Tax Policy, Investment and Profit Shifting

1 November 2024

Katarzyna Bilicka (Utah State University, NBER, and Oxford University Centre for Business Taxation)

Michael Devereux (Oxford University Centre for Business Taxation)

İrem Güçeri (Oxford University Blavatnik School of Government and Centre for Business Taxation)

Tax Policy, Investment
and Profit Shifting

Katarzyna Bilicka[*]   Michael Devereux[†]

İrem Güçeri[‡]

November 1, 2024

Abstract

Many multinational firms (MNEs) pay low or no corporation tax in high-tax countries because they shift taxable income to tax havens. We incorporate nonconvex costs of profit shifting and unobserved heterogeneity in profit-shifting ability in the MNEs' value maximization problem to study responses of firms to tax policies. We estimate our model using UK corporate tax returns data and quantify: (i) the elasticities of tax base and capital stock with respect to tax rates, (ii) the fixed and variable components of profit-shifting costs for different firm types, and (iii) the government's trade-off between raising tax revenue by reducing profit shifting and attracting investment. Accounting for extensive margin profit-reporting decisions, we reconcile most of the discrepancies between previous micro- and macro-level estimates of tax base elasticities. We test the predictions of the model using a quasi-natural experiment that restricted profit-shifting by Italian MNEs that operated in the UK and evaluate two types of tax policies that can be analyzed using our approach.[1]

1 Introduction

What are the behavioral responses to changes in tax policy for multinationals that do not pay any corporation tax in high-tax countries? The incidence of reporting no taxable profit is much more common for multinationals (MNEs) than it is for domestic companies, especially in high-tax jurisdictions, as MNEs have the resources to shift a substantial amount of their taxable profits to tax havens. Those low- or zero-tax paying MNEs reduce their global tax liability while also reducing their cost of capital for investment, while others have more limited access to such capabilities. As a result, MNEs' profits display sharp bunching at zero taxable income, highlighting the importance of extensive-margin profit-shifting decisions (Bilicka; 2019; Koethenbuerger et al.; 2019). Surprisingly, the micro-level tax avoidance literature has placed a strong emphasis on intensive-margin profit-shifting decisions, generating a gap between macro- and micro-level studies (Dharmapala; 2014; Riedel; 2018).[2] Understanding the responses of these non-tax paying MNEs and the trade-offs they are facing has become more important, as *"leaders representing 80% of the world's GDP [...] made clear their support for a strong global minimum tax"*[3] to combat profit shifting.

In this paper, we develop a model that generates corner solutions in multinationals' choices regarding the proportion of profit that they shift to tax havens. We address the commonly-overlooked issue that a high proportion of multinationals are at extreme positions of profit shifting. We use this model, combined with data from the UK tax returns, to estimate structural profit-shifting and investment elasticities and to provide estimates of the costs of profit shifting. We then test the model predictions in a reduced-form difference-in-differences framework. To do that, we investigate the effects of a tax reform aimed at limiting profit shifting on the intensive- and extensive-margin of declared taxable profits and on investment in the UK by subsidiaries of foreign-owned multinationals. Finally, we conduct counterfactual policy experiments in which we analyze changes in corporate tax rates and the introduction of the global minimum tax (GMT) at varying rates. Our findings provide four distinct contributions to the literature as they allow us to: (i) reconcile the differences across previous estimates of the extent of profit shifting based on macro- and micro-data, (ii) quantify the impact of curbing profit shifting on investment, (iii) quantify the size of profit-shifting costs, and (iv) identify the overall impact of proposed tax reforms on global welfare.

---

[2]Such a discrepancy is akin to the gap between macro- and micro-level labor supply elasticities in the presence of frictions in Chetty (2012).

[3]Joe Biden, President of the United States, 30 October 2021.

We start the paper by introducing our conceptual framework in which an MNE undertakes real productive activity in several possible locations and has a subsidiary in a tax haven, to which it shifts profit. A key novelty is in how we model the accumulation of tax avoidance capability. The MNE invests in a tax avoidance asset, which represents a public good for all subsidiaries of the MNE worldwide and reduces the marginal cost of shifting every additional dollar of profit to the tax haven. This is different from prior literature that typically uses the Hines and Rice's 1994 convex profit-shifting cost model and either disregards the subsidiaries of MNEs at corner solutions of taxable income reporting or applies an ad hoc empirical adjustment to deal with MNE subsidiaries that report zero taxable income which are, arguably, the worst offenders in tax avoidance.[4] In our framework, each non-haven subsidiary also invests in productive capital, and chooses how much profit to shift to the haven given its marginal profit-shifting cost. Each MNE faces an idiosyncratic price of investing in each unit of the tax avoidance capability. This heterogeneity mimics the real-world differences across companies that are not captured by standard sectoral or firm size variation: for example, some firms are more digitalized or have a more international customer base than others. Such firms find it easier to move profits to lower tax jurisdictions.

This tax avoidance capability is crucial in our model and can take many forms. In particular, papers in the literature on profit shifting estimate the changes in MNEs' profit reporting behavior to changes in tax policy through various channels such as debt shifting (Desai et al.; 2004; Huizinga et al.; 2008), transfer pricing (Clausing; 2003; Cristea and Nguyen; 2016; Davies et al.; 2018) and patent location (Bilicka et al.; 2023; Desai et al.; 2006; Dischinger and Riedel; 2011; Griffith et al.; 2014).

Our approach is agnostic about the particular form that the tax avoidance network takes. As an example, consider the case of a business with ownership of intellectual property (IP) that chooses to locate the IP in a tax haven subsidiary (and possibly funding that subsidiary under a "cost contribution arrangement" to pay the costs of research and development undertaken elsewhere). The revenue generated from the IP will be treated as arising in the tax haven, but the cost of setting up such an arrangement may be very high.[5] With that arrangement in place, the costs of shifting profit to the tax haven in any subsequent period are significantly lower than they would be without the arrangement. These

---

[4] Other approaches include contributions by Desai et al. (2006); Dharmapala and Riedel (2013); Dischinger et al. (2014); Dischinger and Riedel (2011); Egger et al. (2010); Grubert and Slemrod (1998); Gumpert et al. (2016); Langenmayr and Liu (2023); Slemrod and Wilson (2009); Tørsløv et al. (2022).

[5] See, for example, the descriptions of Apple's and Microsoft's cost sharing arrangements in Chamberlain and Curtis (2021) and Curtis and Avi-Yonah (2023).

costs can include direct payments to in-house staff and for professional advice, but may also include the expected value of potential penalties or reputation costs. This kind of tax planning would be within the range of activities represented by our model.

Differently from the earlier literature, we split the profit-shifting cost into a 'fixed tax avoidance investment cost' and a variable cost of profit shifting. This means that tax reforms induce both an extensive- and an intensive-margin effect on profit shifting. For example, a rise in the effective tax rate of the tax haven country reduces the benefit of profit shifting. However, a company that shifts all of its profit to the tax haven may continue to do so despite the higher tax haven tax rate if the benefits accruing from shifting all profit from a high-tax country continue to exceed its costs. Consequently, the overall macro-level effect of tax reforms will aggregate the response of the companies that, prior to the reform, shift some profit and the companies that shift all their profit – at the intensive and extensive margins of profit-shifting. Micro-level studies that focus only on intensive margins are likely to understate the impact of tax reforms on the proportion of true profit shifted. A key contribution of our paper, therefore, is to highlight the importance of unobserved heterogeneity across different types of firms with access to varying degrees of tax avoidance capabilities in estimating the magnitudes of firm responses to tax policies.

We estimate our model using corporate tax returns data from the UK. We distinguish between short-run and long-run semi-elasticities of the tax base with respect to the tax rates in the home country or the tax haven. In the short run, firms can change the proportion of profit that they shift to the tax haven, but they cannot change their overall tax avoidance capability or their investment in productive capital. In the long run, we relax this constraint and allow firms to adjust all three choices, namely, their investment in productive capital, the tax avoidance asset and the share of profit shifted out of countries with real activity to tax havens.

In aggregate, in the short run, the declared pre-tax profit in the home country declines by 1.5 percent in response to a one percentage point rise in the home country (from 30% to 31%). In the long-run, companies reduce their declared pre-tax profit in the home country by 4.7 percent. Our semi-elasticity estimates capture the overall aggregate responses of tax bases to changes in tax rates. The difference between these short-run and long-run responses constitutes one of the factors that drive a discrepancy between micro- and macro-level estimates in the existing empirical literature. Further, in this paper, we highlight the role of movements into and out of positive taxable income reporting in high-tax jurisdictions (extensive margin responses). To quantify the importance of extensive margin responses in explaining the discrepancy between previous micro- and macro-level studies,

3

we compare our estimates with a version of the semi-elasticity calculation that ignores the responses of firms at the extensive margin. We show that, based on our model, these semi-elasticity estimates are around 15 percent lower in the short-run and 39 percent lower in the long-run.

We also show that investing in tax avoidance capability may reduce the cost of capital for productive assets. This highlights a key relationship between profit shifting and investment, in which reforms that increase the costs of profit shifting for MNEs may reduce their investment in productive assets. This occurs directly in high-tax jurisdictions and has an indirect effect on investment in lower-tax jurisdictions where an affected MNE also has real activity. We estimate that a percentage point rise in the tax rate of a high-tax country results in a reduction in capital accumulation in that country of 0.4 percent in the presence of profit-shifting.[6] If profit shifting did not take place, we find a much larger response; capital accumulation drops by 1.9 percent, as the dampening effect of profit-shifting on cost of capital is removed. A similar type of trade-off has also recently been highlighted in a macro model by Dyrda et al. (2022), who show that policies that reduce profit shifting are also likely to reduce output. Note that this differs from the recent work by Chodorow-Reich et al. (2024) and Kennedy et al. (2024) who both examine the impact of TCJA on global investment and do not take profit shifting into account. Finally, the profit shifting and investment trade-off that our model quantifies is also in line with the recent empirical literature that shows that anti-tax avoidance regulations have real effects on investment and employment (Bilicka et al.; 2022; Suárez Serrato; 2018).

We then move on to quantify the fixed and variable costs of profit shifting. We find that for the average MNE that shifts profit, the fixed cost of setting up a profit-shifting strategy is around 3.3% of the total true tax base, while variable costs amount to twice as much, at 6.9% of the total true tax base. The literature so far has provided few estimates of the size of these costs. An exception is the recent work by McClure (2023), who estimates the costs of avoidance, to be a little lower, at 6.4% of pre-tax income.

We proceed to test the empirical implications of our model using UK tax return data. Our quasi-experimental variation arises from the 2002 Italian Controlled Foreign Company (CFC) reform. These CFC rules stipulated that the income of low-tax subsidiaries in black-listed tax havens should be subject to Italian tax. This reduced or removed incentives to shift profit to such countries (Clifford; 2019). In the context of our model, the CFC rule is akin to an increase in the tax haven's tax rate for affected UK subsidiaries of Italian-

---

[6]Estimate for a country with a baseline 30% tax rate.

headquartered MNEs. We compare a treatment group composed of all UK subsidiaries of Italian-headquartered MNEs with a control group composed of UK subsidiaries of non-UK, non-Italian headquartered MNEs. We expect that the rise in the tax haven tax rate for affected subsidiaries would induce an increase in profit reported in other (i.e. non-haven) subsidiaries of an MNE, including the UK. We estimate a large and significant overall effect of the reform on taxable profits of the treatment group. Consistent with the predictions of our model, for the treatment group we find a significant reduction in the probability of reporting zero taxable income, and a significant increase in the average reported taxable profit. We find no average overall effect of the CFC reform on investment. However, we find suggestive evidence that firms that change their taxable profit reporting behavior from a non-taxable position to a taxable position reduce their investment.

Finally, we carry out counterfactual policy experiments. We consider two types of policy experiments; a reduction in the high-tax country tax rate and an increase in the tax rate in tax haven countries, both marginally and using the 15% tax rate of the Global Minimum Tax. We find that the share of profit shifted from a high-tax country to a tax haven declines in response to a reduction in the high-tax rate country tax rate from 30%, as a lower home country tax rate reduces benefits from costly profit shifting. This change in tax rate also reduces the costs of profit shifting and the fraction of firms that shift *all* of their profits out of the high-tax country. The reduction in tax rate also reduces the cost of capital and increases capital accumulation in high-tax country. In turn, an increase in the tax haven tax rate from zero also reduces the share of profits shifted from high tax country and the costs of shifting profits. However, this comes at the cost of lower capital accumulation. We evaluate the welfare implications of these two types of policies using the marginal value of public funds (MVPF) framework (see Hendren and Sprung-Keyser (2020, 2022) and show that the MVPF of a tax rate reduction in high tax country is higher than that of the tax haven tax rate increase.

We then demonstrate the impact of a global minimum tax at varying threshold rates. A 15% minimum tax threshold implies a large decline in profit-shifting for MNEs, which creates substantial efficiency gains that arise from the reduction in wasteful investment in profit-shifting capability. Offset against that is the cost of reduced investment in productive assets. Overall, for countries with sufficiently high baseline tax rates, we find that the optimal global minimum tax rate (conditional on other tax rates) can be well above zero.

Our paper contributes to two broad strands of literature. The first strand relates to the debate over the magnitudes of profit shifting and its costs. The micro literature that uses firm-level data to estimate elasticities and semi-elasticities of tax base with respect to tax

rate changes estimates much smaller elasticities than the macro literature that uses aggregate data. Our contribution to this literature is in building a model that allows for both the extensive and intensive margin responses to tax policies and which reconciles these differences between micro- and macro-level elasticity estimates. We highlight the role of extensive margin profit-shifting responses in reconciling these elasticity differences. In our model, when we do not allow for extensive margin responses, we estimate the absolute value of the semi-elasticity to be 1.3. This is very close to the consensus from the literature that has been summarized in meta studies by Beer et al. (2019); Heckemeyer and Overesch (2017). For example, in Beer et al. (2019) the average micro-level semi-elasticity in the literature is 1.4 (in absolute value), while in Heckemeyer and Overesch (2017) it is 1.1. The literature proposes much larger macro-level semi-elasticity estimates that range (in absolute value) from an average of 2.9 in Clausing (2016) to an average of 5.2 in Hines and Rice (1994). In a recent paper, Ferrari et al. (2022) build a model of multi-national production and using macro-level data obtain a semi-elasticity estimate of 3.6. Our semi-elasticity estimate of 4.7 is within the range of these macro-level estimates.

Second, profit shifting has been shown to have real consequences on firm operations that feed through the economy to estimates of GDP and productivity. At the micro-level, there is growing evidence that anti-tax avoidance regulations reduce not only the extent of tax avoidance, as intended, but also curb real business operations of MNEs. Suárez Serrato (2018) shows effects on investment and employment and consequences on local labor markets in the US, while Bilicka et al. (2022) show effects for real business operations in the UK and in foreign countries of MNE operations. Bustos et al. (2022) complement this evidence by emphasizing the role that local tax advisors play in enabling profit shifting, while Bilicka and Scur (2024) highlight the role of local organizational capacity. On the macro level, Guvenen et al. (2022) find that profit-shifting reduces US GDP and productivity estimates in the official statistics; Coppola et al. (2021) find that offshore issuance reduces the scale of portfolio investment from developed countries to emerging market companies; and Dyrda et al. (2022) model the consequences of profit shifting for output and growth. Consequently, these papers suggest that profit shifting by MNEs affects output, measurement of GDP, production and international capital flows. Our paper highlights that profit-shifting may reduce the cost of investment in productive assets and our model allows us to quantify the trade-off between investment and tax revenue more systematically.

Our results provide timely evidence to evaluate the possible impact of ongoing international tax reform efforts. In recent years, there has been growing concern over the ability of multinational corporations to shift profit from high-tax jurisdictions to tax havens

in order to reduce their aggregate tax liabilities. Especially following the global financial crisis in 2008-9, governments seeking additional tax revenue have sought to combat such profit-shifting. This led to the OECD/G20 Base Erosion and Profit Shifting (BEPS) project in 2013-5, with sweeping measures aimed at protecting the tax base in high-tax countries. More recently, in 2021, over 140 members of the OECD's Inclusive Framework have agreed the most far-reaching reform to the international taxation of profit in a century: the introduction of a Global Minimum Tax (GMT), set at 15% of "excess profit".[7] There is a growing theoretical literature that analyzes the implications of global minimum taxes for welfare and revenues of high- and low-tax countries (Hebous and Keen; 2021; Hines Jr; 2022; Janeba and Schjelderup; 2022; Johannesen; 2022) and policy simulation exercises (Bares et al.; 2023; Grubert and Altshuler; 2013; Hanappi and Cabral; 2020). While our paper also allows us to examine the implications of GMT on global welfare, tax revenue and investment, the model that we build has a much broader application to how we model the behavior of multinational firms and allows us to estimate the responses to tax policies for firms that pay little or no corporate tax. Our model is easily applicable to a variety of country contexts and reforms.

The rest of the paper is structured as follows. In Section 2, we develop the theoretical framework. In Section 3, we describe the data and empirical evidence. In Section 4, we present our structural approach and in Section 5, we provide reduced-form evidence supporting the structural estimates. In Section 6, we carry out counterfactual policy experiments and we conclude in Section 7.

2 Conceptual framework

2.1 A model of capital accumulation with profit-shifting

We model the behavior of a multinational enterprise (MNE) in a single period. The MNE has subsidiaries in N jurisdictions that have different tax rates on profit. The novelty of the model lies in the MNE investing in an asset which we call the "tax avoidance capability", or "tax avoidance asset", Y. The cost of accumulating the tax avoidance asset varies across MNEs, reflecting the ease with which different businesses can access the legal framework and accumulate know-how to shift profits to tax havens.

Each subsidiary has access to the global tax avoidance capability of the MNE, which

---

[7]See OECD agreements in July 2021 and October 2021.

makes Y a public good within the MNE. Each subsidiary that is not located in a pure tax haven also invests in productive capital, K. In this structure, two otherwise identical businesses in a jurisdiction may behave differently with respect to profit shifting: one may be part of a large multinational which has already invested substantially in the tax avoidance asset, while another may be a subsidiary of an MNE with very little investment in the tax avoidance asset (or a firm with solely domestic activities). The simplest possible version of the MNE in this model operates in three jurisdictions; one with a high tax rate, one with a lower tax rate, and one that acts as a pure tax haven only used for shifting profit into it.

The decision timing for the MNE is as follows:

1. At the beginning of the period, each government j announces its tax rate, τ_j, and any anti-avoidance measures. The anti-avoidance measures determine the extent to which the home government may levy tax on income declared in a subsidiary located in a tax haven X; the key features of the haven X are that it has a very low tax rate, labeled τ_X, and that there is no investment in productive capital. Each MNE i invests in the tax avoidance asset Y_i that serves its subsidiaries globally. The MNE-specific cost to the multinational of purchasing a unit of the tax avoidance asset is denoted p_i. This cost reflects both characteristics of the MNE's activities - for example, its reliance on hard-to-value intangible assets - and also the anti-avoidance measures in place in the countries in which it operates.[8]

2. Still at the beginning of the period, each MNE i receives a draw in a distribution of jurisdiction-specific productivity, ε_{ij}, which introduces heterogeneity in production both across MNEs and across countries. With knowledge of τ_j and ε_{ij} for all possible countries of operation j, as well as τ_X and p_i, each MNE i chooses investment in productive capital K_{ij} in each non-haven country j.

3. At the end of the period, each subsidiary generates output of $F(K_{ij})$, with $F'(K_{ij}) > 0$ and $F''(K_{ij}) < 0$, and sells the remaining productive capital for $(1 - \delta)K_{ij}$, where δ is the depreciation rate of productive capital.

4. Also at the end of the period, the multinational: (i) incurs an exogenous cost of operation Π_i, observed only after it has taken the decisions on capital K_{ij} and tax avoidance asset Y_i – this introduces a third element of heterogeneity, which allows a subsidiary

---

[8]It is possible to add a separate fixed entry cost of investing in the tax avoidance asset. In our simulations and empirical approach, we do not include such a fixed cost and in simulations, we account for the behavior of firms without the tax avoidance asset separately.

to make a loss ex-post; (ii) based on prior choices and on Π_{ij} in each subsidiary, the MNE chooses the proportion α_{ij} of the "true" tax base B_{ij} to shift from each subsidiary j to a tax haven. We assume that it is not possible to shift more than 100% of the tax base.

We assume that the variable costs of shifting profit out of jurisdiction i to a tax haven (conditional on $Y_i > 0$), are:

$$C_{ij} = c\left(\alpha_{ij}, Y_i, B_{ij}\right) B_{ij} \tag{1}$$

with $c_\alpha > 0$, $c_Y < 0$, and $c_B > 0$. That is, variable profit-shifting costs rise with the size of the tax base, B_{ij}, and the proportion of that base shifted, α_{ij}, but fall with higher investment in the tax avoidance asset, Y_i.

The true tax base in country j for MNE i has tax depreciation equal to true economic depreciation at rate δ and there is no relief for any financing costs:

$$B_{ij} = F(K_{ij}) - \delta K_{ij} - \Pi_{ij} \tag{2}$$

There is no tax deduction for expenditure on the tax avoidance asset, $p_i Y_i$. Since this is determined at the level of the MNE as a whole, we assume that it is undertaken by the parent company. Any deduction that may be available in the country of the parent can be thought of as affecting the MNE-specific price, p_i.

At the end of the period, a proportion α_{ij} of the tax base is shifted to the tax haven where it is liable to tax at rate τ_X. The remaining profit is taxed in country j at rate τ_j. The overall tax liability for each MNE in each jurisdiction is therefore:

$$T_{ij} = \widehat{\tau_{ij}} B_{ij} = \left[\tau_j\left(1 - \alpha_{ij}\right) + \alpha_{ij}\tau_X\right] B_{ij} \tag{3}$$

where $\widehat{\tau_{ij}}$ can be thought of as an "effective statutory rate" on profit generated by MNE i in the subsidiary in j.

The MNE chooses Y_i, and for each subsidiary j, it chooses K_{ij} and α_{ij} to maximize its beginning-of-period value. The following value maximization problem summarizes the

MNE's decision-making process at the beginning of the period:

$$\max_{Y_i, K_{ij} \forall j} V_i = -p_i Y_i - \sum_{j=1}^{N} K_{ij} + \beta \sum_{j=1}^{N} \left[F(K_{ij}) - \mathbb{E}[\Pi_{ij} - T_{ij} - c\left(\alpha_{ij}, Y_i, B_{ij}\right) B_{ij}] + (1 - \delta) K_{ij} \right]$$

(4)

subject to the constraint:

$$0 \leq \alpha_{ij} \leq 1$$

where $\beta = 1/(1 + r)$ is the discount factor. This formulation of the maximand reflects the timing considerations that we have described earlier in this section, namely the fact that the MNE chooses K_{ij} and Y_i before it observes the cost of operation, Π_{ij}. These decisions are therefore made based on the expected cost of operation Π_{ij}.

Conditional on some positive investment in $Y_i > 0$, the first order conditions for Y_i, K_{ij} and α_{ij} are as follows:

$$Y_i: \qquad p_i = -\beta \sum_{j=1}^{N} \mathbb{E}[c_Y(\alpha_{ij}, Y_i, B_{ij}) B_{ij}] \tag{5}$$

$$K_{ij}: \qquad F_K(K_{ij}) = \frac{r}{1 - \mathbb{E}[\widehat{\tau_{ij}} + c_B(\alpha_{ij}, Y_i, B_{ij}) B_{ij} + c(\alpha_{ij}, Y_i, B_{ij})]} + \delta \tag{6}$$

$$\alpha_{ij}: \qquad \{\tau_j - \tau_X - c_\alpha(\alpha_{ij}, Y_i, B_{ij})\} B_{ij} + \eta_{ij} = 0 \tag{7}$$

$$\eta_{ij}(\alpha_{ij} - 1) = 0$$

Equation 5 shows that the marginal cost of an additional unit of investment in the tax avoidance capability Y_i, p_i, equals the marginal benefit. The marginal benefit is the sum of the expected marginal reductions in variable profit shifting costs, aggregating over all the subsidiaries.

The condition for the optimal stock of K_{ij} in each subsidiary (Equation 6) resembles the standard, single-country expression for the marginal product of capital being equal to the cost of capital as in Hall and Jorgenson (1967), with adjustments for profit shifting. First, the "effective statutory tax rate", $\widehat{\tau_{ij}}$, replaces the statutory tax rate in country j. Higher profit shifting therefore reduces the cost of capital, and raises K. Second, the expression for cost of capital in Equation (6) depends on K_{ij} through B_{ij}. As K_{ij}, and consequently B_{ij}, rise, the costs of profit shifting increase, which also raises the cost of capital.

10

In Equation 7, for a subsidiary j that shifts less than 100% of its profit, then $\eta_{ij} = 0$, and α_{ij} is determined at the point where the marginal benefit of shifting is equal to the difference between the tax rate differential and the marginal cost:

$$\tau_j - \tau_X = c_\alpha(\alpha_{ij}, Y_i, B_{ij}) \tag{8}$$

This is a conventional expression for the choice of what proportion of profit to shift. Unlike the conventional expression, our model also allows the subsidiary to shift all its profits. Specifically, if the benefits of shifting outweigh the costs for the values of $\alpha_{ij} < 1$, then α_{ij} will be driven to the corner solution of full shifting, with $\alpha_{ij} = 1$.

Welfare. We evaluate the welfare implications of actual and possible reforms to the tax system. To do so, we use a simple measure of welfare, equal to the sum of the net present value of firm cash flows, V_i, and tax revenue, which is appropriate when tax revenue is returned to the private sector as a lump sum. We evaluate welfare from a global perspective, as additional assumptions would be needed to attribute either private income or tax revenue to individual countries. We calculate global welfare numerically, computing for example for each value of τ_X the private income (V_i) and tax revenue (T_i) associated with each firm, and aggregating across firms and countries.

To compare the impact of reform options on global welfare, we use the marginal value of public funds (MVPF) framework. We compare the MVPF of a small tax reduction in the high-tax country with a small tax rise in the tax haven. We calculate changes in private income and tax revenues in response to a percentage point change in tax rates (following the approach of Hendren and Sprung-Keyser (2020, 2022)). The MVPF is the ratio of change in welfare of all firms to the change in cost to the government, which captures both the mechanical and behavioral responses to tax policies. We numerically calculate the following:

$$MVPF = \frac{\sum_{i=1} \Delta V_i}{\sum_{i=1} \Delta T_i} \tag{9}$$

In Appendix C, we also set out an approach that provides a closed-form solution to calculate changes in welfare and MVPF and shows that, similar to Risch (2024), the elasticity of taxable income is no longer the sufficient statistic for welfare analysis. Instead, we require (1) the elasticity of the true tax base with respect to the effective statutory tax rate

11

and (2) the elasticity of the proportion of the tax base shifted to the tax haven with respect to the difference in statutory tax rates between the high tax country and a tax haven.

Functional forms. We follow the conventional functional forms used in the existing literature for $c\,(\alpha_{ij}, Y_i, B_{ij})$ and $F(K_{ij})$, namely, a convex profit-shifting cost function and a Cobb-Douglas production function. The main distinction in our framework is the role of Y_i in the cost function for profit shifting. Specifically, we assume a functional form for the cost of profit-shifting in Equation (10) which exhibits convexity in α_{ij}, along the lines of conventional models (Dharmapala; 2014; Hines and Rice; 1994; Riedel; 2018):

$$c\,(\alpha_{ij}, Y_i, B_{ij}) = \frac{\gamma}{2} \left(\frac{B_{ij}}{Y_i} \right)^m \alpha_{ij}^2 \tag{10}$$

We assume that costs increase with the size of profit available to be shifted from the subsidiary in j, B_{ij}, and fall with the size of the multinational's tax avoidance asset, Y_i. We include the ratio of these two factors, and assume costs to be concave in this ratio, depending on the parameter m, where $m \in (0,1)$.

We also use a simple functional form for $F(K_{ij})$:

$$F(K_{ij}) = \theta_{ij} K_{ij}{}^A \tag{11}$$

where productivity draw $\theta_{ij} = \theta \exp\left(\varepsilon_{ij}\right)$ may vary amongst subsidiaries following the process $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$. We assume that the cost of operation has the distribution $\Pi_{ij} \sim \Gamma(1, \beta_\Pi)$. Chen et al. (2022) derive a dynamic version of a production function with labor optimized out, in a similar spirit to our approach in this paper.

2.2   Implications of the model

We solve the model numerically for the three choices of the firm: investment in the tax avoidance asset for the MNE, Y_i, the optimal productive capital, K_{ij}, and the extent of profit shifting – the share of tax base shifted to the haven, α_{ij} – for each subsidiary. In Figure 1, we illustrate the implications of the model for K_{ij} and α_{ij} and show the heterogeneity across MNEs due to variation in the price of the tax avoidance asset, p_i. The analysis here is based on an MNE with three subsidiaries; one in a high-tax country, H, with a tax rate of 30%, one in a low-tax country, L, with a tax rate of 12.5% and where the MNE also has real activity, and the third subsidiary without any real operations located in a tax haven, X, with a zero

12

tax rate. An MNE with subsidiaries in France (high-tax), Ireland (low-tax) and Bermuda (tax haven) is an example of such a setup. Note that the structural parameters used to generate Figure 1 are those that we estimate in Section 4. As we vary the ease of access (or unit cost) of tax avoidance capability p_i in the horizontal axis of Figure 1, we hold the productivity draw, θ_{ij}, and the cost of operation, Π_{ij}, constant at their mid-point values across all subsidiaries in high-tax and low-tax countries.

In Figure 1a, the black solid line shows the share of profit that the MNE shifts out of H and the blue dashed line shows the share of profit shifted out of L. For MNEs that have very low p_i, i.e., very easy access to tax avoidance, it is optimal to shift the entire tax base in both countries to the tax haven. As p_i increases and the cost of shifting rises, the "full-shifting" behavior becomes less prevalent, first in L and then, as p_i increases further, also in H. This generates the two kinks in the share of shifted profits. Except at very low values of p_i, a larger share of profits is shifted out of H than out of L, reflecting the greater benefit of shifting from the higher-tax country.

In Figure 1b we show capital accumulation in H and L, again with the solid black line and the blue dashed line, respectively. For very small values of p_i, as the MNE shifts all profit out of both countries, capital accumulation in the two countries is the same. However, as p_i increases capital accumulation trajectories diverge, with capital accumulation falling in both countries. Note that capital accumulation falls in H with a rise in p_i even when profit is fully shifted to X so that no tax is paid. This is because profit shifting costs rise, increasing the cost of capital. Capital accumulation also falls in L, but by a smaller amount, reflecting the fact that the sum of taxation and profit shifting costs are lower in L than in H, even for MNEs that face a higher unit cost of tax avoidance.

3 Data and descriptive patterns in profit shifting and investment

We use detailed administrative tax returns data from the HM Revenue and Customs (HMRC; the UK tax authority) matched with financial accounts information and ownership links provided by the FAME dataset to test the predictions of the model. Our tax returns dataset starts in the fiscal year 2000 and comprises all items that are submitted on the corporation tax return form (CT600 form). The unit of observation is an unconsolidated statement of a company in each of the years. Each subsidiary of a company operating in the UK files a separate tax return. This data is matched with financial accounts data that contains in-

formation on firm assets, employment, and other balance sheet items (Bureau van Dijk's FAME dataset). The ownership data from FAME allows us to identify the global ultimate parent companies that own UK subsidiaries and separate firms into domestic and multinational, as well as differentiate by country where the ultimate owner is headquartered. We focus on UK subsidiaries that belong to multinational companies headquartered in foreign countries and that report positive turnover. We use these firms in both our structural model and reduced form analysis because, unlike domestic firms, they have the ability to shift profits to foreign jurisdictions, in line with the MNE behavior outlined in our conceptual framework.[9] Our sample comprises 8,922 subsidiaries, of which 248 are headquartered in Italy and are affected by the Italian reform we use for identification in Section 5 and 8,674 are headquartered across all other countries in the world. We use information for these subsidiaries for years 2000 - 2005 in the reduced form estimations and we use the subsidiaries of all non-Italian MNEs in 2001 to identify the structural parameters in our model. The reason we choose this particular group of firms in 2001 to identify structural parameters is to avoid any influence of the 2002 Italian reform in structural estimates. In what follows, we use these data to present three pieces of descriptive evidence that highlight the margins our model focuses on.

In Table 1, we show the mean and median values for some of the key variables in our balanced panel of MNEs operating in the UK that we use in our structural estimation. 42% of MNEs report zero taxable income. These zero-profit-reporters include those with no tax liability for other reasons than profit shifting, such as genuine operational losses, investment allowances and other reliefs available in the tax system that reduce taxable income for firms. However, many subsidiaries of MNE groups report positive operating profit while reporting zero taxable income, providing the first suggestive descriptive evidence on the importance of the extensive margin of profit shifting that we model in this paper. We depict the prevalence of this behavior in Figure 2, where we plot the distribution of the difference between operating profit and taxable profit scaled by operating profit – the 'operating profit deviation ratio' – for all subsidiaries of MNEs that report a positive operating profit in their financial statements. A striking degree of bunching occurs at the point at which the operating profit deviation ratio is equal to one, indicating that 25% of all subsidiaries that report a positive operating profit also report zero taxable profit. A much smaller proportion of subsidiaries bunch at the zero deviation ratio, which indicates

---

[9]Note that following Bilicka (2019), we do not consider multinational firms that are headquartered in the UK, domestic MNEs, because a majority of them do not report unconsolidated balance sheet information in FAME.

no difference between operating and taxable profit.

In Figure 3, we plot asset accumulation for firms at different taxable income reporting percentiles, grouping firms into 5% bins. Because of large bunching of observations at zero taxable profit, a large share of observations are clustered in the first percentile value, or the leftmost bin in Figure 3. Asset size rises monotonically with taxable income for companies reporting positive taxable income. In turn, companies that report zero taxable profit have a much larger average asset size than the majority of those reporting a positive taxable profit. This indicates a negative relationship between firm size and taxable income reporting. Our model predicts that aggressive profit-shifting MNEs tend to face lower costs of capital, allowing them to grow larger. We explore this relationship in the next section.

4   Estimating profit shifting and investment behavior

We analyze firms' responses to tax policy changes in two steps. First, we use moments from the distribution of key variables, such as taxable profit and capital, from the UK corporate tax returns to match in a simulated method of moments estimation procedure and estimate the key structural parameters of our model. We use these estimates to recover elasticities of tax base and capital with respect to the tax rate. We also calculate the magnitudes of fixed and variable profit-shifting costs as a share of the pre-shifting (true) tax base.

Second, in Section 5, we present results from the reduced-form difference-in-differences regressions that demonstrate MNEs' responses to a reform aimed at curbing tax avoidance for a particular group of firms.

4.1   Estimating key model parameters

In this section, we present structural estimates for the key parameters of our model that microfounds the profit-shifting and investment behavior of MNEs. We use an indirect inference approach (Gallant and Tauchen; 1996; Gourieroux et al.; 1993) and a method of simulated moments (MSM) procedure, in which we simulate firms over the distributions of unobserved relative unit price of the tax avoidance capability, p_i, productivity draws with $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$, and cost of operation that has an exponential distribution with scale parameter β_Π. We minimize the MSM criterion function, which takes the form:

$$L(\Theta) = h(\Theta)^{'} W_N h(\Theta) \tag{12}$$

where Θ is the vector of structural parameters of interest. The two key parameters related to profit-shifting behavior are $\bar{p}$, which is the upper bound of the uniformly distributed relative unit price of the tax avoidance capability, and the convex profit-shifting cost parameter γ. We also estimate the production function parameters within the MSM procedure. These production function parameters are θ, the total factor productivity parameter, a, the output elasticity with respect to K, and σ, the standard deviation of the productivity draw as well as the scale parameter of the linear cost term β_{Π}.

$h(\Theta)$ is the vector of M moment conditions constructed as the difference between simulated moments computed over 100,000 simulated firms following Michaelides and Ng (2000) and empirical moments obtained from the population of subsidiaries of non-Italian foreign multinational firms operating in the UK. As the weight matrix, we use the diagonal elements of the inverse variance-covariance matrix of empirical moments. To identify our key parameters, we use information on taxable income and assets of MNE subsidiaries operating in the UK. In Table 2, we show the correspondence between the specific moments and the parameters that each moment helps to identify in the estimation. Crucially, the share of zero taxable income reporters is a key moment that identifies our profit-shifting 'fixed cost parameter', $\bar{p}$, and the quartiles of taxable profit (in log) are key in identifying the 'variable cost parameter', γ.

The policy environment consists of a high-tax location, a low-tax location and a tax haven. All real investment takes place in the high-tax and the low-tax countries, but profit is then shifted to the tax haven. The high-tax location in our case is the United Kingdom with a 30% statutory corporate income tax rate over the relevant period. We run our simulations using a low-tax location with 12.5% rate, but the availability of this alternative investment location in the model does not have a material impact on our estimates. The tax haven applies a tax rate of zero percent. We assume that the cost of setting up a network of subsidiaries for tax avoidance purposes is absorbed into the cost of the tax avoidance capability captured at the MNE level by p_i, for which the distribution over MNEs is uniform between zero and $\bar{p}$.

In Table 3, we present the estimates from our MSM procedure. We calculate the standard errors of these estimates using the Delta method. All estimates are highly significant and responsive to changes in moments that we use to identify them. We estimate that the elasticity a of output with respect to productive capital K is 0.553, the standard deviation of the productivity shock is 0.926 and the total factor productivity θ is 131.005. We use all our structural estimates in Section 6, where we demonstrate counterfactual policy simulations. Finally, in Figure 4, we present the fit of the model against 21 moments from the

data listed in Table 2 and show that the estimated moments closely match their simulated counterparts.

4.2   Elasticity of taxable income reporting in high-tax jurisdictions

We present implied tax base semi-elasticity estimates based on our structural parameters in Table 4. We present four estimates of semi-elasticities indicating the percent change in the declared tax base in the home country with respect to a one percentage point rise in the home country tax rate, distinguishing across two dimensions: First, we present estimates that account for the extensive margin profit-shifting responses as well as estimates that our model would generate if we did not account for the extensive margin. The latter are equivalent to many of the micro-level estimates from the literature, obtained from a regression analysis that ignores zero-taxable income reporters. Second, we present long-run semi-elasticities which correspond to the case where capital and tax avoidance capability adjust fully in equilibrium, as well as short-run semi-elasticities which capture the case before the adjustment in capital and tax avoidance capability, just after the MNE incurs the cost of operation Π_i. In all these estimates, we present the weighted average of the response of subsidiaries of MNEs and other firms, where we use the Office for National Statistics (ONS) gross value added number for foreign firms operating in the UK as the weight for MNE subsidiaries. In our case, only the subsidiaries of MNEs shift profit to the tax haven, but in the long run all firms adjust their capital. [10]

Table 4 highlights that there are large differences both between long-run and short-run estimates, and between estimates that incorporate zero-taxable income reporters and the estimates that would be equivalent to those obtained from a regression analysis that ignored the zero-taxable income reporters. We find that, in the short-run, the aggregate tax base drops by 1.5 percent in response to a percentage point rise in the high-tax country tax rate. This semi-elasticity estimate is 15 percent higher in comparison to the 'micro-level' estimate that does not take into account zero taxable income reporting MNEs. In the long-run, the tax base drops by 4.7 percent in response to a percentage point rise in the high-tax country tax rate, and this aggregate semi-elasticity estimate is 39% larger in absolute value than the naive micro-level estimate that ignores zero taxable income reporters. This difference highlights the importance of taking the extensive margin responses into account when estimating the effects of tax rate changes on tax bases.

---

[10]Using information from ONS, we use the information that foreign MNEs in the UK comprise 28.5% of gross value added.

In Table 5, we demonstrate that the level of the home country tax rate matters in estimating the response of the declared tax base to changes in the home country tax rate. To do that, in our structural model we change the high tax country rate from 30% to 20% or to 40% and show the resulting semi-elasticity estimates. The average declared tax base response increases as the high-tax country tax rate increases. For example, the implied long-run semi-elasticity for a country with a 20% tax rate is 2.3 in absolute value. As expected, the long-run responses are larger in magnitude than short-run responses, and the overall responses exceed the corresponding micro-level estimates that ignore zero taxable income reporters.[11]

In Figure 5, we demonstrate the heterogeneity in the declared tax base responses in high-tax-rate countries for firms that face different unit prices of tax avoidance capability and have different productivity draws. The left-hand side panel shows the size of the declared tax base in certain states, while the right-hand side panel shows the changes in the declared tax base in response to the one percentage point rise in the high-tax country tax rate. In Figure 5, unbroken lines indicate the status quo before any tax rate change. In black, we plot responses for firms with the values for productivity draws at the mid-point level, in blue for firms with higher productivity than the mid-point, while in red for firms with lower productivity than mid-point. Dotted lines show short-run responses that keep K and Y constant, only allowing the share of profit α to change, and dashed lines allow all three choices to adjust.

Both the short-run and the long-run figures show that for sufficiently low unit prices of the tax avoidance asset, firms shift all their profit. At a value of p_i just over 0.5, firms begin to reduce their profit shifting; the effect of the rise in p_i is strongest for high-productivity firms and weakest for low-productivity firms. As expected, there is also a larger impact in the long-run, compared to the short-run. For values of p_i between 0.504 and 0.527, the long-run response is at the extensive margin: within that range firms respond to a rise in the tax rate by starting to pay zero tax; this is also reflected in the right-hand side panel, where a tax base reaction of -1 shows a firm moving from a positive declared tax base to a zero declared tax base. For higher values of p_i, the magnitude of the short-run response in the declared tax base shrinks rapidly to stabilize to a reduction below 10%. In the long-run, the gradient of the tax base reaction is smaller, indicating a larger tax base reaction for the same unit price of the tax avoidance capability relative to the short-run.

---

[11]This is different than the heterogeneity explored in Dowd et al. (2017), who look at different non-home tax rates and show that the elasticity for tax haven subsidiaries in higher than that for high tax countries.

Comparison with related literature. Two recent papers use meta-regressions to analyze estimates of the semi-elasticity of reported profit with respect to the tax rate differential across countries and provide very comprehensive overviews of the literature and discussions of all potential determinants of the differences between estimates. Heckemeyer and Overesch (2017) find the average semi-elasticity to be -0.8., while Beer et al. (2019) find that this semi-elasticity has increased (in absolute value) in recent years to -1.5. These papers imply that a 10 percentage point increase in the tax rate differential between two countries would reduce the pre-tax income reported by the subsidiary in the low-tax country by 8-15 per cent. However, these aggregates mask a wide range of existing estimates and heterogeneity. As one of the main objectives of our paper is to build a model which reconciles the differences between existing micro- and macro-estimates of tax base, we focus on those differences specifically.

Both meta-studies by Beer et al. (2019); Heckemeyer and Overesch (2017) emphasize in their regressions differences between macro semi-elasticity estimates that use aggregate data and micro semi-elasticity estimates that use detailed firm-level data. Beer et al. (2019) estimate that the implied semi-elasticity for studies that use aggregate data is 48% higher than for studies that do not use aggregate data, while Heckemeyer and Overesch (2017) estimate this difference to be 73%.[12] Our semi-elasticity estimates from Table 4 show that the differences between estimates that ignore the zero-profit reporters and the ones that include these firms are 39% (and as high as 47% if we consider the response for the subsidiaries of MNEs only). This would suggest that the extensive margin response that we have in our model is able to explain most of the difference between the macro- and micro-elasticity estimates.

4.3 Elasticity of capital in high-tax jurisdictions

We next investigate how the elasticity of tangible capital with respect to the tax rate depends on profit shifting. As a higher tax rate induces both a rise in profit shifting and a fall in capital accumulation due to the higher cost of capital, the magnitude of the tangible capital response depends on the extent to which firms engage in profit shifting.

---

[12]Beer et al. (2019) estimate that the implied semi-elasticity for studies that use aggregate data is 2.29. Taking their estimates of the coefficient on aggregate data, which is 1.1, we obtain 1.1/2.29=48%. Heckemeyer and Overesch (2017) estimate that the micro-level studies have semi-elasticity which is 1.65 lower in absolute terms than the average of the aggregate ones. Taking the average aggregate semi-elasticity of 2.77 that Heckemeyer and Overesch (2017) calculate from the papers that use tax returns data from the IRS, suggests a 1.65/2.27=73% difference.

In Table 6, we summarize the semi-elasticities of K_{ij} with respect to the tax rate τ_j. In the first row, we show semi-elasticities for our sample of subsidiaries of foreign-owned MNEs, labeled "Shifters". Using the baseline 30% tax rate in the high-tax country, we find this semi-elasticity to be a modest -0.39. In the second row, we disallow profit shifting by setting $\alpha_{ij} = 0$ and find a considerably higher semi-elasticity of -1.88 for firms we label "Non-Shifters". This difference highlights the importance of profit shifting in how the capital stock responds to changes in tax rates – those firms that shift profit respond considerably less to a change in the tax rate than those that do not. The lower capital response of "Shifters" reflects a high semi-elasticity of the share of the tax base shifted, α_i, with respect to the tax rate. This suggests that MNEs respond to a change in the tax rate far more on the profit-shifting margin than on the capital stock margin and explains large differences in capital responses between "Shifters" and "Non-Shifters".[13]

In the last row of Table 6, we construct a weighted average of the two elasticities, assigning weights based on aggregate data on the share of MNEs in the UK, following our approach in Section 4.2. This aggregate allows us to make a comparison to the existing literature. We convert our estimate of -1.46 to an estimate of the elasticity of the capital stock with respect to the cost of capital; the equivalent user cost elasticity is -1.73.[14] While this estimate is high relative to the range of capital stock elasticities from older literature that are between -0.5 and -1 (Bond and Van Reenen; 2007; Hassett and Hubbard; 2002), our estimates are more closely aligned with the more recent literature that found larger responses.[15]

In Table 6 we present estimates of semi-elasticities of K_{ij} with respect to τ_j for cases where $\tau_j = 20\%$, $\tau_j = 30\%$ and $\tau_j = 40\%$. For "Non-Shifters", the semi-elasticity rises in absolute terms with the tax rate, indicating a rise in the marginal impact of the tax. But for "Shifters", the opposite is true: the marginal impact of the tax on K_{ij} falls as the tax rate increases. This indirectly indicates the marginal effect of a tax rise on profit shifting, which rises with the tax rate. This effect suggests that, at high tax rates, the response of profit shifting becomes so great that the "effective" tax rate rises much less, with consequently a small effect on capital accumulation. At a tax rate of 40% there is effectively no impact on

---

[13]Chodorow-Reich et al. (2024) estimate consolidated capital stock elasticities for domestic and MNEs and find these to be 0.42 and 0.61 respectively. Their approach, however, focuses on the implications of TCJA as a whole and abstracts away from the role of profit shifting. Further, they focus on the implications of TCJA for consolidated capital stock of US-based MNEs, while we consider capital stock of individual subsidiaries.

[14]In Appendix D we outline the calculation behind that transformation.

[15]Some examples of which include Chen et al. (2022); Giroud and Rauh (2019); Maffini et al. (2019); Moon (2022); Ohrn (2018); Suárez Serrato (2018); Zwick and Mahon (2017).

capital accumulation.

4.4 Profit shifting costs

We now consider the level of profit-shifting costs implied by our model. These are indirect and comprehensive estimates of profit shifting costs, inferred from the degree of profit shifting and the estimated parameters. They include the direct costs of any in-house tax department as well as professional and legal fees. They may also include more general anticipated costs, such as any additional tax liabilities or penalties that may arise if risky avoidance schemes turn out ex-post not be consistent with complex international tax law, and any perceived reputational costs of being considered by shareholders or third parties to be aggressive tax avoiders.

In Figure 6 we scale these costs by the total true tax base (before shifting) in both H and L.[16] In general, for the average value of p_i, the fixed costs are 3.3% of the true tax base, while the variable costs are twice as much, around 6.9% of the true tax base.[17] Several papers in the accounting literature estimate the size of specific shifting costs, but very few provide an estimate of total costs. The closest approach to ours is McClure (2023) who also infers the size of costs from a structural model. In that model, firms choose a level of tax risk - the higher the risk, the greater the scrutiny from tax authorities, and the greater the proportion of any tax savings that need to be repaid after an audit. McClure (2023) uses the US FIN 48 provision to estimate risk for consolidated groups, and finds average shifting costs to be 6.4% of pre-avoidance profit, of which 2% is fixed costs. Our estimates, using a very different approach, are comparable to this, although a little higher.

We also investigate heterogeneity of profit shifting costs according to the price of the tax avoidance asset, p_i. The two kinks in the schedule for profit shifting from Figure 1a are also present in Figure 6 . The fixed cost, $p_i Y_i$, rises while at least one of the subsidiaries continues to shift its entire tax base. However, within this region Y_i falls as p_i increases, implying that the variable costs required to shift all profit also rise. This increase is steeper for both costs when an MNE shifts all profits out of both high and low tax countries. As p_i increases further, the higher costs prove prohibitive and the subsidiary in H no longer shifts all its profit. Beyond this point, both fixed and variable costs fall as a proportion of the total true tax base. We find that fixed costs of profit shifting reach a maximum of

---

[16]In Figure A1 in the Appendix we show these costs scaled by the tax base in high tax country or by capital.

[17]Note that to calculate results in this section, we keep ε_{ij} and Π_{ij} constant at their mid-point values.

4.12% of the total true tax base at the point where p_i is the highest value that allows MNEs to still shift all profits away from high tax countries. At that same point, the variable costs of profit shifting are 8.51% of the true tax base. For MNEs that find it easiest to accumulate tax avoidance capability (facing the p_i values at the lower end of the distribution), when MNEs fully shift from both H and L, the fixed costs are at most 3% and the variable costs are at most 6% of the true tax base.

5 Reduced-form estimates

In this section, we evaluate the intensive- and extensive-margin profit-shifting responses to changes in tax rate differentials for model verification. Changes in a high-tax country's tax rate impact all companies operating in that country at the same time, making it difficult to find a valid control group against which to benchmark the MNEs' profit-shifting and investment behavior.

To empirically isolate the profit-shifting responses, we leverage exogenous variation in profit-shifting incentives introduced by a change in the controlled foreign company (CFC) legislation in Italy.[18] Under a CFC rule, subsidiaries of MNEs that are wholly- or partly-owned by a multinational parent that pays tax at an effective rate below a certain threshold (set by the home country) become liable to pay extra tax to the revenue authority of the home jurisdiction (Clifford; 2019).

Before the changes to CFC rules, Italian-owned multinationals operating in the UK could shift profit to jurisdictions that are considered to be tax havens and reduce their effective tax rate to close to zero. After 2002, Italy began to impose additional tax on certain types of income of blacklisted tax haven subsidiaries that are financially controlled by parent companies located in Italy. A simple MNE structure with a parent company in Italy, a high-tax subsidiary in the UK and a tax haven subsidiary in Bermuda illustrates the relevant corporate structure for our empirical analysis. For the period that we study, the UK had a main corporate tax rate of 30% and was considered a high-tax country. The introduction of the CFC regime in Italy increased the tax haven tax rate on passive income of affected Italian MNEs to the Italian tax rate of 34% in 2003 (possibly from an effective rate of zero).

According to the theoretical framework in Section 2, we expect the Italian parent com-

---

[18]Clifford (2019) provides further information on CFC legislation around the world and studies their impact on companies' behavior.

pany's UK 'high-tax' subsidiary to experience a change in its corporate tax liability and investment in the UK. High-tax subsidiaries may become less likely to report zero taxable profit. If tax haven subsidiaries are taxed more intensively, less profit should be shifted out of any of the high-tax subsidiaries of Italian MNEs. This includes subsidiaries in the UK. This is the extensive margin profit-shifting effect. In addition, high-tax subsidiaries with existing profit in high-tax jurisdictions may increase the amount of profit reported in high-tax jurisdictions. This is the intensive margin effect. Depending on the size of the structural parameters $\bar{p}$ and γ, the reform may cause more pronounced effects on the base through the extensive- or the intensive-margin. Finally, if capital adjusts immediately, we may expect a small drop in average investment due to the higher cost of capital driven by the CFC rule for MNEs that shift profits. There are some caveats to this prediction, which we discuss along with the results.

To causally identify the effects of the change in the CFC regime in Italy, we use a difference-in-differences approach and compare the reporting behavior of subsidiaries of Italian MNEs – the treatment group – with subsidiaries of other foreign MNEs – the control group – in the UK over the period 2000 - 2005. Note that we cannot identify directly which Italian MNEs were affected by the reform; by including all Italian MNEs in our treatment group, our estimates therefore reflect the behavioral responses of affected Italian MNEs and their importance within our dataset. The control group is composed of the entire sample of foreign MNE subsidiaries that operate in the UK and we use entropy balancing weights to obtain the most comparable control group to the treatment group. In the paper, we present results using our preferred set of entropy balancing weights, while in the Appendix B, we run the analysis using the unweighted sample of control firms and a series of alternative entropy balancing weights.

Empirical specification   We run three sets of regressions to assess: (i) the change in the probability of reporting zero taxable profit in the UK (using a linear probability model), (ii) the change in the average profit reported in the UK, and (iii) the change in capital in the UK:

$$\mathbb{1}(\text{Taxable Profit} \leq 0)_{it} = \alpha_0 + \beta_0 \text{Treated}_i \times \text{Post-reform}_t + \chi_{i0} + \eta_{tj0} + \varepsilon_{it0} \qquad (13)$$

$$\text{Taxable Profit}_{it} = \alpha_1 + \beta_1 \text{Treated}_i \times \text{Post-reform}_t + \chi_{i1} + \eta_{tj1} + \varepsilon_{it1} \qquad (14)$$

$$\ln(K_{it}) = \alpha_2 + \beta_2 \text{Treated}_i \times \text{Post-reform}_t + \chi_{i2} + \eta_{tj2} + \varepsilon_{it2} \qquad (15)$$

In Equations 13 and 14, the outcome variables relate to the extensive-margin and the

intensive-margin profit-shifting effects, respectively. $\mathbb{1}(\text{Taxable Profit} \leq 0)_{i,t}$ represents a dummy equal to one when a firm reports zero taxable profits in a given year. Taxable Profit$_{it}$ is the level of taxable profit reported in a given year. In Equation 15, the dependent variable is the natural logarithm of total assets which we use to proxy for firm capital.[19] Treated$_i$ is a dummy variable that equals one, if a subsidiary belongs to an MNE that is headquartered in Italy and zero otherwise; Post-reform$_t$ is a dummy variable that equals one from 2003 onward for the Italian CFC reform. χ_is are firm fixed effects, η_{tj}s capture sector-specific time effects, with subscripts j representing two-digit sectors, and ε_{it}s are the error terms.

β_0 captures the effect of the reform on the propensity of the firm to report zero taxable profit in the high-tax jurisdiction (UK), i.e. the extensive margin. Under the model in Section 2, β_0 is negative for the case of the Italian anti-tax avoidance reform. Italian subsidiaries are more likely to report positive profits in the UK after Italy taxes profits located in lower tax jurisdictions in Italy. β_1 represents the overall taxable profit response to the introduction of the CFC legislation. Finally, the impact of taxable position and tax system changes on the cost of capital, reflected in each subsidiary's decision to invest in capital in each jurisdiction is captured by β_2.

Descriptives In Table 7, we present descriptive data on key variables for the pre-reform period for the treatment group (subsidiaries of MNEs with parent company resident in Italy) and the control group (subsidiaries of MNEs with a parent company resident outside the UK and Italy). In this table, we report descriptive statistics before applying entropy-balancing weights. On average, UK subsidiaries of Italian MNEs report £0.89 million and subsidiaries of other MNEs report around £1.15 million of average taxable profit.

Strikingly, in both the treatment and the control groups, 42% of all MNE subsidiaries in the UK report zero taxable profit in the pre-reform period. This is key, as our extensive-margin response to tax reform traces the changes in the prevalence of reporting zero taxable profit. Average asset size of the subsidiaries in the treated group is slightly larger than the average for the control group. When constructing our weighted entropy balance samples, we therefore include size measures to improve comparability between the two groups.

Graphical evidence and reduced-form regression results For the validity of the difference-in-differences approach, treatment and control groups should satisfy the assumption of

---

[19]In our static model, capital stock is accumulated as a result of the beginning-of-period investment.

common counterfactual trends. In the absence of treatment, the change in the average out-come variables of interest for the control and treatment group firms should be similar. We assess whether common counterfactual trends is a plausible assumption in our matched sample by exploring the trajectory of the average outcome variables in pre-reform years. In panel (a) of Figure 7, we show the average probability to report zero taxable income. After the reform, the average probability to report zero taxable profit drops significantly for the treatment group, but not for the control group. The drop in the average probability to report zero taxable profit is consistent with our prediction that the CFC reform in Italy leads to a drop in the likelihood of shifting profit out of the UK for Italian subsidiaries. In panels (b) and (c), we show the evolution of the average taxable income in levels and the inverse hyperbolic sine transformation of that variable respectively. These figures high-light the overall positive response to the CFC reform in terms of reported taxable profits. In panel (d), we show the evolution of the logarithm of total assets, which suggests a small rise after the reform.

In Table 8, we quantify the magnitudes of these responses and present the results of the estimation equations 13, 14, and 15. In columns (1) - (3), we consider the effect of the CFC reform on taxable income. We start by presenting results using OLS and find that subsidiaries of Italian MNEs report £1.33 million more taxable income, on average, after the CFC reform relative to the control group (column 1). Given that the average pre-reform taxable income of the weighted control group (for the weighted sample used in Table 8) is £1.15 million and the CFC reform implied a change in the tax rate of 34 percentage points, the implied semi-elasticity of taxable income with respect to 1 minus tax rate is: 4.4=(1333/893)/0.34. As Table 7 shows that there is a high prevalence of zero taxable income in our sample, we further estimate the general average response of taxable income using the inverse hyperbolic sine and Poisson transformations in columns (2) and (3), respectively.

In column (4), we estimate the effect of the reform on the extensive margin of tax-able income reporting. We find that after the reform, treatment group firms reduce their probability to report zero taxable income in the UK by 5.6 percentage points, which is a 13.3% increase relative to the control group mean of 0.42. In Figure B4 in the Appendix, we show how our empirical and simulated difference-in-difference moments match both for the level of profits and the incidence of reporting zero taxable profits. We find that the reduced-form response is consistent with the case where a fraction of firms responds to the CFC reform, as noted above. We find the best fit when 45% of firms respond to the Italian CFC and the remaining MNEs either respond by moving their tax haven activities

to non-blacklisted alternatives or did not have any affiliates in the blacklisted tax havens in the first place. This is plausible, as the reform only applied to passive income and had a very limited range of tax havens included in the blacklist.[20]

In column (5), we consider the effect of the reform on capital. In contrast to our prediction of a small reduction in capital for aggressive profit shifters, we find no statistically significant average effect of the CFC reform on capital. This is nevertheless consistent with the small elasticities of capital for 'Shifters' that we estimate using our model. It is also possible that capital adjusts more slowly and the impact is harder to pick up during the data period.[21] We complement these results with Figure 8, in which we plot the evolution of capital according to whether a firm switched their status from zero taxable profit reporting before 2003 to positive profit in at least one period after (and including) 2003 (in red) or did not switch their status (in green). These figures suggest that a change in taxable income status is a stronger driver of a user cost change than the CFC reform, as we observe very similar responses for both treated and control group firms. We conclude that the change in total assets appears to be linked with the change in profit reporting, rather than being the direct outcome of the reform.

6   Counterfactual policy experiments and welfare implications

We analyze the impact of three types of tax reforms: (i) a change in the tax rate in the country of production, H or L; (ii) a change in the tax rate applied to profit declared in the haven; and (iii) strengthening anti-avoidance measures, which may imply changes in the tax haven tax rate or may increase the cost of shifting profit by altering the cost of accumulating profit-shifting capability. Governments have used tax rate changes, controlled foreign company rules, and other anti-tax avoidance policies to deter profit shifting and to increase tax revenue. In this section, we analyze the effects of changing the home country tax rate and the rate levied in the haven. In the context of our model, raising the tax rate in the haven is equivalent to introducing a global minimum tax (GMT), which sets a lower bound on the tax rate levied on all profit wherever it is declared. As several countries have already started implementing the OECD's GMT in 2024, we also consider the welfare im-

---

[20]Troiano (2002) discusses the criteria for the black list and this briefing from Clifford Chance discuss the updates to the CFC regime in Italy in 2010, emphasizing that the earlier reform only targeted subsidiaries in blacklisted countries.

[21]Because of subsequent reforms after 2005, we cannot extend the analysis period to later years.

plications of the GMT. Throughout this section, we analyze the responses of multinational firms only and do not include the responses of domestic companies.

6.1 Tax rate changes in high tax country

We start by demonstrating the effects of a small tax rate change for the high-tax country, initially with the tax rate of 30%. In Figure 9 we show, based on our model and estimates, how a one percentage point reduction in the high-tax country tax rate, to 29%, would affect the share of profit shifted from the high-tax country to the tax haven, capital accumulation in the high-tax country, and MNE's profit shifting costs scaled by the global true tax base. The solid lines in each subfigure represent the pre-reform evolution of each variable across the distribution of p_i. The dashed lines show what happens after the tax rate reduction by 1ppt.

The lower tax rate reduces the benefit from shifting profit from high tax country to the tax haven. Firms with very low p_i, which shift all their profit before the reform, continue doing so, as the benefit of it continues to outweigh the costs. For higher values of p_i, some firms respond on the extensive margin by reducing their shifting below 100% of their true profit, reflected in the full-shifting kink moving to the left. As p_i increases further, firms that are already partial shifters reduce their shifting since the marginal benefits are reduced.

The cost of capital does not change for firms that continue to shift all of their profit. They pay no tax, and their costs of shifting are also unaffected by the change in the tax rate. Consequently, as we show in panel (b), their capital accumulation is unaffected by the tax rate change. However, for firms that respond on the extensive profit-shifting margin, and for firms that were partial shifters, the cost of capital falls. This is reflected in slightly higher capital accumulation in the high-tax country for higher values of p_i.

Panel (c) illustrates the impact of the tax rate change on profit-shifting costs. Since there is no change in profit shifting for firms that continue to fully shift, the costs of shifting for these firms are also unaffected. However, the kink in the cost of shifting across p_i moves to the left, consistent with the extensive margin response in panel (a). As p_i increases, the lower shifting induced by the lower tax rate in the high tax country reduces both the fixed and the variable costs of profit shifting.[22]

---

[22]In Figure A2 we also demonstrate the effects of a much larger, 10 percentage points, reduction in the high-tax country tax rate.

6.2 Tax rate changes in tax havens

In Figure 10, we examine the effects of a small tax rate increase in the tax haven country, from 0% to 1%, following the exposition from Figure 9. The effects on profit shifting, plotted in panel (a) are similar to those of a reduction in the tax rate in the high tax country, since the rise in the haven rate also reduces the benefit of profit shifting. We can again observe both the extensive-margin and the intensive-margin profit-shifting effects. However, we find the opposite effect on capital accumulation. In panel (b) we show that as the increase in the tax haven tax rate increases the effective tax rate, $\widehat{\tau_{ij}}$, even for firms that shift all their profit out of the high-tax country, this leads to an increase in the cost of capital and a reduction in K_{ij} for all firms.

Further, the higher tax haven tax rate reduces the benefits of shifting from both the high- and the low-tax countries, leading to a larger decline in both fixed and variable costs of shifting than in the case of tax rate reduction in high-tax country. This happens because the higher tax haven tax rate reduces investment in the tax avoidance asset, Y_i, for all firms, including those in the low-tax country and those that continue to fully shift profits out of the high-tax country. The extensive margin effect in panel (a) reduces the value of p_i at which the kink point in the shifting costs occurs. Therefore, the change in shifting costs after the rise in the tax haven rate is larger for values of p_i higher than at the kink point. MNEs with p_i higher than the value at this kink point are the 'partial' shifters.

6.3 Marginal value of public funds

How can we compare the welfare implications of a home country tax rate reduction and a tax haven tax rate increase that both reduce profit shifting? To answer this question, we use the MVFP framework (Hendren and Sprung-Keyser; 2020, 2022) and compute the ratio of a change in the welfare of firms (benefits) to the change in the net cost of these policies to governments globally. The cost of each of these policies is the change in tax revenues in response to the one percentage point change in the tax rate. This captures both the mechanical and behavioral effects of tax policies through changes in capital accumulation and profit shifting.[23]

For the increase in the tax haven tax rate, we find the MVPF for our sample of MNEs to be 0.38. Because the increase in tax haven tax rate is a tax revenue gain (rather than a cost) to the government in the tax haven country, we interpret this MVPF to say that for a

---

[23]Our approach implicitly assumes that all profit-shifting costs represent a welfare loss to society.

$1 gain to the government from increasing the tax rate, the taxpayer – the firm – is worse off by $0.38. As such, this policy imposes a welfare cost of only $0.38 per dollar of revenue raised. This is lower than $1 since the resulting gain in income from a reduction in profit-shifting costs outweighs the loss in private income due to a reduction in K_i. We calculate the MVPF for our sample of MNEs of a 1 percentage point tax rate reduction in the high tax country moving from 30% to 29% to be infinitely large. That is because this rate reduction not only increases NPV of firms, but also generates an increase in tax revenues, resulting from the combined behavioral response of lower profit shifting and higher investment. This suggests that the 30% tax rate in the UK in 2002 may have been on the right hand side of the Laffer curve for MNEs.

The evidence from Figures 9 and 10 suggests that both the tax rate increase in a tax haven and tax rate reduction in high tax country reduce profit shifting by a similar amount. However, the former reduces investment while the latter increases investment. This implies that that a marginal tax rate reduction in the home country would be more welfare improving than a marginal tax rate increase in tax haven.[24]

6.4 Global Minimum Tax

We then turn to analyzing the implications of setting a range of different tax rates in the tax haven. In the context of our model, with no real activity in the haven, and abstracting from differences in the measure of profitability for tax and accounting purposes, this allows us to mimic the introduction of the new Global Minimum Tax.[25] We first focus on the 15% rate, currently implemented in the EU, UK, Japan and others,[26] on profit-shifting, capital accumulation and shifting costs. In principle, the GMT sets a floor on tax competition between countries, but we do not model any responses in tax rates in the non-haven countries.

First, in Figure 11, we show that our findings are directionally similar to the effects of a small change in tax haven tax rate in Figure 10, but much stronger in magnitude. Our model predicts that a GMT introduced at this rate almost completely eliminates full-shifting out of the high-tax country, accompanied by very large intensive margin responses.

---

[24]Note that the MVPF calculation depends on the starting point for the tax rate. For example, the MVPF of a home country tax rate reduction from 20 to 19% is still infinite, but the MVPF of home country tax rate reduction from 15% to 14% is 4.5.

[25]The OECD's GMT has some features that we do not model. Chief amongst these is that since we assume there is no real activity in the haven, the "substance-based income inclusion" is zero.

[26]Australia, Canada, South Korea, South Africa are among the countries that have also adopted legislation to start the implementation of the OECD's GMT in 2024.

These responses are also reflected in a large reduction of both the variable and fixed costs of profit shifting. The introduction of GMT increases the cost of capital for all firms, adversely affecting capital accumulation.

While the GMT was adopted by many countries in 2024 using the 15% tax rate, we can also use our model to analyze the implications of introducing GMT at varying rates between 0% and 25%. In Figure 12 we first explore the impact of this range of GMT rates on the key factors in the model. We start by showing in the top panel the changes in capital accumulation in the high-tax and low-tax countries, and in aggregate. For each GMT threshold level, we assume that tax haven tax rate (and when the GMT rate reaches the low-tax country tax rate, also the low-tax country's tax rate) automatically increases to the threshold level. Capital accumulation falls in both countries as the minimum tax rate rises. The fall is steeper in the low tax country after the minimum tax rate exceeds 12.5%, as the effective tax rate in that country rises more steeply in line with the minimum tax rate.

In the bottom panel of Figure 12, we demonstrate the trajectories of tax revenue in the high-tax country (black smooth line), in the low-tax country (black dashed line), and in the tax haven (black dotted line)[27]. Tax revenues rise steeply in both the high-tax and low-tax countries as a higher GMT reduces profit shifting; this is more pronounced initially in the high tax country as the rate of profit shifting falls more sharply there, but the revenue in the low-tax country rises more steeply for rates above 12.5%. Revenues in the haven initially rise (from zero) as firms continue to shift profits for low values of GMT, but begin to decline even at relatively low haven tax rates as higher GMT reduces profit shifting to tax havens.

Also in this panel, we demonstrate the impact on fixed and variable profit-shifting costs. We depict the aggregate fixed cost of profit-shifting ($p_i Y_i$) with the red smooth line and the aggregate variable cost of profit-shifting with the red dashed line. Before any minimum tax is introduced, variable costs are roughly twice as large as fixed costs, consistent with evidence from Figure 1. Both types of costs gradually decline, as profit-shifting also declines.

In Figure 13, we summarize these effects using a simple measure of global welfare which aggregates total private value and tax revenue: $W = \sum_{i=1}(V_i + T_i)$. We also show the evolution of the two components separately. In the top panel, total tax revenue rises with the minimum tax rate, and total private income declines. But, consistent with the MVPF framework discussed above, revenue initially rises more quickly than private income de-

---

[27]We assume in this Figure that the additional revenue from the rise in the haven tax rate is captured by the haven itself.

clines. That is because the reduction in profit shifting costs outweighs the reduction in private income due to lower capital accumulation. Aggregating these components into the welfare measure at different GMT tax rates, initially raises global welfare. However, as the rate of decline in profit shifting falls with a higher minimum tax rate, then eventually the reduction in capital accumulation becomes more important and welfare reaches a maximum and then declines. In the top panel, we show that, where the high-tax country has a 30% tax rate, the welfare-maximizing GMT rate is 20%. In the bottom panel, we demonstrate that this optimum depends on the tax rate in the high tax country; if that rate is 20%, the welfare-maximizing GMT rate is 13%.

7   Conclusion

In this paper we explore the joint decisions of MNEs with respect to profit shifting and capital investment. An innovative feature of our approach is a careful modeling of profit-shifting costs, specifically introducing the possibility of investing in a tax avoidance asset, intended to capture the costs of complex structures set up by MNEs to facilitate profit shifting to tax havens. We make four distinct contributions to the literature.

First, we find considerable heterogeneity in profit shifting across MNEs, depending on their characteristics. In particular, we account for the fact that a substantial proportion of MNES pay no tax on profit in high tax countries. We show that, by ignoring extensive margin effects, the traditional convex cost model underestimates the average profit-shifting semi-elasticity with respect to tax rate differentials between high-tax jurisdictions and tax havens. Second, we find considerable interdependence between profit shifting and capital investment. Some MNEs - with high elasticities of profit shifting - respond to changes in tax rates primarily by adjusting their profit shifting, with only a small impact on capital investment. Firms that do not shift profit have much higher elasticities of capital investment with respect to tax rates. Differences in profit shifting can therefore account for large disparities in the elasticity of capital investment with respect to tax rates. Third, we estimate the size of the costs of profit shifting to be substantial, with average fixed costs at 3.3% of the true tax base, and average variable costs to be 6.9% of the true tax base.

Fourth, we use our model to identify the behavioral responses and welfare implications of possible reforms to taxation. The sensitivity of profit shifting and its related costs to changes in the tax rate in the high tax country and the tax haven leads to the possibility of welfare enhancing reforms. For example, we find evidence that, for MNEs, the tax system

31

in the high tax country may be on the wrong side of the Laffer curve, with a reduction in the tax rate increasing tax revenue as profit shifting declines and capital investment increases. A rise in the tax haven rate, as envisaged as part of the Global Minimum Tax, can also raise welfare by reducing wasteful profit-shifting costs, despite having a negative impact on capital investment. Our model results therefore also carry significant policy implications, especially for the GMT. More generally, our counterfactual policy experiments highlight that the optimal GMT tax rate may need to balance the benefits of reduced profit shifting and its associated costs against reduced investment and economic activity.

Our framework can be adapted to analyze alternative policy proposals, offering a versatile tool for evaluating the impact of ongoing and proposed international tax reforms. Future work could extend our model to analyze labor market complementarities and examine implications of profit shifting for employment and wage setting. This could shed light on broader welfare implications of corporate tax reforms beyond firm value and tax revenues.

References

Bares, F., Devereux, M. P., Guceri, I. and Patil, V. (2023). Business location and investment decisions with a global minimum tax, *Centre for Business Taxaiton Working Paper* .

Beer, S., Mooij, R. and Liu, L. (2019). International corporate tax avoidance: A review of the channels, magnitudes, and blind spots, *Journal of Economic Surveys* **34**(3): 660688.

Bilicka, K. A. (2019). Comparing uk tax returns of foreign multinationals to matched domestic firms, *American Economic Review* **109**(8): 2921–53.

Bilicka, K., Devereux, M. and Güçeri, I. (2023). Tax-avoidance networks and the push for a "historic" global tax reform, *Tax Policy and the economy* **37**(1): 57–108.

Bilicka, K., Qi, Y. and Xing, J. (2022). Real responses to anti-tax avoidance: Evidence from the uk worldwide debt cap, *Journal of Public Economics* **214**: 104742.

Bilicka, K. and Scur, D. (2024). Organizational capacity and profit shifting, *Journal of Public Economics* **238**: 105179.

Bond, S. and Van Reenen, J. (2007). Microeconometric models of investment and employment, *Handbook of econometrics* **6**: 4417–4498.

Bustos, S., Pomeranz, D., Serrato, J. C. S., Vila-Belda, J. and Zucman, G. (2022). The race between tax enforcement and tax planning: Evidence from a natural experiment in chile, *Technical report*, National Bureau of Economic Research.

Chamberlain, D. G. and Curtis, S. L. (2021). Apple's cost-sharing arrangement: Frankenstein's monster, *Tax Notes Federal* **172**: 1049.

Chen, Z., Jiang, X., Liu, Z., Serrato, J. C. S. and Xu, D. Y. (2022). Tax Policy and Lumpy Investment Behaviour: Evidence from China's VAT Reform, *The Review of Economic Studies* **90**(2): 634–674.

Chetty, R. (2012). Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply, *Econometrica* **80**(3): 969–1018.

Chodorow-Reich, G., Smith, M., Zidar, O. M. and Zwick, E. (2024). Tax policy and investment in a global economy, *Technical report*, National Bureau of Economic Research.

Clausing, K. (2016). The effect of Profit Shifting on the Corporate Tax Base in the United States and Beyond, *National Tax Journal* **68**(4): 905–934.

Clausing, K. A. (2003). Tax-motivated transfer pricing and us intrafirm trade prices, *Journal of public economics* **87**(9-10): 2207–2223.

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

Coppola, A., Maggiori, M., Neiman, B. and Schreger, J. (2021). Redrawing the map of global capital flows: The role of cross-border financing and tax havens, *The Quarterly Journal of Economics* **136**(3): 1499–1556.

Cristea, A. D. and Nguyen, D. X. (2016). Transfer Pricing by Multinational Firms: New Evidence from Foreign Firm Ownerships, *American Economic Journal: Economic Policy* **8**(3): 170–202.

Curtis, S. L. and Avi-Yonah, R. S. (2023). Microsoft's cost sharing arrangement: Frankenstein strikes again, *Tax Notes Federal* **178**: 23–030.

Davies, R. B., Martin, J., Parenti, M. and Toubal, F. (2018). Knocking on tax haven's door: Multinational firms and transfer pricing, *The Review of Economics and Statistics* **100**(1): 120–134.

Desai, M. A., Foley, C. F. and Hines, J. J. (2006). The demand for tax haven operations, *Journal of Public Economics* **90**(3): 513–531.

Desai, M. A., Foley, C. F. and Hines, J. R. (2004). A Multinational Perspective on Capital Structure Choice and Internal Capital Markets, *Journal of Finance* **59**(6): 2451–2487.

Devereux, M. P., Liu, L. and Loretz, S. (2014). The elasticity of corporate taxable income: New evidence from UK tax records, *American Economic Journal: Economic Policy* **6**(2): 19–53.

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

Dharmapala, D. and Riedel, N. (2013). Earnings shocks and tax-motivated income-shifting: Evidence from European multinationals, *Journal of Public Economics* **97**(C): 95–107.

Dischinger, M., Knoll, B. and Riedel, N. (2014). There's no place like home: The profitability gap between headquarters and their foreign subsidiaries, *Journal of Economics & Management Strategy* **23**(2): 369–395.

Dischinger, M. and Riedel, N. (2011). Corporate taxes and the location of intangible assets within multinational firms, *Journal of Public Economics* **95**(7-8): 691707.

Dowd, T., Landefeld, P. and Moore, A. (2017). Profit shifting of us multinationals, *Journal of Public Economics* **148**: 1–13.

Dyrda, S., Hong, G., Steinberg, J. B. et al. (2022). A macroeconomic perspective on taxing multinational enterprises, *Technical report*.

Egger, P., Eggert, W. and Winner, H. (2010). Saving taxes through foreign plant ownership, *Journal of International Economics* **81**(1): 99 – 108.

Feldstein, M. (1995). The effect of marginal tax rates on taxable income: A panel study of the 1986 Tax Reform Act, *Journal of Political Economy* **103**(3): 551–72.

Feldstein, M. (1999). Tax avoidance and the deadweight loss of the income tax, *Review of Economics and Statistics* **81**(4): 674–80.

Ferrari, A., Laffitte, S., Parenti, M. and Toubal, F. (2022). Profit shifting frictions and the geography of multinational activity, *arXiv preprint arXiv:2211.04388* .

Finkelstein, A. and Hendren, N. (2020). Welfare analysis meets causal inference, *Journal of Economic Perspectives* **34**(4): 146–167.

Gallant, A. R. and Tauchen, G. (1996). Which moments to match?, *Econometric Theory* **12**(4): 657–681.

Giroud, X. and Rauh, J. (2019). State taxation and the reallocation of business activity: Evidence from establishment-level data, *Journal of Political Economy* **127**(3): 1262–1316.

Gourieroux, C., Monfort, A. and Renault, E. (1993). Indirect inference, *Journal of Applied Econometrics* **8**: S85–S118.

Griffith, R., Miller, H. and O'Connell, M. (2014). Ownership of intellectual property and corporate taxation, *Journal of Public Economics* **112**: 1223.

Grubert, H. and Altshuler, R. (2013). Fixing the system: an analysis of alternative proposals for the reform of international tax, *National Tax Journal* **66**(3): 671–711.

Grubert, H. and Slemrod, J. (1998). The Effect Of Taxes On Investment And Income Shifting To Puerto Rico, *The Review of Economics and Statistics* **80**(3): 365–373.

Gumpert, A., Hines, J. R. and Schnitzer, M. (2016). Multinational firms and tax havens, *The Review of Economics and Statistics* **98**(4): 713–727.

Guvenen, F., Mataloni Jr, R. J., Rassier, D. G. and Ruhl, K. J. (2022). Offshore profit shifting and aggregate measurement: Balance of payments, foreign investment, productivity, and the labor share, *American Economic Review* **112**(6): 1848–84.

Hall, R. E. and Jorgenson, D. W. (1967). Tax Policy and Investment Behavior, *The American Economic Review* **57**(3): pp. 391–414.

Hanappi, T. and Cabral, A. C. G. (2020). The impact of the pillar one and pillar two proposals on mne's investment costs, (50).

Hassett, K. A. and Hubbard, R. G. (2002). Chapter 20 - Tax Policy and Business Investment, Vol. 3 of *Handbook of Public Economics*, Elsevier, pp. 1293 – 1343.

Hebous, M. S. and Keen, M. M. (2021). *Pareto-improving minimum corporate taxation*, number 9633, International Monetary Fund.

Heckemeyer, J. H. and Overesch, M. (2017). Multinationals' profit response to tax differentials: Effect size and shifting channels, *Canadian Journal of Economics/Revue canadienne d'économique* **50**(4): 965–994.

Hendren, N. and Sprung-Keyser, B. (2020). A unified welfare analysis of government policies, *Quarterly Journal of Economics* **135**(3): 1209–1318.

Hendren, N. and Sprung-Keyser, B. (2022). The case for using the MVPF in empirical welfare analysis, *NBER Working Paper 30029* .

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

Hines Jr, J. R. (2022). Evaluating tax harmonization, *Ann Arbor* **1001**: 48109–1220.

Huizinga, H., Laeven, L. and Nicodeme, G. (2008). Capital structure and international debt shifting, *Journal of Financial Economics* **88**(1): 80–118.

Janeba, E. and Schjelderup, G. (2022). The global minimum tax raises more revenues than you think, or much less.

Johannesen, N. (2022). The global minimum tax, *Journal of Public Economics* **212**: 104709.

Kennedy, P. J., Dobridge, C., Landefeld, P. and Mortenson, J. (2024). The efficiency-equity tradeoff of the corporate income tax: Evidence from the tax cuts and jobs act, *Unpublished manuscript* .

Koethenbuerger, M., Mardan, M. and Stimmelmayr, M. (2019). Profit shifting and investment effects: The implications of zero-taxable profits, *Journal of Public Economics* **173**: 96 – 112.

Langenmayr, D. and Liu, L. (2023). Home or away? profit shifting with territorial taxation, *Journal of Public Economics* **217**: 104776.

Maffini, G., Xing, J. and Devereux, M. P. (2019). The impact of investment incentives: evidence from uk corporation tax returns, *American Economic Journal: Economic Policy* **11**(3): 361–389.

McClure, C. G. (2023). How costly is tax avoidance? evidence from structural estimation., *Accounting Review* **98**(6).

Michaelides, A. and Ng, S. (2000). Estimating the rational expectations model of speculative storage: A monte carlo comparison of three simulation estimators, *Journal of Econometrics* **96**(2): 231–266.

Moon, T. S. (2022). Capital gains taxes and real corporate investment: Evidence from korea, *American Economic Review* **112**(8): 2669–2700.

Ohrn, E. (2018). The effect of corporate taxation on investment and financial policy, *American Economic Journal: Economic Policy* **2**: 272–301.

Riedel, N. (2018). Quantifying international tax avoidance: A review of the academic literature, *Review of Economics* **69(2)**: 169–181.

Risch, M. (2024). Does taxing business owners affect employees? evidence from a change in the top marginal tax rate, *The Quarterly Journal of Economics* **139**(1): 637–692.

Slemrod, J. and Wilson, J. D. (2009). Tax competition with parasitic tax havens, *Journal of Public Economics* **93**(11-12): 1261–1270.

Suárez Serrato, J. C. (2018). Unintended consequences of eliminating tax havens, *NBER Working Paper* (w24850).

Suárez Serrato, J. C. (2018). Unintended Consequences of Eliminating Tax Havens, *NBER Working Papers 24850*, National Bureau of Economic Research, Inc.

Troiano, P. (2002). Italian controlled foreign company legislation: Publication of theblack list', *Intertax* **30**(3).

Tørsløv, T., Wier, L. and Zucman, G. (2022). The Missing Profits of Nations, *The Review of Economic Studies* .

Zwick, E. and Mahon, J. (2017). Tax Policy and Heterogeneous Investment Behavior, *American Economic Review* **107**(1): 217–248.

Figure 1: Profit-shifting and investment.
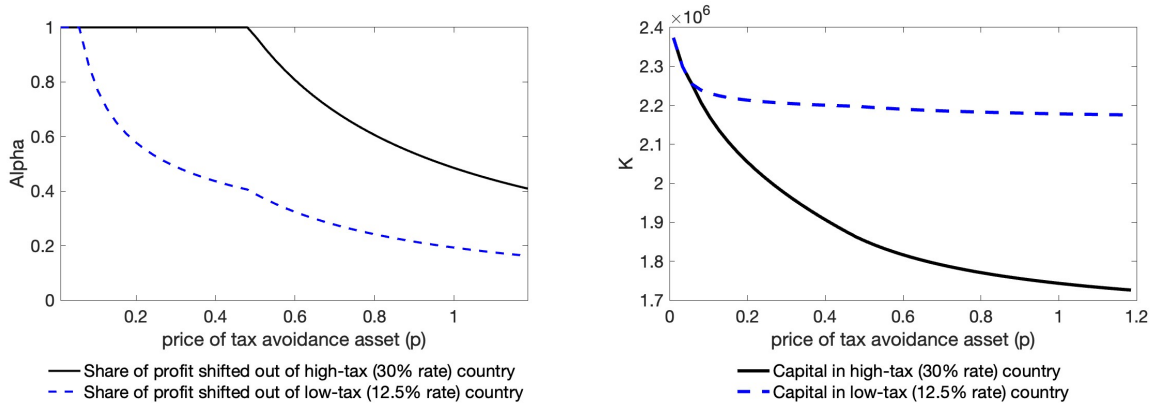
(a) Share of profit shifted, base case

(b) Choice of productive capital, base case

Note: The first panel shows the share of profit shifted out of each of the high-tax and low-tax countries as p_i varies. The second panel traces the total capital in each of the two countries. In both cases, the productivity draw ε and cost of operation Π are fixed at their mid-point values using the parameters that we estimate in Section 4.

Table 1: Descriptive statistics

	Mean	Median
Zero taxable profit (share)	0.421	0
Taxable income (in log)	11.544	11.748
Total assets (in log)	14.409	14.391
Turnover (in log)	14.697	14.785
Observations	8674	8674

Note: This table shows the mean and median for the set of key descriptive variables, namely, the incidence of zero taxable income reporting, taxable income measures and asset values. In constructing the means and medians for this table, we only include the observations that are used to generate the moments in our simulated method of moments estimation. These are the cross-sectional snapshots of non-Italian MNEs in the final pre-Italian reform period of 2001.

Figure 2: Deviation of operating profit from taxable profit for company groups, pooled data

Note: This figure shows the histogram (by fraction) of the ratio of the difference between operating profit and taxable profit as a share of operating profit for all groups that have reported positive operating profit in a given tax year. We pool all years for MNEs available in the data to maintain observation frequencies above the data provider's requirements in each bin. For the same reason, we restrict the sample to contain group level operating profit deviation values between -0.5 and 1.

Figure 3: Average company asset size by percentiles of taxable income reporting

Note: This figure shows the average total asset size (in million GBP) along the distribution of taxable income. We generate percentiles of taxable incomes in our sample, then show average asset values for clusters of bins. Almost half the observations are grouped in the lowest percentile of taxable income reported, that is, at zero taxable profit, generating the gap between the leftmost bin and the next quantile bins. Asset sizes include values between the 5th and 95th percentiles in the overall pooled data.

Table 2: List of moments that identify key structural parameters

Description	Moments	$\bar{p}$	γ	β_Π	θ	a	σ
Zero taxable income	Mean	✓	✓	✓			
Taxable profit (in log)	Uncond. Mean	✓	✓	✓	✓	✓	✓
	Uncond. Variance						✓
	Quartiles	✓	✓	✓	✓	✓	✓
Total assets (in log)	Uncond. Mean	✓	✓		✓	✓	✓
	Uncond. Variance						✓
	Mean, cond. on positive taxable profit				✓	✓	
	Var, cond. on positive taxable profit						✓
	Quartiles, unconditional				✓	✓	
	Quartiles, cond. on positive taxable profit	✓	✓				

Note: This table shows the list of moments used in the MSM procedure. The tickmarks in red indicate the moments that are theoretically the key moments in identifying the specified parameters. For example, the mean of zero taxable income identify $\bar{p}$, γ and β_Π, but this moment is particularly important in the identification of $\bar{p}$.

Table 3: Structural estimates

Assumed and Estimated Parameters			Method
δ:	Depreciation Rate	0.05	Assumed
β:	Discount Factor	0.95	Assumed
$\hat{\bar{p}}$:	Upper bound, unit price of tax avoidance asset	1.185 (0.021)	MSM
$\hat{\gamma}$:	Convex cost of shifting	0.123 (0.002)	MSM
$\hat{\theta}$:	Total factor productivity (in log)	131.005 (18.912)	MSM
$\hat{a}$:	Output elasticity wrt K	0.553 (0.009)	MSM
$\hat{\sigma}$	Std.dev of productivity draw	0.926 (0.019)	MSM
$\hat{\beta}_\Pi$	Scale parameter of linear cost term	1085.663 (97.706)	MSM

Note: This table shows the assumed parameters and our estimates for the structural profit-shifting cost parameters using our MSM procedure. We use the moments listed in Table 2 to match to simulated counterparts. To construct the weight matrix, we use the diagonal elements of the inverse variance-covariance matrix of empirical moments. We assume that the relative unit price of the tax avoidance asset is uniformly distributed over the interval $(0, \bar{p}]$. In the discrete choice application, we use 0.01 to be the lowest value for p.

Figure 4: Model fit
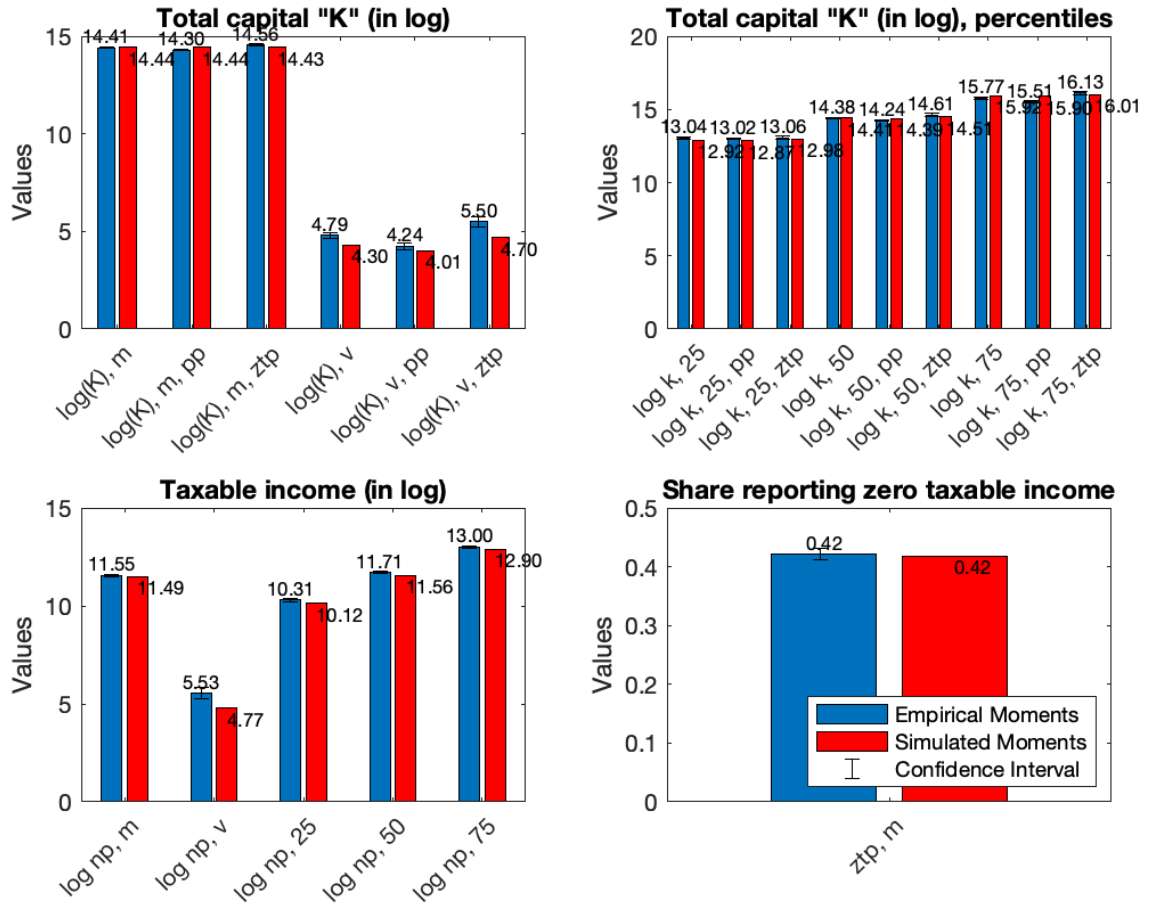
Note: This figure shows the comparison between estimated and structural moments. The blue bars show the empirical moments, with error bars indicating the 95% confidence intervals. The red bars indicate simulated moments constructed using the estimated values for the structural parameters. In the horizontal axes of the figures: 'm' stands for the mean, 'v' stands for the variance, the numbers 25, 50 and 75 stand for the 25-th, 50-th and 75-th percentile values; 'pp' and 'ztp' stand for 'positive taxable profit' and 'zero taxable profit' subsamples (moments conditional on positive profit or zero taxable income reporting). In this figure, the empirical moments are drawn from the 2001 snapshot of corporation tax returns data, and only the observations that are in the control group (of our reduced-form estimation sample) are used. This means that the UK subsidiaries of Italian-headquartered multinationals are left out to align the structural estimation sample with the reduced-form estimation sample discussed in Section 5.

Table 4: Tax base responses to a percentage point rise in the high-tax country tax rate

	Short-run	Long-run
Aggregate	-0.015	-0.047
Ignoring zero-reporters	-0.013	-0.034
Discrepancy	15%	39%

Note: This table shows short-run and long-run average responses to incremental changes in tax rates implied by our structural estimates. Interpretation of the quantities as tax base semi-elasticities is as follows: the aggregate short-run semi-elasticity is -1.5, compared with an aggregate long-run semi-elasticity of -4.7. In short-run estimates, we assume that companies can change their share of profit shifted to the tax haven in response to a percentage point rise in the home (high-tax) country tax rate but they cannot adjust their capital in the high-tax country or their global tax avoidance asset stock. In long-run estimates, we relax this adjustment constraint and allow companies to adjust their capital, tax avoidance capability as well as the share of profit shifted from the high- and low-tax countries to the tax haven. The row labeled 'Discrepancy' shows the differential change in semi-elasticity between the aggregate estimate and the estimate that ignores zero-taxable-income-reporters. The aggregate results are the weighted average of the response of subsidiaries of MNEs and other firms, weighted by gross value added. Domestic firms are assumed to not engage in profit shifting before or after the change in tax rate.

Table 5: Tax base responses to a percentage point rise in high-tax country tax rate: heterogeneity across high tax country tax rate.

High-tax country tax rate:	40%		30%		20%	
Horizon:	Short-run	Long-run	Short-run	Long-run	Short-run	Long-run
Aggregate	-0.110	-0.250	-0.015	-0.047	-0.007	-0.023
Ignoring zero-reporters	-0.050	-0.171	-0.013	-0.034	-0.001	-0.016
Discrepancy	122%	45%	15%	39%	720%	40%

Note: This table shows short-run and long-run average responses to incremental changes in tax rates implied by our structural estimates. The first two columns show the average responses of MNE subsidiaries operating in a high-tax country with a 40% statutory tax rate, the middle two columns show the average responses of MNE subsidiaries operating in a high-tax country with a 30% tax rate (our base case), and the final two columns show the responses of MNE subsidiaries operating in a high-tax country with a 20% tax rate. The interpretation of the quantities is the same as for Table 4. The aggregate results are the weighted average of the response of subsidiaries of MNEs and other firms, weighted by gross value added. Domestic firms are assumed to not engage in profit shifting before or after the change in tax rate.

Figure 5: Heterogeneity in responses to an incremental rise in the home country tax rate
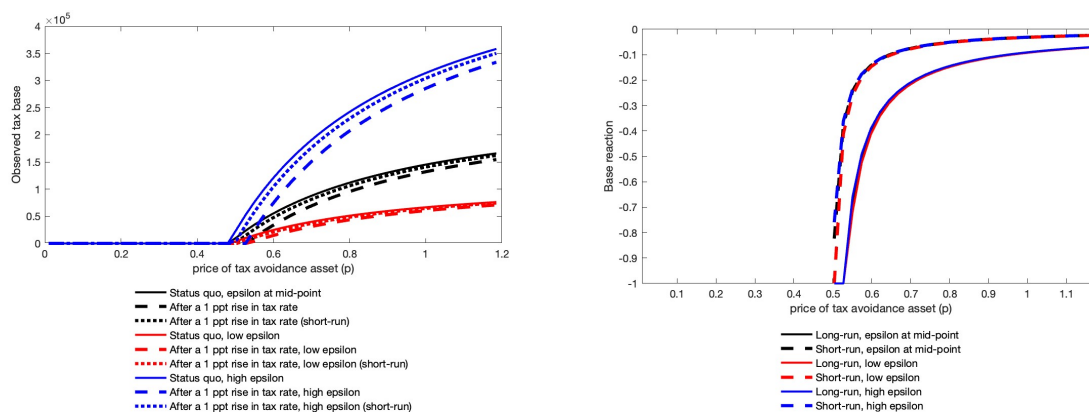
(a) Response of the observed tax base

(b) Relative change in tax base

Note: This figure shows heterogeneous responses to a one percentage point rise in the high-tax country tax rate depending on the unit price of tax avoidance capability at the MNE level. Responses are only for the tax base reported in the high-tax country. The left-hand panel shows the theoretical change in observed tax base for different states along the unit price of the tax avoidance intangible. Black lines represent the values for productivity draws at the mid-point level. Blue lines represent the responses by firms that have slightly higher productivity than the mid-point level. Red lines represent the responses by firms that have slightly lower productivity than the mid-point level. Dashed lines show long-run responses after K and Y adjust as well as α. Dotted lines show short-run responses that keep K and Y constant, only allowing the share of profit α to change. The right hand panel shows the change in base for different levels of the productivity draw.

Table 6: Capital responses to a percentage point rise in high-tax country tax rate.

High-tax country tax rate:	40%	30%	20%
Shifters	-0.0000	-0.0039	-0.0089
Non-Shifters	-0.0233	-0.0188	-0.0156
Aggregate	-0.0167	-0.0146	-0.0137

Note: This table shows capital responses to incremental changes in tax rates implied by our structural estimates. The first column show the average responses of MNE subsidiaries operating in a high-tax country with a 40% statutory tax rate, the second column show the average responses of MNE subsidiaries operating in a high-tax country with a 30% tax rate (our base case), and the last column show the responses of MNE subsidiaries operating in a high-tax country with a 20% tax rate. Shifters are MNEs, non-shifters are firms for which we switch off the profit shifitng channel. The aggregate is the weighted average of shifters and non-shifters using the ONS share of MNEs in the UK as the weight for shifters. The interpretation of the quantities is the same as for Table 4.

Figure 6: Profit-shifting costs.
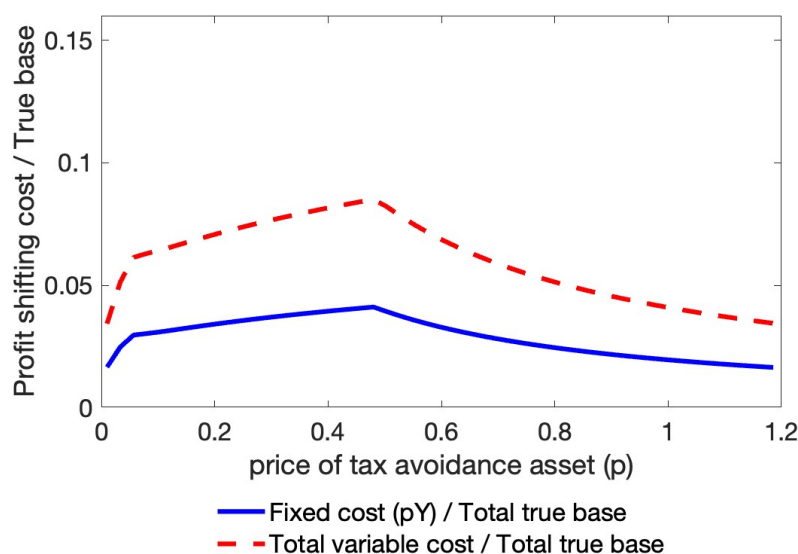
Note: This figure shows the total fixed and variable shifting costs of the MNE as a proportion of the total "true" tax base as p_i varies. The productivity draw ε and cost of operation Π are fixed at their mid-point values using the parameters that we estimate in Section 4.

Table 7: All sampled companies – key descriptive statistics by treatment status, pre-reform period

Group:	Treatment	Control	Treatment-Control Difference	t-test
Zero taxable profit (share)	0.42	0.42	0.006	0.19
Taxable income (thousand GBP)	893	1,146	253.386	0.568
Fixed Assets (in thousand of GBP)	635	654	18.48	0.381
Assets (in thousand of GBP)	2,027	2,018	-9.337	-0.1
Observations	248	8674	8922	

Note: This table shows selected descriptive statistics for means of key variables in the pre-reform period available in the data (2001). Control group companies are the UK subsidiaries of MNEs with parent companies located in countries outside the UK and Italy. Treatment group companies are the UK subsidiaries of MNEs with parent companies located in Italy. Taxable profit data are from the tax return and data on balance sheet size are from company accounts. Units for taxable profit and asset size are nominal British Pounds.

Figure 7: Trends in average propensity to report positive taxable income, average income, and capital by treatment status

(a) Probability to report zero taxable profit

(b) Average taxable income (in level)

(c) Average taxable income (inverse hyperbolic sine)

(d) Average capital (in log)

Note: In Panel (a), we show the evolution of the average probability to report zero taxable income in the United Kingdom, in Panel (b) the evolution of the average taxable income ('Profits chargeable to Corporation Tax' in the UK corporation tax return), in Panel (c) the evolution of the inverse hyperbolic sine transformation of the average taxable income, and in Panel (d) the evolution of the logarithm of total assets. We present the trends separately for the treatment group of Italian-headquartered MNEs (red, dashed line) and for the control group of entropy-weighted subsidiaries of MNEs headquartered outside of the UK and Italy (blue, solid line). We use the following weighting variables: logarithm of total assets in 2001 and 2002, dummy equal to 1 if a firm reported zero taxable profit in 2000, 2001 and 2000, a dummy equal to 1 is a firm has a tax haven subsidiary in its ownership structure.

Table 8: Baseline reduced-form regression results

Outcome variables	(1) np	(2) invhs (np)	(3) Poisson (np)	(4) ztp	(5) log(ta)
treat × post	1333.401*	0.520*	1.197***	-0.056**	0.055
	(715.995)	(0.276)	(0.330)	(0.023)	(0.033)
Observations	53485	53485	45044	53485	53485
Firm FE	Yes	Yes	Yes	Yes	Yes
Sector-year effects	Yes	Yes	Yes	Yes	Yes

Note: This table shows the results of difference-in-difference regression estimates based on equations 13, 14, and 15. Columns (1) - (3) use as a dependent variable np, which is the level of taxable income in 1,000 of GBP ('Profits chargeable to Corporation Tax' in the UK corporation tax return). In Column (1), we estimate the model using OLS, in column (2) using inverse hyperbolic sine transformation, and in column (3) we use Poisson transformation. In column (4) the outcome variable is a dummy equal to 1, if the taxable income is equal to zero, and zero otherwise. In column (5) the outcome variable is the logarithm of total assets. treat is a dummy equal to 1 for subsidiaries of Italian MNEs, 0 for subsidiaries of all other non-UK MNEs. post is equal to 1 is 2003 - 2005, zero before that. We use entropy balancing to weight the control group using the logarithm of total assets in 2000 - 2002, taxable profits in 2000 - 2002, and zero taxable profit share in 2000 - 2002 as matching variables. We do not apply weights to Poisson regressions. In all columns we include firm and sector-year fixed effects. Standard errors are clustered at the subsidiary level.

Figure 8: Capital response for firms that switch reporting status.

(a) Treated group firms

(b) Control group firms

This graph only uses weighted control group means, by eweight3

Note: This figure plots the evolution of the average of the logarithm of total assets between 2000 - 2005. We show responses for both treatment and control group firms to demonstrate that the capital responds to changes in taxable status, but not to the reform. Both the treated group and the control group demonstrate similar patterns. For the years before the reform, we plot the average assets of all firms using blue solid lines. After the reform, we show the differential evolution between firms that change status from reporting zero taxable income in all years 2000 - 2002 to at least one year of positive taxable income in 2003 - 2005 in red dashed line and in green dashed line the evolution of assets for firms that do not switch their taxable income status. We use the following weighting variables in the entropy balancing match procedure: logarithm of total assets in 2001 and 2002, dummy equal to 1 if a firm reported zero taxable profit in 2000, 2001 and 2000, a dummy equal to 1 is a firm has a tax haven subsidiary in its ownership structure.

Figure 9: Effects of reducing the high-tax country tax rate by 1ppt on profit-shifting, capital accumulation, and shifting costs.


(a) Profit shifted

(b) Capital

(c) Shifting costs

Note: In this figure, we compare the trajectories before and after a one percentage point reduction in the high-tax country tax rate along the unit price of accumulating the tax avoidance asset (p_i) for: the share of profit shifted to the tax haven from the high-tax country (α), capital accumulation in the high tax country (K) and profit-shifting costs as a share of the true tax base as p_i varies. Panel (a) shows the share of profit shifted out of the high-tax country in each scenario. Panel (b) traces the total capital in the high-tax country. Panel (c) shows the ratio of fixed and variable profit-shifting costs to the total true tax base. In all cases, the productivity draw ε and cost of operation Π are fixed at their mid-point values using the parameters that we estimate in Section 4. The straight lines trace the status quo and dashed lines trace the values after the change in the tax rate. In Panel (c), the blue lines show the ratio of fixed costs to the total true tax base, black lines show the ratio of variable costs to the total true tax base.

Figure 10: Effects of raising the tax haven tax rate by 1ppt on profit-shifting, capital accumulation, and shifting costs.

(a) Profit shifted

(b) Capital

(c) Shifting costs

Note: In this figure, we compare the trajectories before and after a one percentage point rise in the tax haven tax rate along the unit price of accumulating the tax avoidance asset (p_i) for: the share of profit shifted to the tax haven from the high-tax country (α), capital accumulation in the high tax country (K) and profit-shifting costs as a share of the true tax base as p_i varies. Panel (a) shows the share of profit shifted out of the high-tax country in each scenario. Panel (b) traces the total capital in the high-tax country. Panel (c) shows the ratio of fixed and variable profit-shifting costs to the total true tax base. In all cases, the productivity draw ε and cost of operation Π are fixed at their mid-point values using the parameters that we estimate in Section 4. The straight lines trace the status quo and dashed lines trace the values after the change in the tax rate. In Panel (c), the blue lines show the ratio of fixed costs to the total true tax base, black lines show the ratio of variable costs to the total true tax base.

50

Figure 11: Impact of 15% GMT on profit-shifting, capital accumulation and shifting costs.

(a) Profit shifted

(b) Capital

(c) Shifting costs

Note: In this figure, we compare the trajectories before and after the introduction of a Global Minimum Tax at 15% along the unit price of accumulating the tax avoidance asset (p_i) for: the share of profit shifted to the tax haven from the high-tax country (α), capital accumulation in the high tax country (K) and profit-shifting costs as a share of the true tax base as p_i varies. Panel (a) shows the share of profit shifted out of the high-tax country in each scenario. Panel (b) traces the total capital in the high-tax country. Panel (c) shows the ratio of fixed and variable profit-shifting costs to the total true tax base. In all cases, the productivity draw ε and cost of operation Π are fixed at their mid-point values using the parameters that we estimate in Section 4. The straight lines trace the status quo and dashed lines trace the values after the change in the tax rate. In Panel (c), the blue lines show the ratio of fixed costs to the total true tax base, black lines show the ratio of variable costs to the total true tax base.

Figure 12: Impact of global minimum tax at varying threshold rates on capital accumulation, shifting cost and tax collection

Note: In this figure, we show the effects on our sample of MNEs of a Global Minimum Tax introduced at rates between zero and 25%, assuming that the GMT revenue is collected in the tax haven. Panel (a) shows the impact on capital investment in the high-tax and low-tax countries, as well as in aggregate. Panel (b) shows the effect on revenue collected in each country, as well as the fixed and variable costs of profit shifting in aggregate. The low-tax country tax rate is fixed at 12.5% until the minimum tax threshold reaches this level. It matches the minimum tax rate after this point.

Figure 13: Impact of global minimum tax at varying threshold rates on welfare
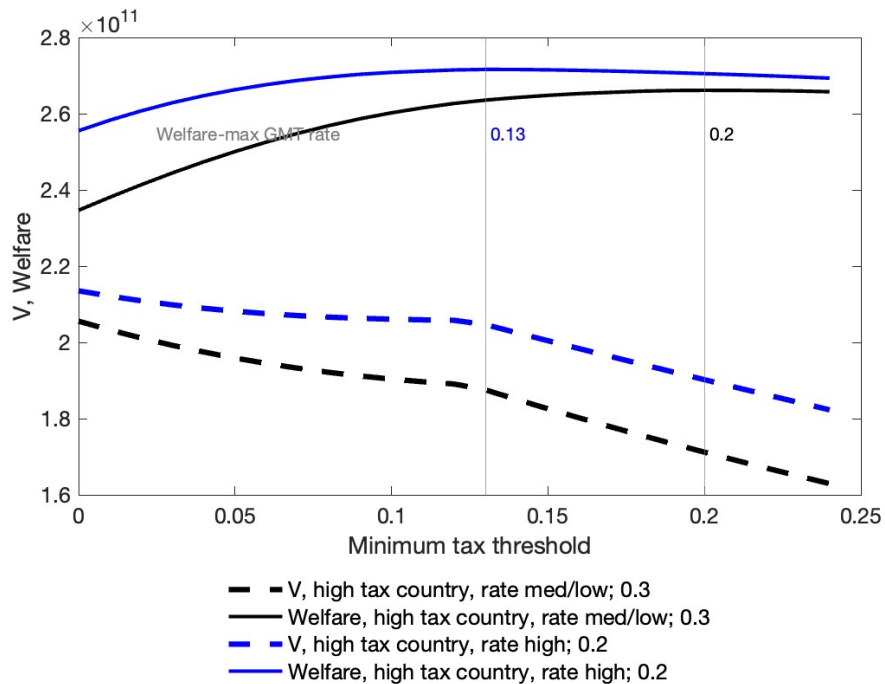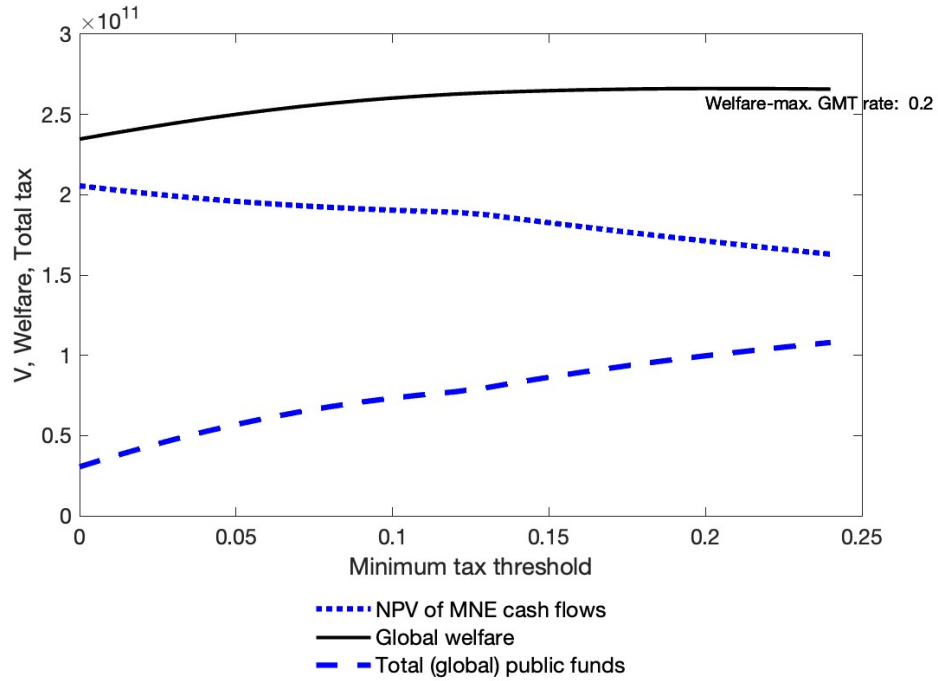
Note: In this figure, we show the effects on global welfare of a Global Minimum Tax introduced at rates between zero and 25%. Panel (a) shows the impact on the sum of NPV of cash flows for the sample of MNES, as well as aggregate tax revenue; welfare is the sum of the two. Panel (b) shows the sum of the NPV of MNE cash flows and welfare under two scenarios: where the high-tax rate country has a rate of 30%, and where it has a rate of 25%. The low-tax country tax rate is fixed at 12.5% until the minimum tax threshold reaches this level. It matches the minimum tax rate after this point.

Appendices

A Structural estimates: supplementary analyses

Figure A1: Fixed and variable costs of shifting profit to tax havens

(a) Profit shifting costs / True base (high-tax country)


(b) Profit shifting costs / capital
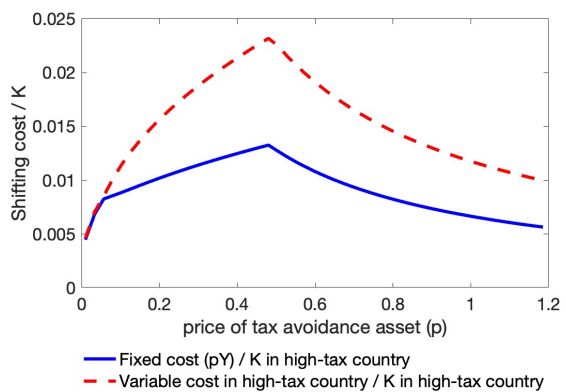
(c) Profit shifting costs / capital in high tax country

Figure A2: Effects of reducing the high-tax country tax rate by 10ppt on profit-shifting, capital accumulation and the costs of profit shifting.

(a) Profit shifted

(b) Capital

(c) Costs of shifting

Note: In this figure, we compare the trajectories before and after a 10 percentage point reduction in the high-tax country tax rate along the unit price of accumulating the tax avoidance asset (p_i) for: the share of profit shifted to the tax haven from the high-tax country (α), capital accumulation in the high tax country (K) and profit-shifting costs as a share of the true tax base as p_i varies. Panel (a) shows the share of profit shifted out of the high-tax country in each scenario. Panel (b) traces the total capital in the high-tax country. Panel (c) shows the ratio of fixed and variable profit-shifting costs to the total true tax base. In all cases, the productivity draw ε and cost of operation Π are fixed at their mid-point values using the parameters that we estimate in Section 4. The straight lines trace the status quo and dashed lines trace the values after the change in the tax rate. In Panel (c), the blue lines show the ratio of fixed costs to the total true tax base, black lines show the ratio of variable costs to the total true tax base.

# B Reduced-form evidence: supplementary analyses

Table B1: All sampled companies – key descriptive statistics by treatment status, pre-reform period

|	Treatment	Control entropy weight 3	Control entropy weight 4
Zero taxable profit (share)	0.415	0.415	0.426
Taxable income (thousand GBP)	892.507	893.209	1340.706
Taxable income (in log)	11.594	11.666	11.601
Fixed Assets (in thousand of GBP)	635.021	668.748	676.494
Fixed assets (in log)	12.427	12.605	12.606
Assets (in thousand of GBP)	2027.081	2065.89	2075.13
Assets (in log)	14.506	14.506	14.503
Observations	248	8669	8669

Note: This table shows selected descriptive statistics for a pre-reform period available in the data (2001). Control group companies are the UK subsidiaries of MNEs with parent companies located in countries outside the UK and Italy. Entropy weight 3 is an entropy balanced sample of control group firms that uses logarithm of total assets in 2000 - 2002, taxable profits in 2000 - 2002, and zero taxable profit share in 2000 - 2002 as matching variables. Entropy weight 4 is an entropy balanced sample of control group firms that uses logarithm of total assets in 2001 and 2002, and a dummy for whether the MNE has a tax haven subsidiary as matching variables. Treatment group companies are the UK subsidiaries of MNEs with parent companies located in Italy. Taxable profit data are from the tax return and data on balance sheet size are from company accounts. Units for taxable profit and asset size are nominal British Pounds.

Figure B1: Trends in average propensity to report positive taxable income, average income, and capital by treatment status: robustness to different weighting scheme.
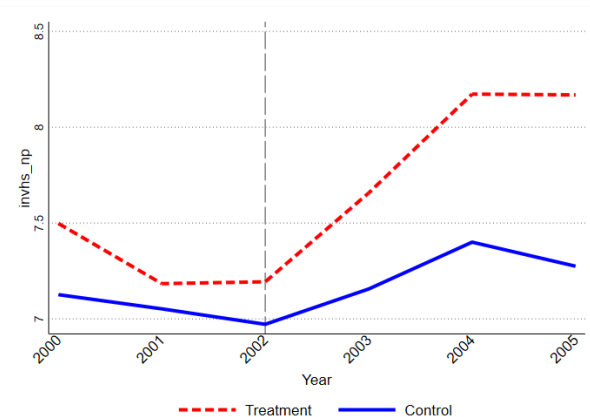
(a) Probability to report zero taxable profit

(b) Average taxable income (in level)

(c) Average taxable income (inverse hyperbolic sine)


(d) Average capital (in log)

Note: In Panel (a), we show the evolution of the average probability to report zero taxable income in the United Kingdom, in Panel (b) the evolution of the average taxable income ('Profits chargeable to Corporation Tax' in the UK corporation tax return), in Panel (c) the evolution of the inverse hyperbolic sine transformation of the average taxable income, and in Panel (d) the evolution of the logarithm of total assets. We present the trends separately for the treatment group of Italian-headquartered MNEs (red, dashed line) and for the control group of entropy-weighted subsidiaries of MNEs headquartered outside of the UK and Italy (blue, solid line). We use the following weights: logarithm of total assets in 2001 and 2002, and a dummy for whether the MNE has a tax haven subsidiary as matching variables. We demean all observations to remove individual effects and rescale the two trends to overlap in the last pre-reform period for ease of comparison. To do this, we subtract from each dot the group mean in the last pre-reform year and add back the pooled mean from the same year.

Figure B2: Trends in average propensity to report positive taxable income, average income, and capital by treatment status: against unweighted control group.

(a) Probability to report zero taxable profit

(b) Average taxable income (in level)

(c) Average taxable income (inverse hyperbolic sine)


(d) Average capital (in log)

Note: In Panel (a), we show the evolution of the average probability to report zero taxable income in the United Kingdom, in Panel (b) the evolution of the average taxable income ('Profits chargeable to Corporation Tax' in the UK corporation tax return), in Panel (c) the evolution of the inverse hyperbolic sine transformation of the average taxable income, and in Panel (d) the evolution of the logarithm of total assets. We present the trends separately for the treatment group of Italian-headquartered MNEs (red, dashed line) and for the control group of un-weighted subsidiaries of MNEs headquartered outside of the UK and Italy (blue, solid line). We demean all observations to remove individual effects and rescale the two trends to overlap in the last pre-reform period for ease of comparison. To do this, we subtract from each dot the group mean in the last pre-reform year and add back the pooled mean from the same year.

Table B2: Baseline reduced-form regression results: robustness to different weighting schemes.

| Outcome variables	(1) np	(2) invhs (np)	(3) Poisson (np)	(4) ztp	(5) log(ta)	(6) log(ta)
Panel A: Entropy balancing alternative weight						
treat × post	1260.450*	0.399	1.197***	-0.044*	0.063*	0.065*
	(683.797)	(0.270)	(0.330)	(0.023)	(0.033)	(0.037)
post × ztp pre-2002						-0.085***
						(0.015)
treat × post × ztp pre-2002						-0.018
						(0.075)
Observations	53485	53485	45044	53485	53485	53485
Panel B: unweighted control group						
treat × post	2275.633*	0.451*	1.197***	-0.051**	0.035	0.042
	(1370.390)	(0.271)	(0.330)	(0.023)	(0.033)	(0.036)
post × ztp pre-2002						-0.075***
						(0.014)
treat × post × ztp pre-2002						-0.029
						(0.078)
Observations	53558	53558	45044	53558	53558	53558
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes
Sector-year FEs	Yes	Yes	Yes	Yes	Yes	Yes

Note: This table shows the results of difference-in-difference regression estimates based on equations 13, 14, and 15. Columns (1) - (3) use as a dependent variable np, which is the level of taxable income in 1,000 of GBP ('Profits chargeable to Corporation Tax' in the UK corporation tax return). In Column (1), we estimate the model using OLS, in column (2) using inverse hyperbolic sine transformation, and in column (3) we use Poisson transformation. In column (4) the outcome variable is a dummy equal to 1, if the taxable income is equal to zero, and zero otherwise. In columns (5) and (6) the outcome variable is the logarithm of total assets. treat is a dummy equal to 1 for subsidiaries of Italian MNEs, 0 for subsidiaries of all other non-UK MNEs. post is equal to 1 is 2003 - 2005, zero before that. ztp pre-2002 is a dummy equal to 1 if a subsidiary reports zero taxable income in all years pre 2002. In Panel A, we use entropy balancing to weight the control group using the logarithm of total assets in 2001 and 2002, and a dummy for whether the MNE has a tax haven subsidiary as matching variables. In Panel B, we use unweighted control group. In all columns we include firm and sector-year fixed effects. Standard errors are clustered at the subsidiary level. We do not apply weights to Poisson regressions.

Figure B3: Capital response for firms that switch reporting status: robustness to different weighting scheme.

(a) Treated group firms (b) Control group firms

Note: This figure plots the evolution of the average of the logarithm of total assets between 2000 - 2005. In Panel (a) we show responses for treated firms and in panel (b) for control group firms. Before the reform we plot the average assets of all firms using blue solid lines. After the reform, we show the differential evolution between firms that change status from reporting zero taxable income in all years 2000 - 2002 to at least one year of positive taxable income in 2003 - 2005 in red dashed line and in green dashed line the evolution of assets for firms that do not switch their taxable income status. In Panel (b), we use the entropy weighted sample of control group firms. We use the following weighting variables: difference between total assets between 2001 and 2000 and between 2002 and 2001, difference in taxable income between 2001 and 2000 and 2002 and 2001, the level of taxable income in 2000, 2001 and 2002, the level of total assets in 2000, 2001 and 2002, the difference in the zero taxable profit reporting status between 2001 and 2000 and 2002 and 2001, the likelihood od reporting zero taxable profits in 2000, 2001 and 2002. We demean all observations to remove individual effects and rescale the two trends to overlap in the last pre-reform period for ease of comparison. To do this, we subtract from each dot the group mean in the last pre-reform year and add back the pooled mean from the same year.

Figure B4: Replication of the reduced-form difference-in-difference analysis in our structural model

Note: In this figure, we show the model fit vis-a-vis the reduced-form difference-in-difference results. The blue bars represent the difference-in-difference estimates in the data, with 95% confidence intervals. The red bars represent the model prediction for the case with 45% of MNEs

C Welfare

There are two elements of welfare in the model: the private value of the MNE and tax revenue. Consider a tax reform, such as the introduction of a global minimum tax. We model that as raising the rate of tax on income declared in the tax haven (without identifying which country receives the tax revenue). For ease of exposition, consider initially the effect of the change in haven tax rate on a subsidiary j of the MNE.

The impact on V_i of this change is:

$$\frac{dV_i}{d\tau_X} = \frac{\partial V_i}{\partial K_{ij}}\frac{\partial K_{ij}}{\partial \tau_X} + \frac{\partial V_i}{\partial Y_i}\frac{\partial Y_i}{\partial \tau_X} + \frac{\partial V_i}{\partial \alpha_{ij}}\frac{\partial \alpha_{ij}}{\partial \tau_X} - \frac{\partial T_{ij}}{\partial \tau_X} \tag{C.1}$$

For interior solutions, the first three terms of this expression are zero by the envelope theorem. If the MNE subsidiary in country H reports zero taxable income, then $\frac{\partial V_i}{\partial \alpha_{ij}} \neq 0$. However, in this case, α_{ij} will not respond to a infinitesimal change in the tax rate, so $\frac{\partial \alpha_{ij}}{\partial \tau_X} = 0$. The impact on private value is therefore just the "mechanical" effect of a change in the tax rate, given business decisions:

$$\frac{dV_i}{d\tau_X} = -\frac{\partial T_{ij}}{\partial \tau_X} = -\alpha_{ij}B_{ij} \tag{C.2}$$

The impact on the tax liability $T_{iH} = \widehat{\tau_{ij}}B_{ij}$ of this change is:

$$\frac{dT_{ij}}{d\tau_X} = \frac{\partial T_{ij}}{\partial \tau_X} + \widehat{\tau_{ij}}\frac{\partial B_{ij}}{\partial \tau_X} - (\tau_j - \tau_X)\frac{\partial \alpha_{ij}}{\partial \tau_X}B_{ij} \tag{C.3}$$

Noting that $\partial B_{ij}/\partial \tau_X = \alpha_{ij}\partial B_{ij}/\partial\widehat{\tau_{ij}}$ and $\partial \alpha_{ij}/\partial \tau_X = -\partial \alpha_{ij}/\partial(\tau_j - \tau_X)$, then in the spirit of the literature relating welfare to the elasticity of taxable income, Equation (C.3) can be rearrangedas

$$\frac{dT_{ij}}{d\tau_X} = \alpha_{ij}B_{ij}(1 + \epsilon_{ij} + z_{ij}) \tag{C.4}$$

where ϵ_{ij} is the elasticity of the true tax base, B_{ij} with respect to the effective statutory rate, $\widehat{\tau_{ij}}$ (conditional on α_{ij}), and z_{ij} is the elasticity of the proportion of the tax base shifted to the tax haven, α_{ij}, with respect to the difference in statutory rates between country j and the haven, $\tau_j - \tau_X$ (conditional on B_{ij}). This case is therefore an extension - along the lines of Devereux et al. (2014) - of the standard result that the elasticity of taxable income is a sufficient statistic for welfare analysis, see Feldstein (1995, 1999). In this case, a change in

the haven's tax rate affects not only the tax base, but also the "choice" of the overall tax rate. through the proportion of profit shifted. Welfare therefore depends on the elasticity of both the base and "effective" statutory rate, determined in this case by the extent of profit shifting.

Considering the overall effect on welfare of a change in the tax haven rate, aggregating over subsidiaries in different countries, different approaches are possible. First, a conventional option would be to assign welfare as a concave function of total tax revenue from the MNE, T_i, where $T_i = \sum_{j=1}^{N}(T_{ij})$, so that:

$$W = V_i + \beta\mu(T_i) \tag{C.5}$$

In the case in which revenue is simply distributed lump-sum back to the taxpayer, then $W = V_i + \beta T_i$. In this case transfers between the private sector and the government wash out, and welfare is equal to the weighted sum of the two elasticities identified above, where the weight is the income from each subsidiary declared in the tax haven:

$$\frac{dW}{d\tau_X} = \frac{dV_i}{d\tau_X} + \beta\sum_{j=1}^{N}\frac{dT_{ij}}{d\tau_X} = \beta\sum_{j=1}^{N}\alpha_{ij}B_{ij}(\epsilon_{ij} + z_{ij}) \tag{C.6}$$

Note that, conditional on other tax rates, there are two offsetting effects of the change in the tax haven rate. On the one hand it raises the cost of capital and hence reduces investment and output, $\epsilon_{ij} < 0$. On the other, it also reduces profit-shifting and its associated costs, $z_{ij} > 0$. The optimal rate would balance these two effects.[28]

This approach can also be used to calculate the "marginal value of public funds" (MVPF) of (Hendren and Sprung-Keyser; 2020, 2022),[29], defined in this case as:

$$MVPF = -\frac{\frac{dV_i}{d\tau_X}}{\frac{dT_{iH}}{d\tau_X}} = \frac{\sum_{j=1}^{N}\alpha_{ij}B_{ij}}{\sum_{j=1}^{N}\alpha_{ij}B_{ij}(1 + \epsilon_{ij} + z_{ij})} \tag{C.7}$$

Note that the inverse of the value of the MVPF is 1 plus the sum of the means of the two elasticities, each weighted by the tax base in the haven. The MVPF exceeds 1 when the average absolute value of ϵ_{ij} exceeds the average value of z_{ij}, and vice versa.

---

[28]Note that this approach gives no additional value to tax revenue raised from multinational profit, which may seem at odds with much of the public discourse on the global minimum tax rate.

[29]See also Finkelstein and Hendren (2020) among others

D   User cost elasticity transformation

In the case of $\alpha_{ij} = Y_i = 0$, our cost of capital is

$$v_{ij} = \frac{r}{1 - \tau_j} + \delta \tag{D.1}$$

Our semi-elasticity of K_{ij} with respect to τ_j is

$$e_{ij} = \frac{\partial K_{ij}/K_{ij}}{\partial \tau_j} \tag{D.2}$$

This implies:

$$e_{ij} = \frac{\partial K_{ij}/K_{ij}}{\partial \tau_j} = \frac{\partial K_{ij}/K_{ij}}{\partial v_i}\frac{\partial v_i}{\partial \tau_j} = \frac{\partial K_{ij}/K_{ij}}{\partial v_i}\frac{r}{(1 - \tau_j)^2} \tag{D.3}$$

And the elasticity of K_{ij} with respect to the cost of capital is x_{ij}:

$$x_{ij} = \frac{\partial K_{ij}/K_{ij}}{\partial v_i/v_i} = e_{ij}v_i\frac{(1 - \tau_j)^2}{r} \tag{D.4}$$

Using the parameters from our data, $r = 0.05$, $\delta = 0.05$, and $\tau_j = 0.3$, we obtain $v_{ij} = 0.1214$, hence $\frac{(1-\tau)^2}{r} = 9.8$ and $x_{ij} = 1.19e_{ij}$. Fore capital stock elasticity with respect to tax rate of 1.456, we obtain elasticity with respect to cost of capital of 1.73=(1.19×1.456).