The Macroeconomic Effects of Revenue-Neutral Corporate Tax Reforms*

Riccardo Silvestrini† Elliott Weder‡

August 26, 2026

[Latest Version]

Abstract

We study how revenue-neutral corporate tax reforms affect aggregate investment, productivity, and welfare in a heterogeneous firm model. Along the revenue-neutral schedule of rates and deductions, we show that a high-rate, high-deduction regime generates more aggregate investment and higher welfare than a low-rate, low-deduction one. More generous deductions act as an investment subsidy that benefits small to medium-large firms, which invest intensively relative to their size, and thereby foster capital accumulation despite the higher tax rate. However, the same mechanism reallocates capital away from the largest and most productive firms, reducing aggregate productivity. An event study exploiting firm-level variation in exposure to a 2005–2007 Dutch reform supports our model's key mechanism.

Keywords: Corporate Taxation, Firm Dynamics, Investment, Distribution, Allocation
JEL Codes: E22, E62, L11

*Previously circulated as 'Winners and Losers: How Corporate Tax Reforms Reshape the Firm Distribution'. We wish to thank Rui Albuquerque, Isaac Baley, Matteo Bizzarri, Timo Boppart, Antoine Bozio, Murillo Campello, Jeff Campbell, Chiara Lacava, Agnieszka Markiewicz, Julien Sauvagnat and the participants of seminars at the University of Naples Federico II, the Erasmus University Rotterdam, the 7^{th} Annual CefES Conference, 5^{th} Sailing the Macro Workshop in Ortygia, 4^{th} XAmsterdam Macroeconomic Workshop, and Theories and Methods in Macroeconomics 2025 for their helpful comments and suggestions.

†University of Naples Federico II and CSEF, riccardo.silvestrini@unina.it
‡Erasmus University Rotterdam and Tinbergen Institute, weder@ese.eur.nl

1 Introduction

Two central features of corporate tax systems are the statutory tax rate and investment deduction allowances. Several recent reforms aimed at stimulating investment have simultaneously reduced the tax rate and expanded deductions, at the expense of foregone tax revenue. For example, Chodorow-Reich, Smith, Zidar, and Zwick (2024) estimate that the 2017 Tax Cuts and Jobs Act boosted investment by roughly 10%, yet reduced US corporate tax revenue by around 40%. Policymakers who want to encourage capital accumulation without sacrificing tax revenue face a trade-off: any reduction in the tax rate must be offset by narrower investment deductions, and vice versa. Whether a high-rate, high-deduction or a low-rate, low-deduction tax regime is more conducive to aggregate investment and productivity is an open and first-order question for fiscal policy design.

We answer the question by tracing out the frontier of revenue-neutral corporate tax reforms in a calibrated heterogeneous firm model in the tradition of Hopenhayn and Rogerson (1993). Along the frontier, we study which combination of tax rate and deduction generates the highest aggregate investment, capital, and productivity. We discipline the model with firm-level balance sheet data from the Netherlands, also exploiting a 2005–2007 reform which was implemented to stimulate investment while preserving budget neutrality. The Dutch reform entailed a simultaneous cut to the statutory corporate tax rate from 34.5% to 25.5% and a reduction to the annual investment deductibility from 100% to 20%. Since the two tax policies push firm investment incentives in opposite directions, understanding the net effect requires a quantitative framework that accounts for realistic firm heterogeneity in size, productivity, and investment behavior.

We find that a high-rate, high-deduction regime stimulates higher investment and output than a low-rate, low-deduction one, holding corporate tax revenue fixed. Along the frontier of balanced reforms in the neighborhood of the Dutch post-reform regime, moving from a 24% tax rate (paired with a 2% annual deduction) to a 32% rate (paired with a 93% deduction) increases aggregate output by 12%, the capital stock by 30%, and raises welfare modestly, by 0.3%. However, aggregate productivity declines by roughly 1% along the same frontier, revealing a trade-off between capital accumulation and allocative efficiency.

The mechanism behind our results operates through heterogeneity in investment intensity across the firm-size distribution. In the balance sheet data, the investment-to-capital ratio is about twice as high for firms in the bottom 95% of the size distribution as for firms in the top 5%, and the ratio decreases monotonically with size. Our model replicates this pattern. The heterogeneity implies that, even in a balanced reform, changes to the rate and the deduction affect firms differently across the distribution. The mechanism operates as follows.

On the one hand, a lower tax rate increases after-tax profits for all firms, but the gain is directly proportional to the level of profit and therefore accrues mainly to the largest and most profitable firms. Since these firms do not invest intensively relative to their size, they do not expand their capital stock substantially in response to the lower tax rate.

On the other hand, a lower deduction delays tax benefits of investment for all firms by essentially raising investment costs. The low-deduction regime therefore depresses aggregate investment and output relative to the high-deduction one. Small to medium-large firms, whose growth strategies rely more heavily on capital investment relative to their size, are disproportionately burdened by the deduction cut. Thus, a low-deduction reform reallocates capital toward top firms, which raises aggregate productivity at the cost of a smaller overall economy.

Together, our results highlight an important policy-relevant trade-off. Along the frontier of revenue-neutral reforms, policymakers can effectively substitute the distortion from a positive tax rate for the distortion from lower investment deductions. Our quantitative analysis suggests that the latter dominates: diluting the tax benefit of investment generates larger aggregate costs than raising the statutory rate.

In the final section of the paper, we complement the model analysis with difference-in-differences evidence for the investment-intensity mechanism at the heart of our quantitative results. Using firm-level balance sheet data, we calculate each firm's exposure to the 2005–2007 Dutch reform as its average pre-reform investment-to-capital ratio. Our model predicts that more exposed firms contract, in relative terms, following the shift to a low-rate, low-deduction regime. The data support the prediction: a one-unit higher pre-reform ratio is associated with a capital stock around 5.2 thousand 2015 euros lower after the reform. Moreover, the differential is absent prior to the reform, emerges as the reform is phased in, and grows through to 2010.

Related literature

Our theoretical model is most closely related to Hopenhayn and Rogerson (1993). We add capital, investment, and corporate fiscal policy to their model, since we focus on frictions in the allocation of capital rather than of labor. Furthermore, we include both adjustment costs and irreversibility frictions for capital investment, echoing the framework presented in Winberry (2021). Cooper, Haltiwanger, and Power (1999), Cooper and Haltiwanger (2006), and Clementi and Palazzo (2016) show that both model features are necessary to improve the match with the data and to characterize realistic investment decisions.

Erosa and González (2019), Sedláček and Sterk (2019), Baley and Blanco (2022), and Colciago, Lewis, and Matyska (2024) all study the long-run effects of corporate taxation in heterogeneous firm models. In contrast to these papers, we focus on both aggregate and distributional effects of balanced corporate tax reforms. Furthermore, we combine detailed microdata with the 2005–2007 Dutch setting, in which a revenue-neutral tax reform was (attempted to be) actually implemented.

We also relate to a large empirical literature that examines how firms investment responds to tax policy changes. For a comprehensive survey of both the theory and the empirics behind this topic, see Chodorow-Reich (2025). Numerous papers in this literature study the response of firm investment to tax rate and investment deductibility changes, the two tax instruments we analyze in this paper.

Zwick and Mahon (2017) find that increased expensing raises investment in eligible capital. They also document that the response of larger firms is less elastic to the policy. Similarly, the reallocation effects generated in our model in the low-deduction regime arise because *non-large* firms respond the strongest. Ohrn (2018) finds that lower effective tax rates increase firm investment, in line with our model's predictions. Chodorow-Reich, Smith, et al. (2024) and Chodorow-Reich, Zidar, and Zwick (2024) study the effects of the 2017 Tax Cuts and Jobs Act (TCJA) in the US, in which the statutory tax rate was cut and immediate expensing was introduced. The authors find that firms more exposed to the reform increased their domestic investment.[1]

We complement the empirical literature in two ways. First, we examine how rate and deduction changes unfold in the cross-section in a general equilibrium model, analyzing both aggregate and distributional effects. Second, we focus on revenue-neutral tax reforms, which differ from several recent reforms, such as the TCJA. Along the revenue-neutral frontier, the implications for aggregate investment, productivity, and welfare are ambiguous. We shed light on these outcomes.

[1]Other papers that study the effect of corporate taxation on firm investment include Becker, Jacob, and Jacob (2013), Campbell, Chyz, Dhaliwal, and Schwartz Jr (2013), Link, Menkhoff, Peichl, and Schüle (2023), Lichter et al. (2025), and Bilicka, Güçeri, and Koumanakos (2025).

Outline

The remainder of the paper proceeds as follows. Section 2 introduces our model, which we calibrate in Section 3, while briefly discussing the Dutch reform we anchor our analysis in. Section 4 describes how we derive the frontier of revenue-neutral reforms and presents the aggregate and distributional effects. Section 5 provides empirical support for the model's main mechanism, while Section 6 concludes.

2 The Model

Our model builds on the framework from Winberry (2021), to which we add endogenous entry and exogenous exit of firms in the spirit of Hopenhayn and Rogerson (1993).[2] The supply side of the economy incorporates an endogenous mass of heterogeneous firms that produce using capital and labor. Firms pay taxes on their profits and can deduct a portion of their capital investment costs for tax purposes. Furthermore, investment is subject to capital adjustment costs and an irreversibility constraint. The model is closed by a representative household, who consumes, supplies labor, and finances the creation of new firms, and by a government, which levies taxes on firms and transfers the revenue back to the household.

2.1 Firms

2.1.1 Production, Investment, and Profits

Time is discrete and indexed by t. A mass of length $\mathcal{M}_t$ of heterogeneous firms populates the economy, where each firm is indexed by the subscript j. At time t, firm j produces output y_{jt} using two inputs, labor l_{jt} and capital k_{jt}. Firms own their capital stock rather than renting it from the household. The production function is given by a constant-returns-to-scale Cobb-Douglas technology:

$$y_{jt} = z_{jt} l_{jt}^{1-\alpha} k_{jt}^{\alpha}, \tag{1}$$

in which $\alpha \in (0,1)$ represents the output elasticity of capital in the production function. z_{jt} denotes the idiosyncratic productivity level, which follows an AR(1) process in logs

[2] There is a key distinction between our model and Hopenhayn and Rogerson (1993) regarding the notion of optimal firm size: in the latter, the non-degenerate firm size distribution comes from the decreasing returns to scale embedded in production. In our model, since we have constant returns to scale in capital and labor, the firm size distribution is pinned down by the convexity of adjustment costs. Our results are robust to the inclusion of a decreasing return to scale technology.

with mean reversion:

$$\log z_{jt+1} = \rho^z \log z_{jt} + (1-\rho^z) \log \bar{z}_{jt} + \omega^z_{jt}, \quad \text{with } \omega^z_{jt} \sim N(0, \sigma^z) \text{ and } 0 < \rho^z < 1. \tag{2}$$

The stationary distribution of the productivity process in levels is given by a log-normal distribution, which we call $G(z)$. Upon entering a period, firms observe a new productivity draw, their pre-existing stock of capital k_{jt}, and their tax deduction allowance d_{jt}, which we define below. Within each period, firms hire l_{jt} workers at the aggregate wage w_t to produce. The labor choice is fully static. Firms also choose investment i_{jt} into new capital.[3] The investment choice is dynamic. Given the investment choice i_{jt}, firm j's capital evolves according to the following law of motion:

$$k_{jt+1} = (1-\delta)k_{jt} + i_{jt}. \tag{3}$$

Whenever investment is non-zero, firms incur a convex adjustment cost $\mathcal{C}_{jt}$. The cost is measured in terms of output and is defined in more detail below.

We move on to corporate taxation. Labor and investment are subject to differential tax treatment. While labor costs are fully deductible in the period in which they occur, only a fraction $\hat{\delta}$ of investment costs is deductible. The partial deductibility of investment costs creates dispersion in the effective tax rate paid by firms. Conditional on pre-tax profitability, firms that invest more pay fewer taxes and see their after-tax profits rise. We assume a geometric deduction scheme: the fraction of investment that is not deducted in period t is carried over to the following period, together with any remaining pre-existing allowance d_{jt}.[4] Therefore, the deduction allowance evolves with the following law of motion:

$$d_{jt+1} = (1-\hat{\delta}) \left(d_{jt} + i_{jt} \right). \tag{4}$$

Firm profits π_{jt} are equal to cash flow $\mathcal{I}_{jt}$ (i.e., *before-tax* profits) minus taxes paid on taxable income $\mathcal{T}_{jt}$ and adjustment costs $\mathcal{C}_{jt}$:

$$\pi_{jt} = \mathcal{I}_{jt} - \tau \mathcal{T}_{jt} - \mathcal{C}_{jt}. \tag{5}$$

Cash flow $\mathcal{I}_{jt}$ is defined as revenues net of period-t costs:

$$\mathcal{I}_{jt} = y_{jt} - w_t l_{jt} - i_{jt}.$$

[3] When simulating the model, we assume that investment is irreversible, i.e. $i_{jt} \geq 0$, although results are similar if we allow for direct disinvestment.

[4] Although depreciation schemes are often linear in the data, working with a geometric scheme comes with computational advantages while remaining quantitatively consistent. For a further discussion, see Furno (2021) and Winberry (2021).

$\mathcal{T}_{jt}$ represents the taxable income, which is taxed at the linear corporate tax rate τ:

$$\mathcal{T}_{jt} = y_{jt} - w_t l_{jt} - \hat{\delta}\left(i_{jt} + d_{jt}\right).$$

Taxable income is equal to output net of fully deductible labor costs $w_t l_{jt}$ and partially deductible investment costs $\hat{\delta}\left(i_{jt} + d_{jt}\right)$. We assume the adjustment costs are convex and given by the following functional form, as in Winberry (2021):

$$\mathcal{C}_{jt} = \frac{\psi}{2}\left(\frac{i_{jt}}{k_{jt}}\right)^2 k_{jt} \equiv \mathcal{C}\left(i_{jt}, k_{jt}\right).$$

We assume that the price of output (consumption) is the numeraire, so revenues $p_t y_{jt}$ are equivalent to output y_{jt}. Even though we assume perfect competition, firms can still make positive profits because of entry costs and technological rents due to productivity differences and investment frictions.

For the remainder of this subsection, it is easier to formulate the firm maximization problem recursively. Let x' indicate the next-period value of a variable x. The state variables for a firm are given by current (idiosyncratic) productivity z, the existing capital stock k, and the deduction allowance d, which together fully identify each firm j.[5] For that reason, we also drop the jt subscript for firm-level variables, unless required. Therefore, we represent the state variables for a firm with the triplet $\{z, k, d\}$ and express the Bellman equation as:

$$V(z,k,d;\mu) = \max_{i,l}\{\pi(\mu) + (1-\gamma)\,\mathbb{E}\left[\Lambda(\mu)V(z',k',d';\mu')|z,k,d,\mu\right]\},$$

where μ represents the aggregate distribution of firms over the individual states, $\Lambda(\mu)$ stands for the stochastic discount factor, and γ is the probability of exit. We discuss firm entry and exit further in the next subsection. Substituting in the production function, the adjustment cost function, and the two laws of motion, the Bellman equation can be rewritten as:

$$\begin{aligned} V(z,k,d;\mu) = {} & \tau\hat{\delta}d + (1-\tau)\max_{l}\{\left[zl^{1-\alpha}k^{\alpha} - w(\mu)l\right]\} + \\ & + \max_{i}\{-\left(1-\tau\hat{\delta}\right)i - \frac{\psi}{2k}i^2 + (1-\gamma)\,\mathbb{E}\left[\Lambda(\mu)V(z',k',d';\mu')|z,k,d,\mu\right]\}, \end{aligned} \tag{6}$$

subject to:

$$\log z' = \rho^z \log z + (1-\rho^z)\log\bar{z} + \omega^z, \quad k' = (1-\delta)k + i \quad \text{and} \quad d' = (1-\hat{\delta})(d+i).$$

[5]For readability, we keep d as a separate state variable in the main text; Winberry (2021) shows that only the present value of deductions per unit of investment matters. This allows us to reduce the state space to just productivity and capital, a property we exploit for our numerical exercises below.

The rewritten Bellman equation indicates that the firm maximization problem leads to two policy functions. The first maximization shows that the optimal labor demand choice $l^* (z, k, d; \mu)$ is static and chosen such that the marginal product of labor $(1 - \alpha) z (l/k)^{-\alpha}$ equals the aggregate wage $w(\mu)$. Second, the dynamic optimal investment choice $i^* (z, k, d; \mu)$ equates the current marginal costs to future marginal benefits of investment. Investment today raises current costs due to forgone profits and adjustment frictions, but raises firm value by increasing future capital and tax deduction allowances. More future capital increases future profits through higher revenues and reduced adjustment costs, while larger deductions raise after-tax profits by expanding the tax shield.

2.1.2 Entry, Exit, and the Distribution of Firms

Entry is endogenous. An unbounded mass of potential entrants is prepared to join the economy in every period t. Entry occurs up to the point where the expected value of the entry equals entry costs $w_t c_e$, which we assume are paid in terms of labor. We denote the mass of entrants in period t by M_t. Entrants learn their productivity only after paying the entry cost, which they draw from the stationary distribution $G(z)$. Entrants immediately start producing with an initial level of capital k_0 and zero deduction allowance.[6] Hence, the free entry condition is given by:

$$\int V(z, k_0, 0; \mu_t) dG(z) = w_t c_e. \tag{7}$$

In contrast, we assume exit is exogenous and occurs between periods with a constant probability γ. Our assumption is consistent with Kindsgrab (2022), who finds that corporate tax rate changes only affect entry and not exit of firms.

Given entry, exit, and firms' optimal policy functions, the evolution of the distribution $\mu(z_t, k_t, d_t)$ satisfies:[7]

$$\mu(z_{t+1}, k^*(z, k, d; \mu_t), d^*(z, k, d; \mu_t)) = (1 - \gamma) \mu(z_t, k_t, d_t) F(z_{t+1} | z_t), \tag{8}$$

for $d_t > 0$ and $k_t \neq k_0$. $F(z_{t+1}|z_t)$ represents the transition probability from productivity z_t to z_{t+1}, and $k^*(z, k, d; \mu_t)$ and $d^*(z, k, d; \mu_t)$ represent, respectively, the optimal

[6] Entrants start with a low but non-zero level of capital to avoid numerical complications that would arise when capital is zero. Two interpretations of this assumption are that a fraction of entry costs can be immediately turned into capital or that paying the entry cost raises equity to begin operating.

[7] For ease of notation, here we represent the transition in a discretized state-space, in which transition probabilities for productivity are described by a discrete Markov process, similarly to our numerical solution method. The logic is the same in continuous space.

capital and deduction chosen for period $t+1$ by a firm with type (z,k,d) in period t.[8] Moreover, for entrants:

$$\mu(z_{t+1}, k^*(z,k_0,0;\mu_t), d^*(z,k_0,0;\mu_t)) = (1-\gamma)\left[\mu(z_t,k_0,0) + M_t G(z)\right] F(z_{t+1}|z_t), \quad (9)$$

for $d_t = 0$ and $k_t = k_0$, as entrants must be included in the transition as well. To close the model, we introduce the household and government actors.

2.2 Households

The representative household consumes, supplies labor, and finances the creation of new firms. By choosing labor supply N_t^s and consumption C_t, the household maximizes expected, discounted lifetime utility:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\log C_t - \nu N_t^s\right), \quad (10)$$

subject to the budget constraint:

$$C_t + M_t w_t c_e = w_t N_t^s + \Pi_t + T_t, \quad (11)$$

in which Π_t represents total profits earned by firms and T_t are government transfers. Combining the first-order conditions of labor and consumption, we obtain the standard optimal intra- temporal condition between consumption and labor supply:

$$\nu C_t = w_t. \quad (12)$$

2.3 Government

The government levies taxes on firms and transfers the proceeds T_t lump-sum to the household. T_t is given by:

$$T_t = \tau \left(\underbrace{Y_t - w_t N_t}_{\text{Net revenues}} - \underbrace{\hat{\delta}\left(D_t + I_t\right)}_{\text{Tax deductions}} \right). \quad (13)$$

τ represents the corporate tax rate, Y_t is aggregate production, N_t stands for the total labor used in production, and D_t and I_t indicate the aggregate stock of deduction allowances and investment, respectively.

[8] This implies $k^*(z,k,d;\mu_t) = (1-\delta)k_t + i^*(z,k,d;\mu_t)$ and $d^*(z,k,d;\mu_t) = (1-\hat{\delta})(d_t + i^*(z,k,d;\mu_t))$.

2.4 Aggregation and Market Clearing

Total labor demanded from households by firms is either used in production or to cover entry costs:

$$N_t^d = M_t c_e + N_t = M_t c_e + \int l^*(z,k,d;\mu_t)\mu(dz,dk,dd).$$

In equilibrium, labor demanded equals labor supplied:

$$N_t^d = N_t^s.$$

The household consumes total production net of investment and the resources dissipated through adjustment costs:

$$C_t = Y_t - I_t - AC_t = \int \left[y^*(z,k,d;\mu_t) - i^*(z,k,d;\mu_t) - \mathcal{C}^*(z,k,d;\mu_t)\right] \mu(dz,dk,dd). \tag{14}$$

Finally, total profits Π_t earned in the economy are:

$$\Pi_t = \int \pi^*(z,k,d;\mu_t)\mu(dz,dk,dd),$$

in which $\pi^*(z,k,d;\mu_t)$ refers to the profit a firm earns under the optimal policy functions.

2.5 Equilibrium

In the exercises we present below, we solve for and compare stationary equilibria that differ in the corporate tax rate τ and the deduction rate $\hat{\delta}$. Given a collection of exogenous parameters, the stationary equilibrium is defined as:

(1) A distribution of firms $\bar{\mu}$ such that $\bar{\mu}(z_{t+1}, k_{t+1}, d_{t+1}) = \bar{\mu}(z_t, k_t, d_t)$, for all z, k, d.

(2) An equilibrium wage $\bar{w}$ which solves the free entry condition in equation (7), given $\bar{\mu}$.

(3) A mass of entrants $\bar{M}$ that solves market clearing from equation (14).

(4) A steady-state stochastic discount factor $\bar{\Lambda} = \beta$.

(5) A collection of policy functions $\bar{l^*}(z,k,d)$ and $\bar{i^*}(z,k,d)$ that solve the firm maximization problem, given $\bar{\Lambda}$ and $\bar{w}$,

(6) Aggregates $\bar{T}$ from (13), and $\bar{\Pi}$, $\bar{I}$, $\bar{C}$, $\bar{N}$, $\bar{N^s}$, $\bar{N^d}$ that solve the aggregation conditions in section 2.4.

3 Calibration

Our quantitative analysis is anchored to the 2005–2007 Dutch *Working on Profit Act* for two reasons. The first reason is that we use the reform as a laboratory for our policy analyses because it was introduced explicitly to stimulate growth, yet aimed to be revenue-neutral. The core feature of the reform was a reduction in the statutory corporate tax rate from 34.5% to 25.5%, and a simultaneous broadening of the tax base through a reduction in the tax deductibility of fixed asset investment. Specifically, before the reform, firms were able to immediately and fully expense eligible investments. Afterwards, only 20% of the costs could be deducted annually, which spread out the tax benefits of the investment. We focus solely on the headline rate cut and fixed asset investment deductibility parts of the reform, although the reform featured other tax policy changes.[9] The second reason we focus on the Dutch context is that we can calibrate and compare our model with firm-level microdata from 2000–2015 in the Netherlands.

In the following, we refer to the 'pre-reform' economy as the one featuring the tax rate τ of 34.5% and a full tax deduction $\hat{\delta} = 1$. The 'post-reform' economy refers to the equilibrium characterized by $\tau = 25.5\%$ and $\hat{\delta} = 0.2$.

3.1 Data

To calibrate our model to reproduce key features of the Dutch economy, we use administrative firm-level microdata from Statistics Netherlands, covering the period 2000–2015, which allows us to analyze firms both before and after the Dutch fiscal reform. First, data from the business registry (*ABR*) contains basic background statistics such as firm birth date, sector and employment (full-time-equivalent number). Second, the balance sheet dataset (*NFO*) covers non-financial firms and reports revenues, profits, material inputs, and tangible capital stocks. Finally, we use the *Production Statistics* to construct a measure of fixed costs relative to total costs in the economy.[10] We link all three datasets via consistent firm identifiers. The final dataset covers all non-financial firms that pay corporate tax in the Netherlands.

3.2 Parameters

We set several parameter values externally. We choose values for $\{\alpha, \beta, \gamma\}$ to match the labor share, real interest rate, and firm exit rate in the Dutch economy. The top panel of

[9] Thus, our analysis should not be interpreted as an evaluation of the Dutch policy change as a whole.
[10] Fixed costs cover variables such as research and development, automation, or advertising costs.

Table 1 reports the values. In particular, (i) α is chosen to target an aggregate wage bill to value added ratio of 68%, which we computed using Dutch microdata, (ii) (the inverse) of β targets a real interest rate of 0.88%, which is the average for the Netherlands in World Bank tables over the period 2000–2004, and (iii) γ matches the 10% exit rate from Freeman, Bettendorf, van Heuvelen, and Meijerink (2021).

We assign the economic depreciation rate δ and the persistence of the productivity process ρ^z standard values from the literature, such as the calibration in Winberry (2021).[11] We then use the balance sheet data of firm-level value-added, capital stock, and employment to estimate the mean and the median of the productivity distribution, following Ackerberg, Caves, and Frazer (2015) (henceforth, ACF). The estimation method is applied in numerous studies of firm-level productivity in the Netherlands and other countries, such as Bartelsman, Dobbelaere, and Mattioli (2024) and De Loecker, Eeckhout, and Unger (2020). Given ρ^z, we back out the two remaining AR(1) parameters, the innovation of the shock σ^z and the unconditional mean $\log \bar{z}$, by calibrating them such that the implied stationary distribution matches the observed median and mean.

The three remaining parameters $\{\psi, c_e, \nu\}$ are calibrated internally by matching two targets and imposing one normalization. Recall, ψ represents the coefficient on the capital quadratic adjustment cost $\mathcal{C}$. We calibrate ψ to match the pre-reform investment to GDP ratio of 0.16. Next, to calibrate entry costs c_e, we exploit the fact that the data allows us to differentiate between total production and non-production (fixed) costs. While the model equivalent is production and non-production labor costs, we acknowledge the mapping from the data to the model is imperfect, but still target the 7% ratio of fixed to total production costs to calibrate c_e. Note that c_e and ψ are jointly calibrated to match the two targets above. Lastly, by combining the household's first-order condition in equation (12) and aggregate market clearing in equation (14), the labor disutility parameter ν is set such that the mass of entrants in the pre-reform steady-state is normalized to one. The bottom panel of Table 1 reports the calibrated values of $\{\psi, c_e, \nu\}$.

[11] To calibrate the productivity process, we need to pin down the parameters of the long-run log-normal distribution implied by our AR(1) process, i.e. log-normal$\left(\log \bar{z}, \frac{\sigma^2}{1-\rho^2}\right)$. We decide to fix ρ^z to a standard annual value of 0.9 (estimating autocorrelation of TFP in the Dutch data, we find ρ^z between 0.89 and 0.91) and calibrate the two remaining parameters with two moments.

Table 1: **Calibration (internal and external).**

Parameter	Description	Value	Target	Model	Data
External					
α	K output elasticity	0.32	Labor share 68%		
γ	Exit rate	0.10	Annual exit rate 10%		
β	Discount factor	0.99	Real interest rate $\approx$ 1%		
δ	Depreciation	0.03	Winberry (2021)		
ρ^z	AR(1) persistence	0.90	Winberry (2021)		
σ^z	AR(1) volatility	0.36	ACF estimation		
$\log \bar{z}$	AR(1) unc. mean	-0.24	ACF estimation		
Parameter	Description	Value	Target	Model	Data
Internal					
ψ	Adjustment cost	0.05	Investment / GDP	0.15	0.16
c_e	Entry cost	8.93	Fixed / total costs	0.07	0.07
ν	Labor disutility	0.01	Mass of entrants	$M = 1$	Normalization

Notes: This table exhibits externally and internally set parameters and indicates the moments we target in our estimation. The 'External' panel shows parameters we set externally, following values from the literature or our own productivity estimation. The 'Internal' panel displays the parameter values that result from our own calibration using the targets from the data. All data moments are measured using data from 2000 to 2004, before the Dutch tax reform.

3.3 Untargeted Moments

To validate the fit of our model, we check how well our calibrated model reproduces several untargeted moments from the data. Panel 1 of Table 2 shows that the model replicates the empirical corporate tax to GDP ratio and the effective tax rate that firms pay.

The model also matches the cross-sectional distribution of capital across firms. We reproduce the high Gini coefficient and share of aggregate capital held by the 0–50th, 50–95th, and 95th+ percentiles of the firm capital distribution, as shown in panels 2 and 3 of Table 2. As expected, our framework results in an employment distribution that is more unequal than the empirical counterpart. This discrepancy is likely due to the lack of labor market frictions in our model. Lastly, our model replicates the empirical pattern that the investment rate (investment over capital) decreases with firm size. In the model, this pattern arises from the mean-reversion of the productivity process: smaller firms expect to receive more favorable productivity shocks and invest to grow, while larger firms invest mainly to maintain their size, not to grow.

Table 2: **Untargeted moments of the pre-reform economy.**

Panel 1	Corp. tax to GDP ratio	Effective tax rate	
Data	0.05	0.18	
Model	0.04	0.15	
Panel 2	Capital	Employment	
Gini			
Data	0.95	0.84	
Model	0.92	0.91	
Panel 3	Share of agg. capital	Share of agg. empl.	Average inv. rate
Top (95 pct.)			
Data	88.47	69.25	0.11
Model	81.43	91.78	0.08
Middle (50–95 pct.)			
Data	10.82	26.41	0.20
Model	17.52	4.19	0.35
Bottom (0–50 pct.)			
Data	0.71	4.35	0.26
Model	0.01	4.02	0.98

Notes: *Panel 1* exhibits the data and model moments of the tax-to-GDP ratio (taxes paid / sum of value-added) and the effective tax rate (taxes paid / taxable income before investment deductions). *Panel 2* exhibits the data and model moments for the cross-sectional Gini measure of inequality based on capital, first column, and employment, second column. *Panel 3* exhibits several measures of the firm distribution for capital, labor, and the investment rate (investment over capital). The 'Top (95 pct.)' sub-panel exhibits the share of aggregate capital and employment held by and the average investment rate (i/k) for the largest 5% of firms. The same applies to the 'Middle (50-95 pct.)' and 'Bottom (0-50 pct.)' sub-panels. The firm size bins are based on capital. All data moments are measured using data from 2000 to 2004, before the Dutch tax reform.

4 Quantitative Results

This section presents our main results. Section 4.1 characterizes the aggregate and distributional effects of that *Working on Profit Act* Dutch reform through the lenses of our quantitative framework. In section 4.2, we trace out a frontier of revenue-neutral reform around the post-reform Dutch economy to document the trade-off between aggregate productivity and investment. In section 4.3 we discuss the key mechanism behind our quantitative results.

4.1 The Aggregate and Distributional Effect of the Dutch Reform

As a first exercise, we use our model to characterize the aggregate and distributional effects of the *Working on Profit Act*. The exercise already reveals the trade-off between the size of the economy, measured by output and the capital stock, and allocative efficiency, measured by aggregate productivity, as the rate and deductions change.

Column 1 of Table 3 shows that the post-reform economy features 6.9% less aggregate capital and 1.5% less production than the pre-reform economy. Columns 2 and 3 isolate the *direct* effect of each tax instrument by changing only τ or only $\hat{\delta}$. The decline in capital and production under the full reform is driven by the lower $\hat{\delta}$, as expected. The deduction cut alone reduces capital by 9.2% and production by 1.4%, while the rate cut alone raises them by 9.2% and 3.0%, respectively.

Table 3: **Baseline exercise; percentage change relative to the pre-reform economy.**

	Post-reform	Tax-only	Deduction-only
Aggregate capital	−6.9	9.2	−9.2
Aggregate production	−1.5	3.0	−1.4
Aggregate productivity	1.3	−0.0	1.1

Notes: This table shows the percentage change in aggregate capital, production, and productivity relative to the pre-reform economy ($\tau = 0.345$, $\hat{\delta} = 1$), for each of the other long-run equilibria, indicated by the column name. Post-reform: $\tau = 0.255$, $\hat{\delta} = 0.2$; Tax-only: $\tau = 0.255$, $\hat{\delta} = 1$; Deduction-only: $\tau = 0.345$, $\hat{\delta} = 0.2$. Aggregate productivity is measured as the Solow residual.

The final row of Table 3 shows that aggregate productivity, as measured by the Solow residual, increases by 1.3% under the full reform.[12] The deduction-only counterfactual generates a 1.1% increase, while the tax-only counterfactual leaves productivity essentially unchanged. Thus, the majority of the post-reform increase in aggregate productivity is driven by the direct effect of the deduction change.

Since we assume the stationary productivity distribution $G(z)$ is fixed across equilibria, any changes in aggregate productivity can only reflect resources moving toward more or less productive firms. To see the reallocation more clearly, panels (a) and (b) of Figure 1 plot the percentage-point change in the share of aggregate capital and investment held by the smallest 50% of firms, medium-large firms (50th –95th percentiles), and the remaining largest firms, all relative to the pre-reform economy shares. Panel (c)

[12] The Solow residual Z is given by $Z \equiv \frac{Y}{K^{\alpha}L^{1-\alpha}}$, in which $Y = \int y^{*}(z,k,d;\mu)\,\mu(dz,dk,dd)$, $K = \int k^{*}(z,k,d;\mu)\,\mu(dz,dk,dd)$, and $L = \int l^{*}(z,k,d;\mu)\,\mu(dz,dk,dd)$.

repeats the exercise for productivity bins. In all three tax regimes, capital and investment reallocate toward the largest and most productive 5% of firms. Furthermore, the reallocation effect is strongest under the deduction cut and for investment.

Figure 1: **Reallocation effects of Dutch reform.**

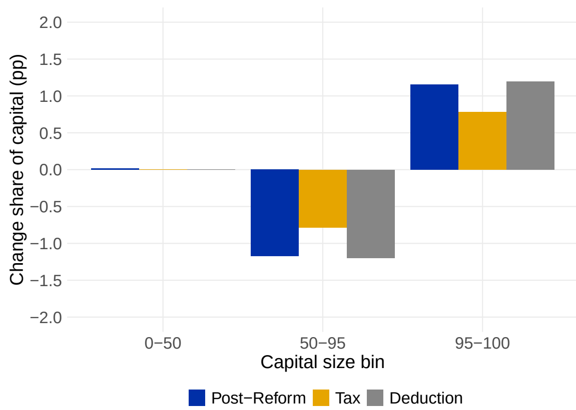

(a) Capital reallocation by size bin.

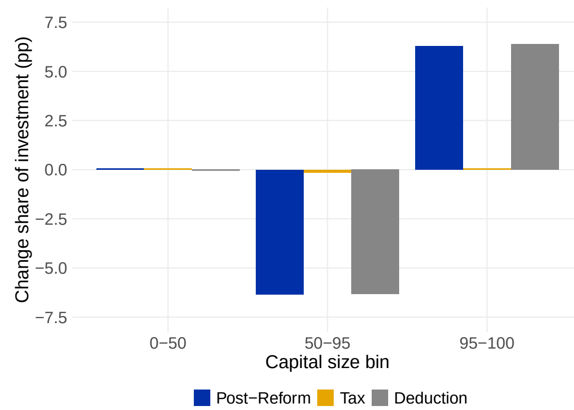

(b) Investment reallocation by size bin.

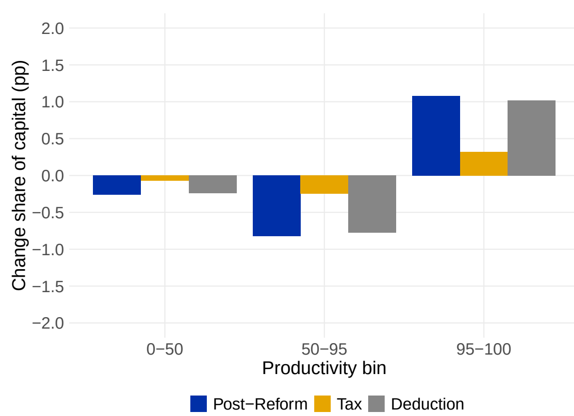

(c) Capital reallocation by productivity bin.

Notes: Panel (a) shows the percentage point change in the share of aggregate capital held by each firm size bin (bottom 50%, 50th–95th percentiles, and top 5%, based on capital stock), relative to the pre-reform economy. Panel (b) shows the corresponding change for the share of aggregate investment. Panel (c) shows capital reallocation across three analogous productivity bins.

4.2 The Frontier of Revenue-Neutral Reforms

Although the investment deduction part of the *Working on Profit Act* was intended to offset the revenue loss from the rate cut, our model suggests that the reform did not achieve budget neutrality. To find the revenue-neutral reform that coincides with the post-reform economy, we fix the post-reform tax rate, denote it $\tau^L = 0.255$, and search through values for $\hat{\delta}$ that keep corporate tax revenue at its pre-reform level. Formally, from the government budget constraint, we back out a new value $\hat{\delta}^L$ such that:

$$T_t^L = \tau^L \left(Y_t - w_t N_t - \hat{\delta}^L \left(D_t + I_t \right) \right), \tag{15}$$

in which T_t^L is the pre-reform corporate tax revenue. When we implement this experiment, we find that a value of $\hat{\delta}^L = 0.074$ satisfies Equation (15) with $\tau^L = 0.255$, which is smaller than the actual $\hat{\delta} = 0.20$ which was implemented. The result suggests that, given the lower tax rate, the implemented deduction cut was not deep enough to balance pre-reform revenue.

Following the logic in Equation (15), we next collect the sequence of revenue-neutral $\{\tau, \hat{\delta}\}$ combinations for $\tau \in [0.24, 0.32]$.[13] Panel (a) of Figure 2 shows the frontier of $\{\tau, \hat{\delta}\}$ combinations that generate the pre-reform level of corporate tax revenue. The triangle marks the actual Dutch reform, clearly not on the frontier. The shape of the frontier shows that a low tax rate requires a wide tax base to generate a balanced reform, as expected. Furthermore, the trade-off between τ and $\hat{\delta}$ is not linear, becoming very steep for higher values of τ.

Panels (b) and (c) of Figure 2 show how aggregate capital and output vary along the frontier. There is a clear positive relationship between the tax rate τ and the size of the economy. Moving from a tax rate of $\tau = 0.24$ to $\tau = 0.32$ increases total capital by 30% and total production by 12.5%. Along the balanced reform frontier, a high-rate, high-deduction is more conducive to economic growth. Furthermore, panels (b) and (c) Figure 3 show that aggregate consumption and household utility also increase along the frontier, 7.5% and 0.3%, respectively.

Nevertheless, panel (d) in Figure 2 shows that aggregate productivity, measured by the Solow residual, declines by roughly 1% as the tax rate increases along the frontier. The result echoes our analysis of the *Working on Profit Act*, in which the low-rate, low-deduction tax regime increases aggregate productivity relative to the pre-reform economy. Moreover, the negative slope highlights the trade-off a policymaker faces between allocative efficiency and the size of the economy. We discuss this trade-off further in the next section, in which we explain the key mechanism behind our results.

[13] Values above or below cannot be matched by $\hat{\delta} \in [0, 1]$ to generate a balanced reform.

Figure 2: **The balanced reform frontier: capital, production, and productivity.**

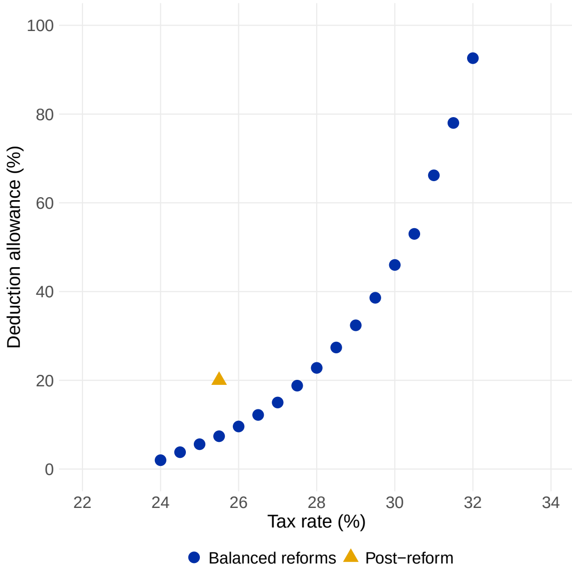

(a) Tax parameter combinations.

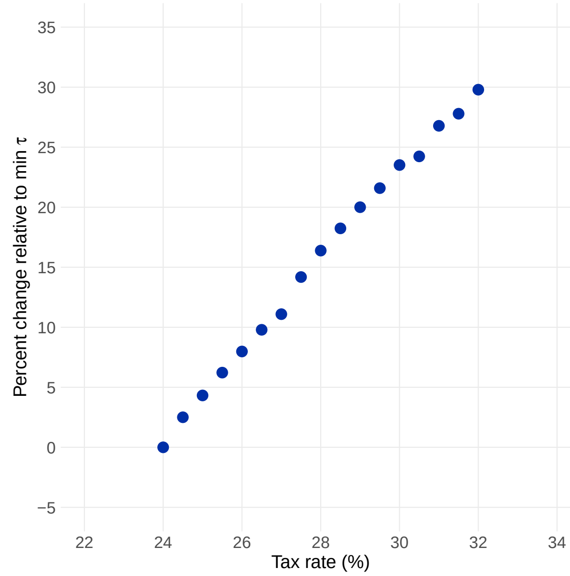

(b) Aggregate capital stock.

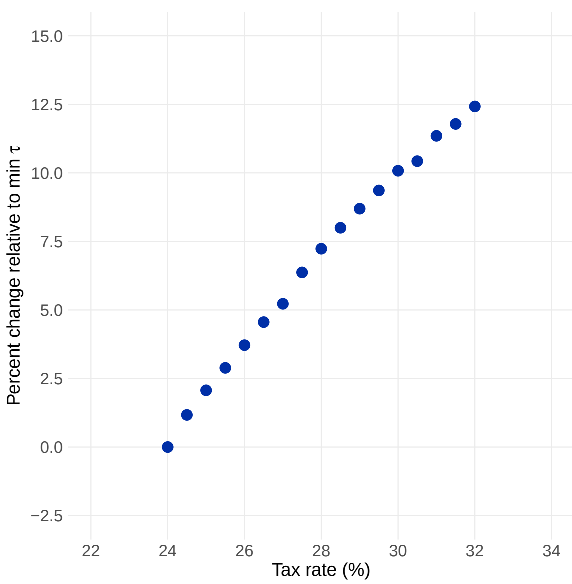

(c) Total production.

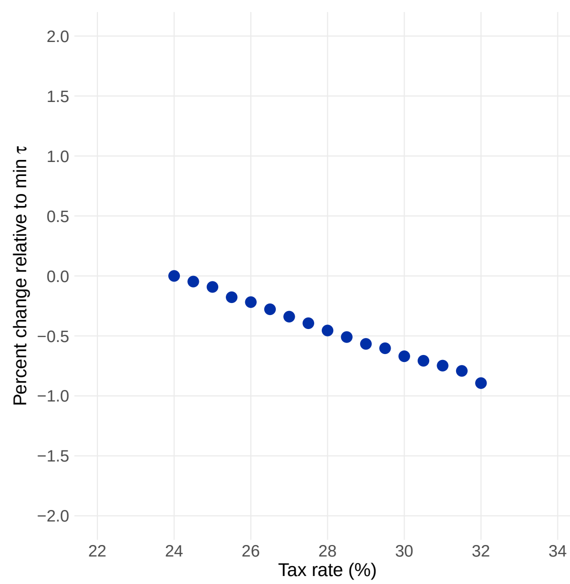

(d) Aggregate productivity.

Notes: Panel (a) shows the rate-base frontier in $\{\tau, \hat{\delta}\}$-space. The triangle marks the implemented Dutch reform. The other panels each show the percentage change in a variable relative to the balanced reform with $\tau = 0.24$, the lowest rate on the frontier we measure.

Figure 3: **The balanced reform frontier: consumption and utility.**

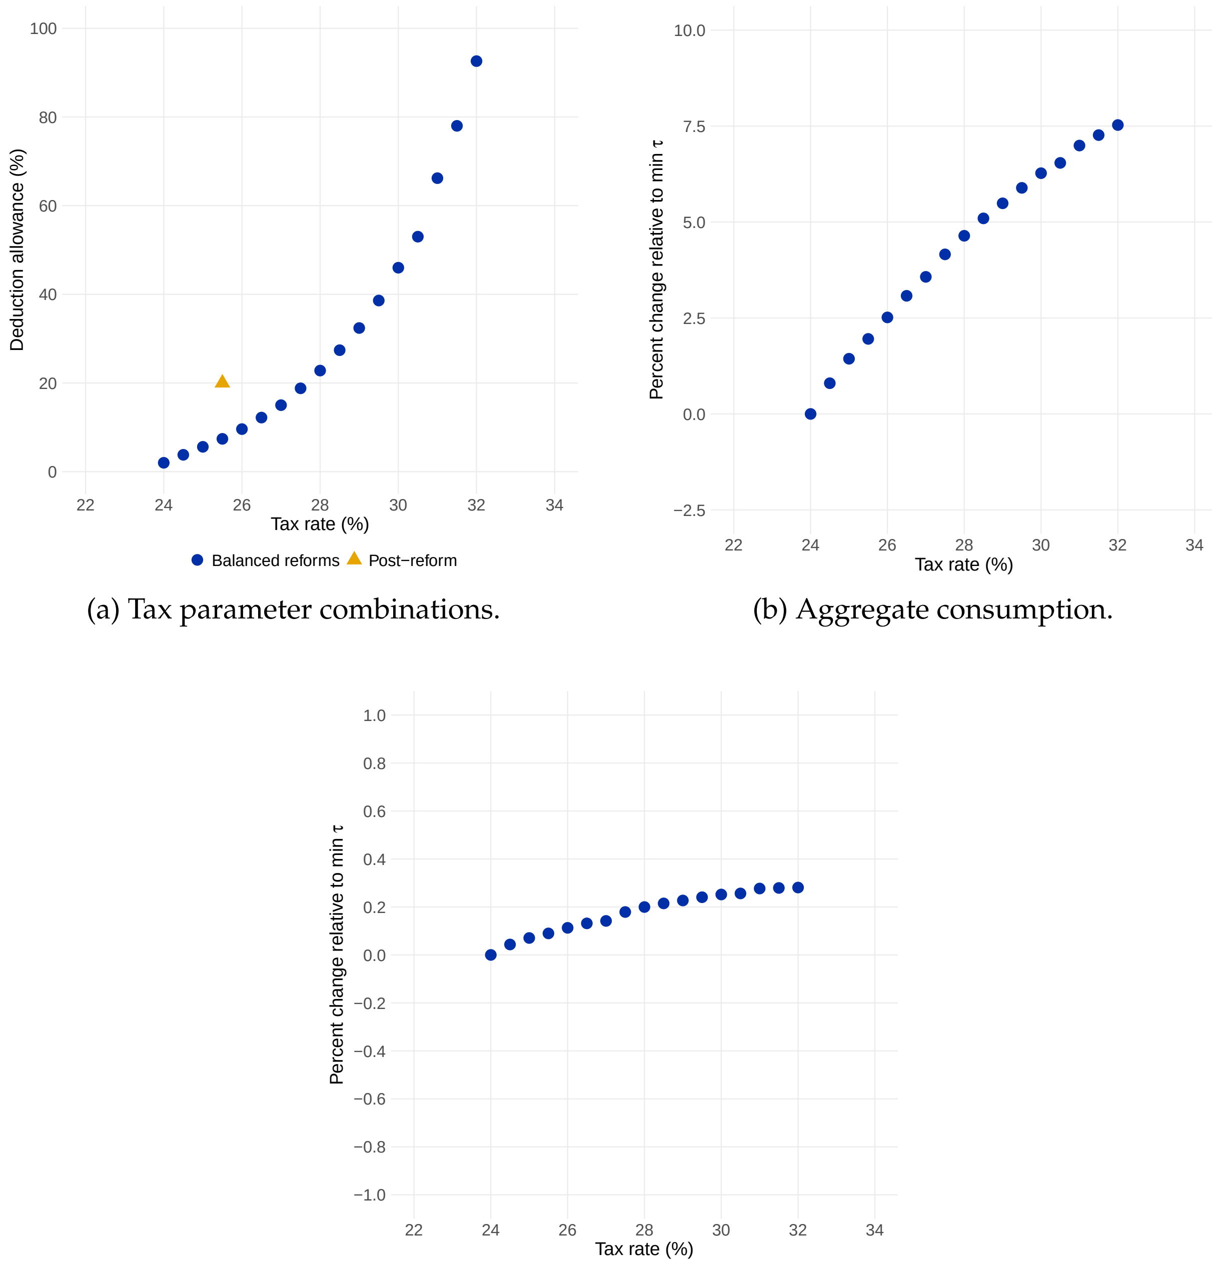

(a) Tax parameter combinations.

(b) Aggregate consumption.

(c) Total utility.

Notes: Panel (a) repeats the rate-base frontier. The other panels each show the percentage change in a variable relative to the balanced reform with $\tau = 0.24$, the lowest rate on the frontier.

4.3 Key Mechanism

Why is a high-rate, high-deduction tax regime more conducive to investment, but not to productivity, than a low-tax, low-deduction regime? The key mechanism works through heterogeneous investment intensity across the firm-size distribution. Recall that in Section 3, the mean investment-to-capital ratio is 0.98 for the bottom 50% of firms (0.26 in the data), 0.35 for the middle-to-large 50th-95th percentiles (0.20 in the data), and 0.08 for the top 5% of firms (0.11 in the data). The heterogeneity implies that, even in a revenue-neutral reform, the effects of a rate and base change differ across the firm distribution.

First, consider a lower tax rate. A rate cut increases after-tax profits and, thus, investment for all firms. However, the windfall gain is directly proportional to profit, so the largest and most profitable firms stand to benefit disproportionately.[14] As a result, the rate cut targets firms with relatively lower investment intensity, who will not respond strongly. The reallocation patterns of the *Working on Profit Act* in Figure 1 illustrate the same point: the reallocation toward large and productive firms is considerably weaker under the direct effect of the lower tax rate than the lower deduction allowance. The same logic holds along the revenue-neutral tax reform frontier.

Next, consider a lower tax deductibility of investment. The lower deduction delays the tax benefits of investment for all firms by effectively raising investment costs. Small to medium-large firms invest more intensely than top firms, and so they are disproportionately burdened. The uneven burden generates stronger reallocation of capital toward the top firms, which raises aggregate productivity at the cost of a smaller overall economy, since the deduction cut falls hardest on the investment-intensive firms. Taken together, the two forces explain why moving along the frontier toward a high-rate, high-deduction regime increases investment and output while lowering measured productivity.

5 Empirical Support for the Key Mechanism

The key mechanism behind our quantitative results is that investment intensity declines with firm size. Therefore, a reduction in investment deductibility burdens investment-intensive, smaller firms the most. In this section, we test the mechanism in the data using the balance sheet data and the 2005–2007 Dutch reform.

Our model predicts that firms with a higher pre-reform investment-intensity were more exposed to the Dutch reform, which shifted the corporate tax regime to a low-tax, low-deduction one. For each firm j we observe in the data from 2002 to 2004, we

[14] $\frac{d(1-\tau)\pi}{d\tau} = -\pi$.

measure exposure E_j as the average pre-reform investment-to-capital ratio:

$$E_j = \frac{1}{4} \sum_{t=2002}^{2004} \left(\frac{i_{jt}}{k_{jt}} \right).$$

Both investment and capital are measured in thousands of (2015) euros. The model predicts that after the Dutch government reduced the tax rate and the tax deduction, firms with higher values of E_j should accumulate less capital: these firms are disproportionately burdened by the deduction cut, even though the tax rate decreased at the same time.

To test the prediction, we estimate a continuous difference-in-differences specification by regressing the capital k_{jt} on the exposure measure interacted with year indicators:

$$k_{jt} = \alpha_j + \lambda_t + \sum_{t \neq 2004} \beta_t \times E_j \times I_t + \varepsilon_{jt}. \tag{16}$$

The terms α_j and λ_t denote firm and year fixed effects, and I_t represents an indicator equal to one in year t and zero otherwise. We omit the indicator for $t = 2004$, the pre-reform year. The coefficients β_t trace out the differential path of capital for fully exposed firms ($E_j = 1$), relative to non-exposed firms ($E_j = 0$). We cluster standard errors at the firm level.

The coefficients β_t identify the causal effect of exposure under two assumptions. First, absent the reform, the capital stocks of high- and low-exposure firms would have evolved in parallel. Second, firms did not adjust their investment in anticipation of the reform. Although neither assumption is directly testable, the pre-reform estimates $\hat{\beta}_t, \;\; t \in \{2002, 2003\}$ provide a standard diagnostic check. If high- and low-exposure firms were on different trends before the reform, or if firms adjusted in anticipation, the estimates would differ from zero.

Figure 4 plots the estimates $\hat{\beta}_t$ for $t \in \{2002, 2010\}$, together with 95% confidence intervals. We can see that the pre-reform coefficients are statistically indistinguishable from zero, supporting the identifying assumptions. As the reform is phased in, the estimates turn negative and significant. In 2010, a one-unit increase in exposure is associated with a capital stock around 10.2 thousand euros lower than before the reform. The pooled difference-in-differences coefficient estimate equals -5.19 (SE: 2.63). Overall, the sign of the difference-in-difference coefficients is consistent with our model's prediction: the deduction cut disproportionately depresses capital accumulation for investment-intensive firms, despite a lower tax rate.

Figure 4: **Event study estimates from 2005–2007 Dutch tax reform.**

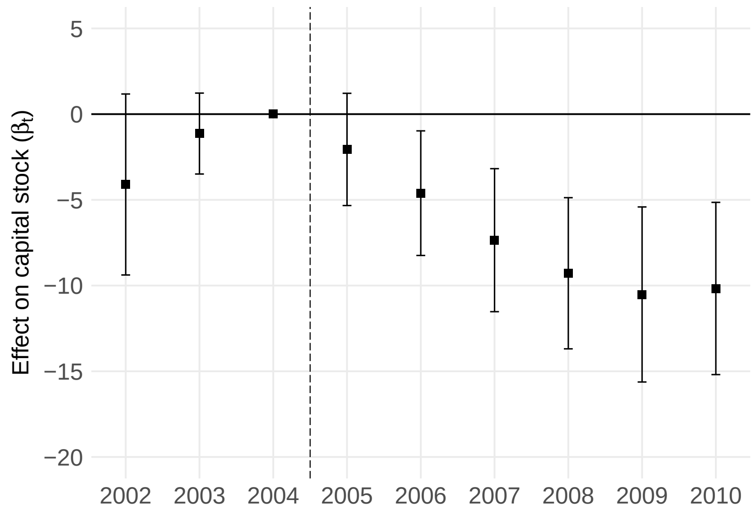

Notes: This figure plots the estimated coefficients $\hat{\beta}_t$ from Equation (16), together with 95% confidence intervals. Standard errors are clustered at the firm level. The dependent variable is the firm-level capital stock k_{jt}. Exposure E_j is the average investment-to-capital ratio of firm j over 2002–2004. Each estimate $\hat{\beta}_t$ measures the differential capital stock in year t associated with a one-unit higher pre-reform investment rate. All estimates are relative to 2004, which is the final pre-reform year. The specification includes firm and year fixed effects α_j and λ_t, respectively. The sample covers non-financial firms observed in the years from 2002 to 2004, required to construct the exposure measure. Number of observations: 356k. Number of unique firms: 47k.

The difference-in-differences estimates capture the differential capital response of more exposed firms. The Global Financial Crisis hit the Dutch economy shortly after the reform was fully phased in. We therefore interpret our results as suggestive evidence, which is nonetheless consistent with the investment-intensity mechanism in our model.

6 Conclusion

Along the frontier of revenue-neutral tax reforms, which combination of tax rate and tax deductions is most conducive to economic growth and aggregate productivity? We answer this question in a heterogeneous firm model, calibrated to reproduce key features of the Dutch economy. We trace out the sequence of revenue-neutral reforms in the neighborhood of the post-reform Dutch tax regime.

Our model shows that a high-rate, high-deduction regime stimulates investment and output more effectively than a low-rate, low-deduction one, holding corporate tax revenue fixed. Along the frontier, moving from a 24% tax rate (paired with a 2% annual deduction) to a 32% rate (paired with a 93% deduction) increases aggregate output by 12%, the capital stock by 30%, and raises welfare marginally, by 0.3%. Aggregate productivity, however, declines by 1% due to the reallocation of investment away from the largest and most productive firms. An event study exploiting firm-level variation in exposure to the 2005–2007 Dutch reform supports our model's mechanism.

Our findings speak to the broader debate on the design of revenue-neutral corporate tax policy (OECD, 2010). The familiar prescription to broaden the base and lower the rate delivers a smaller but better-allocated capital stock than the reverse design. Which side of the trade-off a government takes depends on the weight it places on the size of the economy relative to allocative efficiency.

References

Ackerberg, D. A., Caves, K., & Frazer, G. (2015). Identification Properties of Recent Production Function Estimators. *Econometrica*, *83*(6), 2411–2451.

Baley, I., & Blanco, A. (2022). The Long-Run Effects of Corporate Tax Reforms. *CEPR Discussion Paper No. DP16936*.

Bartelsman, E., Dobbelaere, S., & Mattioli, A. Z. (2024). Non-compete Agreements, Tacit Knowledge and Market Imperfections. *IZA Discussion Paper Series*, *17260*.

Becker, B., Jacob, M., & Jacob, M. (2013). Payout Taxes and the Allocation of Investment. *Journal of Financial Economics*, *107*(1), 1–24.

Bilicka, K., Güçeri, I., & Koumanakos, E. (2025). Dividend Taxation and Firm Performance with Heterogeneous Payout Responses. *American Economic Journal: Economic Policy*, *17*(2), 1–29.

Campbell, J. L., Chyz, J. A., Dhaliwal, D. S., & Schwartz Jr, W. C. (2013). Did the 2003 Tax Act Increase Capital Investments by Corporations? *The Journal of the American Taxation Association*, *35*(2), 33–63.

Chodorow-Reich, G. (2025). *The Neoclassical Theory of Firm Investment and Taxes: A Reassessment* (Tech. Rep.). National Bureau of Economic Research.

Chodorow-Reich, G., Smith, M., Zidar, O. M., & Zwick, E. (2024). *Tax Policy and Investment in a Global Economy* (Tech. Rep.). National Bureau of Economic Research.

Chodorow-Reich, G., Zidar, O., & Zwick, E. (2024). Lessons from the Biggest Business Tax Cut in US History. *Journal of Economic Perspectives*, *38*(3), 61–88.

Clementi, G. L., & Palazzo, B. (2016). Entry, Exit, Firm Dynamics, and Aggregate Fluctuations. *American Economic Journal: Macroeconomics*, *8*(3), 1–41.

Colciago, A., Lewis, V., & Matyska, B. (2024). Corporate Taxes, Productivity, and Business Dynamism. *CEPR Discussion Paper*.

Cooper, R., & Haltiwanger, J. (2006). On the Nature of Capital Adjustment Costs. *The Review of Economic Studies*, *73*(3), 611–633.

Cooper, R., Haltiwanger, J., & Power, L. (1999). Machine Replacement and the Business Cycle: Lumps and Bumps. *American Economic Review*, *89*(4), 921–946.

De Loecker, J., Eeckhout, J., & Unger, G. (2020). The Rise of Market Power and the Macroeconomic Implications. *The Quarterly Journal of Economics*, *135*(2), 561–644.

Erosa, A., & González, B. (2019). Taxation and the Life Cycle of Firms. *Journal of Monetary Economics*, *105*, 114–130.

Freeman, D., Bettendorf, L., van Heuvelen, G. H., & Meijerink, G. (2021). *The contribution of business dynamics to productivity growth in the netherlands* (Tech. Rep.). CPB Netherlands Bureau for Economic Policy Analysis.

Furno, F. (2021). The Macroeconomic Effects of Corporate Tax Reforms. *Unpublished Working Paper*.

Hopenhayn, H., & Rogerson, R. (1993). Job Turnover and Policy Evaluation: A General Equilibrium Analysis. *Journal of political Economy*, *101*(5), 915–938.

Kindsgrab, P. M. (2022). Empirical Evidence on the Effect of Local Business Taxes on Business Entry. *Unpublished Working Paper*.

Lichter, A., Löffler, M., Isphording, I. E., Nguyen, T.-V., Poege, F., & Siegloch, S. (2025). Profit Taxation, R&D Spending, and Innovation. *American Economic Journal: Economic Policy*, *17*(1), 432–463.

Link, S., Menkhoff, M., Peichl, A., & Schüle, P. (2023). Downward Revision of Investment Decisions after Corporate Tax Hikes. *IZA Discussion Paper*.

OECD. (2010). *Choosing a Broad Base - Low Rate Approach to Taxation* (Tech. Rep. No. 19). Paris: OECD Publishing.

Ohrn, E. (2018). The Effect of Corporate Taxation on Investment and Financial Policy: Evidence from the DPAD. *American Economic Journal: Economic Policy*, *10*(2), 272–301.

Sedláček, P., & Sterk, V. (2019). Reviving American Entrepreneurship? Tax Reform and Business Dynamism. *Journal of Monetary Economics*, *105*, 94–108.

Winberry, T. (2021). Lumpy Investment, Business Cycles, and Stimulus Policy. *American Economic Review*, *111*(1), 364–396.

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