Multinational Activity and the Global Minimum Tax

Paul-Emmanuel Chouc*

February 6, 2025

Abstract

Since 2021, over 130 jurisdictions have committed to a 15% global minimum tax on the profits of large multinational companies. This reform is expected to raise additional tax revenues, but little is known about potential behavioral responses. The rise in corporate tax rates and changes in tax differentials may spur targeted firms to scale down operations or adapt location choices. We examine these effects in a general equilibrium model of multinational production with corporate income taxation. Firms jointly choose production locations and whether to adopt a costly profit shifting technology. These discrete choices shape firms' responses to taxation and create heterogeneous exposure to the global minimum tax. Preliminary simulations suggest that the reform may lead to significant real adjustments, reducing private income across all countries. However, behavioral responses seem to affect the revenue potential of the reform only marginally. Net welfare effects depend on the relative valuation of private income and government revenues. Further work will expand the quantification to a richer setting and explore additional features of the global minimum tax.

*CREST, Palaiseau, France. Contact: paul-emmanuel.chouc@ensae.fr. I am grateful to my advisors, Clément Malgouyres and Gabriel Zucman, for their constant eagerness to discuss this ongoing project. I also thank Mathieu Parenti, as well as participants to the CREST Internal Seminar, HEC-CREST PhD Day, and EUTax Internal Workshop, for rich comments. I could make important progress thanks to Rowan Shi's generous help. I acknowledge a full doctoral fellowship from Institut polytechnique de Paris (2023-2026).

1 Introduction

The expansion of multinational firms in recent decades has posed major challenges for tax administrations. Greater economic integration has intensified competition among countries to attract real multinational activity. Scholars and other observers often cite tax competition to explain the sustained decline in corporate income tax rates. Typically, between 1985 and 2020, the global average statutory tax rate decreased from 49% to 23.5%.[1] Profit shifting also enables multinational companies to book taxable income in low-tax jurisdictions, even though they maintain little or no real presence there. These practices have also become a key determinant of the effective taxation of corporate profits over the 2000s and have been widely documented in the literature.[2]

Since October 2021, after years of OECD and G20-led negotiations, 139 jurisdictions have agreed on the "two-pillar solution to address the tax challenges arising from the digitalisation of the economy". In particular, Pillar Two introduces a global minimum tax of 15% on the corporate profits of large multinational companies—above a consolidated revenue threshold of 750 million EUR. It applies on a country-by-country basis such that low-taxed profits in a single jurisdiction may trigger a top-up payment, even though the firm-wide average effective tax rate is above 15%. The distribution of these tax revenues depends on how the agreement is implemented in practice. As of January 2025, 58 jurisdictions have transposed Pillar Two in national law, though a few EU Member-States have delayed application, and US adoption remains unlikely.[3] In principle, the adopters can impose the new rules onto all the multinationals that originate from a signatory jurisdiction, but whether this may extend to US firms without retaliation remains unclear.

This reform is expected to increase tax revenues, both through direct top-up payments and by reducing profit shifting. Indeed, by curtailing tax rate differentials between high-tax and lower-tax jurisdictions, the minimum rate makes profit shifting less attractive. Hugger et al. (2024), who try and account for both mechanisms, estimate that Pillar Two should generate annual gains of 155-192 billion USD globally. However, three key uncertainties remain. First, most revenue gain estimates ignore the potential repercussions on multinationals' real economic activity. If Pillar Two substantially modifies real decisions, this will shape revenue gains and the overall impact on welfare. Second, the reform specifically targets the largest firms globally, and general equilibrium effects may not be negligible.[4]

1. Average statutory corporate income tax rate weighted by GDP, based on the OECD's Corporate Tax Statistics and cited by Clausing, Saez, and Zucman (2020) or Janeba and Schjelderup (2022). Beyond headline rates, the US effective average tax rate (EATR) on corporate profits went down from close to 50% in the 1950s to 17% in 2018, and it declined from 30% to around 20% over the last two decades in the European Union. See Piketty, Saez, and Zucman (2018) for the US and Spengel et al. (2020) for the EU.

2. See Heckemeyer and Overesch (2017) or Beer, Mooij, and Liu (2020) for related literature reviews.

3. Based on the EY BEPS 2.0 Pillar Two Developments Tracker last consulted in January 2025. The EU directive, which was adopted in 2022, allows Member-States to delay implementation until 2030.

4. For instance, in France, the domestic and foreign multinationals in the scope of Pillar Two gathered 4.8 million employees in 2019. This figure comes from the OECD's aggregated tabulation of country-by-country

These effects, in turn, influence tax revenues and overall welfare. Third, partial adoption will likely influence the outcome of the reform. It critically hinges on whether implementing jurisdictions can collect top-up taxes from US-headquartered multinationals, and on possible retaliation by the upcoming administration.

Given these uncertainties, this paper examines how multinational companies respond to a global minimum tax modeled after Pillar Two. What are the expected effects on firms' location choices and real activity? How should aggregate activity respond and how would its geographical distribution be affected? Can we benchmark these responses with the effects of other policy instruments and characterize potential interactions?

Related empirical results—primarily estimates of the elasticity of profits or real variables to tax rates—cannot apply directly. Indeed, we consider a non-marginal change in tax rates that affects many jurisdictions simultaneously, by as much as 15 percentage points for some. We thus develop a quantifiable, general equilibrium theory of firms' relevant decisions. Our model draws from Tintelnot (2016) and Arkolakis, Eckert, and Shi (2023). The key specificity of this approach within the multinational production literature lies in the fixed cost for setting up a foreign affiliate. This creates an extensive margin of firms' multinational expansion, which allows to study the impact of a given reform on location choices but comes at the cost of non-tractability over a part of the profit maximization program. Compared to previous studies, our model captures how heterogeneous production networks shape firms' responses to tax changes and create varying exposure to the global minimum tax.

We extend the theory along three directions. First, we add corporate income taxation. Even a tax on pure profits distorts firms' decisions. Any positive tax rate reduces variable profits, makes it more difficult for the firm to afford the fixed cost of a foreign affiliate, and thus hinders multinational expansion. Given a production network, tax rate differentials also matter for the distribution of real activities across plants. Second, we introduce profit shifting. Our approach makes real and shifting responses to taxation substitutes, similar to Slemrod (1992)'s hierarchy of responses: When profit shifting is cheap, firms primarily respond to a tax hike via accounting manipulations; when profit shifting is less profitable, firms' real activities become more sensitive to taxation. Third, we include a second factor of production, tangible capital, to model specific features of Pillar Two.

We calibrate the model in a simplified setting with five locations. The initial equilibrium tracks well the patterns observed in multinational production statistics, and the estimation yields an exact match with the taxation-related moments that we target. We start by illustrating the main theoretical results. Then, we introduce a global minimum tax of 15% as a counterfactual scenario. The reform affects the multinationals engaged in profit shifting and / or active in Hungary, whose statutory tax rate lies below the minimum. Before general equilibrium adjustments, global multinational activity contracts, and the

reports. It represents roughly 18% of total employment.

effect is concentrated in the US and in Hungary. At the new general equilibrium, changes in factor prices partly offset the direct impacts. The reform reduces real private income in all countries, especially in the Netherlands where it drops by close to 1%. But it also brings revenues, as global corporate income tax collection increases by 5.6%. The net effect on welfare thus depends on the relative valuation of private income and government revenues.

Section 2 reviews the literature to which this work is contributing. Section 3 describes our conceptual framework. Section 4 states our methodology to estimate and simulate the model, and provides the main results. Section 5 sketches ongoing work to build a credible quantitative model of multinational production and taxation. Section 6 concludes.

2 Literature

Global minimum tax. This work is related to previous research on the global minimum tax on corporate income. The two closest analyses by Ferrari et al. (2023) and Dyrda, Hong, and Steinberg (2023) jointly consider real and nominal responses in a quantitative model of firms' behavior. This paper extends the set of real responses to taxation by integrating recent advances from the multinational production literature. First, we allow firms to structure production through multiple affiliates, which not only serve local consumption but also act as export platforms to foreign markets. The location of these affiliates and the allocation of real activities across them depend on corporate income tax rates and profit shifting opportunities. Second, we incorporate fixed costs for establishing foreign affiliates. These fixed costs create an extensive margin in the response to tax changes and result in heterogeneous production networks across firms. This heterogeneity shapes the elasticity of real activities to taxation and leads to differential exposure to the global minimum tax. Parallel to these refinements on the multinational production side, we opt for a reduced-form view of profit shifting that can capture salient empirical patterns.

Public finance. This paper also contributes to the broader public finance literature on the taxation of multinational companies. First, we formalize the substitutability of real and nominal responses to taxation, as articulated by Slemrod (1992), in a quantifiable framework. The elasticity of real activities to the corporate income tax rate increases in the cost of profit shifting. Second, we show that the magnitude of a multinational company's real responses to taxation in a given location depends on the size and distribution of its affiliate network. Third, the estimation of the model may refine certain benchmark quantitative findings about the prevalence and heterogeneity of profit shifting practices.

Multinational production. Finally, we contribute to the international trade literature on multinational production. Our extended model allows us to benchmark the effects of

corporate income taxes and profit shifting regulations against other policy instruments, such as trade barriers. Their interaction is also of high policy relevance. The ability of implementing countries to impose the Pillar Two rules on US multinational companies will likely depend on tariff-based threats of retaliation, similar to the approach used against Digital Service Taxes (DSTs) in prior years. We are also building an updated database of bilateral multinational production flows to follow up on Ramondo, Rodríguez-Clare, and Tintelnot (2015) or Alviarez (2019).

3 Conceptual Framework

3.1 Environment

Preferences. The world economy comprises N countries. Each country j is endowed with a mass L_j of immobile workers who inelastically supply one unit of labor and one unit of tangible capital. Following Tintelnot (2016), workers derive their utility by consuming varieties of differentiated final goods. Preferences are identical across countries with a constant elasticity of substitution (CES) between varieties and goods, denoted σ. Formally, the utility from consuming a quantity $q_j(\omega, \nu)$ of each variety ν of each good ω writes as the following nested aggregator:

$$U^j = [\int_{\Omega_j} (\int_0^1 [q_j(\omega,\nu)]^{\frac{\sigma-1}{\sigma}} d\nu) d\omega]^{\frac{\sigma}{\sigma-1}}$$

Where Ω_j denotes the set of goods available for consumption in country j. These preferences give the usual CES demand function. With Y_j the aggregate income and P_j the ideal price index in country j, the demand for variety ν of good ω is given by:

$$q_j(\omega,\nu) = [p_j(\omega,\nu)]^{-\sigma} \frac{Y_j}{P_j^{1-\sigma}} \tag{1}$$

However, because we have two "layers" of differentiation—a continuum of goods and, for each good, a continuum of varieties, the price index is a composite object:

$$P_j = (\int_{\Omega_j} [p_j(\omega)]^{1-\sigma} d\omega)^{\frac{1}{1-\sigma}}$$
$$\text{where } p_j(\omega) = (\int_0^1 [p_j(\omega,\nu)]^{1-\sigma} d\nu)^{\frac{1}{1-\sigma}}$$

Firm characteristics. An exogenous mass M_i of firms originates from each headquarter country i. Each firm from i is characterized by a core productivity scalar. It is associated with a unique differentiated good ω, and a unit continuum of varieties indexed by ν.

Production. Producing one unit of variety ν requires two inputs. In production location l, labor is paid the wage rate w_l, and the rental rate of tangible capital is r_l. For now, we assume a Cobb-Douglas production function with headquarter-specific input shares α_i for labor and $\beta_i = 1 - \alpha_i$ for capital. The variety-location-specific productivity factor $z_l(\nu)$ is stochastic, with a distribution specified below. By employing m units of labor and k units of capital, the firm produces the following quantity of variety ν:

$$F_{il\nu}(m,k) = z_l(\nu)\, m^{\alpha_i} k^{\beta_i} = z_l(\nu)\, m^{\alpha_i} k^{1-\alpha_i}$$

The function $F_{il\nu}$ is homogeneous of degree 1, which implies constant returns to scale. This assumption is critical because, given a set of foreign affiliates, it allows the firm to choose a production location, a price, and input quantities separately for each final destination market.[5] For this reason, constant returns to scale are common in the multinational production literature where, most often, only labor is used in production. We isolate the deterministic components of the unit cost of production in c_{il}:

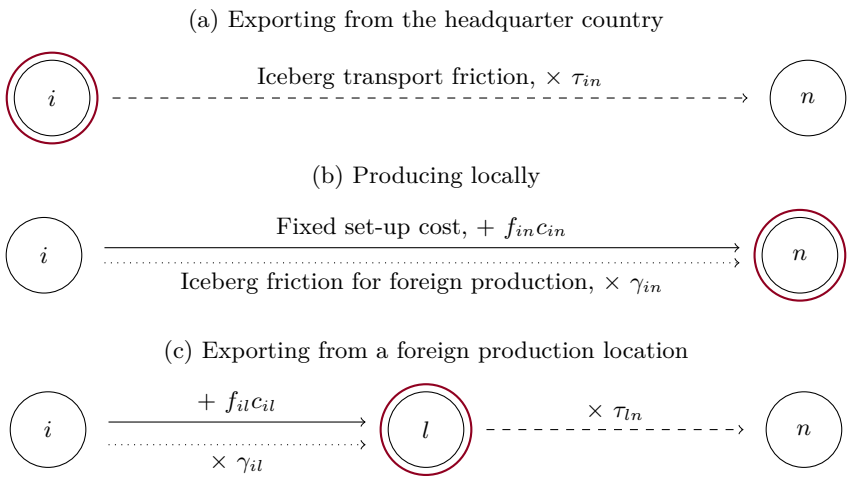
$$c_{il} = \left(\frac{w_l}{\alpha_i}\right)^{\alpha_i} \left(\frac{r_l}{1-\alpha_i}\right)^{1-\alpha_i} \tag{2}$$

We allow for multinational production. Firms are endowed with a production facility in their headquarter country, and they can establish foreign affiliates. All affiliates can sell locally or export into foreign consumption markets. Figure 1 illustrates the three ways in which a firm originating from i can serve final destination market n. Panel 1a shows the usual assumption of international trade models, in which the firm produces in its headquarter country and exports towards the final destination market. The only cost involved is the iceberg friction $\tau_{in} \geq 1$ meaning that, to sell one unit in country n, the firm must produce τ_{in} units in i. Panel 1b describes that firms can also set up a production affiliate in the final destination market to produce and sell the goods locally. The firm must then incur two costs. First, establishing the plant imposes a fixed cost, expressed as a mix of labor and tangible capital in the location of production.[6] Second, production outside of the headquarter country implies iceberg frictions $\gamma_{in} \geq 1$ reflecting increased monitoring duties, frictions in technology transfers, etc. Panel 1c illustrates the scenario where the firm establishes a foreign export platform, from which it serves market n remotely. All types of fixed and variable costs then arise.

5. Oppositely, consider diminishing returns to scale (i.e., $\alpha + \beta < 1$) for instance. Serving market n from production location l would raise the stock of capital and labor employed by the firm in l. This would in turn reduce the profits obtained by serving market n' from production location l. Sourcing decisions could not be made separately from one another.

6. The mix of inputs that must be employed in the production location to set up the plant follows the input shares in the firm's Cobb-Douglas production function, which themselves depend on the firm's headquarter country. With this assumption, the monetary value of the fixed cost is equal to the required quantity of inputs, f_{in}, times the deterministic component of the firm's unit cost in n, c_{in}.

Figure 1: Illustrating multinational production



Notes: The red circle indicates the location where production occurs. The dashed lines indicate (variable) iceberg transportation costs; the dotted lines indicate (variable) iceberg frictions for production outside of the headquarter country; the solid lines indicate the fixed costs for setting up foreign production affiliates. The following assumption remains implicit in the schema: for any country j, we assume that $\tau_{jj} = 1$ and $\gamma_{jj} = 1$.

Taxation. Consider the general case of a firm that originates from country i and serves final destination market n with variety $\nu \in [0, 1]$ through its affiliate in production location l. By default, taxable income is booked where production occurs and is subject to a linear corporate income tax with rate t_l. For now, we assume that input costs are fully deductible so that taxable income and the pre-tax profits perceived by the firm are confounded.

A profit shifting technology allows to relocate taxable income into a tax haven, where it is taxed at rate t^{PS}. This technology involves two types of costs. First, a fixed cost for adoption rationalizes the apparent concentration of complex avoidance practices among the largest multinationals. We denote by ${f_i}^{\text{PS}}$ the quantity of labor and tangible capital that must be employed in the firm's headquarter country to adopt the technology.[7] Second, profit shifting also involves variable costs, linear in total taxable income π^{PT} and quadratic in the share s being relocated:

$$\frac{\delta_i}{2}\pi^{\text{PT}}s^2$$

This convexity assumption, common in the literature, formalizes the idea that the probability of detection increases with tax planning aggressiveness. The parameter δ_i governs the magnitude of variable costs.[8] While the fixed cost for profit shifting is expressed as a

7. Like for the fixed cost of setting up foreign affiliates, the mix of inputs that must be employed in the headquarter country to adopt the profit shifting technology follows the input shares in the firm's Cobb-Douglas production function. With this assumption, the monetary value of the fixed cost is equal to the required quantity of inputs, f_{ii}^{PS}, times the deterministic component of the firm's unit cost in i, c_{ii}.

8. Both fixed and variable costs are parametrized to depend on the firm's headquarter country. This is particularly helpful to replicate Tørsløv, Wier, and Zucman (2022)'s finding that multinationals originating

quantity of inputs employed to adopt the technology and thus accrues to aggregate income, variable costs are modeled as a pure loss of resources.

The optimal share of income shifted depends on the tax rate differential between the production location and the tax haven, and on the headquarter-specific parameter:

$$s^* = \frac{t_l - t^{\text{PS}}}{\delta_i} \text{ if } t^{\text{PS}} < t_l, \text{ 0 otherwise}$$

Then, the firm's effective net-of-tax rate in l writes as:

$$\kappa_{il}^{\text{PS}} = \begin{cases} 1 - t_l + \frac{\left(t_l - t^{\text{PS}}\right)^2}{2\delta_i} & \text{if } t^{\text{PS}} < t_l \\ 1 - t_l & \text{otherwise} \end{cases} \tag{3}$$

Instead, if the firm does not adopt the profit shifting technology, its net-of-tax rate is simply given by $\kappa_{il}^{\text{no PS}} = (1 - t_l)$ in every production location l. As we show below, this effective net-of-tax rate is the relevant tax parameter that affects firms' choices.

3.2 Firms' Program

Timing. Figure 2 summarizes the timing of the firm's problem. Steps (2), (4), and (5) involve discrete or continuous choices; Steps (1) and (3) only mark changes in the firm's information set and endowments. We solve this problem backwards in the rest of this section.

Step (5): Continuous choices. At this step, the firm has full information and has already made discrete choices. Consider variety ν served into market n. Production takes place in $l^*(\nu, n) \equiv l^*$, with total factor productivity $z_{l^*}(\nu)$. We write the firm's effective net-of-tax rate in l^* as κ_{il^*}. If the firm has adopted the profit shifting technology, then it is given by Equation 3; otherwise, it is simply equal to $1 - t_{l^*}$. The final step in the profit maximization program consists in choosing a price and inputs to maximize after-tax profits:

$$\max_{p,m,k} \left(\kappa_{il^*} \left[pD_n(p) - w_{l^*} m - r_{l^*} k\right]\right)$$
$$\text{subject to } F_{il^*\nu}(m,k) = z_{l^*}(\nu)\, m^{\alpha_i} k^{1-\alpha_i} \geq D_n(p)\, \tau_{ln}$$

Where the demand function is given by Equation 1. We raise two remarks. First, maximizing after-tax profits comes down to maximizing the pre-tax ones. Corporate income taxation affects neither the choice of inputs, nor the pricing decision. This derives from our assumptions that (i) input costs are fully deductible from taxable income and (ii) the variable costs of profit shifting are linear in taxable income. Second, constant returns to scale allow

from the US tend to shift more aggressively that their non-US counterparts. We discuss this approach in more details in Section 4.

1. Entry and initial characteristics: *The firm originates from headquarter country i and draws a core productivity scalar φ. It is endowed with a plant in i.*
2. Discrete choices: The firm jointly decides on foreign affiliates and profit shifting:
   (a) It chooses any sub-set Z of the set of locations $[\![1; N]\!] \setminus \{i\}$. For each l in Z, it incurs the fixed cost f_{il}, expressed as a quantity of labor and tangible capital employed in l. The corresponding monetary cost is $f_{il}c_{il}$.
   (b) To adopt the profit shifting, the firm employs a fixed quantity f_i^{PS} of labor and tangible capital in i. It then benefits from reduced effective net-of-tax rates $\{\kappa_{il}\}_{l\in Z\cup\{i\}}$. Otherwise, net-of-tax rates are given by $\{1-t_l\}_{l\in Z\cup\{i\}}$.
3. Location-variety-specific productivity draws: *For each variety ν and each production location l in Z, the firm draws a total factor productivity $z_l(\nu)$.*
4. Mapping production locations and markets: In the absence of fixed costs for exporting, the firm serves all final destination markets. For each variety ν and each market n, the firm selects a production location $l^*(\nu, n)$ in its set of domestic and foreign plants $Z \cup \{i\}$ from which it will sell.
5. Continuous choices: The firm maximizes the after-tax variable profits from selling each variety ν in market n from production location $l^*(\nu, n)$. It chooses employment, tangible capital, the price, and—conditionally on having adopted the technology —a share of taxable income shifted into the tax haven.

Figure 2: Overview of timing in the firm's program

to rewrite the input costs at optimum as proportional to a unit cost. With c_{il^*} defined in Equation 2, we rewrite the profit maximization program as:

$$\max_p \left[pD_n(p) - \frac{c_{il^*}\tau_{l^*n}}{z_{l^*}(\nu)} D_n(p) \right]$$

The optimal price writes as a constant mark-up over the marginal cost of production:

$$p^* = \frac{\sigma}{\sigma-1} \frac{c_{il^*}\tau_{l^*n}}{z_{l^*}(\nu)} \tag{4}$$

Given this price, the firm's after-tax profits at optimum write as:

$$\pi^* = \frac{\kappa_{il}}{\sigma} \left[\frac{\sigma}{\sigma-1} \frac{c_{il^*}\tau_{l^*n}}{z_{l^*}(\nu)} \frac{1}{P_n} \right]^{1-\sigma} Y_n$$

Step (4): Mapping production locations and markets. We move one step back in the profit maximization program. At this stage, for each variety ν, the firm matches market n with a production location $l^*(\nu, n)$. It chooses this production location to maximize the

after-tax profits that we have derived right above:

$$l^{*}(\nu, n) = \arg\max_{l} \left(\frac{\kappa_{il}}{\sigma} \left[\frac{\sigma}{\sigma - 1} \frac{c_{il}\tau_{ln}}{z_{l}(\nu)} \frac{1}{P_{n}} \right]^{1-\sigma} Y_{n} \right) = \arg\min_{l} \left[\frac{\kappa_{il}^{\frac{1}{1-\sigma}} c_{il}\tau_{ln}}{z_{l}(\nu)} \right]$$

Step (3): Location-variety-specific productivity draws. We now characterize the distribution of location-variety-specific productivity draws. Suppose a firm has opened a set Z of foreign production locations, adding to the domestic facility with which it is endowed upon entry. For each production location l in $Z \cup \{i\}$ and each variety ν, the firm draws a productivity parameter $z_{l}(\nu)$ from the following cumulative distribution function:

$$\mathbb{P}\left[z_{l}(\nu) \leq z\right] = exp\left[-\left(\varphi\epsilon_{l}\right)^{\theta} \left(\gamma_{il} z\right)^{-\theta} \right] \tag{5}$$

This Fréchet law follows the one laid out by Tintelnot (2016). For now, we assume that the productivity draws are independently distributed across affiliates and varieties. The shape parameter θ governs the dispersion of the productivity draws across countries and varieties. A higher value leads to less heterogeneous productivity draws, and thus to more substitutable production locations. A higher scale parameter $\varphi\epsilon_{l}/\gamma_{il}$ translates into better productivity draws overall. It can be decomposed into core productivity φ and the location-specific shifter ϵ_{l}, that both yield better productivity draws, and the iceberg friction γ_{il} for foreign production which weighs down on productivity draws.

Tintelnot (2016) imposes $\theta > \sigma - 1$ as a technical assumption. Arkolakis, Eckert, and Shi (2023) highlight this inequality, satisfied in their benchmark calibration. They describe two types of complementarities in firms' choice of production locations. Negative complementarities result from cannibalization effects between affiliates, and the parameter θ governs this substitutability. Positive complementarities arise due to "scale" effects. Opening a new production location essentially comes down to trading off a fixed cost with reduced variable costs. The latter outweigh the former when applied to a sufficiently large volume of sales. Therefore, a firm with an extensive network of affiliates, which benefits from relatively low marginal costs and high sales, has an incentive to open further production locations. The strength of the scale effects depends on the elasticity of revenues to marginal cost, $\sigma - 1$. Assuming $\theta > \sigma - 1$ means that substitutability forces dominate complementarities.

Consider a given variety ν served to final destination market n. We derive the probability that l is selected as the production location for these transactions. Applying the law of large numbers to the continuum of varieties, the probability can also be interpreted as the share

of production location l in the varieties served to market n. Denoting it as ζ_{iln}, we have:

$$\zeta_{iln} = \frac{\epsilon_l^\theta \left(\kappa_{il}^{\frac{1}{1-\sigma}} c_{il} \gamma_{il} \tau_{ln} \right)^{-\theta}}{\sum_{j \in Z \cup \{i\}} \epsilon_j^\theta \left(\kappa_{ij}^{\frac{1}{1-\sigma}} c_{ij} \gamma_{ij} \tau_{jn} \right)^{-\theta}} \tag{6}$$

We then derive the distribution of the marginal costs of production with which the firm will serve final destination market n. With a constant mark-up, we deduce the distribution of the price at which variety ν is sold in n. Once we have this distribution, we can compute the local firm-level price index and we obtain the firm's sales. Among those, sales x_{iln} are operated from production location l:

$$\begin{aligned} x_{iln} &= A \varphi^{\sigma-1} Y_n P_n^{\sigma-1} \Psi_{i,n}^{\frac{\sigma-1}{\theta}} \zeta_{iln} \kappa_{il}^{-1} \\ \text{with } \Psi_{i,n} &= \sum_{j \in Z \cup \{i\}} \epsilon_j^\theta \left(\kappa_{ij}^{\frac{1}{1-\sigma}} c_{ij} \gamma_{ij} \tau_{jn} \right)^{-\theta} \end{aligned} \tag{7}$$

The term $\Psi_{i,n}$ may be interpreted as the firm's capacity to serve market n, and A is a constant that depends on parameters common to all locations. In our model, the value share of production location l in the sales to market n differs from its share in varieties. Corporate income taxation drives a wedge between the quantity that the firm minimizes when selecting a production location and the marginal cost of production. Then, the conditional distribution of the marginal cost varies across affiliates. If l has a relatively low corporate income tax rate, it may indeed be selected as production location despite a relatively high marginal cost of production or transportation.

The firm's variable profits are then given by:

$$\pi(i, \varphi, Z, \mathbb{1}^{\text{PS}}) = \frac{1}{\sigma} A \varphi^{\sigma-1} \sum_{n=1}^{N} P_n^{\sigma-1} Y_n \Psi_{i,n}^{\frac{\sigma-1}{\theta}} \tag{8}$$

Discrete choices affect variable profits through the market access term $\Psi_{i,n}$, explicit in Equation 7. First, as it sums a positive term for each production location, this object is unambiguously increasing in the cardinal of Z. With a wider set of affiliates, the firm has a greater capacity to serve each market, and draws larger variable profits. Second, adopting the profit shifting technology, which we indicate with the dummy $\mathbb{1}^{\text{PS}}$, weakly raises all effective net-of-tax rates κ_{il}. This limits the distortive effect of corporate income taxation in the mapping between production locations and markets, so that the firm "focuses on" minimizing the marginal cost of production. This results in a lower marginal cost, higher sales, and ultimately higher variable profits.

Step (2): Discrete choices. The firm decides on a set of foreign affiliates and whether to adopt profit shifting to maximize total profits net of all fixed costs. This step writes as:

$$\max_{Z \subset [\![1;N]\!] \setminus \{i\}, \mathbb{1}^{\text{PS}}} \left[\pi(i, \varphi, Z, \mathbb{1}^{\text{PS}}) - \sum_{j \in Z} f_{ij} c_{ij} - \mathbb{1}^{\text{PS}} f_i^{\text{PS}} c_{ii} \right] \tag{9}$$

The choice of production locations forms a combinatorial discrete choice problem. The profit shifting decision further interacts with this choice. The brute force approach would compare net profits for the 2^{N+1} possible combinations. However, for given effective net-of-tax rates, the algorithm by Arkolakis, Eckert, and Shi (2023) exactly solves for the optimal set of foreign affiliates while remaining much more computationally efficient. A straightforward extension of their algorithm derives the optimal set of production locations with and without profit shifting, and retains the option that yields the highest net profits.

Step (1): Entry and initial characteristics. So far, we have reasoned conditionally on core productivity and a headquarter country. In their analysis of the choice of production locations, Arkolakis, Eckert, and Shi (2023) go one step further and solve for the "policy function": $\mathcal{L}$ maps firm characteristics, i and φ, into the optimal set of affiliates. They show that it is piece-wise constant with respect to core productivity, and their algorithm computes the cutoff productivity levels alongside the corresponding choices. In our case, there are two functions. In addition to $\mathcal{L}$, $\mathcal{PS}$ maps firm characteristics into an indicator variable for profit shifting adoption. Again, we extend the algorithm by Arkolakis, Eckert, and Shi (2023) to account for the profit shifting decision.

3.3 General Equilibrium

In this section, we aggregate firms' choices to define general equilibrium outcomes and the corresponding conditions. To lighten the notations, we ignore that $\mathcal{L}$ and $\mathcal{PS}$ are functions of the firm's headquarter country and core productivity.

Aggregate sales. It is useful to write down the total sales operated between production location l and market n by firms that originate from headquarter country i. Integrating Equation 7 over the distribution G^i of core productivity yields the expected sales of a single firm. Because a continuum of firms originates from each headquarter country, the law of large numbers applies and we simply multiply by the mass of entrants, M_i, to obtain aggregate sales.

$$X_{iln} = AM_i \frac{Y_n}{P_n^{1-\sigma}} \int_{\varphi} \varphi^{\sigma-1} \left[\Psi_{i,n}\left(\mathcal{L}, \mathcal{PS}\right)\right]^{\frac{\sigma-1}{\theta}} \zeta_{iln}\left(\mathcal{L}, \mathcal{PS}\right) \left[\kappa_{il}\left(\mathcal{PS}\right)\right]^{-1} dG^i\left(\varphi\right) \tag{10}$$

Labor market clearing. In each production location l, the labor market clearing condition equates workers' aggregate labor income and firms' total labor expenditures. Labor income is the product of the wage rate w_l and total labor supply L_l. Labor expenditures sum three components.

$$\begin{aligned} w_l L_l = & \frac{\sigma-1}{\sigma} \sum_{i=1}^{N} \sum_{n=1}^{N} \alpha_i X_{iln} \\ & + \sum_{\substack{i=1 \\ i \neq l}}^{N} \alpha_i M_i \int_{\varphi} \mathbb{1}^{l \in \mathcal{L}(i,\varphi)} \; f_{il} c_{il} \; dG^i(\varphi) \\ & + \alpha_l M_l \int_{\varphi} \mathcal{PS}(l,\varphi) \; f_l^{\mathrm{PS}} c_{ll} \; dG^l(\varphi) \end{aligned} \tag{11}$$

First, labor is employed in production. For a firm headquartered in i, labor represents a share α_i of variable costs, and variable costs are themselves a share $\frac{\sigma-1}{\sigma}$ of the revenues—which corresponds to the inverse of the mark-up. Additionally, labor is employed by foreign multinationals to establish their affiliates in l. Each firm from i entering the production location incurs a fixed cost $f_{il}c_{il}$, with a share α_i of labor expenditures. The second term in the right-hand side of Equation 11 aggregates these expenditures over all firms from foreign headquarter countries. Finally, labor is employed by the domestic multinationals that adopt the profit shifting technology. These incur a fixed cost $f_{ll}^{\mathrm{PS}}c_{ll}$, with a share α_l of labor expenditures. The third term aggregates these expenditures over all the firms originating from l that adopt the profit shifting technology.

Capital market clearing. The market clearing condition for tangible capital equates workers' aggregate capital income and firms' total rental expenditures. Capital income is the product of the rental rate r_l and the total supply L_l of tangible capital. Rental expenditures have the same structure as for labor, with input shares $(1-\alpha_j)$.

$$\begin{aligned} r_l L_l = & \frac{\sigma-1}{\sigma} \sum_{i=1}^{N} \sum_{n=1}^{N} (1-\alpha_i) X_{iln} \\ & + \sum_{\substack{i=1 \\ i \neq l}}^{N} (1-\alpha_i) M_i \int_{\varphi} \mathbb{1}^{l \in \mathcal{L}(i,\varphi)} \; f_{il} c_{il} \; dG^i(\varphi) \\ & + (1-\alpha_l) M_l \int_{\varphi} \mathcal{PS}(l,\varphi) \; f_l^{\mathrm{PS}} c_{ll} \; dG^l(\varphi) \end{aligned} \tag{12}$$

Price index. The price index in each final destination market n aggregates all firm-level price indices. The firm-level price index depends not only on the firm's capacity to serve market n, but also on the wedge that corporate income taxation drives between the marginal cost of production and the objective of the firm when mapping production locations. We integrate it over the headquarter country's core productivity distribution, multiply by the mass of firms, and sum over headquarter countries.

$$\begin{aligned} P_n^{1-\sigma} &= \sum_{i=1}^{N} M_i \int_{\varphi} \left[p_n\left(i,\varphi,\mathcal{L},\mathcal{PS}\right)\right]^{1-\sigma} dG^i\left(\varphi\right) \\ \text{where } \left[p_n\left(i,\varphi,\mathcal{L},\mathcal{PS}\right)\right]^{1-\sigma} &= A\left[\Psi_{in}\left(\mathcal{L},\mathcal{PS}\right)\right]^{\frac{\sigma-1}{\theta}} \sum_{l\in\mathcal{L}(i,\varphi)\cup\{i\}} \frac{\zeta_{iln}\left(\mathcal{L},\mathcal{PS}\right)}{\kappa_{il}\left(\mathcal{PS}\right)} \end{aligned} \tag{13}$$

Balance of payments. Eventually, the balance of payments equates total consumption expenditures with aggregate private income. Aggregate private income in country i sums local workers' labor and capital income with the total after-tax profits, net of all fixed costs, of firms originating from i. We thereby assume that firms are fully and uniformly owned by their headquarter country's residents. We also assume that government revenues are valued as such by households, and not spent on consumption goods.

$$Y_i = w_i L_i + r_i K_i + M_i \int_{\varphi} \Pi\left[i,\varphi,\mathcal{L}\left(i,\varphi\right)\right] dG^i(\varphi) \tag{14}$$

Definition. Similarly to Arkolakis, Eckert, and Shi (2023), a general equilibrium is formed from two elements. First, policy functions, $\{\mathcal{L},\mathcal{PS}\}$, map firm characteristics into production locations and a profit shifting decision. Second, four vectors of aggregates, $\{w_i, r_i, P_i, Y_i\}_i$, respectively stand for wages, rental rates of tangible capital, price indices, and aggregate consumption expenditures. These two elements must be so that:

- Given the aggregates, policy functions solve the optimization problem of Equation 9;
- Given policy functions, the aggregates satisfy Equations 11, 12, 13, and 14.

4 Simulations

4.1 Methodology

Setting. We simulate the model in a simplified environment with three large high-tax countries—the US, Germany, and France —, one tax haven—the Netherlands —, and one low-tax non-haven country—Hungary. We assume that the Netherlands is the only profit shifting destination, and shifted profits are initially not taxed. That is, $t^{\text{PS}} = 0$.

Data. We mainly rely on data shared by Arkolakis, Eckert, and Shi (2023). Most statistics related to multinational production come from Alviarez (2019)'s database on manufacturing multinational production and trade flows. Three matrices are particularly important when estimating the model: the count of affiliates, multinational production sales, and trade flows. For the first two, one observation corresponds to a pair of headquarter country and production location; for the latter, one observation corresponds to a pair of production location and final destination market. All values are averaged over the 2003-2012 period. We also use information on GDP, employment, and total population gathered by Arkolakis, Eckert, and Shi (2023) from the Penn World Tables. We complement this dataset with 2023 statutory corporate income tax rates from the Tax Foundation.

Parametrization. Simulations require to further parametrize the model. First, we assume that core productivity is drawn in each headquarter country i from a Pareto distribution with scale φ_i and shape ξ. This specification is common to Tintelnot (2016) and Arkolakis, Eckert, and Shi (2023).

Another approach common to both studies consists in writing bilateral frictions as functions of gravity variables. We follow Arkolakis, Eckert, and Shi (2023) who add and estimate destination-specific fixed effects. For any triplet of locations (i, l, n), we assume:

$$\begin{cases} \ln(\tau_{ln}) = \bar{\tau}_n \mathbb{1}^{l \neq n} + \kappa_\tau^d \ln(d_{ln}) + \kappa_\tau^{\text{COL}} \text{COL}_{ln} + \kappa_\tau^{\text{COM}} \text{COM}_{ln} + \kappa_\tau^{\text{BOR}} \text{BOR}_{ln} + \ln\left(1 + t_{ln}^{\text{Tariff}}\right) \\ \ln(\gamma_{il}) = \bar{\gamma}_l \mathbb{1}^{l \neq n} + \kappa_\gamma^d \ln(d_{il}) + \kappa_\gamma^{\text{COL}} \text{COL}_{il} + \kappa_\gamma^{\text{COM}} \text{COM}_{il} + \kappa_\gamma^{\text{BOR}} \text{BOR}_{il} \\ \ln(f_{il}) = \bar{f}_l \mathbb{1}^{l \neq n} + \kappa_f^d \ln(d_{il}) + \kappa_f^{\text{COL}} \text{COL}_{il} + \kappa_f^{\text{COM}} \text{COM}_{il} + \kappa_f^{\text{BOR}} \text{BOR}_{il} \end{cases}$$

The same gravity variables determine the three bilateral frictions. d_{jk} measures distance between countries j and k, COL_{jk} indicates whether they share any historical colonial tie, COM_{jk} stands for common languague(s), and BOR_{jk} is a dummy for contiguity. These expressions yield 12 elasticities, specific to each friction and gravity variable but common across locations. Tariffs only affect the iceberg trade costs, multiplicatively.

The destination-specific fixed effects $\bar{\tau}_n$, $\bar{\gamma}_l$, and $\bar{f}_l$ respectively reflect how costly it is on average to export into market n, to produce in location l from any foreign headquarter country, and to set up a foreign affiliate in location l from any headquarter country. We re-estimate the three vectors of fixed effects as part of the calibration effort described below.

Calibration. Some quantities, like population or enterprise counts, are directly observed in the data. From there, we can distinguish three types of unknown quantities: (i) parameters common to all locations—e.g., the price elasticity of demand σ; (ii) country-specific parameters, the so-called "fundamentals" of the model—e.g., location-specific productivity shifters $\{\epsilon_l\}$; (iii) general equilibrium outcomes—e.g., wage rates $\{w_l\}$.

The parameters in (i) are all calibrated as per Arkolakis, Eckert, and Shi (2023). The parameter θ that governs the dispersion of the Fréchet draws is set to 4.5, and the price elasticity of demand σ is equal to 4. With this calibration, the $\theta > \sigma - 1$ condition is verified, and substitutability forces thus dominate complementarities in the choice of production locations. The dispersion parameter for core productivity, ξ, takes value 4.95, which Arkolakis, Eckert, and Shi (2023) choose to match estimates in Arkolakis (2010). Finally, we build bilateral frictions from the same elasticities to gravity variables as those estimated by Arkolakis, Eckert, and Shi (2023).

Given this calibration, a unified iterative procedure estimates the model fundamentals (ii) and solves for general equilibrium outcomes (iii). As we summarize in Table A1, each fundamental is associated with one empirical moment and each outcome is associated with one equilibrium condition. At each iteration, we make a guess about general equilibrium outcomes and fundamentals. Based on this guess, we update general equilibrium outcomes. Then, we compute values of the fundamentals consistent with the new general equilibrium outcomes and the targeted moments. We form our next guess and stop the iteration when the relative change in each outcome or fundamental is sufficiently small. The procedure yields an exact match for all targeted moments.

4.2 Model Fit and Initial Patterns

Prevalence of multinational expansion. The bilateral count of affiliates in the data shows that expansion outside of the headquarter country is very rare. In the estimation, we rationalize this observation via large fixed costs for setting up foreign affiliates, and the model thus produces very few multinationals.

Figure 3 plots in red the share of multinationals in the mass of firms originating from each headquarter country. This share always lies below 1%. The Netherlands displays a substantially larger concentration of multinationals than all the other locations; on the contrary, there are almost no Hungarian multinationals.

We cannot observe each firm's multinational status in the data, which we would need to compute the corresponding shares. However, we can deduce bounds from the bilateral counts of foreign affiliates: The upper bound in green assumes that each foreign affiliate belongs to a different multinational, while the lower bound in blue only counts multinationals present in the most attractive production location. The model tracks well the lower bound.

Foreign affiliate counts. We turn to the distribution of foreign affiliates, between headquarter countries and production locations. Two moments that we target are closely related to this distribution. On the production location side, the fixed effect $\bar{f}_l$ ensures an exact match between l's own share of affiliates in the model and that measured in the data. On the headquarter country side, the scale of the core productivity distribution φ_i adjusts to track

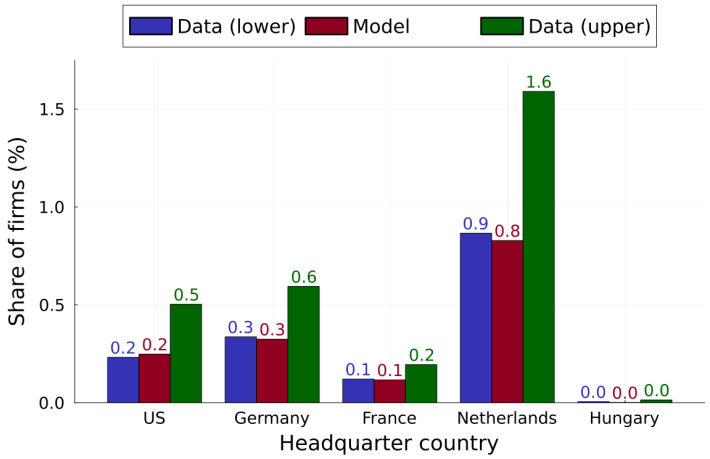
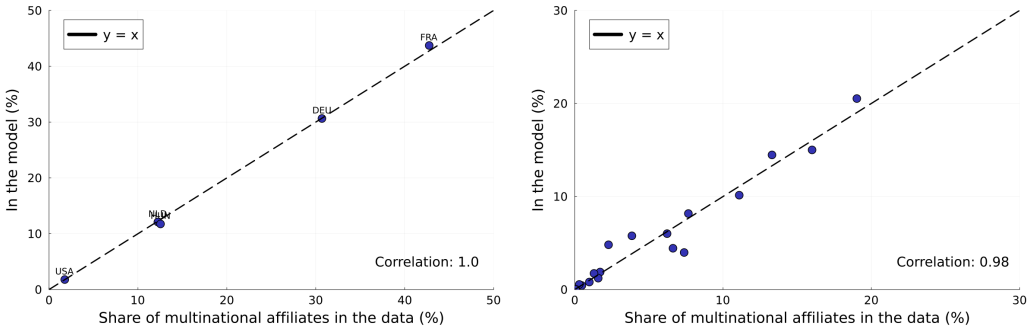


Figure 3: Share of multinationals by headquarter country

i's multinational sales relatively to US ones. We obtain a quasi-exact match in the weights of headquarter countries and production locations in the count of multinational affiliates.

The left-hand panel of Figure 4 illustrates this match for production locations. Strikingly, the US is associated with very few foreign affiliates, accounting for a few percents of the total. The model rationalizes this observation from the data via a large fixed effect $\bar{f}_{\text{US}}$ in the fixed cost for setting up a foreign affiliate there. The Netherlands and Hungary attract very similar numbers of foreign affiliates, while France hosts more than 40% of the total.

Our estimation procedure does not target the bilateral split of these affiliate counts. The right-hand panel of Figure 4 compares the weight of each pair of headquarter country and production location in the data and the model. The two series are very close to one another.



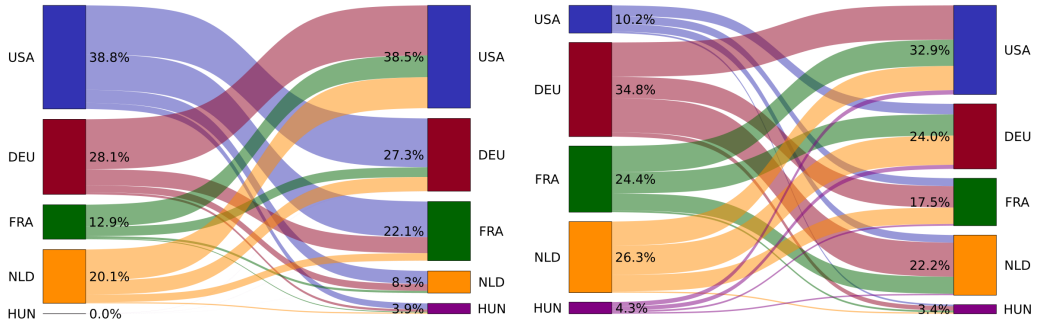
(a) At the production location level

(b) At the bilateral level

Figure 4: Multinational expansion in the model and the data

Multinational sales and trade flows. We obtain similar results when we include the intensive margin and look at the distribution of multinational sales. Here too, several targeted moments contribute to matching both sides with the data. On the headquarter country side, the link with the estimation of the core productivity scales is even more direct. On the production location side, the technology shifter ϵ_l adjusts to fit l's per-capita output—by domestic and foreign affiliates—relative to the US, and the fixed effect $\bar{\gamma}_l$ ensures an exact match between l's own share of local sales in the model and that measured in the data. With a correlation of 0.91, the relation between the data and the model at the bilateral level is satisfying although slightly noisier than for affiliate counts.

The left-hand panel of Figure 5 represents foreign affiliates' sales by headquarter country on the left and production location on the right. We saw that only a few European multinationals enter the US, but they are the most productive ones and get direct access to a massive consumption market. This explains why the US appears as the largest host of multinational activity. France, which was the most attractive location at the extensive margin, comes in third after we account for the intensive margin.



(a) Multinational sales

(b) Exports

Figure 5: Bilateral flows at the initial equilibrium

For trade flows, the estimation procedure leaves more freedom. The main related moment that we target is country n's own share of trade, which determines the fixed effect $\bar{\tau}_n$. This results in a rather close match for the weights of final destination markets in total trade flows, with a correlation of 0.93. We find more variation on the production location side, where the correlation is reduced to 0.86 and the model underestimates the importance of German exports. That being said, the bilateral split still displays a correlation of 0.87, and the model remains rather consistent with the data.

The right-hand panel of Figure 5 plots trade flows between production locations and final destination markets. It considers the exports of purely local firms as well as multinationals. The excessively low weight of the US among exporters highlights a limitation of our simplified setting. With our selected sample of locations, we ignore many of the US' largest trade

partners. Early simulations can hardly be seen as a scaled down version of the ideal, more comprehensive study since qualitative patterns would also change drastically.

Taxation. So far, we have discussed the fit to multinational production data. Regarding taxation, the estimation procedure ensures an exact match with the targeted moments. We use the headquarter-specific fixed cost to match the share of multinationals adopting the profit shifting technology. We assume that this share is 10% across all headquarters, and the initial equilibrium exactly reflects this proportion. We also target profit shifting intensities broadly inspired from Tørsløv, Wier, and Zucman (2022) and computed as the share of shifted profits in the total profits booked outside of the headquarter country. We assume that they are respectively 60% for the US, 35% for France and Germany, and 25% for Hungary. Then, we adjust the parameter that governs variable costs and obtain exactly the same shares in the initial equilibrium.

4.3 Illustrating Theoretical Results

Responses to taxation. From the initial equilibrium, we can simulate several theoretical results. We start with the relation of substitutability between real and profit shifting responses to taxation. Consider a firm originating from i with an affiliate in l and access to the profit shifting technology. We measure real responses as the semi-elasticity of the total output of its affiliate in l to the local corporate income tax rate. Starting from Equation 7, we derive:

$$\frac{\partial \ln\left(\sum_n x_{iln}\right)}{\partial t_l} = -\left(\frac{\theta}{\sigma-1}-1\right)\frac{1}{\kappa_{il}}\left(1-\frac{t_l}{\delta_i}\right)\left(1-\sum_{n=1}^{N}\zeta_{iln}\,\xi_{iln}\right)$$

$$\text{with } \xi_{iln} = \frac{x_{iln}}{\sum_{k=1}^{N} x_{ilk}} = \text{Share of } n \text{ in total sales from } l$$

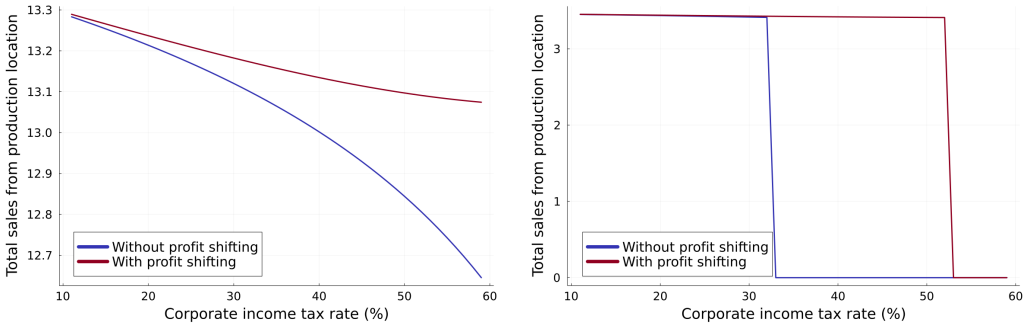
Remember our assumption that $\theta > \sigma - 1$, discussed above. Then, the semi-elasticity is unambiguously negative. This illustrates the intuitive idea that a higher tax rate makes the production location less attractive at the intensive margin. But what happens to the magnitude of these real responses when δ_i increases and profit shifting becomes more costly?

- The effective net-of-tax rate κ_{il} is reduced, and its inverse increases;
- The share of profits that remains taxed in the production location, $\left(1-\frac{t_l}{\delta_i}\right)$, increases;
- The last term, $\left(1-\sum_{n=1}^{N}\zeta_{iln}\,\xi_{iln}\right)$, changes ambiguously.

The first two effects tend to raise the magnitude of the semi-elasticity. In the specific case where conventional corporate income tax rates are equal across production locations, the last term does not depend on the cost of profit shifting anymore, and the third effect

disappears. Under this illustrative assumption, real responses to taxation intensify as the cost of pure tax planning increases. This induces a substitutability relation between real and profit shifting responses to taxation. Since we cannot conclude in the general case, we further rely on a simulation.

In Figure 6, we isolate a multinational and let the conventional corporate income tax rate of one of its affiliates vary between 10% and 60%. We then plot the affiliate's output as a function of the corporate income tax rate. The left-hand panel keeps the firm's location choices unchanged, while the right-hand panel allows for extensive margin responses.



(a) Intensive margin only

(b) Including the extensive margin

Figure 6: Effect of corporate income taxation on a foreign multinational's operations

Consistently with the sign of the semi-elasticity above, both curves in Panel 6a are downward sloping. A higher tax rate makes the production location less attractive. Furthermore, when we move from a finite cost of profit shifting with the red curve to an infinite one with the blue curve, we find that real activities respond more sharply to the local tax rate. That is, curtailing tax planning opportunities intensifies real responses to taxation, which corroborates the substitutability relation hypothesized above.

Interestingly, Panel 6b suggests that this substitutability relation also holds along the extensive margin of location choices, even though we lack analytical results. When the firm is prevented from shifting profits, it cannot sustain a range of relatively high corporate income tax rates and exits the location "earlier" than with tax planning capacities.

This echoes Slemrod (1992)'s hierarchy of responses to taxation. According to this hierarchy, firms preferably engage in reporting responses, instead of costly adjustments to real activities. If profit shifting is cheap, tax planners do not feel the need to move employment or capital in response to changes in the tax environment. This gives rise to a substitutability relation between real and reporting responses to taxation very similar to the one that we lay out in our model. What about the global minimum tax then? The reform is expected to curtail profit shifting incentives, which should raise revenues for relatively high-tax countries. However, we see here that it may also make real activities more sensitive to

conventional tax rates, putting the same high-tax countries at risk.

Efficiency. Simulations also highlight that the impact of a global minimum tax on efficiency is *a priori* ambiguous. In our benchmark conceptual framework, we can distinguish three types of distortions related to taxation. First, holding everything else constant, an increase in the average conventional corporate income tax rate hinders multinational expansion and reduces aggregate output. The global minimum tax should aggravate distortions through this channel. Second, tax rate differentials alter firms' choice of production locations, at the extensive and intensive margins. Since the reform is expected to limit differentials in effective tax rates, it would be beneficial to efficiency. Third, we model the variable costs of profit shifting as a pure loss of resources, which could be drastically reduced by a global minimum tax.

4.4 Counterfactual

The reform. We introduce a global minimum tax of 15%. Similarly to Pillar Two, top-up taxes apply on a country-by-country basis. Besides, in the Netherlands, they distinctly apply to income generated locally and to artificially relocated income. The taxes paid there on conventional profits cannot shield shifted profits from the minimum tax rate.

The reform applies to *all* multinationals, ignoring the revenue threshold included in the original agreement. However, *only* multinationals are concerned. As a consequence, some firms may find it optimal to close their foreign affiliates in order to avoid the global minimum tax. In practice, only relatively small Hungarian multinationals have such an incentive, as they avoid an increase in their domestic tax rate by 6 percentage points.

Table 1 shows the initial effective tax rates and their changes with the reform respectively for purely local firms, non-tax-planning multinationals, and multinationals adopting the profit shifting technology. Purely local firms initially face statutory corporate income tax rates, and they do not experience any change with the reform. Multinationals not engaged in profit shifting also face statutory rates before the reform. However, after the reform, their perceived tax rate in Hungary increases from 9% to 15%. Profit shifting multinationals initially benefit from lower effective tax rates, which account for their capacity to relocate income into the tax haven. Because the pre-reform tax rate on shifted profits is null, profit shifting occurs from all production locations. The reform raises this tax rate to 15%. So, shifting becomes less intensive from the US, France, Germany, and the Netherlands, which results in higher effective tax rates. Shifting from Hungary completely stops, and top-up taxes raise the effective rate to 15%.

Location	Purely local		MNC, not shifting		MNC, shifting	
	Initial (%)	Change (ppt.)	Initial	Change	Initial	Change
United States	26	—	26	—	$26 - \frac{6.76}{\delta_i}$	$+ \frac{1.65}{\delta_i}$
Germany	30	—	30	—	$30 - \frac{9}{\delta_i}$	$+ \frac{2.25}{\delta_i}$
France	20	—	20	—	$20 - \frac{4}{\delta_i}$	$+ \frac{0.75}{\delta_i}$
Netherlands	20	—	20	—	$20 - \frac{4}{\delta_i}$	$+ \frac{0.75}{\delta_i}$
Hungary	9	—	9	$+6$	$9 - \frac{0.81}{\delta_i}$	$+6 + \frac{0.81}{\delta_i}$

Table 1: Initial effective tax rates and changes due to the reform

Notes: The initial tax rates for firms that do not adopt the profit shifting technology are based on the 2023 statutory corporate income tax rates compiled by the Tax Foundation. For the other firms, they account for profits relocated to the tax haven and taxed at rate $t^{\text{PS}} = 0$. The formula is: $t_l - t_l^2 / (100\delta_i)$. Purely local firms are not concerned with the reform. Non-shifting multinationals see their effective tax rate in each location increase by $\max(15 - t_l, 0)$ percentage points. Eventually, two changes occur for shifting multinationals. Not only do they face the same increase as non-shifting multinationals, but they also relocate a lower fraction of their income into the tax haven. This raises their effective tax rate by $t_l^2 / (100\delta_i) - [\max(t_l - 15, 0)]^2 / (100\delta_i)$. This table ignores changes in multinational or shifting status.

Direct effects. We update all of firms' decisions in response to the reform, including their discrete choice of production locations and their decision with respect to the profit shifting technology. However, for now, we keep general equilibrium outcomes unchanged. We do not analyze an equilibrium, but we instead isolate the immediate effects of the global minimum tax on firms' incentives.

Figure 7 shows the impact of the reform on the mean effective tax rates of three groups of firms. As expected, purely local firms are never concerned. The largest changes are found in Hungary, where both domestic and foreign multinationals suffer a tax hike by at least 6 percentage points. Foreign multinationals in the US are also part of the main losers. To explain this sharp increase, remember that only the largest European multinationals expand into the US and most of these firms also adopt the profit shifting technology.

Changes in effective tax rates translate into firms' expansion and the distribution of their activities. Figure 8 shows that the US and Hungary lose the most foreign activities in relative terms. At the extensive margin, in the left-hand panel, the number of foreign affiliates drops by 12% in both locations. Because stayers are more productive than leavers, the intensive margin mitigates this effect. Still, in the right-hand panel, the output of foreign multinationals is reduced by 5% in the US and in Hungary. Not only does multinational activity contract as a whole, but the reform also induces some substitutions between production locations. In particular, in our simplified setting, the Netherlands seems to benefit from Hungary's loss of attractivity as a production location.

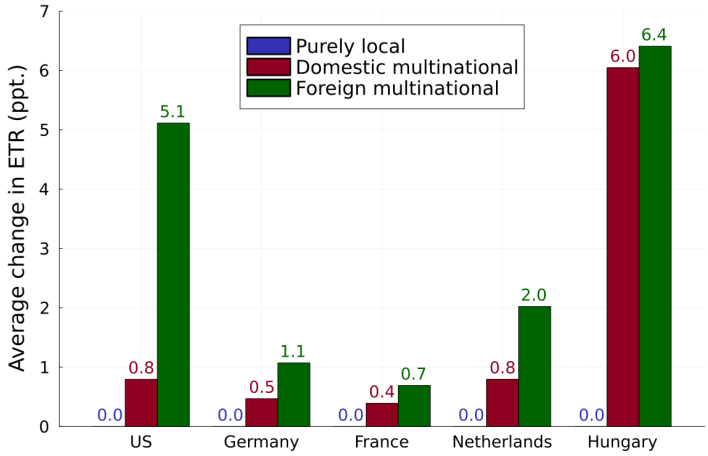
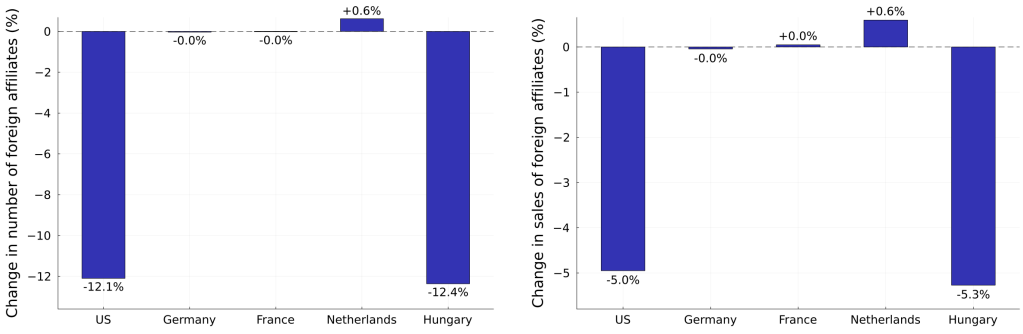


Figure 7: Direct change in effective tax rates by location and firm type



(a) Extensive margin only

(b) Including the intensive margin

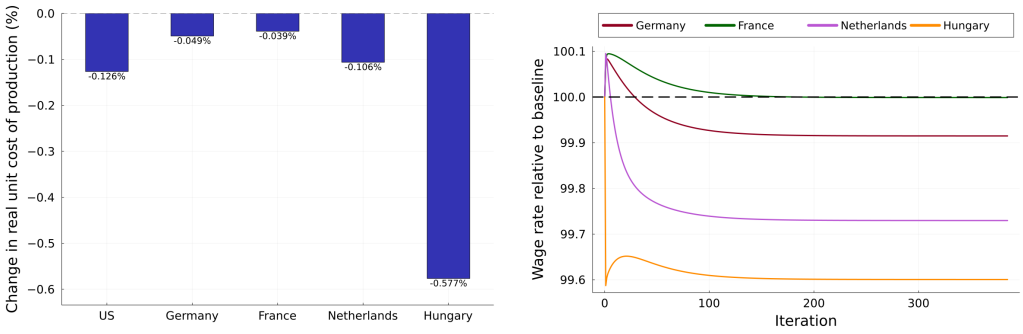
Figure 8: Direct change in multinational activity by location

Further results suggest that multinationals tend to relocate foreign activities into their headquarter country. Domestic output increases across all locations, especially in the Netherlands that displays a larger share of multinationals. We can attribute this effect to multinationals because we focus on the direct consequences of the reform and purely domestic firms face do not face any change in effective tax rates.

General equilibrium. We turn to the new general equilibrium with the 15% global minimum tax. We update wage rates, rental rates for tangible capital, price indices, and aggregate expenditures. The iterative procedure focus on aggregates while keeping fundamentals unchanged. The previous results could be interpreted as stopping in the middle of the first iteration, before updating our guess for general equilibrium outcomes.

The left-hand panel of Figure 9 highlights how real input costs adjust following the reform. All countries display noticeable drops in the unit cost of production. For Hungary, this decrease is consistent with previous direct effects. As it becomes less attractive as a production location, the reduced demand for labor or capital by foreign multinationals translates into lower real input costs. The graph corroborates a similar narrative for the US, with a muted impact since foreign affiliates account for a smaller share of total activities. However, the substantial drop in the Dutch real cost of production seems to contradict the immediate consequences of the reform. Typically, we saw in the direct effects of Figure 8 that the Netherlands tends to attract more foreign multinational activity. This should raise the demand for labor and capital, eventually pushing the unit cost of production upwards.

The right-hand panel of Figure 9 decomposes the path of our guesses for nominal wage rates over the iterative procedure that determines the new equilibrium with the global minimum tax. The first iteration clearly reflects the direct effects of the reform. The wage rate declines sharply in Hungary and increases in the other production locations, including the Netherlands. But from the second iteration, the lower unit cost of production in Hungary counteracts the tax hike, spurs some firms to substitute Hungary for the other production locations, and drives their wage rates downwards. The strength of this mechanism depends on how substitutable production locations are with Hungary in our simplified setting, and the Netherlands appears as the most exposed.



(a) Change in the unit cost of production

(b) Path of the nominal wage rate

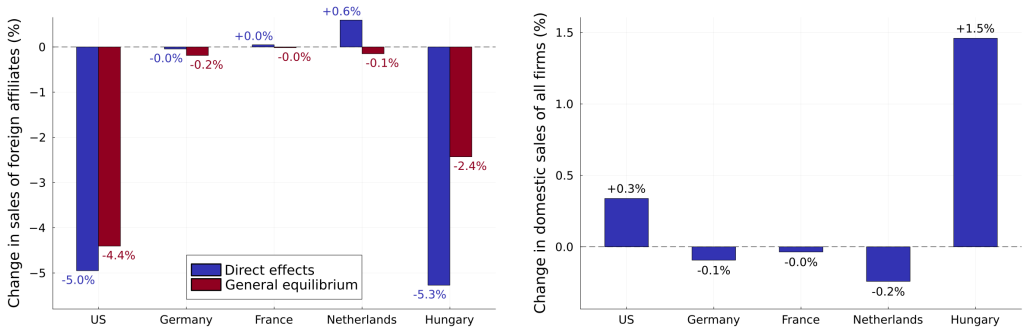
Figure 9: Change in input costs

Notes: The right-hand panel displays the path of our guesses for nominal wage rates over the iterative procedure used to compute the counterfactual general equilibrium. We only display four countries because the US wage rate is systematically normalized to 1 to anchor nominal values. The guesses include some "dampening" to facilitate convergence. From iteration k and for country i, the next guess, $\hat{w}_{i,k+1}$, is a weighted average of the previous guess, $\hat{w}_{i,k}$, and what is obtained from the labor market clearing condition, $w_i\left(\hat{w}_{i,k},...\right)$. We give a weight of 80% to the previous guess, which contributes to the smooth aspect of the curves.

How do these changes in the unit cost of production translate into multinational and domestic activity? In the left-hand panel of Figure 10, we show in red the change in foreign affiliates' sales by location in the new general equilibrium. We benchmark it against the

direct effects of Figure 8, in blue. In Hungary, the decrease in the unit cost of production visibly moderates the direct loss of attractivity as a production location. The drop in multinational activity in the US is also slightly muted. Conversely, in the Netherlands, the positive direct effect on foreign affiliates' output is more than offset by general equilibrium adjustments, despite the reduction in the Dutch real cost of production. We again interpret this as resulting from substitution with Hungary, where input prices are cut much more drastically than in any other location.

The right-hand panel of Figure 10 turns to domestic affiliates' sales, showing how they vary in the new general equilibrium for each location. Although the reform initially targets multinationals, general equilibrium effects transfer the shock onto purely local firms. In Hungary, and to a lower extent in the US, these firms face unchanged tax rates and benefit from lower costs of production, so they expand. This adds to the direct repatriation incentive described above and fuels the output of domestic affiliates.



(a) Change in multinational activity by location (b) Change in domestic activity by location

Figure 10: Change in multinational and domestic activity

Welfare. The reform reduces real private income in all countries. Table 2 provides a detailed decomposition of this change. In the Netherlands, where real private income is reduced by close to 1%, the impact primarily relates to the lower after-tax profits that accrue to consumers after the reform: Dutch firms have the strongest tendency to expand abroad in our simplified setting, which exposes them relatively more to the global minimum tax. Similar effects seem at play for Germany and France, although the reduction in real private income is more limited. Oppositely, in Hungary, purely local firms benefit from lower input prices and grow. The gains in after-tax profits partly offset losses in labor and capital income. The higher price level aggravates the drop in real private income, up to 0.3%. Eventually, the change in US real private income flows almost only through prices.

But the desirability of the global minimum tax also depends on government revenues. Figure 11 shows the percentage change in global corporate income tax collection under var-

Country	Change in $\frac{Y}{P}$ (%)	Change in Y (%)	Decomposing the change in Y		
			Labor income (%)	Capital income (%)	Profits (%)
United States	−0.15	−0.02	+0.00	+0.00	−0.02
Germany	−0.22	−0.25	−0.03	−0.03	−0.18
France	−0.15	−0.11	+0.00	+0.00	−0.11
Netherlands	−0.96	−1.13	−0.11	−0.11	−0.91
Hungary	−0.3	−0.12	−0.17	−0.17	+0.22

Table 2: Impact on real private income

Notes: The last three columns sum to the third column, "Change in Y".

ious scenarios. First, without any behavioral response, total corporate income tax revenues increase by roughly 4%. This approach echoes the mechanical estimates in the literature that keep profits and effective tax rates unchanged, and directly apply top-up taxes—e.g., Baraké et al. (2022). Then, revenue gains increase when firms adjust profit shifting, only at the intensive margin in the second bar and with the extensive margin in the third one. Once certain firms abandon the profit shifting technology, the global revenue gains reach almost 6%. Eventually, firms' real responses and general equilibrium adjustments affect the revenue potential of the reform only marginally, as shown in the last two bars.

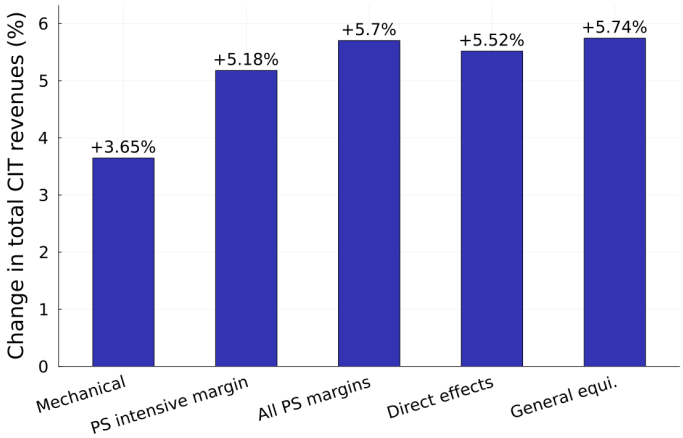


Figure 11: Change in global tax revenues in different scenarios

Notes: This graph shows the percentage change in global tax revenues induced by the reform under various scenarios. The "Mechanical" scenario applies top-up taxes without any behavioral response by firms. "PS intensive margin" allows firms to adjust their profit shifting behavior at the intensive margin only. In addition, "All PS margins" allows firms to review their discrete profit shifting decision, and some do not find the technology profitable anymore. "Direct effects" incorporates all of firms' choices but keeps general equilibrium outcomes unchanged. "General equi." finally adjusts general equilibrium outcomes.

The net impact of the global minimum tax on welfare in each location depends on two additional ingredients. First, we must add an assumption to distribute additional tax revenues across countries. We consider two scenarios. On the one hand, top-up taxes are collected where low-taxed profits are booked. On the other hand, top-up taxes accrue to headquarter countries. In technical terms, both scenarios respectively correspond to a global implementation of Qualified Domestic Minimum Top-up Taxes—"QDMTT scenario"—or the Income Inclusion Rule—"IIR scenario".

Second, welfare effects depend on the relative valuation of private income and government revenues. We write welfare in country j as a linear combination of real consumption and real government revenues, with a parameter μ_j governing the weight of the latter:

$$W_j = \frac{Y_j}{P_j} + \mu_j \frac{G_j}{P_j} \tag{15}$$

Calibrating μ_j seems difficult. Instead, we compute the minimum value of the parameter that makes the global minimum tax welfare-improving for country j. We denote this threshold by $\underline{\mu}_j^s$, with $s \in \{\text{QDMTT}, \text{IIR}\}$ designating the scenario of interest. Table 3 shows the result for each country and implementation scenario. For instance, under the QDMTT scenario, the reform improves welfare in France if and only €1 of government revenues is more valuable than $\underline{\mu}_{\text{FR}}^{\text{QDMTT}} = €0.46$ of private income. Under the IIR scenario, because French revenue gains are larger, welfare is improved for a slightly wider range of values of the parameter—$\underline{\mu}_{\text{FR}}^{\text{IIR}} = 0.39$.

Country	**Change in $\frac{Y}{P}$**	**QDMTT scenario**		**IIR scenario**	
		Change in $\frac{G}{P}$	$\underline{\mu}^{\text{QDMTT}}$	Change in $\frac{G}{P}$	$\underline{\mu}^{\text{QDMTT}}$
	(%)	(%)		(%)	
US	−0.15	+4	0.64	+4	0.57
Germany	−0.22	+7	0.45	+8	0.37
France	−0.15	+5	0.46	+6	0.39
Netherlands	−0.96	+41	0.42	+22	0.78
Hungary	−0.3	+33	0.41	+7	2.09

Table 3: Minimum weight of government revenues for welfare improvement

5 Quantitative Model

The next steps for this work consist in developing and estimating a quantitative version of the model. First, the effects of a global minimum tax critically hinge on the magnitude of real responses to taxation. This magnitude depends on the degree of substitutability

between production locations, captured by the $[\theta/(\sigma-1)]$ ratio in Equation 4.3. The calibration proposed above fixes this ratio, and we cannot adjust it to empirical estimates. By introducing some correlation between countries in the Fréchet productivity draws as in Lind and Ramondo (2023), we can further parametrize the ratio of interest, leaving more freedom in the calibration. Real responses to taxation are also driven by the imperfect deductibility of capital costs, fully absent from the above conceptual framework.

Second, recent empirical studies point at potential complementarities between real and nominal responses to taxation. For instance, De Simone and Olbert (2022) find that firms respond to new disclosure requirements by increasing real activities in low-tax jurisdictions in the EU. This effect is concentrated within *ex ante* more aggressive tax planners, and they interpret it as firms moving employment and / or assets to substantiate their profit shifting practices. The interaction of real and nominal responses critically shapes the effects of a global minimum tax. Hence the need to model and estimate such complementarities, in addition to the substitutability channel described so far.

Third, the literature highlights large variations in effective tax rates across firms. This results in substantial pockets of under-taxed profits in otherwise high-tax jurisdictions. Hugger, Cabral, and O'Reilly (2023) estimate that the jurisdictions with an average effective tax rate above 15% host more than half of the global profits taxed below 15%. Our early simulations based on statutory rates completely ignore those. Firm-level country-by-country reports appear as a valuable data source to introduce and quantify this heterogeneity.

6 Conclusion

We study the real responses of multinational companies to a global minimum tax on corporate income. We leverage a quantifiable general equilibrium model of multinational production, that builds upon Tintelnot (2016) and Arkolakis, Eckert, and Shi (2023). The fixed costs for setting up foreign affiliates imply that firms establish heterogeneous production networks. Certain firms also adopt a profit shifting technology, thanks to which they can relocate taxable income into a tax haven. These discrete choices shape real responses to taxation and create differential exposure to the reform.

In the current version, we propose a simplified simulation to illustrate some theoretical results and the main mechanisms at play with a global minimum tax. We sketch the next steps to move from our baseline framework to a credible quantitative model. Progress would come from the modeling and calibration of real responses to taxation, refined complementarities between real and profit shifting responses, and accounting for the large heterogeneity in effective tax rates across firms. Ongoing efforts also aim at updating the multinational production statistics and improving the taxation-related moments used in estimation.

References

Alviarez, Vanessa. 2019. "Multinational production and comparative advantage." *Journal of International Economics* 119:1–54. ISSN: 0022-1996. https://doi.org/https://doi.org/10.1016/j.jinteco.2019.03.004.

Arkolakis, Costas. 2010. "Market Penetration Costs and the New Consumers Margin in International Trade." *Journal of Political Economy* 118 (6): 1151–1199. ISSN: 00223808, 1537534X, accessed January 21, 2025. http://www.jstor.org/stable/10.1086/657949.

Arkolakis, Costas, Fabian Eckert, and Rowan Shi. 2023. *Combinatorial Discrete Choice: A Quantitative Model of Multinational Location Decisions.* Technical report. Yale University.

Baraké, Mona, Paul-Emmanuel Chouc, Theresa Neef, and Gabriel Zucman. 2022. "Revenue Effects of the Global Minimum Tax Under Pillar Two." *Intertax* 50 (10): 689–710.

Beer, Sebastian, Ruud de Mooij, and Li Liu. 2020. "International Corporate Tax Avoidance: A Review of the Channels, Magnitudes, and Blind Spots." *Journal of Economic Surveys* 34 (3): 660–688. https://doi.org/https://doi.org/10.1111/joes.12305.

Clausing, Kimberly, Emmanuel Saez, and Gabriel Zucman. 2020. "Ending Corporate Tax Avoidance and Tax Competition: A Plan to Collect the Tax Deficit of Multinationals." *SSRN Electronic Journal* (January). https://doi.org/10.2139/ssrn.3655850.

De Simone, Lisa, and Marcel Olbert. 2022. "Real Effects of Private Country-by-Country Disclosure." *The Accounting Review* 97, no. 6 (October): 201–232. ISSN: 0001-4826. https://doi.org/10.2308/TAR-2020-0714.

Dyrda, Sebastian, Guangbin Hong, and Joseph B. Steinberg. 2023. *A Macroeconomic Perspective on Taxing Multinational Enterprises.* Working Papers. University of Toronto, Department of Economics.

Ferrari, Alessandro, Sébastien Laffitte, Mathieu Parenti, and Farid Toubal. 2023. *Profit Shifting Frictions and the Geography of Multinational Activity.* Technical report. arXiv: 2211.04388 [econ.GN].

Gumpert, Anna, Jr. Hines James R., and Monika Schnitzer. 2016. "Multinational Firms and Tax Havens." *The Review of Economics and Statistics* 98, no. 4 (October): 713–727. ISSN: 0034-6535. https://doi.org/10.1162/REST_a_00591.

Heckemeyer, Jost H., and Michael Overesch. 2017. "Multinationals' profit response to tax differentials: Effect size and shifting channels." *Canadian Journal of Economics* 50 (4): 965–994. ISSN: 00084085. http://search.ebscohost.com.ezproxy.hec.fr/login.aspx?direct=true&db=bth&AN=125995512&site=ehost-live.

Hugger, Felix, Ana Cinta González Cabral, Massimo Bucci, Maria Gesualdo, and Pierce O'Reilly. 2024. "The Global Minimum Tax and the taxation of MNE profit," no. 68, https://www.oecd-ilibrary.org/content/paper/9a815d6b-en.

Hugger, Felix, Ana Cinta González Cabral, and Pierce O'Reilly. 2023. "Effective tax rates of MNEs," no. 67, https://www.oecd-ilibrary.org/content/paper/4a494083-en.

Janeba, Eckhard, and Guttorm Schjelderup. 2022. *The Global Minimum Tax Raises More Revenues than You Think, or Much Less* [in en]. SSRN Scholarly Paper 4059715. Rochester, NY: Social Science Research Network. Accessed June 13, 2022. https://doi.org/10.2139/ssrn.4059715. https://papers.ssrn.com/abstract=4059715.

Lind, Nelson, and Natalia Ramondo. 2023. "Trade with Correlation." *American Economic Review* 113, no. 2 (February): 317–53. https://doi.org/10.1257/aer.20190781. https://www.aeaweb.org/articles?id=10.1257/aer.20190781.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. "Distributional National Accounts: Methods and Estimates for the United States." *The Quarterly Journal of Economics* 133 (2): 553–609. https://EconPapers.repec.org/RePEc:oup:qjecon:v:133:y:2018:i:2:p:553-609..

Ramondo, Natalia, Andrés Rodríguez-Clare, and Felix Tintelnot. 2015. "Multinational Production: Data and Stylized Facts." *American Economic Review* 105, no. 5 (May): 530–36. https://doi.org/10.1257/aer.p20151046. https://www.aeaweb.org/articles?id=10.1257/aer.p20151046.

Slemrod, Joel. 1992. "Do Taxes Matter? Lessons from the 1980's." *The American Economic Review* 82 (2): 250–256. ISSN: 00028282, accessed August 5, 2024. http://www.jstor.org/stable/2117409.

Spengel, Christoph, Frank Schmidt, Jost Heckemeyer, and Katharina Nicolay. 2020. *Effective Tax Levels Using the Devereux / Griffith Methodology.* Vol. Project for the EU Commission - TAXUD/2020/DE/308. Final Report 2020. European Commission, October.

Tintelnot, Felix. 2016. "Global Production with Export Platforms." *The Quarterly Journal of Economics* 132, no. 1 (October): 157–209. ISSN: 0033-5533. https://doi.org/10.1093/qje/qjw037.

Tørsløv, Thomas, Ludvig Wier, and Gabriel Zucman. 2022. "The Missing Profits of Nations." *The Review of Economic Studies* 90, no. 3 (July): 1499–1534. ISSN: 0034-6527. https://doi.org/10.1093/restud/rdac049.

A Summary of the Quantification Strategy

Outcome	Equilibrium condition	Equation
w_i	Labor market clearing	11
r_i	*Tangible capital market clearing*	12
P_i	Definition of the ideal price index	13
Y_i	Balance of payments	14
Fundamental	**Matched moment**	**Source**
φ_i	Foreign sales of firms from i vs. US: $\frac{\sum_{l\neq i,n} X_{iln}}{\sum_{l\neq \text{US},n} X_{\text{US}ln}}$	AES (2023)
ϵ_l	Output per capita in l vs. US: $\frac{\sum_{i,n} X_{iln}/L_l}{\sum_{i,n} X_{i\text{US}n}/L_{\text{US}}}$	AES (2023)
$\bar{\tau}_n$	Own share of trade: $\sum_i X_{inn}/\sum_{il} X_{iln}$	AES (2023)
$\bar{\gamma}_l$	Own share of local sales: $\sum_n X_{lln}/\sum_{in} X_{iln}$	AES (2023)
$\bar{f}_l$	Own share of local affiliates: $M_l/\sum_i M_{il}$	AES (2023)
δ_i	*Share of shifted profits in multinational profits*	$\approx$ *TWZ (2022)*
f_i^{PS}	*Share of multinationals shifting profits*	*Arbitrary, 10%*
α_i	—	*Arbitrary, 50%*

Table A1: Summary of the quantification strategy

Notes: In the table, X_{iln} denotes the total sales of firms originating from i, producing in l, and serving market n. L_i and M_i respectively denote the mass of workers and firms with which country i is endowed. M_{il} denotes the mass of firms originating from i that set up a foreign affiliate in l. Notice that for any country l, $M_l = M_{ll}$. For the parameter that governs the variable costs of profit shifting, we target headquarter-specific shares of profits shifted in multinational profits broadly consistent with Tørsløv, Wier, and Zucman (2022). They report that, on aggregate, shifted profits represent 36% of global multinational profits, and they argue that US-headquartered firms are more aggressive than the non-US ones. We thus target 60% for US-headquartered firms, 35% for France and Germany, and 25% for Hungary. δ_{NLD} is set equal to the minimum of the parameters obtained for the other, non-haven headquarter countries. To estimate the fixed costs of profit shifting, we assume for all headquarter countries that 10% of multinationals adopt the tax planning technology. For instance, Gumpert, Hines, and Schnitzer (2016) notice that 23% of German multinationals are present in at least one tax haven; we downgrade this benchmark figure as German multinationals may enter large neighboring havens like the Netherlands or Switzerland for motivations unrelated to profit shifting. Eventually, we set all Cobb-Douglas labor shares equal to 50% so that, in preliminary simulations, both inputs are treated in a fully symmetric way. **Link:** Back to Section 4.1.