tax-haven affiliates when those affiliates come within the scope of a substance requirement. This is in line with the idea that the fixed cost of maintaining tax-haven affiliates increases due to substance requirements.

To test the distance-effect in the model in Section 4, I measure, for each group and year, the average distance from its base-year productive affiliates to the nearest tax-haven affiliate still belonging to the MNE group. I then regress the distance on exposure, with group and year fixed effects and standard errors clustered at MNE level.

[32] Each observation is the share of a group's base-year affiliates of one type (treatment versus control) still in the group in a given year, so it is a sample proportion whose variance is inversely proportional to the number of affiliates; weighting by the affiliate count is therefore inverse-variance weighting

[33] Increasing the second threshold to three, six and eleven affiliates moves the substance-rule estimate only within -0.046 to -0.032.

[34] The event-study design is not optimal for a continuous-exposure treatment variable, and I code the event-study indicator as 1 upon first increase in exposure to the rule. The event study allows me to show visual evidence for hypotheses 1a and 1b.

Figure 6: Extensive Margin: The Share of Tax-Haven Affiliates

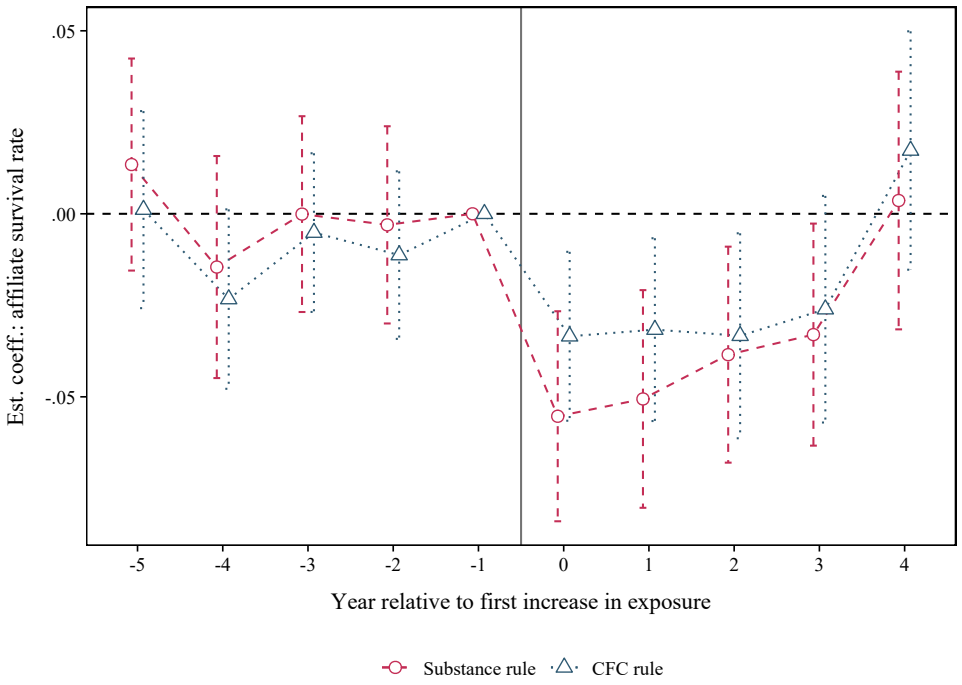

Figure 7: Extensive Margin: Distance to the Nearest Tax-Haven Affiliate

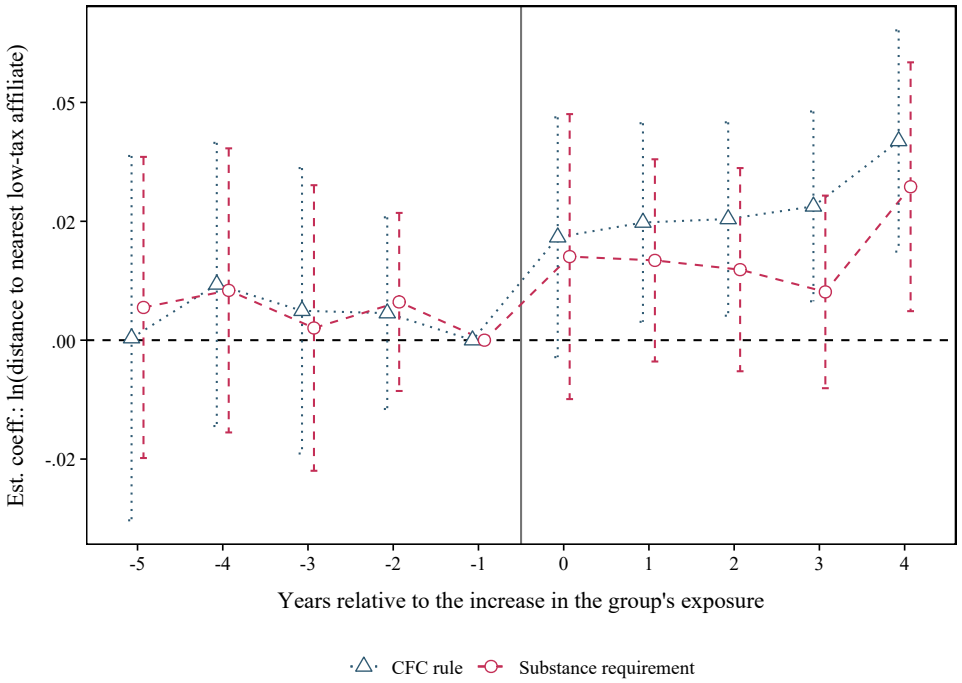

Note: Data from Orbis for 2010–2023. The figure plots event-study estimates for equation (7), in which exposure is interacted with indicators for years relative to the first increase in a group's exposure to a CFC rule. The dependent variable in Figure 6 is the share of a group's base-year tax-haven affiliates that still belong to the group. The dependent variable in Figure 7 is the average distance to the nearest-by tax-haven affiliate. Tax-havens defined as in Dharmapala and Hines (2009) excluding nine conduit countries. Year four pools that year and all later ones. 90% confidence intervals from standard errors clustered at the MNE level. All specifications include MNE × year and affiliate-type fixed effects and are weighted by the number of base-year affiliates of that type.

Table 3: The Number of Low-Tax Affiliates: CFC and Substance Rules

	Share of Low-Tax Affiliates		Distance to Nearest	
	(1) CFC Rule	(2) Substance Rule	(3) CFC Rule	(4) Substance Rule
Exposure × Low-Tax	-0.017	-0.046**		
	(0.016)	(0.019)		
Exposure			0.028**	0.018
			(0.012)	(0.013)
Group × Year FE	Yes	Yes	No	No
Group FE	No	No	Yes	Yes
Year FE	No	No	Yes	Yes
Affiliate-Type FE	Yes	Yes	No	No
Observations	74,160	74,160	32,351	32,351
MNE Groups	3,507	3,507	3,507	3,507
Mean Survival Rate	0.849	0.849		
Mean Distance (km)			5,999	5,999

Note: Group-level data from Orbis for the years 2010–2023. Each group's composition is fixed at its first observed year; the outcome is the share of those base-year affiliates of a given type still belonging to the group. Low-tax affiliates are those located in the tax havens listed by Dharmapala and Hines (2009), excluding Belgium, Cyprus, Hong Kong, Ireland, Luxembourg, Malta, the Netherlands, Singapore and Switzerland. *Exposure* is the share of the group's base-year affiliates of that type that fall within the scope of a CFC rule (column 1) or a substance rule (column 2), that is, whose parent country operates the rule and whose statutory tax rate lies below the parent's threshold. Group × year fixed effects absorb all group-level variation, so the coefficient is identified from the difference between low-tax and other affiliates of the same group in the same year. Observations are weighted by the number of base-year affiliates of that type. In columns 3 and 4 the outcome is instead the logarithm of the average distance from the group's base-year productive affiliates to the nearest low-tax affiliate still belonging to the group, where affiliate locations are country centroids. Production is held at the base-year roster and only the low-tax set is allowed to change, so the measure responds to the closure of low-tax affiliates rather than to the exit of productive ones; it is then a minimum over a shrinking set, and a positive coefficient means the surviving low-tax affiliates are worse placed to serve the group's production. The reported coefficient is the average of the five post-treatment event-study estimates plotted in Figure 7, with a standard error computed from their joint covariance. Distance is undefined once a group's last low-tax affiliate closes, so those group-years leave the sample, which retains 1.3 percent of them; the estimates are therefore conditional on the group still holding a low-tax affiliate and understate the unconditional change. Observations in columns 3 and 4 are weighted by the number of base-year productive affiliates. Standard errors are clustered at the MNE (GUO) level. ***, ** and * indicate significance at the 1%, 5% and 10% levels.

Columns (3) and (4) of Table 3 report the average effect over the five post-treatment years. Exposure to a CFC rule increases the distance to the nearest tax-haven affiliate by 2.8 percent, with a standard error of 0.012 and significance at the five percent level. Exposure to a substance requirement gives 1.8 percent (0.013), of the same sign and magnitude but insignificant.

Figure 7 shows the event-study estimates. The pre-treatment estimates are small and individually insignificant for both rules, and the effect then grows: by the fourth year the distance is 4.2 percent higher (1.4) after exposure to a CFC rule and 3.2 percent (1.6) after exposure to a substance requirement, both significant. Taken together with the decrease in the number of tax-haven affiliates in columns (1) and (2), the estimates are in line with the model's predictions.

7 The Intensive Margin

7.1 Measuring Investment in Tax-Haven Affiliates

Hypothesis 2 predicts that substance requirements increase real activity in tax-haven affiliates. As measures for real activity, I use the natural logarithm of tangible fixed assets, the number of employees and the cost of employees, each winsorised at the first and ninety-ninth percentiles. The sample is based on 21,629 treated affiliates that report tangible fixed assets, 20,536 report employment and 15,951 report employee costs in Orbis.

Treatment is defined as follows: An affiliate is treated in year t when its parent country operates a substance requirement *and* the affiliate's statutory tax rate lies below that country's threshold. Treatment is not absorbing: 11,084 affiliates, 37.1 percent of those ever treated, become untreated at some point. An affiliate can become untreated because the statutory rate crosses above the threshold—plausibly exogenous to the affiliate—or because its ownership changes. I keep the first type and drop the second type to address endogeneity concerns. Treatment changes due to ownership changes drop 1.53 percent of treated affiliate-years.

7.2 Identification

Because substance-treated affiliates and non-treated affiliates likely differ in unobserved ways, I consider three control groups.

The first control group includes affiliates whose parent country operates a CFC rule but no substance requirement, so the coefficient measures the *incremental* effect of the substance requirement over the CFC rule itself. This maps to the model most closely, because the CFC rule alone operates through the tax-rate channel and the substance requirement adds the fixed-cost channel.

The second control group includes never-treated affiliates whose parent country operates neither rule, giving the total effect of being subject to a substance rule.

The third control group restricts the CFC-only controls from the first control group to affiliates matched to the substance affiliates on pre-treatment size. I match each treated affiliate with a CFC-only affiliate in the same affiliate country and the same treatment-cohort year on log total assets in the year before treatment, using one-to-one exact matching without replacement, and a caliper of 0.25 of a pooled standard deviation. The matching yields 8,763 pairs and improves balance.[35]

Using each control group, I estimate

$$Y_{i,t} = \beta \, Substance_{i,t} + \alpha_i + \delta_t + \varepsilon_{i,t}, \tag{8}$$

with i for the affiliate and t for the year. $Y_{i,t}$ is the natural logarithm of tangible fixed assets, number of employees, or costs of employees. $Substance_{i,t}$ indicates that the affiliate is within the scope of a substance requirement. The affiliate fixed effects α_i absorb time-invariant differences between affiliates, including jurisdiction, sector and any average size difference. The year effects δ_t absorb common trends. Standard errors are clustered at the MNE group level.

Treatment is staggered over thirteen cohorts in 2010-2023, so already-treated affiliates act as controls for later ones, and treatment is non-absorbing. Both are known to bias the two-way estimator when effects are heterogeneous. I therefore also report estimates using De Chaisemartin and D'Haultfœuille (2024).

7.3 Investment in Tax-Haven Affiliates

Tables 4 through 6 report estimates for equation (8), with one table per outcome. The first four columns show estimates for the two unmatched control groups; the last two show the matched control group.

[35] The standardised difference in log total assets decreases from −0.095 to +0.005, and the median pair differs by 0.05 log points. 69 percent of treated affiliates have no match, but the loss of observations is not selective on size: the mean within-cell size percentile of the unmatched treated affiliates is 0.511 against 0.507 for the matched.

Tangible Fixed Assets. The event study in Figure 8 shows that the non-matched samples have significant pre-trends. The matched specification has no significant pre-trends, and shows a modest significant increase in tangible fixed assets upon becoming subject to substance rules.

Table 4 shows that upon becoming subject to a substance requirement, tangible fixed assets increase significantly under all specifications. Against affiliates whose parents operate a CFC rule but no substance requirement in column (2), tangible fixed assets increase by 6.8 percent. Against the never-treated in column (4), this difference is 9.9 percent. In the matched sample in column (6), tangible fixed assets increase by 5.3 percent.

Figure 8: Intensive Margin: Tangible Fixed Assets

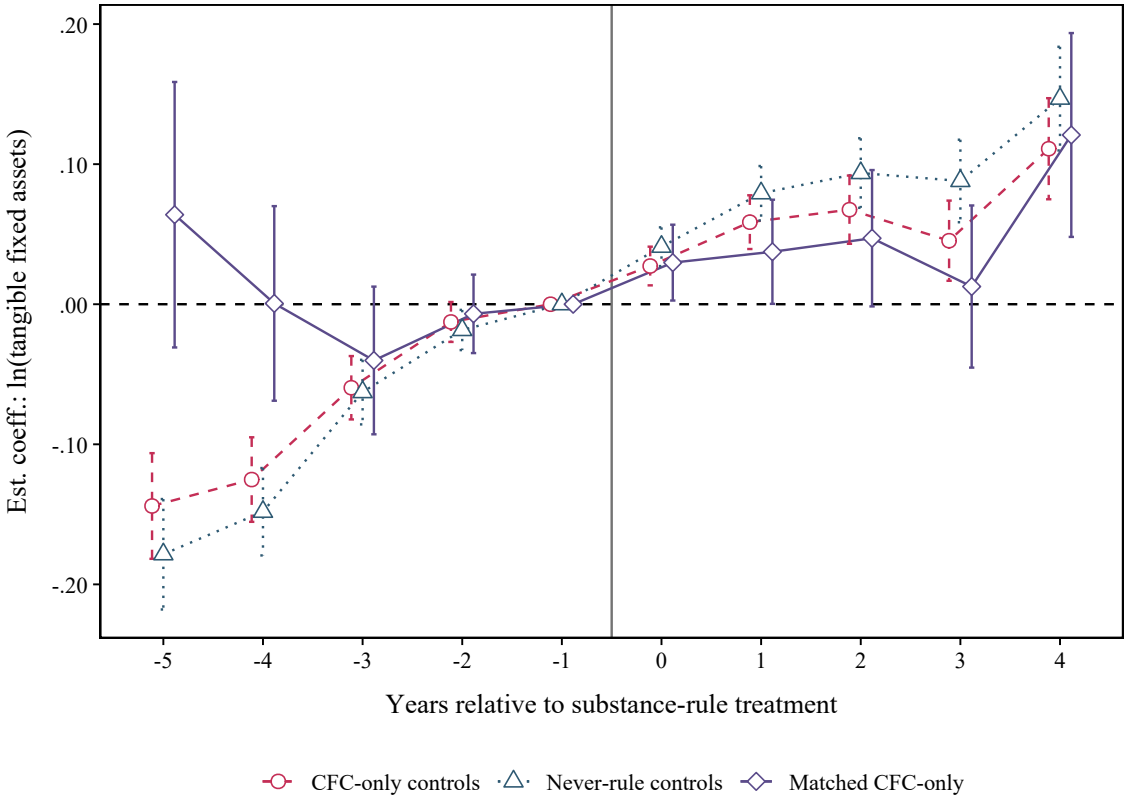

Note: Affiliate-year level data from Orbis over 2010–2023. Dynamic estimates using De Chaisemartin and D'Haultfœuille (2024) for the natural logarithm of tangible fixed assets for each of the three control groups. Includes affiliate and year fixed effects. 90% confidence intervals from standard errors clustered at the MNE level.

Table 4: Investment in Tangible Fixed Assets: Substance Rules

	CFC-Only		Never-Rule		Matched	
	(1) TWFE	(2) DCDH	(3) TWFE	(4) DCDH	(5) TWFE	(6) DCDH
Substance Rule	0.090***	0.068***	0.112***	0.099***	0.027	0.053**
	(0.012)	(0.013)	(0.012)	(0.014)	(0.019)	(0.026)
Affiliate FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,023,317	858,308	691,448	556,062	90,179	70,056
Affiliates	145,436	145,436	122,095	122,095	12,207	12,207
p(joint placebo)	–	0.000	–	0.000	–	0.342

Note: Affiliate-level data from Orbis for the years 2010–2023. *Substance Rule* indicates that the affiliate's parent country operates a substance requirement and the affiliate's statutory tax rate lies below that parent's threshold. Columns (1) and (2) use as controls affiliates whose parent country operates a CFC rule but no substance requirement; columns (3) and (4) use affiliates whose parent country operates neither rule; columns (5) and (6) use CFC-only controls matched to the treated affiliates within cohort year and host country on log total assets in the year before treatment, with a caliper of 0.25 of a pooled standard deviation. TWFE columns report two-way fixed-effects estimates. DCDH columns report the average total effect of the heterogeneity-robust estimator of De Chaisemartin and D'Haultfœuille (2024), estimated with five post-treatment effects and four placebos. All outcome variables are transformed by the natural logarithm and winsorised at the 1st and 99th percentiles. All columns include affiliate and year fixed effects; standard errors are clustered at the MNE (GUO) level. p(joint placebo) is the p-value of the joint test that all four placebo estimates equal zero; it is not defined for the two-way fixed-effects columns. ***, ** and * indicate significance at the 1%, 5% and 10% levels.

Number of Employees. The event study in Figure 9 again shows that the non-matched samples have significant pre-trends. All three specifications however, show a significant increase in the number of employees upon becoming subject to a substance rule.

Table 5 estimates that upon becoming subject to a substance requirement, the number of employees increases significantly under all specifications. Against affiliates that are subject to a CFC rule but not to a substance rule in column (2), the number of employees increases by 10.8 percent. Against the never-treated in column (4), this difference is 13.5 percent. In the matched sample in column (6), the number of employees increases by 6.7 percent.

Figure 9: Intensive Margin: Number of Employees

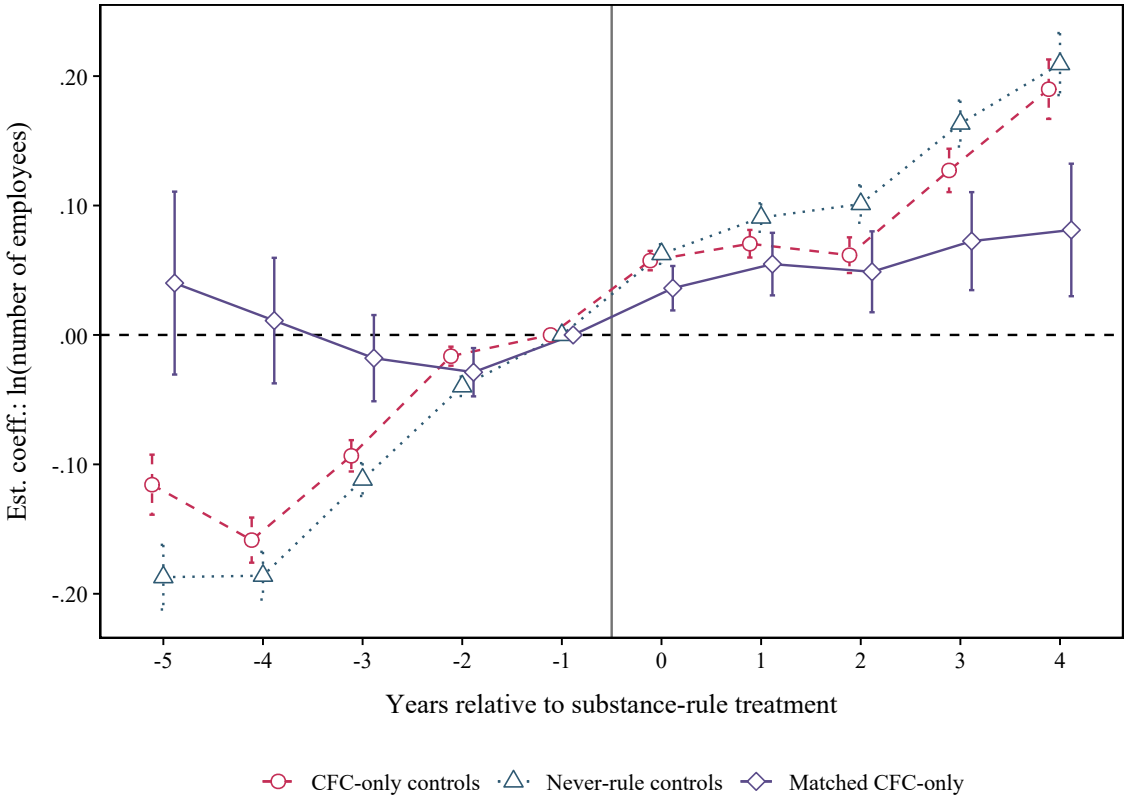

Note: Affiliate-year level data from Orbis over 2010–2023. Dynamic estimates using De Chaisemartin and D'Haultfœuille (2024) for the natural logarithm of the number of employees for each of the three control groups. Includes affiliate and year fixed effects. 90% confidence intervals from standard errors clustered at the MNE level.

Table 5: Investment in Number of Employees: Substance Rules

	CFC-Only		Never-Rule		Matched	
	(1) TWFE	(2) DCDH	(3) TWFE	(4) DCDH	(5) TWFE	(6) DCDH
Substance Rule	0.144***	0.108***	0.148***	0.135***	0.056***	0.067***
	(0.007)	(0.008)	(0.007)	(0.008)	(0.013)	(0.018)
Affiliate FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	827,771	690,365	511,584	405,483	69,638	53,767
Affiliates	126,363	126,363	101,899	101,899	10,353	10,353
p(joint placebo)	–	0.000	–	0.000	–	0.069

Note: Affiliate-level data from Orbis for the years 2010–2023. *Substance Rule* indicates that the affiliate's parent country operates a substance requirement and the affiliate's statutory tax rate lies below that parent's threshold. Columns (1) and (2) use as controls affiliates whose parent country operates a CFC rule but no substance requirement; columns (3) and (4) use affiliates whose parent country operates neither rule; columns (5) and (6) use CFC-only controls matched to the treated affiliates within cohort year and host country on log total assets in the year before treatment, with a caliper of 0.25 of a pooled standard deviation. TWFE columns report two-way fixed-effects estimates. DCDH columns report the average total effect of the heterogeneity-robust estimator of De Chaisemartin and D'Haultfœuille (2024), estimated with five post-treatment effects and four placebos. All outcome variables are transformed by the natural logarithm and winsorised at the 1st and 99th percentiles. All columns include affiliate and year fixed effects; standard errors are clustered at the MNE (GUO) level. p(joint placebo) is the p-value of the joint test that all four placebo estimates equal zero; it is not defined for the two-way fixed-effects columns. ***, ** and * indicate significance at the 1%, 5% and 10% levels.

Costs of Employees. The event study in Figure 10 shows that the non-matched samples have significant pre-trends also in this cae, and the matched sample now has a significant estimate for the last year before treatment. Upon becoming subject to a substance rule, the event-study coefficients show a significant increase in the costs of employees.

Table 6 estimates that upon becoming subject to a substance rule, the costs of employees increase significantly under all specifications. Against the CFC-only control group in column (2), the costs of employees increase by 5.9 percent. Against the never-treated in column (4), this difference is 7.0 percent. In the matched sample in column (6), the costs of employees increase by 5.9 percent.

Figure 10: Intensive Margin: Costs of Employees

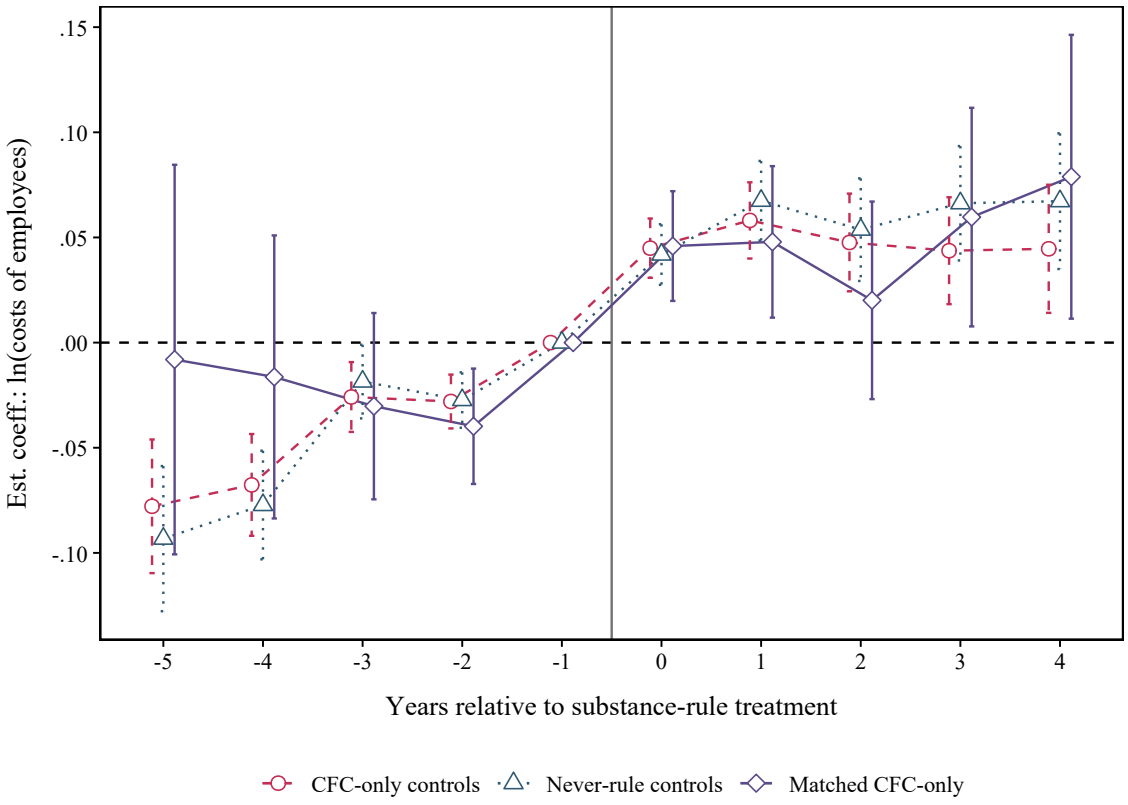

Note: Affiliate-year level data from Orbis over 2010–2023. Dynamic estimates using De Chaisemartin and D'Haultfœuille (2024) for the natural logarithm of the costs of employees for each of the three control groups. Includes affiliate and year fixed effects. 90% confidence intervals from standard errors clustered at the MNE level.

Table 6: Investment in Costs of Employees: Substance Rules

	CFC-Only		Never-Rule		Matched	
	(1) TWFE	(2) DCDH	(3) TWFE	(4) DCDH	(5) TWFE	(6) DCDH
Substance Rule	0.085***	0.059***	0.105***	0.070***	0.080***	0.059**
	(0.010)	(0.012)	(0.010)	(0.013)	(0.018)	(0.025)
Affiliate FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	594,899	497,348	341,222	271,447	57,037	44,476
Affiliates	87,396	87,396	64,366	64,366	8,059	8,059
p(joint placebo)	–	0.000	–	0.000	–	0.181

Note: Affiliate-level data from Orbis for the years 2010–2023. *Substance Rule* indicates that the affiliate's parent country operates a substance requirement and the affiliate's statutory tax rate lies below that parent's threshold. Columns (1) and (2) use as controls affiliates whose parent country operates a CFC rule but no substance requirement; columns (3) and (4) use affiliates whose parent country operates neither rule; columns (5) and (6) use CFC-only controls matched to the treated affiliates within cohort year and host country on log total assets in the year before treatment, with a caliper of 0.25 of a pooled standard deviation. TWFE columns report two-way fixed-effects estimates. DCDH columns report the average total effect of the heterogeneity-robust estimator of De Chaisemartin and D'Haultfœuille (2024), estimated with five post-treatment effects and four placebos. All outcome variables are transformed by the natural logarithm and winsorised at the 1st and 99th percentiles. All columns include affiliate and year fixed effects; standard errors are clustered at the MNE (GUO) level. p(joint placebo) is the p-value of the joint test that all four placebo estimates equal zero; it is not defined for the two-way fixed-effects columns. ***, ** and * indicate significance at the 1%, 5% and 10% levels.

The results in this section are in line with the model's predictions. A substance requirement makes the exemption from the CFC rule conditional on real activity in the tax-haven affiliate. Affiliates that enter the scope of the substance rule increase both capital and labor relative to otherwise similar affiliates facing a CFC rule without such a test. Taken together with the extensive-margin results in Section 6.3, the estimates suggest that MNEs close some of their tax-haven affiliates, and invest more in the ones they keep. That is the mechanism that the model in Section 4 describes, and it shows that substance rules affect MNE group structures.

8 The Spillover Effect

8.1 Measuring Investment in Unaffected Affiliates

Hypothesis 3 asks whether real investment in affiliates not subject to a substance rule changes when a different affiliate of the same multinational group is brought within the scope of a rule. To test this, treatment is defined as follows: an affiliate is exposed in year t when *another* affiliate in the MNE group becomes subject to a substance requirement in that year.

The outcome variables are the same three as in Section 7.1—the natural logarithm of tangible fixed assets, of the number of employees and of the cost of employees, each winsorised at the first and ninety-ninth percentiles. 21,952 tax-haven affiliates and 150,391 other affiliates are exposed to within-MNE substance treatment at some point. Because Orbis does not report every outcome for every affiliate, the samples for analysis rely on smaller numbers.

8.2 Identification

I estimate whether outcomes in unaffected affiliates change once another affiliate within the group becomes subject to substance rules,

$$Y_{i,t} = \gamma \, GroupSubstance_{i,t} + \alpha_i + \delta_t + \varepsilon_{i,t}, \tag{9}$$

with i for the affiliate and t for the year. $GroupSubstance_{i,t}$ indicates that another affiliate in affiliate i's group is within the scope of a substance requirement. The affiliate fixed effects α_i absorb time-invariant differences between affiliates and the year effects δ_t common trends, as in equation (8). Standard errors are clustered at the MNE level.

If the MNE shifts real activity out of its high-tax affiliates, the estimate would be $\gamma < 0$ (Buettner et al., 2018; Suárez Serrato, 2019; Drake et al., 2022; Williams, 2018). Under complementarity of investment between tax-haven affiliates and productive affiliates, investment in treated and untreated affiliates moves together and $\gamma > 0$ in both (Hong and Smart, 2010).

I use two control groups. MNE groups with affiliates that are subject to substance rules are substantially larger than groups that do not have substance-treated affiliates. Therefore, I first report estimates using all never-exposed groups. Second, I report estimates using a matched control group, in which each exposed group is paired with an MNE group subject to a CFC rule but never to a substance requirement. I exact-match on cohort-year and country, as well as log total assets.

As on the intensive margin, exposure is staggered and non-absorbing, so I report the two-way fixed-effects estimator alongside the heterogeneity-robust estimator of De Chaisemartin and D'Haultfœuille (2024), with 5 post-treatment effects and 4 placebos.

8.3 Investment in Unaffected Affiliates

Figure 11 reports estimated coefficients for an event-study based on equation (9). The results for tangible fixed assets are insignificant across specifications. Some of the estimated coefficients for the number of employees are significant in later years.

Table 7 reports the difference-in-differences estimates of equation (9). Panel A compares exposed MNEs with never-exposed MNEs. For tangible fixed assets, the results are insignificant across specifications. The two-way fixed-effects estimates are positive and significant for the number of employees (0.021) and employee costs (0.014), while the heterogeneity-robust estimates are negative and significant (-0.016 and -0.019).

MNEs with substance-treated affiliates are larger, which is why I also obtain results for a matched sample. This reduces the sample to 3,070 matched MNEs, but the pre-trends in Figure 12 are insignificant. Now, all estimated coefficients are insignificant as well.

The difference-in-differences estimates are reported in Panel B of Table 7. The TWFE estimates for the number of employees and the costs of employees are positive and significant, but since treatment is staggered, I prefer to rely on the estimates using De Chaisemartin and D'Haultfœuille (2024). Those estimates are all insignificant.

Taken together, the estimates in this section suggest that there is no significant evidence for an investment-shifting effect.

Figure 11: Investment Shifting, Unmatched Control Group

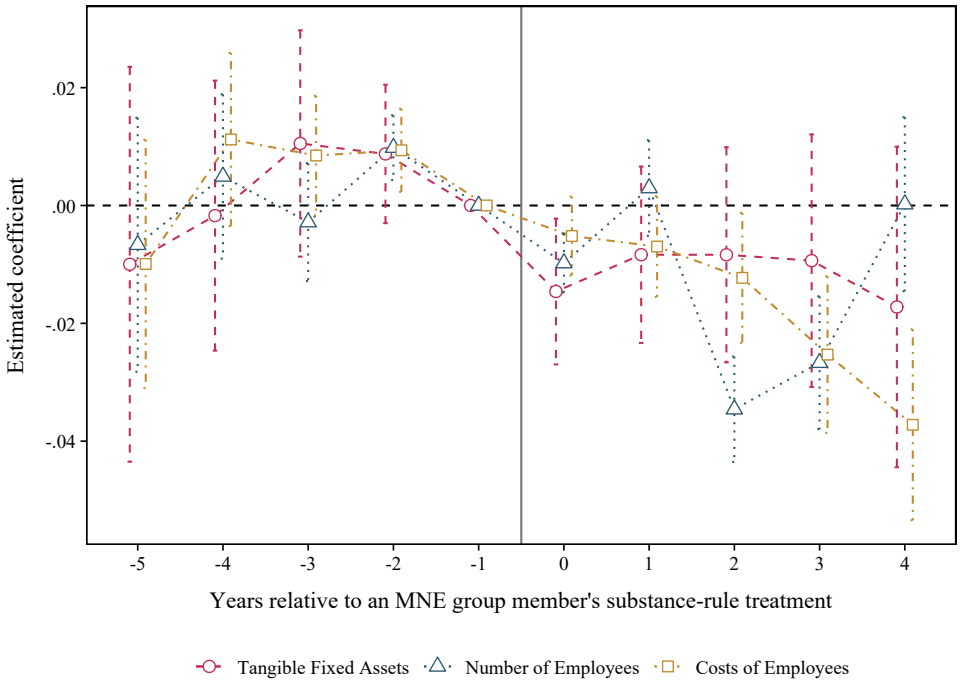

Figure 12: Investment Shifting, Matched Control Group

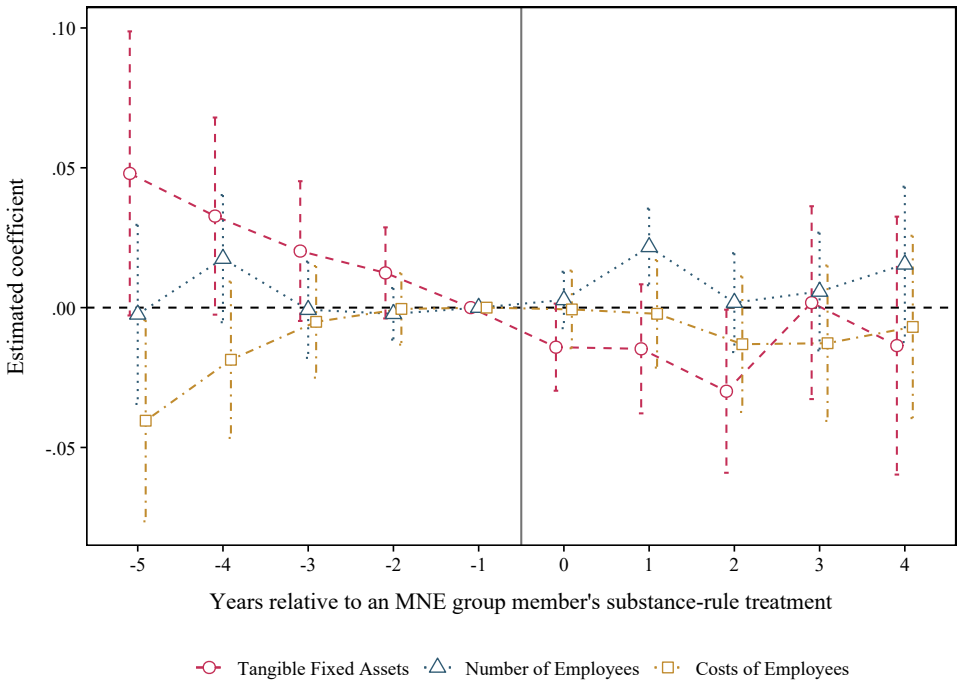

Note: Affiliate-year level data from Orbis over 2010–2023. Dynamic estimates using De Chaisemartin and D'Haultfœuille (2024). Outcomes are the natural logarithm of tangible fixed assets (red, circles), the number of employees (blue, triangles) and the costs of employees (gold, squares). Exposure is another affiliate of the same MNE group becoming subject to a substance requirement, MNE group composition is fixed in the first-observed year. Includes affiliate and year fixed effects. 90% confidence intervals from standard errors clustered at the MNE level.

Table 7: Investment Shifting in High-Tax Affiliates: Substance Rules

	Tangible Fixed Assets		Number of Employees		Costs of Employees	
	(1) TWFE	(2) DCDH	(3) TWFE	(4) DCDH	(5) TWFE	(6) DCDH
Panel A: Unmatched control group						
Group-Substance	0.009	-0.014	0.021***	-0.016***	0.014**	-0.019***
	(0.010)	(0.011)	(0.006)	(0.006)	(0.007)	(0.007)
Observations	1,482,056	1,197,531	1,465,496	1,152,332	1,480,333	1,185,448
Affiliates	212,760	212,760	256,091	256,091	206,381	206,381
p(joint placebo)	–	0.704	–	0.005	–	0.045
Panel B: Matched control group						
Group-Substance	-0.001	-0.019	0.027***	0.011	0.024**	-0.008
	(0.014)	(0.018)	(0.010)	(0.010)	(0.010)	(0.014)
Observations	300,056	238,468	237,056	184,085	200,374	158,854
Affiliates	39,977	39,977	38,216	38,216	26,330	26,330
p(joint placebo)	–	0.497	–	0.447	–	0.451
Affiliate FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Note: Affiliate-level data from Orbis for the years 2010–2023. *Group-Substance* indicates that another affiliate in the same group is within the scope of a substance requirement in that year. MNE group composition is fixed in the first observed year. Panel A compares exposed MNEs with all MNEs never exposed to a substance requirement. Panel B restricts the comparison to 3,070 matched MNEs. I match CFC-treated MNEs on cohort, affiliate country, and total assets within a caliper of 0.25 pooled standard deviations. TWFE columns report two-way fixed-effects estimates and DCDH columns the average total effect of the heterogeneity-robust estimator of De Chaisemartin and D'Haultfœuille (2024), estimated with 5 post-treatment effects and 4 placebos. All outcome variables are transformed by the natural logarithm and winsorised at the 1st and 99th percentiles. All columns include affiliate and year fixed effects; standard errors are clustered at the MNE (GUO) level. ***, ** and * indicate significance at the 1%, 5% and 10% levels.

9 Conclusion

Countries have implemented CFC rules to target MNEs' shifting of passive income to low-tax countries. Many countries provide exemptions to CFC rules in case an MNE proves that an affiliate has sufficient substance to plausibly carry out real economic activity. This substance exemption incentivises MNEs to increase real investment in low-tax countries to facilitate profit shifting.

In this paper, I modeled the effect of substance rules on real investment in fixed assets and employment in tax-haven affiliates. The model predicts that a CFC rule should decrease real investment in tax-haven affiliates, because profit shifting becomes less attractive. However, if there is a substance exemption, investment in tax-haven affiliates should increase. The reason for this effect is that convincing tax authorities that real economic activities are being carried out in a tax-haven affiliate allows the MNE to avoid application of the CFC rule.

Using data from Orbis, I text the model's predictions. On the extensive margin, substance rules decrease the share of affected affiliates in the MNE group. The intensive-margin results then suggest that the MNE then concentrates investment in remaining tax-haven affiliates to satisfy the substance requirement.

The reason to implement substance rules is to ensure that anti-profit shifting rules have a carve-out for genuine economic activities (OECD, 2015, p. 16). The results in this paper indicate that substance rules may increase real investment in tax-haven affiliates. Importantly, it seems that there is no investment shifting from non-affected (high-tax) affiliates to tax-haven affiliates.

A Deriving the Net Profits Shifted per Affiliate

This appendix derives the shifting benefit per tax-haven affiliate, equation (3), from the optimal profits shifted in equation (2). Using the after-tax profit of each productive affiliate in equation (2), the MNE's after-tax profit, before it chooses k_j and h, is

$$\Pi = (1 - t_i)m_i + \sum_{j=1}^{h} \left[\int_{\hat{x}_-}^{\hat{x}_+} \frac{1}{2} \left(t_i - t_j - \frac{\bar{\alpha}}{k_j} \right)^2 dx - \frac{\phi k_j^2}{2} \right], \tag{10}$$

where the first term is after-tax profit from production and the second collects, for each tax-haven affiliate j, the shifting benefits of the productive affiliates it serves net of its fixed cost. The limits $\hat{x}_-$ and $\hat{x}_+$ are the productive affiliates that are just indifferent between two closest-by tax-haven affiliates.

A productive affiliate is indifferent between two tax-haven affiliates where its after-tax profits are equalized, $\pi^*_{i,j} = \pi^*_{i,j+1} = \pi^*_{i,j-1}$. Solving the upper- and lower-bound indifference conditions yields

$$\hat{x}_+ = \frac{k_j^2}{2\bar{\alpha}} \left[t_j - t_{j+1} + \frac{\bar{\alpha}}{k_j} z_j + \frac{\bar{\alpha}}{k_j} z_{j+1} \right], \qquad \hat{x}_- = \frac{k_j^2}{2\bar{\alpha}} \left[-t_j + t_{j-1} + \frac{\bar{\alpha}}{k_j} z_j + \frac{\bar{\alpha}}{k_j} z_{j-1} \right].$$

Assuming the tax-haven affiliates are symmetric, $t_j = t_{j+1} = t_{j-1}$ and using that the Salop circle has measure one, the limits simplify to $\hat{x}_+ = \frac{1}{2h}$ and $\hat{x}_- = -\frac{1}{2h}$. Evaluating the integral in equation (10) over $[\hat{x}_-, \hat{x}_+]$ then yields the shifting benefit per tax-haven affiliate in equation (3).

B Tables

Table A1: CFC Rules

Country	CFC Rule	Year	Substance Rule	Tax Rate Threshold
Argentina	Yes	1999	Subjective	75% of Argentinian ETR
Australia*	Yes	1990	No	no threshold
Austria*	Yes	2019	Subjective	12.5% ETR
Belgium*	Yes	2017	No	50% of Belgian ETR
Brazil*	Yes	1995	No	no threshold
Bulgaria*	Yes	2019	Subjective	lower than Bulgarian ETR
Canada	Yes	1976	Safe Harbour	no threshold
Chile	Yes	2014	Subjective	30% ETR
China*	Yes	2008	Subjective	12.5% ETR
Colombia*	Yes	2017	Subjective	no threshold
Croatia*	Yes	2019	Subjective	lower than Croatian ETR
Cyprus*	Yes	2019	Subjective	50% of Cypriotic ETR
Czech Republic*	Yes	2019	Subjective	50% of Czech ETR
Denmark*	Yes	1995	No	no threshold
Estonia*	Yes	2000	Subjective	no threshold
Finland*	Yes	1995	Subjective	60% of Finnish ETR
France*	Yes	1980	Subjective	50% of French tax rate
Germany*	Yes	1972	Subjective, EEA only	25% of German tax rate
Greece*	Yes	2014	Subjective, EEA only	50% of Greek ETR
Hungary*	Yes	2017	Safe Harbour	50% of Hungarian ETR
Iceland*	Yes	2010	No	66.67% of Icelandic tax rate
Indonesia	Yes	1995	No	no threshold
Ireland*	Yes	2019	Subjective	lower than Irish ETR
Israel	Yes	2003	No	15% ETR
Italy*	Yes	2001	Subjective	50% of Italian ETR
Japan*	Yes	1978	Subjective	20% absolute ETR
Kazakhstan*	Yes	1995	No	10% absolute ETR
Latvia*	Yes	2019	Subjective	lower than Latvian ETR
Lithuania*	Yes	2004	Subjective	50% of Lithuanian ETR
Luxembourg*	Yes	2019	No	lower than Luxembourg ETR
Malta*	Yes	2019	Subjective	lower than Maltese ETR
Mauritius*	Yes	2019	No	50% of Mauritian tax rate
Mexico*	Yes	1997	Subjective	75% of Mexican ETR
Netherlands*	Yes	2019	Subjective	9% absolute tax rate

Table A1 continued from previous page

Country	CFC Rule	Year	Substance Rule	Tax Rate Threshold
New Zealand*	Yes	1988	No	no threshold
Norway*	Yes	1992	Subjective, EEA only	66.67% of Norwegian ETR
Pakistan	Yes	2019	No	60% of Pakistani ETR
Peru	Yes	2013	No	75% of Peruvian ETR
Poland*	Yes	2015	Subjective, EEA only	75% of Polish ETR
Portugal*	Yes	1995	Subjective	50% of Portugese ETR
Romania*	Yes	2018	Subjective, EEA only	lower than Romanian ETR
Russia*	Yes	2015	No	75% of Russian tax rate
Slovak Republic*	Yes	2019	Subjective	lower than Slovakian ETR
Slovenia*	Yes	2019	Subjective	50% of Slovenian ETR
South Africa	Yes	1997	Safe Harbour	75% of South African ETR
South Korea*	Yes	1995	Subjective	15% absolute ETR
Spain*	Yes	1995	Subjective	75% of Spanish ETR
Sweden*	Yes	1989	Subjective, EEA only	55% of Swedish tax rate
Türkiye*	Yes	2006	No	10% ETR
United Kingdom*	Yes	1984	Subjective	75% of UK CIT rate
United States	Yes	1962	No	90% of US ETR

Note: The table shows data on CFC rules and substance rules from OECD (2024), in addition to hand-collected data from national CFC legislation. The data relates to the analysed sample from Orbis for the years 2010–2023 on the observed parent countries for both Global Ultimate Owners and Immediate Shareholders. An asterisk (*) indicates that the sample observes affiliates located in this country; 9 of the 51 countries are observed as parent countries only. *CFC Rule* indicates that a country has implemented CFC rules. *Year* shows the year of introduction of the CFC rule. *Substance Rule* indicates whether a country has implemented a substance rule and what it looks like. Subjective rules require testing the specific facts and circumstances. Safe Harbour rules are fixed thresholds. *Tax Rate Threshold* indicates the threshold below which a Controlled Foreign Company becomes subject to CFC rules. These thresholds can be absolute or relative, and can be based on statutory or effective tax rates.

Table A2: Other Countries with Affiliates in Sample

Country			
Albania	Egypt	Moldova	Switzerland
Bahamas	Gibraltar	Montenegro	Taiwan
Barbados	Guernsey	Morocco	Ukraine
Bermuda	Hong Kong	North Macedonia	United Arab Emirates
Bosnia and Herzegovina	India	Philippines	Uruguay
British Virgin Islands	Jersey	Serbia	
Cayman Islands	Malaysia	Singapore	

Note: Table shows the additional countries in which the sample observes affiliates and which do not have a CFC rule in place. Data from Orbis for the years 2010–2023.

References

Alstadsæter, A., Davies, R. B., Parenti, M., Toubal, F., 2024. "The real effects of tax havens," In Research Handbook on the Economics of Tax Havens, Edward Elgar Publishing.

Becker, J., Riedel, N., 2012. "Cross-border tax effects on affiliate investment—evidence from european multinationals," European Economic Review 56, 436–450.

Beer, S., de Mooij, R., Liu, L., 2020. "International corporate tax avoidance: A review of the channels, magnitudes, and blind spots," Journal of Economic Surveys 34, 660–688.

Brekke, K. R., Garcia Pires, A. J., Schindler, D., Schjelderup, G., 2017. "Capital taxation and imperfect competition: ACE vs. CBIT," Journal of Public Economics 147, 1–15.

Buettner, T., Overesch, M., Wamser, G., 2018. "Anti profit-shifting rules and foreign direct investment," International Tax and Public Finance 25, 553–580.

Clifford, S., 2019. "Taxing multinationals beyond borders: Financial and locational responses to CFC rules," Journal of Public Economics 173, 44–71.

De Chaisemartin, C., D'Haultfœuille, X., 2024. "Difference-in-differences estimators of intertemporal treatment effects," Review of Economics and Statistics 108, 1–18.

De Mooij, R., Liu, L., 2020. "At a cost: The real effects of transfer pricing regulations," IMF Economic Review 68, 268–306.

De Simone, L., Olbert, M., 2022. "Real effects of private country-by-country disclosure," The Accounting Review 97, 201–232.

Desai, M. A., Dharmapala, D., 2009. "Earnings management, corporate tax shelters, and book-tax alignment," National Tax Journal 62, 169–186.

Desai, M. A., Foley, C. F., Hines, J. R., 2004. "Economic effects of regional tax havens," NBER Working Paper Series 10806.

Desai, M. A., Foley, C. F., Hines, J. R., 2006. "Do tax havens divert economic activity?" Economics Letters 90, 219–224.

Dharmapala, D., 2020. "Do Multinational Firms use Tax Havens to the Detriment of Other Countries?" Coase-Sandor Institute for Law & Economics Research Paper No. 910.

Dharmapala, D., Hines, J. R., 2009. "Which countries become tax havens?" Journal of Public Economics 93, 1058–1068.

Drake, K. D., Goldman, N. C., Murphy, F., 2022. "Foreign employment, income shifting, and tax uncertainty," The Accounting Review 97, 183–212.

Egger, P. H., Wamser, G., 2015. "The impact of controlled foreign company legislation on real investments abroad. A multi-dimensional regression discontinuity design," Journal of Public Economics 129, 77–91.

Ferrari, A., Laffitte, S., Parenti, M., Toubal, F., 2023. "Profit shifting frictions and the geography of multinational activity," CEPR DP17801.

Fuest, C., Hugger, F., Neumeier, F., 2022. "Corporate profit shifting and the role of tax havens: Evidence from German country-by-country reporting data," Journal of Economic Behavior & Organization 194, 454–477.

Gresik, T. A., Konrad, K. A., 2017. "Tax havens, accounting experts, and fee-setting rules," CESifo Working Paper Series No. 6774.

Gschossmann, E., Pfrang, A., 2024. "Multinationals' location, financial and real responses to the EU-wide implementation of CFC rules by the ATAD," SSRN Electronic Jounal 4735272.

Gunny, K. A., 2010. "The relation between earnings management using real activities manipulation and future performance: Evidence from meeting earnings benchmarks*," Contemporary Accounting Research 27, 855–888.

Hansen, J., Merlo, V., Wamser, G., 2023. "Taxes, profit shifting, and the real activities of mnes: Evidence from corporate tax notches," CESifo Working Paper Series 10593.

Haufler, A., Mardan, M., Schindler, D., 2018. "Double tax discrimination to attract FDI and fight profit shifting: The role of CFC rules," Journal of International Economics 114, 25–43.

Hong, Q., Smart, M., 2010. "In praise of tax havens: International tax planning and foreign direct investment," European Economic Review 54, 82–95.

Kofler, G., Krever, R., Lang, M., Owens, J., Pistone, P., Rust, A., Schuch, J., Spies, K., Staringer, C., Papulova, A., Ramharter, L., Rosca, I.-F., 2020. Controlled Foreign Company Legislation. WU - Tax Law and Policy Series, IBFD.

Laffitte, S., 2024. "The market for tax havens," Working Paper https://slaffitte1.github.io/website/SL_V4.pdf.

Merlo, V., Wamser, G., 2024. "Profit-shifting elasticities, channels, and the role of tax havens: Evidence from micro-level data," In Research Handbook on the Economics of Tax Havens, Edward Elgar Publishing.

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

OECD, 2022. OECD Transfer Pricing Guidelines for Multinational Enterprises and Tax Administrations 2022. OECD Transfer Pricing Guidelines for Multinational Enterprises and Tax Administrations, OECD.

OECD, 2024. "Controlled Foreign Company (CFC) Rules," https://qdd.oecd.org/data/CFC/.ALL.

Overesch, M., Schindler, D., Wamser, G., 2024. "Design and consequences of CFC and GILTI rules: A review and potential lessons for the global minimum tax," Research Handbook on Tax Havens, A. Lejour and D. Schindler (eds), Mimeo.

Pieretti, P., Pulina, G., 2020. "Does eliminating international profit shifting increase tax revenue in high-tax countries?" Economic Modelling 93, 717–727.

Ruf, M., Weichenrieder, A. J., 2013. "CFC legislation, passive assets and the impact of the ECJ's cadbury-schweppes decision," CESifo Working Paper No. 4461.

Salop, S. C., 1979. "Monopolistic competition with outside goods," The Bell Journal of Economics 10, 141–156.

Schenkelberg, S., 2020. "The cadbury schweppes judgment and its implications on profit shifting activities within europe," International Tax and Public Finance 27, 1–31.

Suárez Serrato, J. C., 2019. "Unintended consequences of eliminating tax havens," American Economic Review 24850.

Tørsløv, T., Wier, L., Zucman, G., 2023. "The missing profits of nations," The Review of Economic Studies 90, 1499–1534.

United Nations Conference on Trade and Development, 2016. World Investment Report 2016: Investor Nationality - Policy Challenges. United Nations Conference on Trade and Development (UNCTAD) World Investment Report (WIR), United Nations.

Voget, J., 2011. "Relocation of headquarters and international taxation," Journal of Public Economics 95, 1067–1081.

Weichenrieder, A. J., 1996. "Anti-tax-avoidance provisions and the size of foreign direct investment," International Tax and Public Finance 3, 67–81.

Whitehead, S., 2007. "Practical implications arising from the European Court's recent decisions concerning CFC legislation and dividend taxation," EC Tax Review 16, 176–183.

Williams, B. M., 2018. "Multinational tax incentives and offshored U.S. jobs," The Accounting Review 93, 293–324.